COF-C02 Dumps

SnowPro Core Certification Exam (COF-C02)

https://www.certleader.com/COF-C02-dumps.html

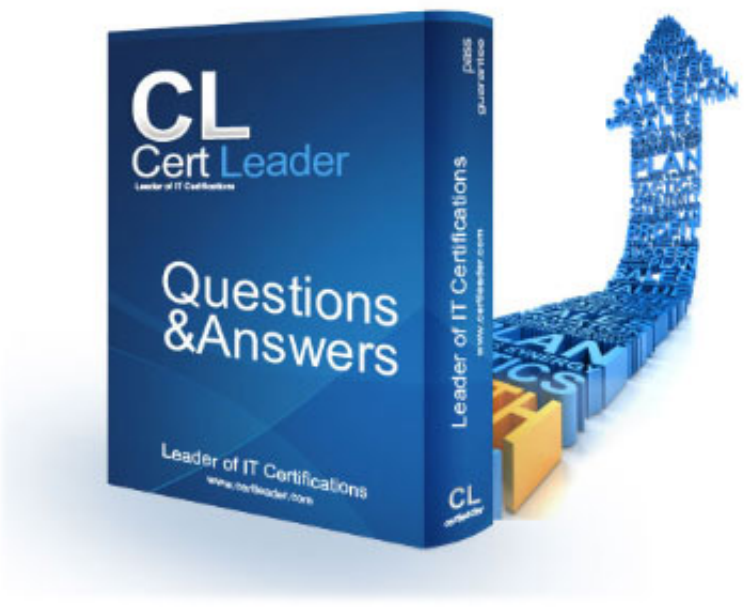

NEW QUESTION 1
- (Topic 1)
A user has 10 files in a stage containing new customer data. The ingest operation completes with no errors, using the following command:
COPY INTO my table FROM @my stage;
The next day the user adds 10 files to the stage so that now the stage contains a mixture of new customer data and updates to the previous data. The user did not remove the 10 original files.
If the user runs the same copy into command what will happen?

A. All data from all of the files on the stage will be appended to the table
B. Only data about new customers from the new files will be appended to the table
C. The operation will fail with the error uncertain files in stage.
D. All data from only the newly-added files will be appended to the table.

Answer: A

Explanation:
When the COPY INTO command is executed in Snowflake, it processes all files present in the specified stage that have not been ingested before or marked as already loaded. Since the user did not remove the original 10 files after the first load, running the same COPY INTO command again will result in all 20 files being processed. This means that the data from the original 10 files will be appended to the table again, along with the data from the new 10 files, potentially leading to duplicate records for the original data set.
References:
? Snowflake Documentation on Data Loading
? SnowPro® Core Certification Study Guide

NEW QUESTION 2
- (Topic 1)
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 3
- (Topic 1)
The Information Schema and Account Usage Share provide storage information for which of the following objects? (Choose three.)

A. Users
B. Tables
C. Databases
D. Internal Stages

Answer: BCD

Explanation:
The Information Schema and Account Usage Share in Snowflake provide metadata and historical usage data for various objects within a Snowflake account. Specifically, they offer storage information for Tables, Databases, and Internal Stages. These schemas contain views and table functions that allow users to query object metadata and usage metrics, such as the amount of data stored and historical activity.
? Tables: The storage information includes data on the daily average amount of data in database tables.
? Databases: For databases, the storage usage is calculated based on all the data contained within the database, including tables and stages.
? Internal Stages: Internal stages are locations within Snowflake for temporarily storing data, and their storage usage is also tracked.
References: The information is verified according to the SnowPro Core Certification Study Guide and Snowflake documentation

NEW QUESTION 4
- (Topic 1)
Which of the following objects can be shared through secure data sharing?

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

Answer: D

Explanation:
Secure data sharing in Snowflake allows users to share various objects between Snowflake accounts without physically copying the data, thus not consuming additional storage. Among the options provided, external tables can be shared through secure data sharing. External tables are used to query data directly from files in a stage without loading the data into Snowflake tables, making them suitable for sharing across different Snowflake accounts.
References:
? Snowflake Documentation on Secure Data Sharing
? SnowPro™ Core Certification Companion: Hands-on Preparation and Practice

NEW QUESTION 5
- (Topic 1)
What features does Snowflake Time Travel enable?

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

Answer: BC

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

NEW QUESTION 6
- (Topic 1)
Which stage type can be altered and dropped?

A. Database stage
B. External stage
C. Table stage
D. User stage

Answer: B

Explanation:
External stages can be altered and dropped in Snowflake. An external stage points to an external location, such as an S3 bucket, where data files are stored. Users can modify the stage??s definition or drop it entirely if it??s no longer needed. This is in contrast to table stages, which are tied to specific tables and cannot be altered or dropped independently.
References:
? [COF-C02] SnowPro Core Certification Exam Study Guide
? Snowflake Documentation on Stages1

NEW QUESTION 7
- (Topic 1)
When is the result set cache no longer available? (Select TWO)

A. When another warehouse is used to execute the query
B. When another user executes the query
C. When the underlying data has changed
D. When the warehouse used to execute the query is suspended
E. When it has been 24 hours since the last query

Answer: CE

Explanation:
The result set cache in Snowflake is invalidated and no longer available when the underlying data of the query results has changed, ensuring that queries return the most current data. Additionally, the cache expires after 24 hours to maintain the efficiency and accuracy of data retrieval1.

NEW QUESTION 8
- (Topic 1)
A company strongly encourages all Snowflake users to self-enroll in Snowflake's default Multi-Factor Authentication (MFA) service to provide increased login security for users connecting to Snowflake.
Which application will the Snowflake users need to install on their devices in order to connect with MFA?

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

Answer: B

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

NEW QUESTION 9
- (Topic 1)
What is a machine learning and data science partner within the Snowflake Partner Ecosystem?

