Snowflake

Exam Questions COF-C02

SnowPro Core Certification Exam (COF-C02)

NEW QUESTION 1
- (Topic 1)
Which of the following Snowflake capabilities are available in all Snowflake editions? (Select TWO)

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

Answer: BD

Explanation:
In all Snowflake editions, two key capabilities are universally available:
? B. Automatic encryption of all data: Snowflake automatically encrypts all data stored in its platform, ensuring security and compliance with various regulations. This encryption is transparent to users and does not require any configuration or management.
? D. Object-level access control: Snowflake provides granular access control mechanisms that allow administrators to define permissions at the object level, including databases, schemas, tables, and views. This ensures that only authorized users can access specific data objects.
These features are part of Snowflake??s commitment to security and governance, and they are included in every edition of the Snowflake Data Cloud.
References:
? Snowflake Documentation on Security Features
? SnowPro® Core Certification Exam Study Guide

NEW QUESTION 2
- (Topic 1)
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 3
- (Topic 1)
What is the default character set used when loading CSV files into Snowflake?

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

Answer: A

Explanation:
https://docs.snowflake.com/en/user-guide/intro-summary-loading.html#:~:text=For%20delimited%20files%20(CSV%2C%20TSV,encoding%20to%20use%20for%20loading.
For delimited files (CSV, TSV, etc.), the default character set is UTF-8. To use any other characters sets, you must explicitly specify the encoding to use for loading. For the list of supported character sets, see Supported Character Sets for Delimited Files (in this topic).

NEW QUESTION 4
- (Topic 1)
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 can be used to view warehouse usage over time? (Select Two).

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

Answer: BD

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

NEW QUESTION 6
- (Topic 1)
Which Snowflake feature is used for both querying and restoring data?

A. Cluster keys
B. Time Travel
C. Fail-safe
D. Cloning

Answer: B

Explanation:
Snowflake??s Time Travel feature is used for both querying historical data in tables and restoring and cloning historical data in databases, schemas, and tables3. It allows users to access historical data within a defined period (1 day by default, up to 90 days for Snowflake Enterprise Edition) and is a key feature for data recovery and management. References: [COF-C02] SnowPro Core Certification Exam Study Guide

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

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

Answer: BDE

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

NEW QUESTION 8
- (Topic 1)
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 9
- (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 10
- (Topic 1)
A user has an application that writes a new Tile to a cloud storage location every 5 minutes.
What would be the MOST efficient way to get the files into Snowflake?

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

Answer: D

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

NEW QUESTION 10
- (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 14
- (Topic 1)
What are value types that a VARIANT column can store? (Select TWO)

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

Answer: BD

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

NEW QUESTION 16
- (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 20
- (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 24
- (Topic 1)
What tasks can be completed using the copy command? (Select TWO)

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

Answer: CD

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

NEW QUESTION 25
- (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 28
- (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 29
- (Topic 1)
True or False: It is possible for a user to run a query against the query result cache without requiring an active Warehouse.

A. True
B. False

Answer: A

Explanation:
Snowflake??s architecture allows for the use of a query result cache that stores the results of queries for a period of time. If the same query is run again and the underlying data has not changed, Snowflake can retrieve the result from this cache without needing to re-run the query on an active warehouse, thus saving on compute resources.

NEW QUESTION 30
- (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 33
- (Topic 1)
What is a best practice after creating a custom role?

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

Answer: B

Explanation:
Assigning the custom role to the SYSADMIN role is considered a best practice because it allows the SYSADMIN role to manage objects created by the custom role. This is important for maintaining proper access control and ensuring that the SYSADMIN can perform necessary administrative tasks on objects created by users with the custom role.
References:
? [COF-C02] SnowPro Core Certification Exam Study Guide
? Section 1.3 - SnowPro Core Certification Study Guide1

NEW QUESTION 34
- (Topic 1)
A developer is granted ownership of a table that has a masking policy. The developer's role is not able to see the masked data. Will the developer be able to modify the table to read the masked data?

A. Yes, because a table owner has full control and can unset masking policies.
B. Yes, because masking policies only apply to cloned tables.
C. No, because masking policies must always reference specific access roles.
D. No, because ownership of a table does not include the ability to change masking policies

Answer: D

Explanation:
Even if a developer is granted ownership of a table with a masking policy, they will not be able to modify the table to read the masked data if their role does not have the necessary permissions. Ownership of a table does not automatically confer the ability
to alter masking policies, which are designed to protect sensitive data. Masking policies are applied at the schema level and require specific privileges to modify12.
References:
? [COF-C02] SnowPro Core Certification Exam Study Guide
? Snowflake Documentation on Masking Policies

NEW QUESTION 37
- (Topic 1)
The fail-safe retention period is how many days?

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

Answer: B

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

NEW QUESTION 42
- (Topic 1)
A sales table FCT_SALES has 100 million records. The following Query was executed
SELECT COUNT (1) FROM FCT SALES;
How did Snowflake fulfill this query?

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

Answer: D

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

NEW QUESTION 47
- (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 49
- (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 50
- (Topic 1)
True or False: Fail-safe can be disabled within a Snowflake account.

A. True
B. False

Answer: B

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

NEW QUESTION 53
- (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 55
- (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 57
- (Topic 1)
Which Snowflake object enables loading data from files as soon as they are available in a cloud storage location?

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

Answer: A

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

NEW QUESTION 60
- (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 62
- (Topic 1)
A marketing co-worker has requested the ability to change a warehouse size on their medium virtual warehouse called mktg WH.
Which of the following statements will accommodate this request?

