Number: SnowPro-Core
Passing Score: 800
Time Limit: 120
File Version: 13.0

Exam Code: SnowPro Core
Exam Name: SnowPro-Core

Exam A

QUESTION 1
How can a user improve the performance of a single large complex query in Snowflake?

A. Scale up the virtual warehouse.
B. Scale out the virtual warehouse.
C. Enable standard warehouse scaling.
D. Enable economy warehouse scaling.

Correct Answer: A
Section:
Explanation:
Scaling up the virtual warehouse in Snowflake involves increasing the compute resources available for a single warehouse, which can improve the performance of large and complex queries by providing more CPU and memory resources.Reference:Based on general cloud data warehousing knowledge as of 2021.

QUESTION 2
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

Correct Answer: C, D
Section:
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.Reference:Based on general security practices in cloud services as of 2021.

QUESTION 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

Correct Answer: C
Section:
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.Reference:Based on general programming and cloud data service knowledge as of 2021.

QUESTION 4
If file format options are specified in multiple locations, the load operation selects which option FIRST to apply in order of precedence?

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

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

QUESTION 5
Which statistics can be used to identify queries that have inefficient pruning? (Select TWO).

A. Bytes scanned
B. Bytes written to result
C. Partitions scanned
D. Partitions total
E. Percentage scanned from cache

Correct Answer: C, D
Section:
Explanation:
The statistics that can be used to identify queries with inefficient pruning are 'Partitions scanned' and 'Partitions total'.These statistics indicate how much of the data was actually needed and scanned versus the total available, which can highlight inefficiencies in data pruning34.

QUESTION 6
Which ACCOUNT_USAGE schema database role provides visibility into policy-related information?

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

Correct Answer: B
Section:
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.

QUESTION 7
Which Snowflake data types can be used to build nested hierarchical data? (Select TWO)

A. INTEGER

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

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

QUESTION 8
What is the MAXIMUM size limit for a record of a VARIANT data type?

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

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

QUESTION 9
What is the name of the SnowSQLfile that can store connection information?

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

Correct Answer: B
Section:
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.Reference:Based on general database knowledge as of 2021.

QUESTION 10
A Snowflake user has been granted the create data EXCHANGE listing privilege with their role.
Which tasks can this user now perform on the Data Exchange? (Select TWO).

A. Rename listings.
B. Delete provider profiles.
C. Modify listings properties.
D. Modify incoming listing access requests.

E. Submit listings for approval/publishing.

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

QUESTION 11
What service is provided as an integrated Snowflake feature to enhance Multi-Factor Authentication (MFA) support?

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

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

QUESTION 12
Which user object property requires contacting Snowflake Support in order to set a value for it?

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

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

QUESTION 13
Which file format will keep floating-point numbers from being truncated when data is unloaded?

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

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

QUESTION 14
In which Snowflake layer does Snowflake reorganize data into its internal optimized, compressed, columnar format?

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

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

QUESTION 15
What can a Snowflake user do in the Admin area of Snowsight?

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

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

QUESTION 16
How can a Snowflake user optimize query performance in Snowflake? (Select TWO).

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

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

QUESTION 17
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

Correct Answer: D
Section:
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.
[COF-C02] SnowPro Core Certification Exam Study Guide
Snowflake Documentation on Masking Policies

QUESTION 18
Which of the following describes how clustering keys work in Snowflake?

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

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

QUESTION 19
What is a machine learning and data science partner within the Snowflake Partner Ecosystem?

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

Correct Answer: D
Section:
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.
[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

QUESTION 20
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

Correct Answer: D
Section:
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 theazure://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.Reference:[COF-C02] SnowPro Core Certification Exam Study Guide

QUESTION 21
Which data type can be used to store geospatial data in Snowflake?

A. Variant
B. Object
C. Geometry
D. Geography

Correct Answer: D
Section:
Explanation:
Snowflake supports two geospatial data types:GEOGRAPHYandGEOMETRY. TheGEOGRAPHYdata type is used to store geospatial data that models the Earth as a perfect sphere, which is suitable for global geospatial data. This data type follows the WGS 84 standard and is used for storing points, lines, and polygons on the Earth's surface. TheGEOMETRYdata type, on the other hand, represents features in a planar (Euclidean, Cartesian) coordinate system and is typically used for local spatial reference systems.Since the question specifically asks about geospatial data, which commonly refers to Earth-related spatial data, the correct answer isGEOGRAPHY3.Reference:[COF-C02] SnowPro Core Certification Exam Study Guide

QUESTION 22
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

Correct Answer: B, D
Section:
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.Reference:[COF-C02] SnowPro Core Certification Exam Study Guide

QUESTION 23
Which Snowflake partner specializes in data catalog solutions?

A. Alation

B. DataRobot
C. dbt
D. Tableau

Correct Answer: A
Section:
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.
[COF-C02] SnowPro Core Certification Exam Study Guide
Snowflake's official documentation and partner listings

QUESTION 24
What is the MOST performant file format for loading data in Snowflake?

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

Correct Answer: B
Section:
Explanation:
Parquet is a columnar storage file format that is optimized for performance in Snowflake. It is designed to be efficient for both storage and query performance, particularly for complex queries on large datasets. Parquet files support efficient compression and encoding schemes, which can lead to significant savings in storage and speed in query processing, making it the most performant file format for loading data into Snowflake.
[COF-C02] SnowPro Core Certification Exam Study Guide
Snowflake Documentation on Data Loading1

QUESTION 25
Which copy INTO command outputs the data into one file?