A. Informatica
B. Power Bl
C. Adobe
D. Data Robot

Answer: D

Explanation:
Data Robot is recognized as a machine learning and data science partner within the Snowflake Partner Ecosystem. It provides an enterprise AI platform that enables users to build and deploy accurate predictive models quickly. As a partner, Data Robot integrates with Snowflake to enhance data science capabilities2.
References:
? [COF-C02] SnowPro Core Certification Exam Study Guide
? Snowflake Documentation on Machine Learning & Data Science Partners
https://docs.snowflake.com/en/user-guide/ecosystem-analytics.html

NEW QUESTION 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 14
- (Topic 1)
What is the recommended file sizing for data loading using Snowpipe?

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

Answer: C

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

NEW QUESTION 19

- (Topic 1)
Which cache type is used to cache data output from SQL queries?

A. Metadata cache
B. Result cache
C. Remote cache
D. Local file cache

Answer: B

Explanation:
TheResult cache is used in Snowflake to cache the data output from SQL queries. This feature is designed to improve performance by storing the results of queries for a period of time. When the same or similar query is executed again, Snowflake can retrieve the result from this cache instead of re-computing the result, which saves time and computational resources.
References:
? Snowflake Documentation on Query Results Cache
? SnowPro® Core Certification Study Guide

NEW QUESTION 22
- (Topic 1)
Which Snowflake objects track DML changes made to tables, like inserts, updates, and deletes?

A. Pipes
B. Streams
C. Tasks
D. Procedures

Answer: B

Explanation:
In Snowflake, Streams are the objects that track Data Manipulation Language (DML) changes made to tables, such as inserts, updates, and deletes. Streams record these changes along with metadata about each change, enabling actions to be taken using the changed data. This process is known as change data capture (CDC)2.

NEW QUESTION 23
- (Topic 1)
What are ways to create and manage data shares in Snowflake? (Select TWO)

A. Through the Snowflake web interface (Ul)
B. Through the DATA_SHARE=TRUE parameter
C. Through SQL commands
D. Through the enable share=true parameter
E. Using the CREATE SHARE AS SELECT * TABLE command

Answer: AC

Explanation:
Data shares in Snowflake can be created and managed through the Snowflake web interface, which provides a user-friendly graphical interface for various operations. Additionally, SQL commands can be used to perform these tasks programmatically, offering flexibility and automation capabilities123.

NEW QUESTION 26
- (Topic 1)
User-level network policies can be created by which of the following roles? (Select TWO).

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

Answer: BD

Explanation:
User-level network policies in Snowflake can be created by roles with the necessary privileges to manage security and account settings. The ACCOUNTADMIN role has the highest level of privileges across the account, including the ability to manage network policies. The SECURITYADMIN role is specifically responsible for managing security objects within Snowflake, which includes the creation and management of network policies.
References:
? [COF-C02] SnowPro Core Certification Exam Study Guide
? Snowflake Documentation on Network Policies1
? Section 1.3 - SnowPro Core Certification Study Guide

NEW QUESTION 31
- (Topic 1)
Which of the following Snowflake objects can be shared using a secure share? (Select TWO).

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

E. Secure User Defined Functions (UDFs)

Answer: DE

Explanation:
Secure sharing in Snowflake allows users to share specific objects with other Snowflake accounts without physically copying the data, thus not consuming additional storage. Tables and Secure User Defined Functions (UDFs) are among the objects that can be shared using this feature. Materialized views, sequences, and procedures are not shareable objects in Snowflake.
References:
? [COF-C02] SnowPro Core Certification Exam Study Guide
? Snowflake Documentation on Secure Data Sharing1

NEW QUESTION 34
- (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 37
- (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 38
- (Topic 1)
What are the default Time Travel and Fail-safe retention periods for transient tables?

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

Answer: C

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

NEW QUESTION 43
- (Topic 1)
Which Snowflake technique can be used to improve the performance of a query?

A. Clustering
B. Indexing
C. Fragmenting
D. Using INDEX HINTS

Answer: A

Explanation:
Clustering is a technique used in Snowflake to improve the performance of queries. It involves organizing the data in a table into micro-partitions based on the values of one or more columns. This organization allows Snowflake to efficiently prune non- relevant micro-partitions during a query, which reduces the amount of data scanned and improves query performance.
References:
? [COF-C02] SnowPro Core Certification Exam Study Guide
? Snowflake Documentation on Clustering

NEW QUESTION 44
- (Topic 1)
Which data types does Snowflake support when querying semi-structured data? (Select TWO)

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

Answer: AB

Explanation:
Snowflake supports querying semi-structured data using specific data types that are capable of handling the flexibility and structure of such data. The data types supported for this purpose are:
? A. VARIANT: This is a universal data type that can store values of any other type,
including structured and semi-structured types. It is particularly useful for handling JSON, Avro, ORC, Parquet, and XML data formats1.
? B. ARRAY: An array is a list of elements that can be of any data type, including
VARIANT, and is used to handle semi-structured data that is naturally represented as a list1.
These data types are part of Snowflake??s built-in support for semi-structured data, allowing for the storage, querying, and analysis of data that does not fit into the traditional row- column format.
References:
? Snowflake Documentation on Semi-Structured Data
? [COF-C02] SnowPro Core Certification Exam Study Guide

NEW QUESTION 45
- (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 46
- (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 51
- (Topic 1)
Which of the following is a valid source for an external stage when the Snowflake account is located on Microsoft Azure?

A. An FTP server with TLS encryption
B. An HTTPS server with WebDAV
C. A Google Cloud storage bucket
D. A Windows server file share on Azure

Answer: D

Explanation:
In Snowflake, when the account is located on Microsoft Azure, a valid source for an external stage can be an Azure container or a folder path within an Azure container. This includes Azure Blob storage which is accessible via the azure:// endpoint. A Windows server file share on Azure, if configured properly, can be a valid source for staging data files for Snowflake. Options A, B, and C are not supported as direct sources for an external stage in Snowflake on Azure12.
References: [COF-C02] SnowPro Core Certification Exam Study Guide

NEW QUESTION 55
- (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 57
- (Topic 1)
How would you determine the size of the virtual warehouse used for a task?