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

Answer: B

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

NEW QUESTION 63
- (Topic 1)
What SQL command would be used to view all roles that were granted to user.1?

A. show grants to user USER1;
B. show grants of user USER1;
C. describe user USER1;
D. show grants on user USER1;

Answer: A

Explanation:
The correct command to view all roles granted to a specific user in Snowflake is SHOW GRANTS TO USER <user_name>;. This command lists all access control privileges that have been explicitly granted to the specified user.
References: SHOW GRANTS | Snowflake Documentation

NEW QUESTION 68
- (Topic 1)
What transformations are supported in a CREATE PIPE ... AS COPY ... FROM (....) statement? (Select TWO.)

A. Data can be filtered by an optional where clause
B. Incoming data can be joined with other tables
C. Columns can be reordered
D. Columns can be omitted
E. Row level access can be defined

Answer: AD

Explanation:
In a CREATE PIPE ... AS COPY ... FROM (....) statement, the supported transformations include filtering data using an optional WHERE clause and omitting columns. The WHERE clause allows for the specification of conditions to filter the data that is being loaded, ensuring only relevant data is inserted into the table. Omitting columns enables the exclusion of certain columns from the data load, which can be useful when the incoming data contains more columns than are needed for the target table.
References:
? [COF-C02] SnowPro Core Certification Exam Study Guide
? Simple Transformations During a Load1

NEW QUESTION 71
- (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 76

- (Topic 1)
A user has unloaded data from Snowflake to a stage
Which SQL command should be used to validate which data was loaded into the stage?

A. list @file stage
B. show @file stage
C. view @file stage
D. verify @file stage

Answer: A

Explanation:
The list command in Snowflake is used to validate and display the list of files in a specified stage. When a user has unloaded data to a stage, running the list @file stage command will show all the files that have been uploaded to that stage, allowing the user to verify the data that was unloaded.
References:
? Snowflake Documentation on Stages
? SnowPro® Core Certification Study Guide

NEW QUESTION 77
- (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 81
- (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 85
- (Topic 1)
Which semi-structured file formats are supported when unloading data from a table? (Select TWO).

A. ORC
B. XML
C. Avro
D. Parquet
E. JSON

Answer: DE

Explanation:
Semi-structured JSON, Parquet Snowflake supports unloading data in several semi-structured file formats, including Parquet and JSON. These formats allow for efficient storage and querying of semi-
structured data, which can be loaded directly into Snowflake tables without requiring a predefined schema12.
https://docs.snowflake.com/en/user-guide/data-unload- prepare.html#:~:text=Supported%20File%20Formats,-The%20following%20file&text=Delimited%20(CSV%2C%20TSV%2C%20etc.)

NEW QUESTION 89

- (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 90
- (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 95
- (Topic 1)
Which command can be used to stage local files from which Snowflake interface?

A. SnowSQL
B. Snowflake classic web interface (Ul)
C. Snowsight
D. .NET driver

Answer: A

Explanation:
SnowSQL is the command-line client for Snowflake that allows users to execute SQL queries and perform all DDL and DML operations, including staging files for bulk data loading. It is specifically designed for scripting and automating tasks. References:
? SnowPro Core Certification Exam Study Guide
? Snowflake Documentation on SnowSQL https://docs.snowflake.com/en/user-guide/snowsql-use.html

NEW QUESTION 96
- (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 99
- (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 103
- (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 107
- (Topic 1)
What happens when an external or an internal stage is dropped? (Select TWO).

A. When dropping an external stage, the files are not removed and only the stage is dropped
B. When dropping an external stage, both the stage and the files within the stage are removed
C. When dropping an internal stage, the files are deleted with the stage and the files are recoverable
D. When dropping an internal stage, the files are deleted with the stage and the files arenot recoverable
E. When dropping an internal stage, only selected files are deleted with the stage and are not recoverable

Answer: AD

Explanation:
When an external stage is dropped in Snowflake, the reference to the external storage location is removed, but the actual files within the external storage (like Amazon S3, Google Cloud Storage, or Microsoft Azure) are not deleted. This means that the data remains intact in the external storage location, and only the stage object in Snowflake is removed.
On the other hand, when an internal stage is dropped, any files that were uploaded to the stage are deleted along with the stage itself. These files are not recoverable once the internal stage is dropped, as they are permanently removed from Snowflake??s storage. References:
? [COF-C02] SnowPro Core Certification Exam Study Guide
? Snowflake Documentation on Stages

NEW QUESTION 108
- (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 109
- (Topic 2)
The following JSON is stored in a VARIANT column called src of the CAR_SALES table:

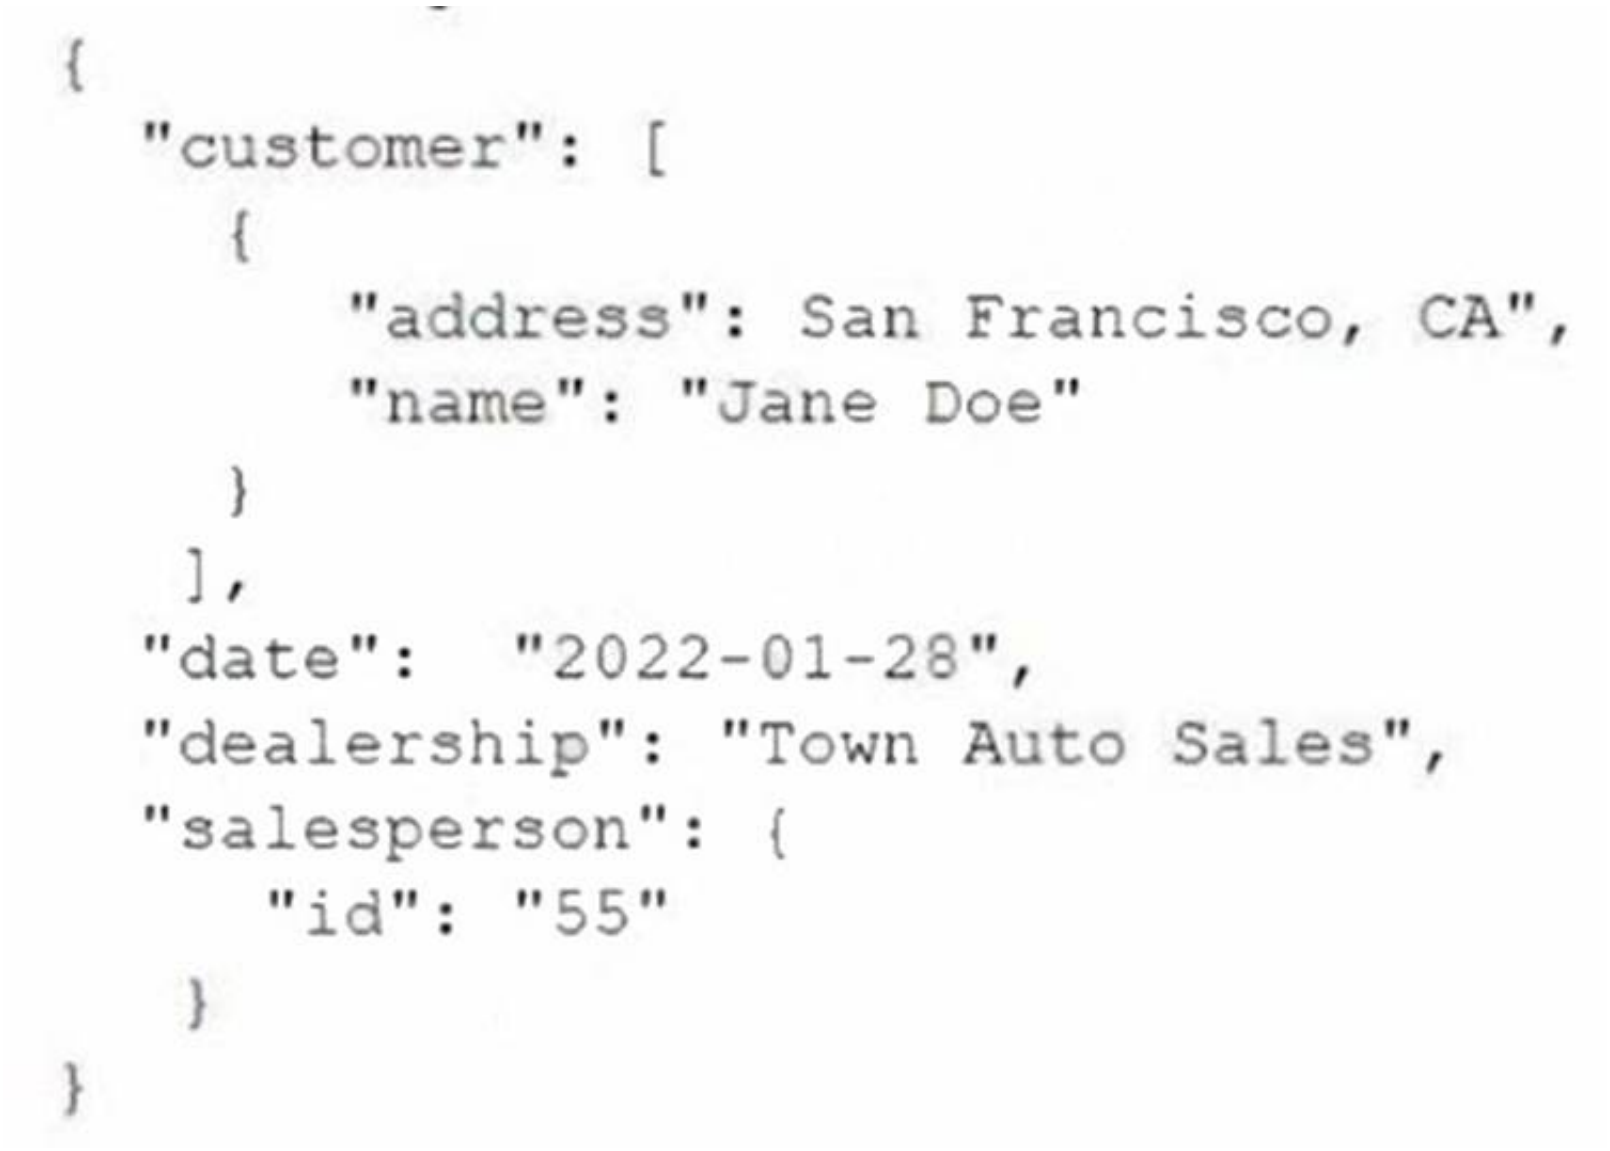

A user needs to extract the dealership information from the JSON. How can this be accomplished?

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 112
- (Topic 2)
How does Snowflake Fail-safe protect data in a permanent table?

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

Answer: C

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

NEW QUESTION 114
- (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 116
- (Topic 2)
When publishing a Snowflake Data Marketplace listing into a remote region what should be taken into consideration? (Choose two.)

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

Answer: BC

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

NEW QUESTION 119
- (Topic 2)
Which file formats are supported for unloading data from Snowflake? (Choose two.)