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

Correct Answer: B
Section:
Explanation:
TheCOPY INTOcommand in Snowflake can be configured to output data into a single file by setting theMAX_FILE_NUMBERoption to 1. This option limits the number of files generated by the command, ensuring that only one file is created regardless of the amount of data being exported.
[COF-C02] SnowPro Core Certification Exam Study Guide
Snowflake Documentation on Data Unloading

QUESTION 26
The fail-safe retention period is how many days?

A. 1 day

B. 7 days
C. 45 days
D. 90 days

Correct Answer: B
Section:
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.

QUESTION 27
True or False: A 4X-Large Warehouse may, at times, take longer to provision than a X-Small Warehouse.

A. True
B. False

Correct Answer: A
Section:
Explanation:
Provisioning time can vary based on the size of the warehouse. A4X-Large Warehousetypically has more resources and may take longer to provision compared to aX-Small Warehouse, which has fewer resources and can generally be provisioned more quickly.Reference:Understanding and viewing Fail-safe | Snowflake Documentation

QUESTION 28
How would you determine the size of the virtual warehouse used for a task?

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

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

QUESTION 29
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

Correct Answer: B, C, D
Section:
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 forTables,Databases, andInternal 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.

QUESTION 30
What is the default File Format used in the COPY command if one is not specified?

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

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

QUESTION 31
True or False: Reader Accounts are able to extract data from shared data objects for use outside of Snowflake.

A. True
B. False

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

QUESTION 32
True or False: Loading data into Snowflake requires that source data files be no larger than 16MB.

A. True
B. False

Correct Answer: B
Section:
Explanation:
Snowflake does not require source data files to be no larger than 16MB. In fact, Snowflake recommends that for optimal load performance, data files should be roughly100-250 MBin size when compressed. However, it is not recommended to load very large files (e.g., 100 GB or larger) due to potential delays and wasted credits if errors occur. Smaller files should be aggregated to minimize processing overhead, and larger files should be split to distribute the load among compute resources in an active warehouse.

QUESTION 33
True or False: A Virtual Warehouse can be resized while suspended.

A. True
B. False

Correct Answer: A
Section:
Explanation:
Virtual Warehouses in Snowflake can indeed be resized while they are suspended. Resizing a warehouse involves changing the number of compute resources (servers) allocated to it, which can be done to adjust performance and cost. When a warehouse is suspended, it is not currently running any queries, but its definition and metadata remain intact, allowing for modifications like resizing.

QUESTION 34
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

Correct Answer: B
Section:
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.

QUESTION 35
What are two ways to create and manage Data Shares in Snowflake? (Choose two.)

A. Via the Snowflake Web Interface (UI)
B. Via the data_share=true parameter
C. Via SQL commands
D. Via Virtual Warehouses

Correct Answer: A, C
Section:
Explanation:
In Snowflake, Data Shares can be created and managed in two primary ways:
Via the Snowflake Web Interface (UI): Users can create and manage shares through the graphical interface provided by Snowflake, which allows for a user-friendly experience.
Via SQL commands: Snowflake also allows the creation and management of shares using SQL commands. This method is more suited for users who prefer scripting or need to automate the process.

QUESTION 36
True or False: Fail-safe can be disabled within a Snowflake account.

A. True
B. False

Correct Answer: B
Section:
Explanation:
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.

QUESTION 37
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

Correct Answer: A
Section:
Explanation:
Snowflake's architecture allows for the use of aquery result cachethat 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.

QUESTION 38
A virtual warehouse's auto-suspend and auto-resume settings apply to which of the following?

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

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

QUESTION 39
Which of the following Snowflake features provide continuous data protection automatically? (Select TWO).

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

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

QUESTION 40
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

Correct Answer: A, E
Section:
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.

QUESTION 41
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

Correct Answer: B, C
Section:
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.
https://docs.snowflake.com/en/user-guide/tables-clustering-micropartitions.html

QUESTION 42
What is the minimum Snowflake edition required to create a materialized view?

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

Correct Answer: B
Section:
Explanation:
Materialized views in Snowflake are a feature that allows for the pre-computation and storage of query results for faster query performance. This feature is available starting from the Enterprise Edition of Snowflake. It is not available in the Standard Edition, and while it is also available in higher editions like Business Critical and Virtual Private Snowflake, the Enterprise Edition is the minimum requirement.

Snowflake Documentation on CREATE MATERIALIZED VIEW1.
Snowflake Documentation on Working with Materialized Views
https://docs.snowflake.com/en/sql-reference/sql/create-materialized-view.html#:~:text=Materialized%20views%20require%20Enterprise%20Edition,upgrading%2C%20please%20contact%20Snowflake%20Support.

QUESTION 43
What happens to the underlying table data when a CLUSTER BY clause is added to a Snowflake table?

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

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

QUESTION 44
Which feature is only available in the Enterprise or higher editions of Snowflake?

A. Column-level security
B. SOC 2 type II certification
C. Multi-factor Authentication (MFA)
D. Object-level access control

Correct Answer: A
Section:
Explanation:
Column-level security is a feature that allows fine-grained control over access to specific columns within a table. This is particularly useful for managing sensitive data and ensuring that only authorized users can view or manipulate certain pieces of information. According to my last update, this feature was available in the Enterprise Edition or higher editions of Snowflake.
https://docs.snowflake.com/en/user-guide/intro-editions.html

QUESTION 45
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

Correct Answer: B, D, E

Section:
Explanation:
Snowflake supports several methods for authenticating users, includingfederated authentication,key-pair authentication, andOAuth. 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.Reference:Authentication policies | Snowflake Documentation,Authenticating to the server | Snowflake Documentation,External API authentication and secrets | Snowflake Documentation.

QUESTION 46
During periods of warehouse contention which parameter controls the maximum length of time a warehouse will hold a query for processing?