A. Root task may be executed concurrently (i.
B. multiple instances), it is recommended to leave some margins in the execution window to avoid missing instances of execution
C. Querying(select)the size of the stream content would help determine the warehouse siz
D. For example, if querying large stream content, use a larger warehouse size
E. If using the stored procedure to execute multiple SQL statements, it's best to test run the stored procedure separately to size the compute resource first
F. Since task infrastructure is based on running the task body on schedule, it's recommended to configure the virtual warehouse for automatic concurrency handling using Multi-cluster warehouse (MCW) to match the task schedule

Answer: D

Explanation:
The size of the virtual warehouse for a task can be configured to handle concurrency automatically using a Multi-cluster warehouse (MCW). This is because tasks are designed to run their body on a schedule, and MCW allows for scaling compute resources to match the task??s execution needs without manual intervention.
References: [COF-C02] SnowPro Core Certification Exam Study Guide

NEW QUESTION 58
- (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 59
- (Topic 1)
When reviewing the load for a warehouse using the load monitoring chart, the chart indicates that a high volume of Queries are always queuing in the warehouse
According to recommended best practice, what should be done to reduce the Queue volume? (Select TWO).

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

Answer: AB

Explanation:
To address a high volume of queries queuing in a warehouse, Snowflake recommends two best practices:
? A. Use multi-clustered warehousing to scale out warehouse capacity: This approach allows for the distribution of queries across multiple clusters within a

warehouse, effectively managing the load and reducing the queue volume.
? B. Scale up the warehouse size to allow Queries to execute faster: Increasing the size of the warehouse provides more compute resources, which can reduce the time it takes for queries to execute and thus decrease the number of queries waiting in the queue.
These strategies help to optimize the performance of the warehouse by ensuring that resources are scaled appropriately to meet demand.
References:
? Snowflake Documentation on Multi-Cluster Warehousing
? SnowPro Core Certification best practices

NEW QUESTION 62
- (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 66
- (Topic 1)
Which command is used to unload data from a Snowflake table into a file in a stage?

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

Answer: A

Explanation:
The COPY INTO command is used in Snowflake to unload data from a table into a file in a stage. This command allows for the export of data from Snowflake tables into flat files, which can then be used for further analysis, processing, or storage in external systems.
References:
? Snowflake Documentation on Unloading Data
? Snowflake SnowPro Core: Copy Into Command to Unload Rows to Files in Named Stage

NEW QUESTION 69
- (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 70
- (Topic 1)
Which of the following conditions must be met in order to return results from the results cache? (Select TWO).

A. The user has the appropriate privileges on the objects associated with the query
B. Micro-partitions have been reclustered since the query was last run
C. The new query is run using the same virtual warehouse as the previous query
D. The query includes a User Defined Function (UDF)

E. The query has been run within 24 hours of the previously-run query

Answer: AE

Explanation:
To return results from the results cache in Snowflake, certain conditions must be met:
? Privileges: The user must have the appropriate privileges on the objects associated with the query. This ensures that only authorized users can access cached data.
? Time Frame: The query must have been run within 24 hours of the previously-run query. Snowflake??s results cache is designed to store the results of queries for a short period, typically 24 hours, to improve performance for repeated queries.

NEW QUESTION 72
- (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 75
- (Topic 1)
Which of the following are benefits of micro-partitioning? (Select TWO)

A. Micro-partitions cannot overlap in their range of values
B. Micro-partitions are immutable objects that support the use of Time Travel.
C. Micro-partitions can reduce the amount of I/O from object storage to virtual warehouses
D. Rows are automatically stored in sorted order within micro-partitions
E. Micro-partitions can be defined on a schema-by-schema basis

Answer: BC

Explanation:
Micro-partitions in Snowflake are immutable objects, which means once they are written, they cannot be modified. This immutability supports the use of Time Travel, allowing users to access historical data within a defined period. Additionally, micro-partitions can significantly reduce the amount of I/O from object storage to virtual warehouses. This is because Snowflake??s query optimizer can skip over micro-partitions that do not contain relevant data for a query, thus reducing the amount of data that needs to be scanned and transferred.
References: [COF-C02] SnowPro Core Certification Exam Study Guide https://docs.snowflake.com/en/user-guide/tables-clustering-micropartitions.html

NEW QUESTION 77
- (Topic 1)
True or False: Loading data into Snowflake requires that source data files be no larger than 16MB.

A. True
B. False

Answer: B

Explanation:
Snowflake does not require source data files to be no larger than 16MB. In fact, Snowflake recommends that for optimal load performance, data files should be roughly 100-250 MB in size when compressed. However, it is not recommended to load very large files (e.g., 100 GB or larger) due to potential delays and wasted credits if errors occur. Smaller files should be aggregated to minimize processing overhead, and larger files should be split to distribute the load among compute resources in an active warehouse. References: Preparing your data files | Snowflake Documentation

NEW QUESTION 79
- (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 81
- (Topic 1)
Which of the following compute resources or features are managed by Snowflake? (Select TWO).

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

Answer: CE

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

NEW QUESTION 84
- (Topic 1)
In the query profiler view for a query, which components represent areas that can be used to help optimize query performance? (Select TWO)

A. Bytes scanned
B. Bytes sent over the network
C. Number of partitions scanned
D. Percentage scanned from cache
E. External bytes scanned

Answer: AC

Explanation:
In the query profiler view, the components that represent areas that can be used to help optimize query performance include ??Bytes scanned?? and ??Number of partitions scanned??. ??Bytes scanned?? indicates the total amount of data the query had to read and is a direct indicator of the query??s efficiency. Reducing the bytes scanned can lead to lower data transfer costs and faster query execution. ??Number of partitions scanned?? reflects how well the data is clustered; fewer partitions scanned typically means better performance because the system can skip irrelevant data more effectively.
References:
? [COF-C02] SnowPro Core Certification Exam Study Guide
? Snowflake Documentation on Query Profiling1

NEW QUESTION 89
- (Topic 1)
What is the default File Format used in the COPY command if one is not specified?