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

Answer: BE

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

NEW QUESTION 122
- (Topic 2)
When should a multi-cluster warehouse be used in auto-scaling mode?

A. When it is unknown how much compute power is needed
B. If the select statement contains a large number of temporary tables or Common Table Expressions (CTEs)
C. If the runtime of the executed query is very slow
D. When a large number of concurrent queries are run on the same warehouse

Answer: D

Explanation:
A multi-cluster warehouse should be used in auto-scaling mode when there is a need to handle a large number of concurrent queries. Auto-scaling allows Snowflake to automatically add or remove compute clusters to balance the load, ensuring that performance remains consistent during varying levels of demand

NEW QUESTION 123
- (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 128
- (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 131
- (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 133
- (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 135
- (Topic 2)
In a Snowflake role hierarchy, what is the top-level role?

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

Answer: C

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

NEW QUESTION 140
- (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 144
- (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 146
- (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 147
- (Topic 2)
Which of the following describes a Snowflake stored procedure?

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

Answer: D

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

NEW QUESTION 149
- (Topic 2)
The Snowflake Cloud Data Platform is described as having which of the following architectures?

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

Answer: C

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

NEW QUESTION 152
- (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 154
- (Topic 2)
When cloning a database containing stored procedures and regular views, that have fully qualified table references, which of the following will occur?

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

Answer: A

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

NEW QUESTION 155
- (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 157
- (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 162
- (Topic 2)
What is the purpose of multi-cluster virtual warehouses?

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

Answer: C

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

NEW QUESTION 163
- (Topic 2)
The Snowflake Search Optimization Services supports improved performance of which kind of query?

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

Answer: C

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

NEW QUESTION 168
- (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 172
- (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 177
- (Topic 2)
A single user of a virtual warehouse has set the warehouse to auto-resume and auto- suspend after 10 minutes. The warehouse is currently suspended and the user performs the following actions:
* 1. Runs a query that takes 3 minutes to complete
* 2. Leaves for 15 minutes
* 3. Returns and runs a query that takes 10 seconds to complete
* 4. Manually suspends the warehouse as soon as the last query was completed When the user returns, how much billable compute time will have been consumed?

A. 4 minutes
B. 10 minutes
C. 14 minutes
D. 24 minutes

Answer: C

Explanation:
The billable compute time includes the time the warehouse is running queries plus the auto-suspend time after the last query if the warehouse is not manually suspended. In this scenario, the warehouse runs for 3 minutes, suspends after 10 minutes of inactivity, resumes for a 10-second query, and then is manually suspended. The total billable time is the sum of the initial 3 minutes, the 10 minutes of auto-suspend time, and the brief period for the 10-second query, which is rounded up to the next full minute due to Snowflake??s billing practices. References: [COF-C02] SnowPro Core Certification Exam Study Guide

NEW QUESTION 181
- (Topic 2)
Which data types are supported by Snowflake when using semi-structured data? (Choose two.)

A. VARIANT
B. VARRAY
C. STRUCT
D. ARRAY
E. QUEUE

Answer: AD

Explanation:
Snowflake supports the VARIANT and ARRAY data types for semi- structured data. VARIANT can store values of any other type, including OBJECT and ARRAY, making it suitable for semi-structured data formats like JSON. ARRAY is used to store an ordered list of elements

NEW QUESTION 185
- (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 189
- (Topic 2)
Which command can be used to load data files into a Snowflake stage?

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

Answer: C

Explanation:
The PUT command is used to load data files into a Snowflake stage. This command uploads data files from a local file system to a specified stage in Snowflake

NEW QUESTION 191
- (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 195
- (Topic 2)
What is an advantage of using an explain plan instead of the query profiler to evaluate the performance of a query?

A. The explain plan output is available graphically.
B. An explain plan can be used to conduct performance analysis without executing a query.
C. An explain plan will handle queries with temporary tables and the query profiler will not.
D. An explain plan's output will display automatic data skew optimization information.

Answer: B

Explanation:
An explain plan is beneficial because it allows for the evaluation of how a query will be processed without the need to actually execute the query. This can help in understanding the query??s performance implications and potential bottlenecks without consuming resources that would be used if the query were run

NEW QUESTION 196
- (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 198
- (Topic 2)
How can a row access policy be applied to a table or a view? (Choose two.)

A. Within the policy DDL
B. Within the create table or create view DDL
C. By future APPLY for all objects in a schema
D. Within a control table
E. Using the command ALTER <object> ADD ROW ACCESS POLICY <policy>;

Answer: AE

Explanation:
A row access policy can be applied to a table or a view within the policy DDL when defining the policy. Additionally, an existing row access policy can be applied to a table or a view using the ALTER <object> ADD ROW ACCESS POLICY <policy> command

NEW QUESTION 203
- (Topic 2)
True or False: Snowpipe via REST API can only reference External Stages as source.

A. True
B. False

Answer: B

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

NEW QUESTION 207
- (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 212
- (Topic 2)
Which statement is true about running tasks in Snowflake?

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

Answer: B

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

NEW QUESTION 217
- (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 221
- (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 223
- (Topic 2)
Which command should be used to download files from a Snowflake stage to a local folder on a client's machine?

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

Answer: B

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

NEW QUESTION 226
- (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 230
- (Topic 2)
In an auto-scaling multi-cluster virtual warehouse with the setting SCALING_POLICY = ECONOMY enabled, when is another cluster started?

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

Answer: A

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

NEW QUESTION 234
- (Topic 2)
In the Snowflake access control model, which entity owns an object by default?