A. STATEMENT_TIMEOUT__IN__SECONDS
B. STATEMENT_QUEUED_TIMEOUT_IN_SECONDS
C. MAX_CONCURRENCY__LEVEL
D. QUERY_TIMEOUT_IN_SECONDS

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

QUESTION 47
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

Correct Answer: D, E
Section:
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 can lead 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.
Snowflake Documentation on Clustering Keys & Clustered Tables
Snowflake Documentation on SYSTEM$CLUSTERING_INFORMATION
Stack Overflow discussion on cluster key selection in Snowflake

QUESTION 48
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

Correct Answer: A
Section:
Explanation:
In Snowflake, aPipeis 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'sCOPYcommand to load data and can be associated with aStageobject to monitor for new files. When new data files appear in the stage, the pipe automatically loads the data into the target table.
Snowflake Documentation on Pipes
SnowPro Core Certification Study Guide
https://docs.snowflake.com/en/user-guide/data-load-snowpipe-intro.html

QUESTION 49
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; CREATE MATERIALIZED VIEW ON SCHEMA MYDB.MYSCHEMA TO ROLE MYROLE;
B. GRANT ROLE MYROLE TO USER USER1; CREATE MATERIALIZED VIEW ON SCHEMA MYDB.MYSCHEMA TO USER USER1;
C. GRANT ROLE MYROLE TO USER USER1; CREATE MATERIALIZED VIEW ON SCHEMA MYDB.MYSCHEMA TO USER1;
D. GRANT ROLE MYROLE TO USER USER1; CREATE MATERIALIZED VIEW ON SCHEMA MYDB.MYSCHEMA TO MYROLE;

Correct Answer: D
Section:
Explanation:
In Snowflake, to create a materialized view, the user must have the necessary privileges on the schema where the view will be created. These privileges are granted through roles, not directly to individual users. Therefore, the correct process is to grant the role to the user and then grant the privilege to create the materialized view to the role itself.
The statementGRANT ROLE MYROLE TO USER USER1;grants the specified role to the user, allowing them to assume that role and exercise its privileges. The subsequent statementCREATE MATERIALIZED VIEW ON SCHEMA MYDB.MYSCHEMA TO MYROLE;grants the privilege to create a materialized view within the specified schema to the roleMYROLE. Any user who has been grantedMYROLEcan then create materialized views inMYDB.MYSCHEMA.
Snowflake Documentation on Roles
Snowflake Documentation on Materialized Views

QUESTION 50
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

Correct Answer: A
Section:
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, seeSupported Character Sets for Delimited Files(in this topic).

QUESTION 51

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

Correct Answer: D
Section:
Explanation:
Snowflake is designed to optimize query performance by utilizing metadata for certain types of queries. When executing aCOUNTquery, 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 likeFCT_SALESwith 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 aCOUNTquery without the need to perform a full scan.
Snowflake Documentation on Metadata Management
SnowPro Core Certification Study Guide

QUESTION 52
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

Correct Answer: B
Section:
Explanation:
TheResult cacheis 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.
Snowflake Documentation on Query Results Cache
SnowPro Core Certification Study Guide

QUESTION 53
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

Correct Answer: C
Section:
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.
Snowflake Documentation on Virtual Warehouses
SnowPro Core Certification Study Guide

QUESTION 54
What is a limitation of a Materialized View?

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

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

QUESTION 55
What features does Snowflake Time Travel enable?

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

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

QUESTION 56
Which statement about billing applies to Snowflake credits?

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

D. Credits are consumed based on the warehouse size and the time the warehouse is running

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

QUESTION 57
What Snowflake features allow virtual warehouses to handle high concurrency workloads? (Select TWO)

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

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

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

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

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

Snowflake Documentation on Multi-Cluster Warehousing
SnowPro Core Certification best practices

QUESTION 59
Which of the following objects can be shared through secure data sharing?

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

Correct Answer: D
Section:
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.
Snowflake Documentation on Secure Data Sharing
SnowPro Core Certification Companion: Hands-on Preparation and Practice

QUESTION 60
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

Correct Answer: A
Section:
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. TheCREATE SHAREcommand 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.
Snowflake Documentation on Reader Accounts
SnowPro Core Certification Study Guide

QUESTION 61
A user unloaded a Snowflake table called mytable to an internal stage called mystage.
Which command can be used to view the list of files that has been uploaded to the staged?

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

Correct Answer: D

Section:
Explanation:
The commandlist @mystage;is used to view the list of files that have been uploaded to an internal stage in Snowflake. Thelistcommand displays the metadata for all files in the specified stage, which in this case ismystage. This command is particularly useful for verifying that files have been successfully unloaded from a Snowflake table to the stage and for managing the files within the stage.
Snowflake Documentation on Stages
SnowPro Core Certification Study Guide

QUESTION 62
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

Correct Answer: B, D
Section:
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.
Snowflake Documentation on Security Features
SnowPro Core Certification Exam Study Guide

QUESTION 63
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

Correct Answer: A
Section:
Explanation:
TheCOPY INTOcommand 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.
Snowflake Documentation on Unloading Data
Snowflake SnowPro Core: Copy Into Command to Unload Rows to Files in Named Stage

QUESTION 64
How often are encryption keys automatically rotated by Snowflake?

A. 30 Days

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

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

QUESTION 65
What are value types that a VARIANT column can store? (Select TWO)

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

Correct Answer: B, D
Section:
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.
Snowflake Documentation on Semi-Structured Data Types
SnowPro Core Certification Study Guide

QUESTION 66
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

Correct Answer: D
Section:
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.
Snowflake Documentation on Snowpipe
SnowPro Core Certification Study Guide

QUESTION 67
Which of the following are best practice recommendations that should be considered when loading data into Snowflake? (Select TWO).