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

Answer: A

Explanation:
The default file format for the COPY command in Snowflake, when not specified, is CSV (Comma-Separated Values). This format is widely used for data exchange because it is simple, easy to read, and supported by many data analysis tools.

NEW QUESTION 91
- (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 92
- (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 97
- (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 102
- (Topic 1)
A virtual warehouse's auto-suspend and auto-resume settings apply to which of the following?

A. The primary cluster in the virtual warehouse
B. The entire virtual warehouse
C. The database in which the virtual warehouse resides
D. The Queries currently being run on the virtual warehouse

Answer: B

Explanation:
The auto-suspend and auto-resume settings in Snowflake apply to the entire virtual warehouse. These settings allow the warehouse to automatically suspend when it??s not in use, helping to save on compute costs. When queries or tasks are submitted to the warehouse, it can automatically resume operation. This functionality is designed to optimize resource usage and cost-efficiency.
References:
? SnowPro Core Certification Exam Study Guide (as of 2021)
? Snowflake documentation on virtual warehouses and their settings (as of 2021)

NEW QUESTION 106
- (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 108
- (Topic 1)
Which of the following commands cannot be used within a reader account?

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

Answer: A

Explanation:

In Snowflake, a reader account is a type of account that is intended for consuming shared data rather than performing any data management or DDL operations. The CREATE SHARE command is used to share data from your account with another account, which is not a capability provided to reader accounts. Reader accounts are typically restricted from creating shares, as their primary purpose is to read shared data rather than to share it themselves.
References:
? Snowflake Documentation on Reader Accounts
? SnowPro® Core Certification Study Guide

NEW QUESTION 109
- (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 114
- (Topic 1)
Which is the MINIMUM required Snowflake edition that a user must have if they want to use AWS/Azure Privatelink or Google Cloud Private Service Connect?

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

Answer: D

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

NEW QUESTION 116
- (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 117
- (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 121
- (Topic 2)
What occurs when a pipe is recreated using the CREATE OR REPLACE PIPE command?

A. The Pipe load history is reset to empty.
B. The REFRESH command is executed.
C. The stage will be purged.
D. The destination table is truncated.

Answer: A

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

NEW QUESTION 126
- (Topic 2)
The following JSON is stored in a VARIANT column called src of the CAR_SALES table:

```
{
  "customer": [
    {
       "address": San Francisco, CA",
       "name": "Jane Doe"
    }
  ],
  "date":  "2022-01-28",
  "dealership": "Town Auto Sales",
  "salesperson": {
     "id": "55"
  }
}
```

A user needs to extract the dealership information from the JSON. How can this be accomplished?

A. select src:dealership from car_sales;
B. select src.dealership from car_sales;
C. select src:Dealership from car_sales;
D. select dealership from car_sales;

Answer: B

Explanation:
In Snowflake, to extract a specific element from a JSON stored in a VARIANT column, the correct syntax is to use the dot notation. Therefore, the query select src.dealership from car_sales; will return the dealership information contained within each JSON object in the src column.
References: For a detailed explanation, please refer to the Snowflake documentation on querying semi-structured data.

NEW QUESTION 131
- (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 132
- (Topic 2)
Which of the following features are available with the Snowflake Enterprise edition? (Choose two.)

A. Database replication and failover
B. Automated index management
C. Customer managed keys (Tri-secret secure)
D. Extended time travel
E. Native support for geospatial data

Answer: AD

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

NEW QUESTION 136
- (Topic 2)
A user is preparing to load data from an external stage
Which practice will provide the MOST efficient loading performance?

A. Organize files into logical paths
B. Store the files on the external stage to ensure caching is maintained
C. Use pattern matching for regular expression execution
D. Load the data in one large file

Answer: A

Explanation:
Organizing files into logical paths can significantly improve the efficiency of data loading from an external stage. This practice helps in managing and locating files easily, which can be particularly beneficial when dealing with large datasets or complex directory structures1.

NEW QUESTION 137
- (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 139
- (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 144
- (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 147
- (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 149
- (Topic 2)
Which command should be used to load data from a file, located in an external stage, into a table in Snowflake?

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

Answer: D

Explanation:
The COPY command is used in Snowflake to load data from files located in an external stage into a table. This command allows for efficient and parallelized data loading from various file formats1.
References = [COF-C02] SnowPro Core Certification Exam Study Guide, Snowflake Documentation

NEW QUESTION 151
- (Topic 2)
What features that are part of the Continuous Data Protection (CDP) feature set in Snowflake do not require additional configuration? (Choose two.)

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

Answer: CD

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

NEW QUESTION 156
- (Topic 2)
A Snowflake Administrator needs to ensure that sensitive corporate data in Snowflake tables is not visible to end users, but is partially visible to functional managers.
How can this requirement be met?

A. Use data encryption.
B. Use dynamic data masking.
C. Use secure materialized views.
D. Revoke all roles for functional managers and end users.

Answer: B

Explanation:
Dynamic data masking is a feature in Snowflake that allows administrators to define masking policies to protect sensitive data. It enables partial visibility of the data to certain roles, such as functional managers, while hiding it from others, like end users

NEW QUESTION 157
- (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 161
- (Topic 2)
What affects whether the query results cache can be used?

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

Answer: C

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

NEW QUESTION 164
- (Topic 2)
Assume there is a table consisting of five micro-partitions with values ranging from A to Z. Which diagram indicates a well-clustered table?

A.

B.