A. The user who created the object
B. The SYSADMIN role
C. Ownership depends on the type of object
D. The role used to create the object

Answer: D

Explanation:
In Snowflake??s access control model, the default owner of an object is the role that was used to create the object. This role has the OWNERSHIP privilege on the object and can grant access to other roles1

NEW QUESTION 237
- (Topic 2)
A running virtual warehouse is suspended.
What is the MINIMUM amount of time that the warehouse will incur charges for when it is restarted?

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

Answer: B

Explanation:

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

NEW QUESTION 242
- (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 247
- (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 252
- (Topic 3)
If a Snowflake user decides a table should be clustered, what should be used as the cluster key?

A. The columns that are queried in the select clause.
B. The columns with very high cardinality.
C. The columns with many different values.
D. The columns most actively used in the select filters.

Answer: D

Explanation:
When deciding on a clustering key for a table, Snowflake recommends using the columns that are most actively used in the select filters. This is because clustering by these columns can improve the performance of queries that filter on these values, leading to more efficient scans and better overall query performance2. References: [COF-C02] SnowPro Core Certification Exam Study Guide

NEW QUESTION 255
- (Topic 3)
What statistical information in a Query Profile indicates that the query is too large to fit in memory? (Select TWO).

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

Answer: AB

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

NEW QUESTION 259
- (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 264
- (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 269
- (Topic 3)
Which role has the ability to create and manage users and roles?

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

Answer: B

Explanation:
The USERADMIN role in Snowflake has the ability to create and manage users and roles within the Snowflake environment. This role is specifically dedicated to user and role management and creation

NEW QUESTION 273
- (Topic 3)
What column type does a Kafka connector store formatted information in a single column?

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

Answer: D

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

NEW QUESTION 275
- (Topic 3)
Which Snowflake URL type allows users or applications to download or access files directly from Snowflake stage without authentication?

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

Answer: C

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

NEW QUESTION 278
- (Topic 3)
For the ALLOWED VALUES tag property, what is the MAXIMUM number of possible string values for a single tag?

A. 10
B. 50
C. 64

D. 256

Answer: D

Explanation:
For the ALLOWED VALUES tag property, the maximum number of possible string values for a single tag is 256. This allows for a wide range of values to be assigned to a tag when it is set on an object

NEW QUESTION 281
- (Topic 3)
What effect does WAIT_FOR_COMPLETION = TRUE have when running an ALTER WAREHOUSE command and changing the warehouse size?

A. The warehouse size does not change until all queries currently running in the warehouse have completed.
B. The warehouse size does not change until all queries currently in the warehouse queue have completed.
C. The warehouse size does not change until the warehouse is suspended and restarted.
D. It does not return from the command until the warehouse has finished changing its size.

Answer: D

Explanation:
The WAIT_FOR_COMPLETION = TRUE parameter in an ALTER WAREHOUSE command ensures that the command does not return until the warehouse has completed resizing. This means that the command will wait until all the necessary compute resources have been provisioned and the warehouse size has been changed. References: [COF-C02] SnowPro Core Certification Exam Study Guide

NEW QUESTION 286
- (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 287
- (Topic 3)
What is the purpose of using the OBJECT_CONSTRUCT function with me COPY INTO command?

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

Answer: B

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

NEW QUESTION 289
- (Topic 3)
Which query profile statistics help determine if efficient pruning is occurring? (Choose two.)

A. Bytes sent over network
B. Percentage scanned from cache
C. Partitions total
D. Bytes spilled to local storage
E. Partitions scanned

Answer: CE

Explanation:
Efficient pruning in Snowflake is indicated by the number of partitions scanned out of the total available. If a small percentage of partitions are scanned, it suggests that the pruning process is effectively narrowing down the data, which can lead to improved query performance

NEW QUESTION 291
- (Topic 3)
What MINIMUM privilege is required on the external stage for any role in the GET REST API to access unstructured data files using a file URL?

A. READ
B. OWNERSHIP

C. USAGK
D. WRTTF

Answer: A

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

NEW QUESTION 294
- (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 299
- (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 304
- (Topic 3)
What is the recommended way to change the existing file format type in my format from CSV to JSON?

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

Answer: A

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

NEW QUESTION 309
- (Topic 3)
Which Snowflake edition enables data sharing only through Snowflake Support?

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

Answer: A

Explanation:
The Snowflake edition that enables data sharing only through Snowflake Support is the Virtual Private Snowflake (VPS). By default, VPS does not permit data sharing outside of the VPS environment, but it can be enabled through Snowflake Support4.

NEW QUESTION 313
- (Topic 3)
Which features could be used to improve the performance of queries that return a small subset of rows from a large table? (Select TWO).

A. Search optimization service
B. Automatic clustering
C. Row access policies

D. Multi-cluster virtual warehouses
E. Secure views

Answer: AB

Explanation:
The search optimization service and automatic clustering are features that can improve the performance of queries returning a small subset of rows from a large table. The search optimization service is designed for low-latency point lookup queries, while automatic clustering organizes data in micro-partitions based on specific dimensions to reduce the amount of data scanned during queries.

NEW QUESTION 316
- (Topic 3)
What happens when a database is cloned?

A. It does not retain any privileges granted on the source object.
B. It replicates all granted privileges on the corresponding source objects.
C. It replicates all granted privileges on the corresponding child objects.
D. It replicates all granted privileges on the corresponding child schema objects.

Answer: A

Explanation:
When a database is cloned in Snowflake, it does not retain any privileges that were granted on the source object. The clone will need to have privileges reassigned as necessary for users to access it. References: [COF-C02] SnowPro Core Certification Exam Study Guide

NEW QUESTION 317
- (Topic 3)
Snowflake's hierarchical key mode includes which keys? (Select TWO).