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

Correct Answer: C, D
Section:
Explanation:
When loading data into Snowflake, it is recommended to:
C) Load files that are approximately 100-250 MB (or larger): This size is optimal for parallel processing and can help to maximize throughput. Smaller files can lead to overhead that outweighs the actual data processing time.
D) Avoid using embedded characters such as commas for numeric data types: Embedded characters can cause issues during data loading as they may be interpreted incorrectly. It's best to clean the data of such characters to ensure accurate and efficient data loading.
These best practices are designed to optimize the data loading process, ensuring that data is loaded quickly and accurately into Snowflake.
Snowflake Documentation on Data Loading Considerations
[COF-C02] SnowPro Core Certification Exam Study Guide

QUESTION 68
A user has 10 files in a stage containing new customer data. The ingest operation completes with no errors, using the following command:
COPY INTO my__table FROM @my__stage;
The next day the user adds 10 files to the stage so that now the stage contains a mixture of new customer data and updates to the previous data. The user did not remove the 10 original files.
If the user runs the same copy into command what will happen?

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

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

QUESTION 69
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;

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

QUESTION 70
Which of the following can be executed/called with Snowpipe?

A. A User Defined Function (UDF)
B. A stored procedure
C. A single copy_into statement
D. A single insert__into statement

Correct Answer: C
Section:
Explanation:
Snowpipe is used for continuous, automated data loading into Snowflake. It uses aCOPY INTO <table>statement within a pipe object to load data from files as soon as they are available in a stage. Snowpipe does not execute UDFs, stored procedures, or insert statements.
Reference:Snowpipe | Snowflake Documentation

QUESTION 71
What Snowflake role must be granted for a user to create and manage accounts?

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

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

QUESTION 72
When unloading to a stage, which of the following is a recommended practice or approach?

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

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

QUESTION 73
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

Correct Answer: C, E
Section:
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.

QUESTION 74
Which of the following describes external functions in Snowflake?

A. They are a type of User-defined Function (UDF).
B. They contain their own SQL code.
C. They call code that is stored inside of Snowflake.
D. They can return multiple rows for each row received

Correct Answer: A
Section:
Explanation:
External functions in Snowflake are a special type of User-Defined Function (UDF) that call code executed outside of Snowflake, typically through a remote service.Unlike traditional UDFs, external functions do not contain SQL code within Snowflake; instead, they interact with external services to process data2.
https://docs.snowflake.com/en/sql-reference/external-functions.html#:~:text=External%20functions%20are%20user%2Ddefined,code%20running%20outside%20of%20Snowflake.

QUESTION 75
What are ways to create and manage data shares in Snowflake? (Select TWO)

A. Through the Snowflake web interface (UI)
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

Correct Answer: A, C
Section:
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.

QUESTION 76
A company's security audit requires generating a report listing all Snowflake logins (e.g.. date and user) within the last 90 days. Which of the following statements will return the required information?

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

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

QUESTION 77
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

Correct Answer: D, E
Section:
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.)

QUESTION 78
What is the purpose of an External Function?

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

Correct Answer: A
Section:
Explanation:
The purpose of an External Function in Snowflake is to call code that executes outside of the Snowflake environment.This allows Snowflake to interact with external services and leverage functionalities that are not natively available within Snowflake, such as calling APIs or running custom code hosted on cloud services3.
https://docs.snowflake.com/en/sql-reference/external-functions.html
Topic 2, Exam pool B

QUESTION 79
A user created a new worksheet within the Snowsight UI and wants to share this with teammates

How can this worksheet be shared?

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

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

QUESTION 80
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

Correct Answer: C
Section:
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.

QUESTION 81
Which statements are true concerning Snowflake's underlying cloud infrastructure? (Select THREE),

A. Snowflake data and services are deployed in a single availability zone within a cloud provider's region.
B. Snowflake data and services are available in a single cloud provider and a single region, the use of multiple cloud providers is not supported.
C. Snowflake can be deployed in a customer's private cloud using the customer's own compute and storage resources for Snowflake compute and storage
D. Snowflake uses the core compute and storage services of each cloud provider for its own compute and storage
E. All three layers of Snowflake's architecture (storage, compute, and cloud services) are deployed and managed entirely on a selected cloud platform
F. Snowflake data and services are deployed in at least three availability zones within a cloud provider's region

Correct Answer: D, E, F
Section:
Explanation:
Snowflake's architecture is designed to operate entirely on cloud infrastructure. It uses the core compute and storage services of each cloud provider, which allows it to leverage the scalability and reliability of cloud resources. Snowflake's services are deployed across multiple availability zones within a cloud provider's region to ensure high availability and fault tolerance.
Reference: [COF-C02] SnowPro Core Certification Exam Study Guide

QUESTION 82
Which snowflake objects will incur both storage and cloud compute charges? (Select TWO)

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

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

QUESTION 83
A user is preparing to load data from an external stage
Which practice will provide the MOST efficient loading performance?

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

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

QUESTION 84
If 3 size Small virtual warehouse is made up of two servers, how many servers make up a Large warehouse?

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

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

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

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

QUESTION 85
Which command sets the Virtual Warehouse for a session?

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

Correct Answer: C
Section:
Explanation:
The commandUSE WAREHOUSE <<warehouse name>>;is used to set the virtual warehouse for the current session in Snowflake.This command specifies which virtual warehouse to use for executing queries in that session1.

QUESTION 86
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.

Correct Answer: A
Section:
Explanation:
When a pipe is recreated using theCREATE OR REPLACE PIPEcommand, 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.

QUESTION 87
True or False: Snowpipe via REST API can only reference External Stages as source.