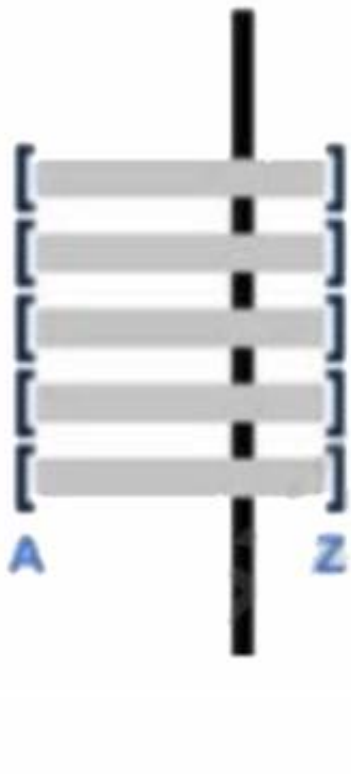

C.

D.

Answer: C

Explanation:
A well-clustered table in Snowflake means that the data is organized in such a way that related data points are stored close to each other within the micro-partitions. This optimizes query performance by reducing the amount of scanned data. The diagram indicated by option C shows a well-clustered table, as it likely represents a more evenly distributed range of values across the micro-partitions1.
References = Snowflake Micro-partitions & Table Clustering

NEW QUESTION 166
- (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 168
- (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 169
- (Topic 2)
Which Snowflake layer is always leveraged when accessing a query from the result cache?

A. Metadata
B. Data Storage
C. Compute
D. Cloud Services

Answer: D

Explanation:
The Cloud Services layer in Snowflake is responsible for managing the result cache. When a query is executed, the results are stored in this cache, and subsequent identical queries can leverage these cached results without re-executing the entire query1.

NEW QUESTION 174
- (Topic 2)

Which of the following statements describe features of Snowflake data caching? (Choose two.)

A. When a virtual warehouse is suspended, the data cache is saved on the remote storage layer.
B. When the data cache is full, the least-recently used data will be cleared to make room.
C. A user can only access their own queries from the query result cache.
D. A user must set USE_METADATA_CACHE to TRUE to use the metadata cache in queries.
E. The RESULT_SCAN table function can access and filter the contents of the query result cache.

Answer: BE

Explanation:
Snowflake??s data caching features include the ability to clear the least- recently used data when the data cache is full to make room for new data. Additionally, the RESULT_SCAN table function can access and filter the contents of the query result cache, allowing users to retrieve and work with the results of previous queries. The other statements are incorrect: the data cache is not saved on the remote storage layer when a virtual warehouse is suspended, users can access queries from the query result cache that were run by other users, and there is no setting calledUSE_METADATA_CACHE in Snowflake. References: Caching in the Snowflake Cloud Data Platform, Optimizing the warehouse cache

NEW QUESTION 178
- (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 181
- (Topic 2)
What COPY INTO SQL command should be used to unload data into multiple files?

A. SINGLE=TRUE
B. MULTIPLE=TRUE
C. MULTIPLE=FALSE
D. SINGLE=FALSE

Answer: D

Explanation:
The COPY INTO SQL command with the option SINGLE=FALSE is used to unload data into multiple files. This option allows the data to be split into multiple files during the unload process. References: SnowPro Core Certification COPY INTO SQL command unload multiple files

NEW QUESTION 183
- (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 184
- (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 187
- (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 190
- (Topic 2)
What actions will prevent leveraging of the ResultSet cache? (Choose two.)

A. Removing a column from the query SELECT list
B. Stopping the virtual warehouse that the query is running against
C. Clustering of the data used by the query
D. Executing the RESULTS_SCAN() table function
E. Changing a column that is not in the cached query

Answer: BD

Explanation:
The ResultSet cache is leveraged to quickly return results for repeated queries. Actions that prevent leveraging this cache include stopping the virtual warehouse that the query is running against (B) and executing the RESULTS_SCAN() table function (D). Stopping the warehouse clears the local disk cache, including the ResultSet cache1. The RESULTS_SCAN() function is used to retrieve the result of a previously executed query, which bypasses the need for the ResultSet cache.

NEW QUESTION 194
- (Topic 2)
Which of the following are characteristics of Snowflake virtual warehouses? (Choose two.)

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

Answer: CE

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

NEW QUESTION 195
- (Topic 2)
Which of the following is a data tokenization integration partner?

A. Protegrity
B. Tableau
C. DBeaver
D. SAP

Answer: A

Explanation:
Protegrity is listed as a data tokenization integration partner for Snowflake. This partnership allows Snowflake users to utilize Protegrity??s tokenization solutions within the Snowflake environment3.
References = [COF-C02] SnowPro Core Certification Exam Study Guide, Snowflake Documentation

NEW QUESTION 198
- (Topic 2)

How are serverless features billed?

A. Per second multiplied by an automatic sizing for the job
B. Per minute multiplied by an automatic sizing for the job, with a minimum of one minute
C. Per second multiplied by the size, as determined by the SERVERLESS_FEATURES_SIZE account parameter
D. Serverless features are not billed, unless the total cost for the month exceeds 10% of the warehouse credits, on the account

Answer: B

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

NEW QUESTION 199
- (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 201
- (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 203
- (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 208
- (Topic 2)
What are best practice recommendations for using the ACCOUNTADMIN system-defined role in Snowflake? (Choose two.)

A. Ensure all ACCOUNTADMIN roles use Multi-factor Authentication (MFA).
B. All users granted ACCOUNTADMIN role must be owned by the ACCOUNTADMIN role.
C. The ACCOUNTADMIN role must be granted to only one user.
D. Assign the ACCOUNTADMIN role to at least two users, but as few as possible.
E. All users granted ACCOUNTADMIN role must also be granted SECURITYADMIN role.

Answer: AD

Explanation:
Best practices for using the ACCOUNTADMIN role include ensuring that all users with this role use Multi-factor Authentication (MFA) for added security. Additionally, it is recommended to assign the ACCOUNTADMIN role to at least two users to avoid delays in case of password recovery issues, but to as few users as possible to maintain strict control over account-level operations4.

NEW QUESTION 210
- (Topic 2)
Which columns are part of the result set of the Snowflake LATERAL FLATTEN command? (Choose two.)