A. Account master keys
B. Database master keys
C. File keys
D. Secure view keys
E. Schema master keys

Answer: AC

Explanation:
Snowflake??s hierarchical key model includes several levels of keys, where Account master keys and File keys are part of this hierarchy. Account master keys are used to encrypt all the data within an account, while File keys are used to encrypt individual files within the database2.

NEW QUESTION 319
- (Topic 3)
Which privilege is required for a role to be able to resume a suspended warehouse if auto- resume is not enabled?

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

Answer: B

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

NEW QUESTION 320
- (Topic 3)
By definition, a secure view is exposed only to users with what privilege?

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

Answer: B

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

NEW QUESTION 325
- (Topic 3)
What is a characteristic of the Snowflake Query Profile?

A. It can provide statistics on a maximum number of 100 queries per week.

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

Answer: B

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

NEW QUESTION 328
- (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 332
- (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 336
- (Topic 3)
Data storage for individual tables can be monitored using which commands and/or objects? (Choose two.)

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

Answer: AE

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

NEW QUESTION 339
- (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 344

- (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)
A company needs to read multiple terabytes of data for an initial load as part of a Snowflake migration. The company can control the number and size of CSV extract files.
How does Snowflake recommend maximizing the load performance?

A. Use auto-ingest Snowpipes to load large files in a serverless model.
B. Produce the largest files possible, reducing the overall number of files to process.
C. Produce a larger number of smaller files and process the ingestion with size Small virtual warehouses.
D. Use an external tool to issue batched row-by-row inserts within BEGIN TRANSACTION and COMMIT commands.

Answer: B

Explanation:
Snowflake??s documentation recommends producing the largest files possible for data loading, as larger files reduce the number of files to process and the overhead associated with handling many small files. This approach can maximize the load performance by leveraging Snowflake??s ability to ingest large files efficiently1. References:
[COF-C02] SnowPro Core Certification Exam Study Guide

NEW QUESTION 349
- (Topic 3)
Query parsing and compilation occurs in which architecture layer of the Snowflake Cloud Data Platform?

A. Cloud services layer
B. Compute layer
C. Storage layer
D. Cloud agnostic layer

Answer: A

Explanation:
Query parsing and compilation in Snowflake occur within the cloud services layer. This layer is responsible for various management tasks, including query compilation and optimization

NEW QUESTION 354
- (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 356
- (Topic 3)
What privilege should a user be granted to change permissions for new objects in a managed access schema?

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

Answer: C

Explanation:
To change permissions for new objects in a managed access schema, a user should be granted the MANAGE GRANTS global privilege. This privilege allows the

user to manage access control through grants on all securable objects within Snowflake2. References: [COF-C02] SnowPro Core Certification Exam Study Guide

NEW QUESTION 358
- (Topic 3)
When should a user consider disabling auto-suspend for a virtual warehouse? (Select TWO).

A. When users will be using compute at different times throughout a 24/7 period
B. When managing a steady workload
C. When the compute must be available with no delay or lag time
D. When the user does not want to have to manually turn on the warehouse each time it is needed
E. When the warehouse is shared across different teams

Answer: BC

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

NEW QUESTION 363
- (Topic 3)
For non-materialized views, what column in Information Schema and Account Usage identifies whether a view is secure or not?

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

Answer: B

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

NEW QUESTION 365
- (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 370
- (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)
Which statement describes pruning?

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

Answer: A

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

NEW QUESTION 380
- (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 383
- (Topic 3)
If queries start to queue in a multi-cluster virtual warehouse, an additional compute cluster starts immediately under what setting?

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

Answer: A

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

NEW QUESTION 387
- (Topic 3)
Which features make up Snowflake's column level security? (Select TWO).

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

Answer: BC

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

NEW QUESTION 389
- (Topic 3)
The bulk data load history that is available upon completion of the COPY statement is stored where and for how long?

A. In the metadata of the target table for 14 days
B. In the metadata of the pipe for 14 days
C. In the metadata of the target table for 64 days
D. In the metadata of the pipe for 64 days

Answer: D

Explanation:
The bulk data load history available after a COPY statement is stored in the metadata of the pipe and is retained for 64 days1.

NEW QUESTION 391
- (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 393
- (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 395
- (Topic 3)
What action can a user take to address query concurrency issues?

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

Answer: C

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

NEW QUESTION 400
- (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 403
- (Topic 3)
When unloading data to an external stage, what is the MAXIMUM file size supported?

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

Answer: B

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

NEW QUESTION 404
- (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 409
- (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 414
- (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 415
- (Topic 3)
How many resource monitors can be assigned at the account level?

A. 1
B. 2
C. 3
D. 4

Answer: A

Explanation:
Snowflake allows for only one resource monitor to be assigned at the account level. This monitor oversees the credit usage of all the warehouses in the account. References: Snowflake Documentation

NEW QUESTION 417
- (Topic 3)
A tabular User-Defined Function (UDF) is defined by specifying a return clause that contains which keyword?

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

Answer: B

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

NEW QUESTION 421
- (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 426
- (Topic 3)
What is the name of the SnowSQLfile that can store connection information?

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

Answer: B

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

NEW QUESTION 427
- (Topic 3)
Which commands should be used to grant the privilege allowing a role to select data from all current tables and any tables that will be created later in a schema? (Choose two.)

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

Answer: CD

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

NEW QUESTION 431
- (Topic 3)
A user has a standard multi-cluster warehouse auto-scaling policy in place.
Which condition will trigger a cluster to shut-down?