A. True
B. False

Correct Answer: B
Section:
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.
Reference = [COF-C02] SnowPro Core Certification Exam Study Guide, Snowflake Documentation1

QUESTION 88
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.

Correct Answer: A, C, D
Section:
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.

QUESTION 89
What do the terms scale up and scale out refer to in Snowflake? (Choose two.)

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

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

QUESTION 90
What is the minimum Snowflake edition that has column-level security enabled?

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

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

QUESTION 91
When cloning a database, what is cloned with the database? (Choose two.)

A. Privileges on the database
B. Existing child objects within the database
C. Future child objects within the database
D. Privileges on the schemas within the database
E. Only schemas and tables within the database

Correct Answer: A, B
Section:
Explanation:
When cloning a database in Snowflake, the clone includes all privileges on the database as well as existing child objects within the database, such as schemas, tables, views, etc.However, it does not include future child objects or privileges on schemas within the database2.
Reference = [COF-C02] SnowPro Core Certification Exam Study Guide, Snowflake Documentation

QUESTION 92
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

Correct Answer: A
Section:
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.
Reference = [COF-C02] SnowPro Core Certification Exam Study Guide, Snowflake Documentation1

QUESTION 93
What is the maximum total Continuous Data Protection (CDP) charges incurred for a temporary table?

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

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

QUESTION 94
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

Correct Answer: C
Section:
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.
Reference = [COF-C02] SnowPro Core Certification Exam Study Guide, Snowflake Documentation

QUESTION 95
Which of the following are characteristics of Snowflake virtual warehouses? (Choose two.)

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

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

QUESTION 96
Which command should be used to load data from a file, located in an external stage, into a table in Snowflake?

A. INSERT
B. PUT

C. GET
D. COPY

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

QUESTION 97
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

Correct Answer: C
Section:
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.
Reference = [COF-C02] SnowPro Core Certification Exam Study Guide, Snowflake Documentation

QUESTION 98
Which of the following is a data tokenization integration partner?

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

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

QUESTION 99
What versions of Snowflake should be used to manage compliance with Personal Identifiable Information (PII) requirements? (Choose two.)

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

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

QUESTION 100
What are supported file formats for unloading data from Snowflake? (Choose three.)

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

Correct Answer: B, C, F
Section:
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

QUESTION 101
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

Correct Answer: B, C
Section:
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.

QUESTION 102
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.

Correct Answer: B, C
Section:

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.

QUESTION 103
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

Correct Answer: B
Section:
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.

QUESTION 104
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)

Correct Answer: A
Section:
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.

QUESTION 105
What is the SNOWFLAKE.ACCOUNT_USAGE view that contains information about which objects were read by queries within the last 365 days (1 year)?

A. VIEWS_HISTORY
B. OBJECT_HISTORY
C. ACCESS_HISTORY
D. LOGIN_HISTORY

Correct Answer: C
Section:
Explanation:
TheACCESS_HISTORYview in theSNOWFLAKE.ACCOUNT_USAGEschema contains information about the access history of Snowflake objects, such as tables and views, within the last 365 days1.

QUESTION 106
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

Correct Answer: B
Section:
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.

QUESTION 107
What are the responsibilities of Snowflake's Cloud Service layer? (Choose three.)

A. Authentication
B. Resource management
C. Virtual warehouse caching
D. Query parsing and optimization
E. Query execution
F. Physical storage of micro-partitions

Correct Answer: A, B, D
Section:
Explanation:
The responsibilities of Snowflake's Cloud Service layer include authentication (A), which ensures secure access to the platform; resource management (B), which involves allocating and managing compute resources; and query parsing and optimization (D), which improves the efficiency and performance of SQL query execution3.

QUESTION 108
How long is the Fail-safe period for temporary and transient tables?

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

Correct Answer: A
Section:
Explanation:
Temporary and transient tables in Snowflake do not have a Fail-safe period.Once the session ends or the tables are dropped, the data is purged and not recoverable1.

QUESTION 109
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

Correct Answer: B
Section:
Explanation:
The GET command is used to download files from a Snowflake stage to a local folder on a client's machine2.

QUESTION 110
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.

Correct Answer: C
Section:
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.

QUESTION 111
A virtual warehouse is created using the following command:
Create warehouse my_WH with
warehouse_size = MEDIUM
min_cluster_count = 1
max_cluster_count = 1
auto_suspend = 60
auto_resume = true;
The image below is a graphical representation of the warehouse utilization across two days.