A. CONTENT
B. PATH
C. BYTE_SIZE
D. INDEX
E. DATATYPE

Answer: BD

Explanation:
The LATERAL FLATTEN command in Snowflake produces a result set that includes several columns, among which PATH and INDEX are includedPATH indicates the path to the element within a data structure that needs to be flattened, and INDEX represents the index of the element if it is an array2.

NEW QUESTION 211
- (Topic 2)
What is the default file size when unloading data from Snowflake using the COPY command?

A. 5 MB
B. 8 GB
C. 16 MB
D. 32 MB

Answer: C

Explanation:
The default file size when unloading data from Snowflake using the COPY command is not explicitly stated in the provided resources. However, Snowflake documentation suggests that the file size can be specified using the MAX_FILE_SIZE option in the COPY INTO <location> command2.

NEW QUESTION 216
- (Topic 2)
What happens to historical data when the retention period for an object ends?

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

Answer: C

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

NEW QUESTION 218
- (Topic 2)
Network policies can be set at which Snowflake levels? (Choose two.)

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

Answer: CE

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

NEW QUESTION 219
- (Topic 2)
Which of the following describes the Snowflake Cloud Services layer?

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

Answer: A

Explanation:
The Snowflake Cloud Services layer is a collection of services that coordinate activities across Snowflake, tying together all the different components to process user requests, from login to query dispatch1.
References = [COF-C02] SnowPro Core Certification Exam Study Guide, Snowflake Documentation1

NEW QUESTION 224
- (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 226
- (Topic 2)
How many days is load history for Snowpipe retained?

A. 1 day
B. 7 days
C. 14 days
D. 64 days

Answer: C

Explanation:
Snowpipe retains load history for 14 days. This allows users to view and audit the data that has been loaded into Snowflake using Snowpipe within this time frame3.

NEW QUESTION 229
- (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 232
- (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 235
- (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 237
- (Topic 2)
What impacts the credit consumption of maintaining a materialized view? (Choose two.)

A. Whether or not it is also a secure view
B. How often the underlying base table is queried
C. How often the base table changes
D. Whether the materialized view has a cluster key defined
E. How often the materialized view is queried

Answer: CD

Explanation:
The credit consumption for maintaining a materialized view is impacted by how often the base table changes © and whether the materialized view has a cluster key defined (D). Changes to the base table can trigger a refresh of the materialized view, consuming credits. Additionally, having a cluster key defined can optimize the performance and credit usage during the materialized view??s maintenance. References: SnowPro Core Certification materialized view credit consumption

NEW QUESTION 240
- (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 242
- (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 244
- (Topic 2)
Which of the following accurately describes shares?

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

Answer: A

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

NEW QUESTION 245
- (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 248
- (Topic 2)
Why does Snowflake recommend file sizes of 100-250 MB compressed when loading data?

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

Answer: D

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

NEW QUESTION 253
- (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 257
- (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 259
- (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 260
- (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 263
- (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 266
- (Topic 2)
Which SQL commands, when committed, will consume a stream and advance the stream offset? (Choose two.)

A. UPDATE TABLE FROM STREAM
B. SELECT FROM STREAM
C. INSERT INTO TABLE SELECT FROM STREAM
D. ALTER TABLE AS SELECT FROM STREAM
E. BEGIN COMMIT

Answer: AC

Explanation:
The SQL commands that consume a stream and advance the stream offset are those that result in changes to the data, such as UPDATE and INSERT operations. Specifically, ??UPDATE TABLE FROM STREAM?? and ??INSERT INTO TABLE SELECT
FROM STREAM?? will consume the stream and move the offset forward, reflecting the changes made to the data.
References: [COF-C02] SnowPro Core Certification Exam Study Guide

NEW QUESTION 271
- (Topic 3)
Using variables in Snowflake is denoted by using which SQL character?

A. @
B. &
C. $
D. #

Answer: C

Explanation:
VeryComprehensiveExplanation=InSnowflake,variablesaredenotedbyadollarsign().Variable s can be used in SQL statements where a literal constant is allowed, and they must be prefixed with a $ sign to distinguish them from bind values and column names.

NEW QUESTION 272
- (Topic 3)
What is a responsibility of Snowflake??s virtual warehouses?

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

Answer: C

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

NEW QUESTION 273
- (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 276
- (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 281
- (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 284
- (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 are considerations when using a directory table when working with unstructured data? (Choose two.)

A. A directory table is a separate database object.
B. Directory tables store data file metadata.
C. A directory table will be automatically added to a stage.
D. Directory tables do not have their own grantable privileges.
E. Directory table data can not be refreshed manually.

Answer: BD

Explanation:
Directory tables in Snowflake are used to store metadata about data files in a stage. They are not separate database objects but are conceptually similar to external tables. Directory tables do not have grantable privileges of their own

NEW QUESTION 288

- (Topic 3)
A user needs to create a materialized view in the schema MYDB.MYSCHEMA. Which statements will provide this access?

A. GRANT ROLE MYROLE TO USER USER1;GRANT CREATE MATERIALIZED VIEW ON SCHEMA MYDB.MYSCHEMA TO ROLE MYROLE;
B. GRANT ROLE MYROLE TO USER USER1;GRANT CREATE MATERIALIZED VIEW ON SCHEMA MYDB.MYSCHEMA TO USER USER1;
C. GRANT ROLE MYROLE TO USER USER1;GRANT CREATE MATERIALIZED VIEW ON SCHEMA MYD
D. K"-'SCHEMA TO USER! ;
E. GRANT ROLE MYROLE TO USER USER1;GRANT CREATE MATERIALIZED VIEW ON SCHEMA MYDB.MYSCHEMA TO MYROLE;

Answer: A

Explanation:
To provide a user with the necessary access to create a materialized view in a schema, the user must be granted a role that has the CREATE MATERIALIZED VIEW privilege on that schema. First, the role is granted to the user, and then the privilege is granted to the role

NEW QUESTION 292
- (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 297
- (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 299
- (Topic 3)
How can a Snowflake user optimize query performance in Snowflake? (Select TWO).