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

Answer: D

Explanation:
In a standard multi-cluster warehouse with auto-scaling, a cluster will shut down when, after 2-3 consecutive checks, the system determines that the load on the least-loaded cluster could be redistributed to other clusters. This ensures efficient resource utilization and cost management. References: [COF-C02] SnowPro Core Certification Exam Study Guide

NEW QUESTION 435
- (Topic 3)
Which parameter can be used to instruct a COPY command to verify data files instead of loading them into a specified table?

A. STRIP_NULL_VALUES
B. SKIP_BYTE_ORDER_MARK
C. REPLACE_INVALID_CHARACTERS
D. VALIDATION_MODE

Answer: D

Explanation:
The VALIDATION_MODE parameter can be used with the COPY command to verify data files without loading them into the specified table. This parameter allows users to check for errors in the files

NEW QUESTION 440
- (Topic 3)
A Snowflake user executed a query and received the results. Another user executed the same query 4 hours later. The data had not changed.
What will occur?

A. No virtual warehouse will be used, data will be read from the result cache.
B. No virtual warehouse will be used, data will be read from the local disk cache.
C. The default virtual warehouse will be used to read all data.
D. The virtual warehouse that is defined at the session level will be used to read all data.

Answer: A

Explanation:
Snowflake maintains a result cache that stores the results of every query for 24 hours. If the same query is executed again within this time frame and the data has not changed, Snowflake will retrieve the data from the result cache instead of using a virtual warehouse to recompute the results2.

NEW QUESTION 445
- (Topic 3)
Which feature is integrated to support Multi-Factor Authentication (MFA) at Snowflake?

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

Answer: B

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

NEW QUESTION 448
- (Topic 3)
How does Snowflake recommend handling the bulk loading of data batches from files already available in cloud storage?

A. Use Snowpipe.
B. Use the INSERT command.
C. Use an external table.
D. Use the COPY command.

Answer: D

Explanation:
Snowflake recommends using the COPY command for bulk loading data batches from files already available in cloud storage. This command allows for efficient and large-scale data loading operations from files staged in cloud storage into Snowflake tables3.

NEW QUESTION 450
- (Topic 3)
How would a user run a multi-cluster warehouse in maximized mode?

A. Configure the maximum clusters setting to "Maximum."
B. Turn on the additional clusters manually after starting the warehouse.
C. Set the minimum Clusters and maximum Clusters settings to the same value.
D. Set the minimum clusters and maximum clusters settings to different values.

Answer: C

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

NEW QUESTION 451
- (Topic 3)
Which REST API can be used with unstructured data?

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

Answer: C

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

NEW QUESTION 455
- (Topic 4)
Which Snowflake feature allows a user to track sensitive data for compliance, discovery, protection, and resource usage?

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

Answer: A

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

NEW QUESTION 457
- (Topic 4)
How can a Snowflake administrator determine which user has accessed a database object that contains sensitive information?

A. Review the granted privileges to the database object.
B. Review the row access policy for the database object.
C. Query the ACCESS_HISTORY view in the ACCOUNT_USAGE schema.
D. Query the REPLICATION USAGE HISTORY view in the ORGANIZATION USAGEschema.

Answer: C

Explanation:
To determine which user has accessed a database object containing sensitive information, a Snowflake administrator can query the ACCESS_HISTORY view in the ACCOUNT_USAGE schema, which provides information about access to database objects3.

NEW QUESTION 460
- (Topic 4)
Which data types can be used in Snowflake to store semi-structured data? (Select TWO)

A. ARRAY
B. BLOB
C. CLOB
D. JSON
E. VARIANT

Answer: AE

Explanation:
Snowflake supports the storage of semi-structured data using the ARRAY and VARIANT data types. The ARRAY data type can directly contain VARIANT, and thus indirectly contain any other data type, including itself. The VARIANT data type can store a value of any other type, including OBJECT and ARRAY, and is often used to represent semi-structured data formats like JSON, Avro, ORC, Parquet, or XML34.
References: [COF-C02] SnowPro Core Certification Exam Study Guide

NEW QUESTION 464
- (Topic 4)
A user with which privileges can create or manage other users in a Snowflake account? (Select TWO).

A. GRANT
B. SELECT
C. MODIFY
D. OWNERSHIP
E. CREATE USER

Answer: DE

Explanation:
A user with the OWNERSHIP privilege on a user object or the CREATE USER privilege on the account can create or manage other users in a Snowflake account56.

NEW QUESTION 469
- (Topic 4)
Which views are included in the DATA SHARING USAGE schema? (Select TWO).

A. ACCESS_HISTORY
B. DATA_TRANSFER_HISTORY
C. WAREHOUSE_METERING_HISTORY
D. MONETIZED_USAGE_DAILY
E. LISTING TELEMETRY DAILY

Answer: DE

Explanation:
The DATA_SHARING_USAGE schema includes views that display information about listings published in the Snowflake Marketplace or a data exchange, which includes DATA_TRANSFER_HISTORY and LISTING_TELEMETRY_DAILY2.

NEW QUESTION 471
- (Topic 4)
A tag object has been assigned to a table (TABLE_A) in a schema within a Snowflake database.
Which CREATE object statement will automatically assign the TABLE_A tag to a target object?

A. CREATE TABLE <table_name> LIKE TABLE_A;
B. CREATE VIEW <view_name> AS SELECT * FROM TABLE_A;
C. CREATE TABLE <table_name> AS SELECT * FROM TABLE_A;
D. CREATE MATERIALIZED VIEW <view name> AS SELECT * FROM TABLE A;

Answer: C

Explanation:
When a tag object is assigned to a table, using the statement CREATE TABLE <table_name> AS SELECT * FROM TABLE_A will automatically assign the TABLE_A tag to the newly created table2.