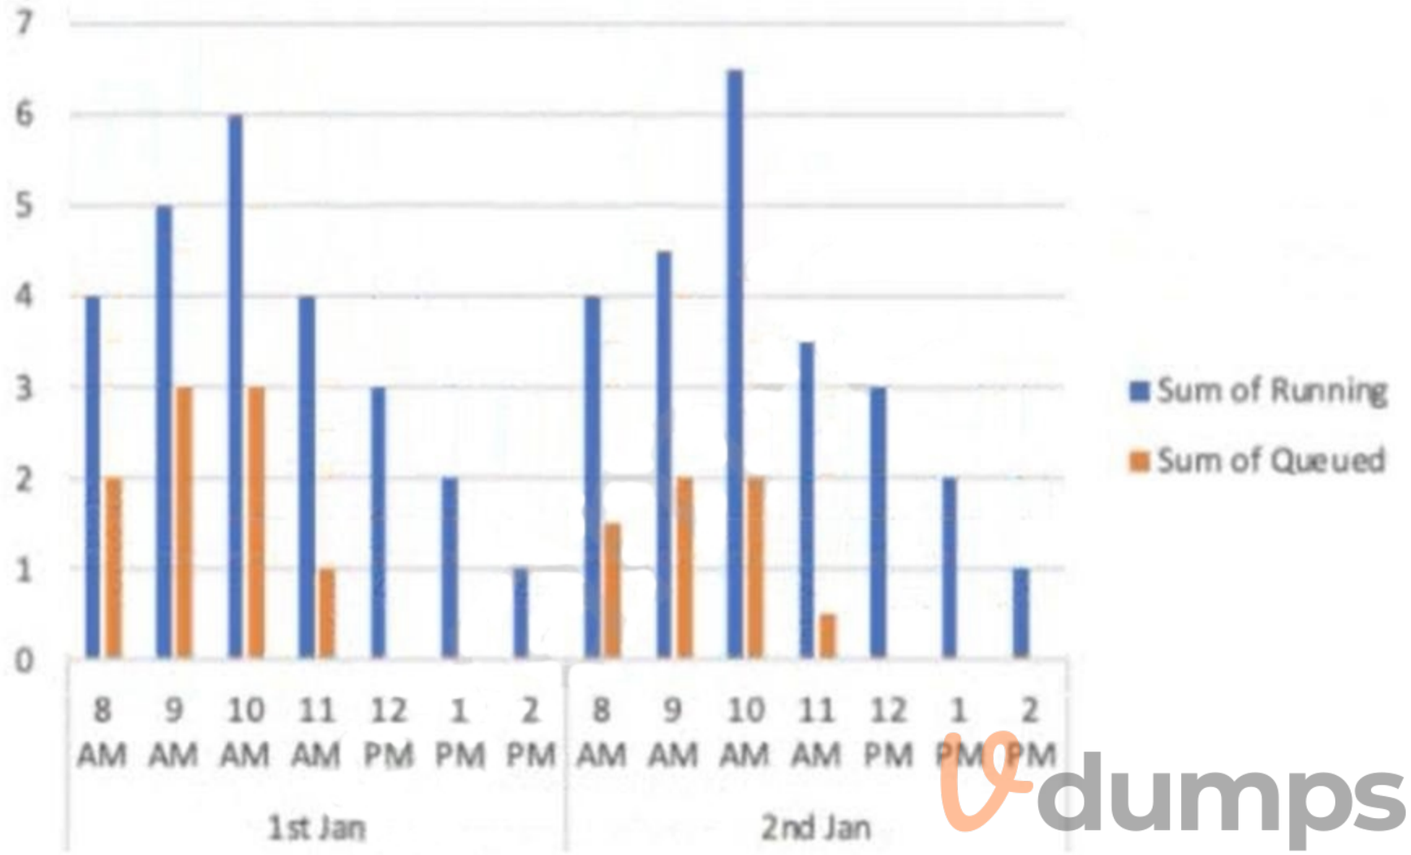

What action should be taken to address this situation?

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

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

QUESTION 112
Which minimum Snowflake edition allows for a dedicated metadata store?

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

Correct Answer: B
Section:
Explanation:
The Enterprise edition of Snowflake allows for a dedicated metadata store, providing additional features designed for large-scale enterprises

QUESTION 113
Network policies can be set at which Snowflake levels? (Choose two.)

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

Correct Answer: C, E
Section:
Explanation:
Network policies in Snowflake can be set at the user level and at the account level2.

QUESTION 114
What are the correct parameters for time travel and fail-safe in the Snowflake Enterprise Edition?

A. Default Time Travel Retention is set to 0 days. Maximum Time Travel Retention is 30 days. Fail Safe retention time is 1 day.
B. Default Time Travel Retention is set to 1 day. Maximum Time Travel Retention is 365 days. Fail Safe retention time is 7 days.
C. Default Time Travel Retention is set to 0 days. Maximum Time Travel Retention is 90 days. Fail Safe retention time is 7 days.
D. Default Time Travel Retention is set to 1 day. Maximum Time Travel Retention is 90 days. Fail Safe retention time is 7 days.
E. Default Time Travel Retention is set to 7 days. Maximum Time Travel Retention is 1 day. Fail Safe retention time is 90 days.
F. Default Time Travel Retention is set to 90 days. Maximum Time Travel Retention is 7 days. Fail Safe retention time is 356 days.

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

QUESTION 115
Which of the following objects are contained within a schema? (Choose two.)

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

Correct Answer: B, D
Section:

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

QUESTION 116
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.

Correct Answer: B, E
Section:
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 called USE_METADATA_CACHE in Snowflake.
Reference:Caching in the Snowflake Cloud Data Platform,Optimizing the warehouse cache

QUESTION 117
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.

Correct Answer: B
Section:
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.
Reference:COPY INTO Error during Structured Data Load: ''Max LOB size (16777216) exceeded...''

QUESTION 118
Which of the following features, associated with Continuous Data Protection (CDP), require additional Snowflake-provided data storage? (Choose two.)

A. Tri-Secret Secure
B. Time Travel
C. Fail-safe

D. Data encryption
E. External stages

Correct Answer: B, C
Section:
Explanation:
The features associated with Continuous Data Protection (CDP) that require additional Snowflake-provided data storage are Time Travel and Fail-safe. Time Travel allows users to access historical data within a defined period, while Fail-safe provides an additional layer of data protection beyond the Time Travel period.Reference:[COF-C02] SnowPro Core Certification Exam Study Guide

QUESTION 119
Where can a user find and review the failed logins of a specific user for the past 30 days?

A. The USERS view in ACCOUNT_USAGE
B. The LOGIN_HISTORY view in ACCOUNT_USAGE
C. The ACCESS_HISTORY view in ACCOUNT_USAGE
D. The SESSIONS view in ACCOUNT_USAGE

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

QUESTION 120
Which of the following statements apply to Snowflake in terms of security? (Choose two.)

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

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

QUESTION 121
Which Snowflake object can be accessed in he FROM clause of a query, returning a set of rows having one or more columns?