A. Create a view.
B. Cluster a table.
C. Enable the search optimization service.
D. Enable Time Travel.
E. Index a table.

Answer: BC

Explanation:
To optimize query performance in Snowflake, users can cluster a table, which organizes the data in a way that minimizes the amount of data scanned during queries. Additionally, enabling the searchoptimization service can improve the performance of selective point lookup queries on large tables34.

NEW QUESTION 303
- (Topic 3)
Which command is used to unload files from an internal or external stage to a local file system?

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

Answer: B

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

NEW QUESTION 305
- (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 310
- (Topic 3)
Which formats does Snowflake store unstructured data in? (Choose two.)

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

Answer: AC

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

NEW QUESTION 315
- (Topic 3)
Which data type can store more than one type of data structure?

A. JSON
B. BINARY
C. VARCHAR
D. VARIANT

Answer: D

Explanation:
The VARIANT data type in Snowflake can store multiple types of data structures, as it is designed to hold semi-structured data. It can contain any other data type, including OBJECT and ARRAY, which allows it to represent various data structures

NEW QUESTION 317
- (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 321
- (Topic 3)
Which Snowflake object can be accessed in he FROM clause of a query, returning a set of rows having one or more columns?

A. A User-Defined Table Function (UDTF)
B. A Scalar User Function (UDF)
C. A stored procedure
D. A task

Answer: A

Explanation:
In Snowflake, a User-Defined Table Function (UDTF) can be accessed in the FROM clause of a query. UDTFs return a set of rows with one or more columns, which can be queried like a regular table

NEW QUESTION 322
- (Topic 3)
Which of the following are characteristics of security in Snowflake?

A. Account and user authentication is only available with the Snowflake Business Critical edition.
B. Support for HIPAA and GDPR compliance is available for UI Snowflake editions.
C. Periodic rekeying of encrypted data is available with the Snowflake Enterprise edition and higher
D. Private communication to internal stages is allowed in the Snowflake Enterprise edition and higher.

Answer: C

Explanation:
One of the security features of Snowflake includes the periodic rekeying of encrypted data, which is available with the Snowflake Enterprise edition and higher2. This ensures that the encryption keys are rotated regularly to maintain a high level of security. References: [COF-C02] SnowPro Core Certification Exam Study Guide

NEW QUESTION 326
- (Topic 3)
Which operations are handled in the Cloud Services layer of Snowflake? (Select TWO).

A. Security
B. Data storage
C. Data visualization
D. Query computation
E. Metadata management

Answer: AE

Explanation:
The Cloud Services layer in Snowflake is responsible for various services, including security (like authentication and authorization) and metadata management (like query parsing and optimization). References: Based on general cloud architecture knowledge as of 2021.

NEW QUESTION 330
- (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 335
- (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 337
- (Topic 3)
Which clients does Snowflake support Multi-Factor Authentication (MFA) token caching for? (Select TWO).

A. GO driver
B. Node.js driver
C. ODBC driver
D. Python connector
E. Spark connector

Answer:

CD

Explanation:
Multi-Factor Authentication (MFA) token caching is typically supported for clients that maintain a persistent connection or session with Snowflake, such as the ODBC driver and Python connector, to reduce the need for repeated MFA challenges. References: Based on general security practices in cloud services as of 2021.

NEW QUESTION 339
- (Topic 3)
Which privilege must be granted to a share to allow secure views the ability to reference data in multiple databases?

A. CREATE_SHARE on the account
B. SHARE on databases and schemas
C. SELECT on tables used by the secure view
D. REFERENCE_USAGE on databases

Answer: D

Explanation:
To allow secure views the ability to reference data in multiple databases, the REFERENCE_USAGE privilege must be granted on each database that contains objects referenced by the secure view2. This privilege is necessary before granting the SELECT privilege on a secure view to a share.

NEW QUESTION 341
- (Topic 3)
Which stages are used with the Snowflake PUT command to upload files from a local file system? (Choose three.)

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

Answer: BDF

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

NEW QUESTION 344
- (Topic 3)
What computer language can be selected when creating User-Defined Functions (UDFs) using the Snowpark API?

A. Swift
B. JavaScript
C. Python
D. SQL

Answer: C

Explanation:
The Snowpark API allows developers to create User-Defined Functions (UDFs) in various languages, including Python, which is known for its ease of use and wide adoption in data-related tasks. References: Based on general programming and cloud data service knowledge as of 2021.

NEW QUESTION 346
- (Topic 3)
Which Snowflake objects can be shared with other Snowflake accounts? (Choose three.)

A. Schemas
B. Roles
C. Secure Views
D. Stored Procedures
E. Tables
F. Secure User-Defined Functions (UDFs)

Answer: ACF

Explanation:
In Snowflake, you can share several types of objects with other Snowflake accounts. These include schemas, secure views, and secure user-defined functions (UDFs). Sharing these objects allows for collaboration and data access across different Snowflake accounts while maintaining security and governance controls4.

NEW QUESTION 348
- (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 351
- (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 356
- (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 359
- (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 361
- (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 366
- (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 371
- (Topic 3)
What is the minimum Snowflake edition needed for database failover and fail-back between Snowflake accounts for business continuity and disaster recovery?

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

Answer: C

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

NEW QUESTION 372
- (Topic 3)
How many network policies can be assigned to an account or specific user at a time?

A. One
B. Two
C. Three
D. Unlimited

Answer: A

Explanation:
According to my knowledge, a security administrator can create multiple network policies, but only one network policy can be active for an account or specific user at any given time. This ensures that there is a clear and consistent policy being applied without conflicts. References: Based on my internal knowledge as of 2021.

NEW QUESTION 373
- (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 375
- (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 378
- (Topic 3)

What does Snowflake recommend regarding database object ownership? (Select TWO).

A. Create objects with ACCOUNTADMIN and do not reassign ownership.
B. Create objects with SYSADMIN.
C. Create objects with SECURITYADMIN to ease granting of privileges later.
D. Create objects with a custom role and grant this role to SYSADMIN.
E. Use only MANAGED ACCESS SCHEMAS for66 objects owned by ACCOUNTADMIN.