NEW QUESTION 476
- (Topic 4)
Which commands are restricted in owner's rights stored procedures? (Select TWO).

A. SHOW
B. MERGE
C. INSERT
D. DELETE
E. DESCRIBE

Answer: AE

Explanation:
In owner??s rights stored procedures, certain commands are restricted to maintain security and integrity. The SHOW and DESCRIBE commands are limited because they can reveal metadata and structure information that may not be intended for all roles.

NEW QUESTION 481
- (Topic 4)
What metadata does Snowflake store for rows in micro-partitions? (Select TWO).

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

Answer: AB

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

NEW QUESTION 483
- (Topic 4)
How can performance be optimized for a query that returns a small amount of data from a very large base table?

A. Use clustering keys
B. Create materialized views
C. Use the search optimization service
D. Use the query acceleration service

Answer: C

Explanation:
The search optimization service in Snowflake is designed to improve the performance of selective point lookup queries on large tables, which is ideal for scenarios where a query returns a small amount of data from a very large base table1. References: [COF-C02] SnowPro Core Certification Exam Study Guide

NEW QUESTION 486
- (Topic 4)
Which VALIDATION_MODE value will return the errors across the files specified in a COPY command, including files that were partially loaded during an earlier load?

A. RETURN_-1_R0WS
B. RETURN_n_ROWS
C. RETURN_ERRORS
D. RETURN ALL ERRORS

Answer: C

Explanation:
The RETURN_ERRORS value in the VALIDATION_MODE option of the COPY command instructs Snowflake to validate the data files and return errors encountered across all specified files, including those that were partially loaded during an earlier load2. References: [COF-C02] SnowPro Core Certification Exam Study Guide

NEW QUESTION 491
- (Topic 4)
Which type of loop requires a BREAK statement to stop executing?

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

Answer: B

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

NEW QUESTION 494
- (Topic 4)
What factors impact storage costs in Snowflake? (Select TWO).

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

Answer: AC

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

NEW QUESTION 495
- (Topic 4)
Which Snowflake function will parse a JSON-null into a SQL-null?

A. TO_CHAR
B. TO_VARIANT
C. TO_VARCHAR
D. STRIP NULL VALUE

Answer: D

Explanation:
The STRIP_NULL_VALUE function in Snowflake is used to convert a JSON null value into a SQL NULL value1.

NEW QUESTION 498
- (Topic 4)
How is unstructured data retrieved from data storage?

A. SQL functions like the GET command can be used to copy the unstructured data to a location on the client.
B. SQL functions can be used to create different types of URLs pointing to the unstructured dat
C. These URLs can be used to download the data to a client.
D. SQL functions can be used to retrieve the data from the query results cach
E. When the query results are output to a client, the unstructured data will be output to the client as files.
F. SQL functions can call on different web extensions designed to display different types of files as a web pag
G. The web extensions will allow the files to be downloaded to the client.

Answer: B

Explanation:
Unstructured data stored in Snowflake can be retrieved by using SQL functions to generate URLs that point to the data. These URLs can then be used to download the data directly to a client

NEW QUESTION 500
- (Topic 4)
What is the minimum Snowflake Edition that supports secure storage of Protected Health Information (PHI) data?

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

Answer: C

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

NEW QUESTION 505
- (Topic 4)
How can a Snowflake user traverse semi-structured data?

A. Insert a colon (:) between the VARIANT column name and any first-level element.
B. Insert a colon (:) between the VARIANT column name and any second-level elemen
C. Insert a double colon (: :) between the VARIANT column name and any first-level element.
D. Insert a double colon (: :) between the VARIANT column name and any second-level element.

Answer: A

Explanation:
To traverse semi-structured data in Snowflake, a user can insert a colon (:) between the VARIANT column name and any first-level element. This path syntax is used to retrieve elements in a VARIANT column4.

NEW QUESTION 509
- (Topic 4)
Which object can be used with Secure Data Sharing?

A. View
B. Materialized view
C. External table
D. User-Defined Function (UDF)

Answer: A

Explanation:
Views can be used with Secure Data Sharing in Snowflake. Materialized views, external tables, and UDFs are not typically shared directly for security and performance reasons2.

NEW QUESTION 512
- (Topic 4)
What is the primary purpose of a directory table in Snowflake?

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

Answer: D

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

NEW QUESTION 513
- (Topic 4)
What feature of Snowflake Continuous Data Protection can be used for maintenance of historical data?

A. Access control
B. Fail-safe
C. Network policies
D. Time Travel

Answer: D

Explanation:
Snowflake??s Time Travel feature is used for the maintenance of historical data, allowing users to access and restore data that has been changed or deleted within a defined period4.

NEW QUESTION 518
- (Topic 4)
Which ACCOUNT_USAGE schema database role provides visibility into policy-related information?

A. USAGE_VIEWER
B. GOVERNANCE_VIEWER
C. OBJECT_VIEWER
D. SECURITY_VIEWER

Answer: B

Explanation:
The GOVERNANCE_VIEWER role in the ACCOUNT_USAGE schema provides visibility into policy-related information within Snowflake. This role is specifically designed to access views that display object metadata and usage metrics related to governance12.

NEW QUESTION 522
.....

Thank You for Trying Our Product

We offer two products:

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

2nd - Questons and Answers in PDF Format

COF-C02 Practice Exam Features:

* COF-C02 Questions and Answers Updated Frequently

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

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

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