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

Correct Answer: A

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

QUESTION 122
Which type of join will list a I rows in the specified table, even if those rows have no match in the other table?

A. Cross join
B. Inner join
C. Natural join
D. Outer join

Correct Answer: D
Section:
Explanation:
An outer join, specifically a left outer join, will list all rows from the left table and match them with rows from the right table. If there is no match, the result will still include the row from the left table, with NULLs for columns from the right table.Reference:Based on general SQL knowledge as of 2021.

QUESTION 123
Which operations are handled in the Cloud Services layer of Snowflake? (Select TWO).

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

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

QUESTION 124
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.

Correct Answer: A
Section:
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.Reference:Based on general data sharing practices in cloud services as of 2021.

QUESTION 125

What can a Snowflake user do with the information included in the details section of a Query Profile?

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

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

QUESTION 126
How many network policies can be assigned to an account or specific user at a time?

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

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

QUESTION 127
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

Correct Answer: B
Section:
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.Reference:Based on my internal knowledge as of 2021.

QUESTION 128
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;

Correct Answer: A
Section:
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.Reference:Based on my internal knowledge as of 2021.

QUESTION 129
What technique does Snowflake use to limit the number of micro-partitions scanned by each query?

A. B-tree
B. Indexing
C. Map reduce
D. Pruning

Correct Answer: D
Section:
Explanation:
Snowflake uses a technique called pruning to limit the number of micro-partitions scanned by each query.Pruning effectively filters out unnecessary micro-partitions based on the query's filter conditions, which can significantly improve query performance by reducing the amount of data scanned1.

QUESTION 130
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

Correct Answer: B
Section:
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.

QUESTION 131
Which feature is integrated to support Multi-Factor Authentication (MFA) at Snowflake?

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

Correct Answer: B
Section:
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.

QUESTION 132
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

Correct Answer: A
Section:
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 allows Snowflake 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

QUESTION 133
Which Snowflake object helps evaluate virtual warehouse performance impacted by query queuing?

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

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

QUESTION 134
A Snowflake user has two tables that contain numeric values and is trying to find out which values are present in both tables. Which set operator should be used?

A. INTERSECT
B. MFRCK
C. MINUS
D. UNION

Correct Answer: A
Section:
Explanation:
To find out which numeric values are present in both tables, the INTERSECT set operator should be used.This operator returns rows from one query's result set which also appear in another query's result set, effectively finding the common elements between the two tables45.

QUESTION 135
What type of columns does Snowflake recommend to be used as clustering keys? (Select TWO).

A. A VARIANT column

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

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

QUESTION 136
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

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

QUESTION 137
What does Snowflake recommend regarding database object ownership? (Select TWO).

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

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

QUESTION 138
If a multi-cluster warehouse is using an economy scaling policy, how long will queries wait in the queue before another cluster is started?

A. 1 minute

B. 2 minutes
C. 6 minutes
D. 8 minutes

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

QUESTION 139
How can a Snowflake user access a JSON object, given the following table? (Select TWO).

SRC
{ "salesPerson": { "id": "1234", "name": "user" } }

A. src:salesperson.name
B. src:sa1esPerson. name
C. src:salesperson.Name
D. SRC:salesperson.name
E. SRC:salesperson.Name

Correct Answer: A, C
Section:
Explanation:
To access a JSON object in Snowflake, dot notation is used where the path to the object is specified after the column name containing the JSON data. Both lowercase and uppercase can be used for attribute names, so both ''name'' and ''Name'' are valid.
Reference: [COF-C02] SnowPro Core Certification Exam Study Guide

QUESTION 140
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.

Correct Answer: A
Section:
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.
Reference: [COF-C02] SnowPro Core Certification Exam Study Guide

QUESTION 141
Which native data types are used for storing semi-structured data in Snowflake? (Select TWO)

A. NUMBER
B. OBJECT
C. STRING
D. VARCHAR
E. VARIANT

Correct Answer: B, E
Section:
Explanation:
Snowflake supports semi-structured data types, which include OBJECT and VARIANT. These data types are capable of storing JSON-like data structures, allowing for flexibility in data representation.OBJECT can directly contain VARIANT, and thus indirectly contain any other data type, including itself1.

QUESTION 142
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

Correct Answer: A, C
Section:
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.

QUESTION 143
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.

Correct Answer: D
Section:
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.

QUESTION 144
What role is required to use Partner Connect?

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

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

QUESTION 145
How does a scoped URL expire?

A. When the data cache clears.
B. When the persisted query result period ends.
C. The encoded URL access is permanent.
D. The length of time is specified in the expiration_time argument.

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

QUESTION 146
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

Correct Answer: B, C
Section:
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.

QUESTION 147
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.

Correct Answer: A
Section:
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.

QUESTION 148
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

Correct Answer: B, C
Section:
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

QUESTION 149
How does Snowflake handle the bulk unloading of data into single or multiple files?

A. It assigns each unloaded data file a unique name.
B. It uses the put command to download the data by default.
C. It uses COPY INTO <location> for bulk unloading where the default option is SINGLE - TRUE.
D. It uses COPY INTO <location> to copy the data from a table into one or more files in an external stage only.

Correct Answer: A
Section:
Explanation:
When unloading data, Snowflake assigns each file a unique name to ensure there is no overlap or confusion between files.This is part of the bulk unloading process where data is exported from Snowflake tables into flat files3.

QUESTION 150
How can a dropped internal stage be restored?

A. Enable Time Travel.
B. Clone the dropped stage.
C. Execute the UNDROP command.
D. Recreate the dropped stage.

Correct Answer: D
Section:
Explanation:
Once an internal stage is dropped in Snowflake, it cannot be recovered or restored using Time Travel or UNDROP commands.The only option is to recreate the dropped stage

QUESTION 151
Who can grant object privileges in a regular schema?