Answer: BD

Explanation:
Snowflake recommends creating objects with a role that has the necessary privileges and is not overly permissive. SYSADMIN is typically used for managing system- level objects and operations. Creating objects with a custom role and granting this role to SYSADMIN allows for more granular control and adherence to the principle of least privilege. References: Based on best practices for database object ownership and role management.

NEW QUESTION 379
- (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 381
- (Topic 3)
User INQUISITIVE_PERSON has been granted the role DATA_SCIENCE. The role DATA_SCIENCE has privileges OWNERSHIP on the schema MARKETING of the database ANALYTICS_DW.
Which command will show all privileges granted to that schema?

A. SHOW GRANTS ON ROLE DATA_SCIENCE
B. SHOW GRANTS ON SCHEMA ANALYTICS_DW.MARKETING
C. SHOW GRANTS TO USER INQUISITIVE_PERSON
D. SHOW GRANTS OF ROLE DATA_SCIENCE

Answer: B

Explanation:
To show all privileges granted to a specific schema, the command SHOW GRANTS ON SCHEMA <schema_name> should be used3. In this case, it would be SHOW GRANTS ON SCHEMA ANALYTICS_DW.MARKETING. References: [COF-C02] SnowPro Core Certification Exam Study Guide

NEW QUESTION 386
- (Topic 3)
Which of the following statements describes a schema in Snowflake?

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

Answer: A

Explanation:
A schema in Snowflake is a logical grouping of database objects, such as tables and views, that belongs to a single database. Each schema is part of a namespace in Snowflake, which is inferred from the current database and schema in use for the session5

NEW QUESTION 391
- (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 395

- (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 397
- (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 401
- (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 406
- (Topic 3)
What can a Snowflake user do in the Admin area of Snowsight?

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

Answer: A

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

NEW QUESTION 409
- (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 412

- (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 416
- (Topic 3)
A data provider wants to share data with a consumer who does not have a Snowflake account. The provider creates a reader account for the consumer following these steps:
* 1. Created a user called "CONSUMER"
* 2. Created a database to hold the share and an extra-small warehouse to query the data
* 3. Granted the role PUBLIC the following privileges: Usage on the warehouse, database, and schema, and SELECT on all the objects in the share
Based on this configuration what is true of the reader account?

A. The reader account will automatically use the Standard edition of Snowflake.
B. The reader account compute will be billed to the provider account.
C. The reader account can clone data the provider has shared, but cannot re-share it.
D. The reader account can create a copy of the shared data using CREATE TABLE AS...

Answer: B

Explanation:
The reader account compute will be billed to the provider account.
Very Comprehensive Explanation
In Snowflake, when a provider creates a reader account for a consumer who does not have a Snowflake account, the compute resources used by the reader account are billed to the provider??s account. This allows the consumer to query the shared data without incurring any costs. References: [COF-C02] SnowPro Core Certification Exam Study Guide

NEW QUESTION 420
- (Topic 3)
Which of the following describes the Snowflake Cloud Services layer?

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

Answer: A

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

NEW QUESTION 425
- (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 428
- (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 433
- (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 436
- (Topic 3)
Which Snowflake URL type is used by directory tables?

A. File
B. Pre-signed
C. Scoped
D. Virtual-hosted style

Answer: C

Explanation:
The Snowflake URL type used by directory tables is the scoped URL. This type of URL provides access to files in a stage with metadata, such as the Snowflake file URL, for each file

NEW QUESTION 441
- (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 442
- (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 447
- (Topic 3)
A view is defined on a permanent table. A temporary table with the same name is created in the same schema as the referenced table. What will the query from the view return?

A. The data from the permanent table.
B. The data from the temporary table.

C. An error stating that the view could not be compiled.
D. An error stating that the referenced object could not be uniquely identified.

Answer: A

Explanation:
When a view is defined on a permanent table, and a temporary table with the same name is created in the same schema, the query from the view will return the data from the permanent table. Temporary tables are session-specific and do not affect the data returned by views defined on permanent tables2.

NEW QUESTION 452
- (Topic 3)
Credit charges for Snowflake virtual warehouses are calculated based on which of the following considerations? (Choose two.)

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

Answer: CD

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

NEW QUESTION 453
- (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 455
- (Topic 3)
What is cached during a query on a virtual warehouse?

A. All columns in a micro-partition
B. Any columns accessed during the query
C. The columns in the result set of the query
D. All rows accessed during the query

Answer: C

Explanation:
During a query on a virtual warehouse, the columns in the result set of the query are cached. This allows for faster retrieval of data if the same or a similar query is run again, as the system can retrieve the data from the cache rather than reprocessing the entire query. References: [COF-C02] SnowPro Core Certification Exam Study Guide

NEW QUESTION 459
- (Topic 3)
A Snowflake user has been granted the create data EXCHANGE listing privilege with their role.
Which tasks can this user now perform on the Data Exchange? (Select TWO).

A. Rename listings.
B. Delete provider profiles.
C. Modify listings properties.
D. Modify incoming listing access requests.
E. Submit listings for approval/publishing.

Answer: CE

Explanation:
With the create data EXCHANGE listing privilege, a Snowflake user can modify the properties of listings and submit them for approval or publishing on the Data Exchange. This allows them to manage and share data sets with consumers effectively. References: Based on general data exchange practices in cloud services as of 2021.

NEW QUESTION 463
.....

Thank You for Trying Our Product

* 100% Pass or Money Back

All our products come with a 90-day Money Back Guarantee.

* One year free update

You can enjoy free update one year. 24x7 online support.

* Trusted by Millions

We currently serve more than 30,000,000 customers.

* Shop Securely

All transactions are protected by VeriSign!

100% Pass Your COF-C02 Exam with Our Prep Materials Via below:

https://www.certleader.com/COF-C02-dumps.html