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

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

QUESTION 152
Which command is used to unload data from a Snowflake database table into one or more files in a Snowflake stage?

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

Correct Answer: C
Section:
Explanation:
TheCOPY INTO <location>command is used to unload data from a Snowflake database table into one or more files in a Snowflake stage1.

QUESTION 153
Which metadata table will store the storage utilization information even for dropped tables?

A. DATABASE_STORAGE_USAGE_HISTORY
B. TABLE_STORAGE_METRICS
C. STORAGE_DAILY_HISTORY
D. STAGE STORAGE USAGE HISTORY

Correct Answer: B
Section:
Explanation:
TheTABLE_STORAGE_METRICSmetadata table stores the storage utilization information, including for tables that have been dropped but are still incurring storage costs2.

QUESTION 154
What does a masking policy consist of in Snowflake?

A. A single data type, with one or more conditions, and one or more masking functions
B. A single data type, with only one condition, and only one masking function
C. Multiple data types, with only one condition, and one or more masking functions
D. Multiple data types, with one or more conditions, and one or more masking functions

Correct Answer: A

Section:
Explanation:
A masking policy in Snowflake consists of a single data type, with one or more conditions, and one or more masking functions.These components define how the data is masked based on the specified conditions3.

QUESTION 155
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

Correct Answer: D
Section:
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.

QUESTION 156
A JSON file, that contains lots of dates and arrays, needs to be processed in Snowflake. The user wants to ensure optimal performance while querying the data.
How can this be achieved?

A. Flatten the data and store it in structured data types in a flattened table. Query the table.
B. Store the data in a table with a variant data type. Query the table.
C. Store the data in a table with a vai : ant data type and include STRIP_NULL_VALUES while loading the table. Query the table.
D. Store the data in an external stage and create views on top of it. Query the views.

Correct Answer: B
Section:
Explanation:
Storing JSON data in a table with a VARIANT data type is optimal for querying because it allows Snowflake to leverage its semi-structured data capabilities.This approach enables efficient storage and querying without the need for flattening the data, which can be performance-intensive1.

QUESTION 157
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

Correct Answer: A
Section:
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.

QUESTION 158
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

Correct Answer: A, B
Section:
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.

QUESTION 159
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

Correct Answer: C
Section:
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.

QUESTION 160
At what levels can a resource monitor be configured? (Select TWO).

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

Correct Answer: A, E
Section:
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.Reference:[COF-C02] SnowPro Core Certification Exam Study Guide

QUESTION 161
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)

Correct Answer: B
Section:
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.Reference:[COF-C02] SnowPro Core Certification Exam Study Guide

QUESTION 162
For which use cases is running a virtual warehouse required? (Select TWO).

A. When creating a table
B. When loading data into a table
C. When unloading data from a table
D. When executing a show command
E. When executing a list command

Correct Answer: B, C
Section:
Explanation:
Running a virtual warehouse is required when loading data into a table and when unloading data from a table because these operations require compute resources that are provided by the virtual warehouse23.

QUESTION 163
What does SnowCD help Snowflake users to do?

A. Copy data into files.
B. Manage different databases and schemas.
C. Troubleshoot network connections to Snowflake.
D. Write SELECT queries to retrieve data from external tables.

Correct Answer: C
Section:
Explanation:
SnowCD is a connectivity diagnostic tool that helps users troubleshoot network connections to Snowflake.It performs a series of checks to evaluate the network connection and provides suggestions for resolving any issues4.

QUESTION 164
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

Correct Answer: D, E
Section:
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.

QUESTION 165
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 data. These URLs can be used to download the data to a client.
C. SQL functions can be used to retrieve the data from the query results cache. When the query results are output to a client, the unstructured data will be output to the client as files.
D. SQL functions can call on different web extensions designed to display different types of files as a web page. The web extensions will allow the files to be downloaded to the client.

Correct Answer: B
Section:
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

QUESTION 166
Which Snowflake view is used to support compliance auditing?

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

Correct Answer: A
Section:
Explanation:
TheACCESS_HISTORYview in Snowflake is utilized to support compliance auditing. It provides detailed information on data access within Snowflake, including reads and writes by user queries.This view is essential for regulatory compliance auditing as it offers insights into the usage of tables and columns, and maintains a direct link between the user, the query, and the accessed data1.

QUESTION 167
What is the purpose of the STRIP NULL_VALUES file format option when loading semi-structured data files into Snowflake?

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

Correct Answer: D
Section:
Explanation:
TheSTRIP NULL_VALUESfile format option, when set toTRUE, removes object or array elements that contain null values during the loading process of semi-structured data files into Snowflake.This ensures that the data loaded into Snowflake tables does not contain these null elements, which can be useful when the ''null'' values in files indicate missing values and have no other special meaning2.

QUESTION 168
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

Correct Answer: A, E
Section:
Explanation:
Snowflake supports the storage of semi-structured data using theARRAYandVARIANTdata types. TheARRAYdata type can directly containVARIANT, and thus indirectly contain any other data type, including itself.TheVARIANTdata type can store a value of any other type, includingOBJECTandARRAY, and is often used to represent semi-structured data formats like JSON, Avro, ORC, Parquet, or XML34.

QUESTION 169
A user wants to access files stored in a stage without authenticating into Snowflake. Which type of URL should be used?

A. File URL
B. Staged URL
C. Scoped URL
D. Pre-signed URL

Correct Answer: D
Section:
Explanation:
APre-signed URLshould be used to access files stored in a Snowflake stage without requiring authentication into Snowflake. Pre-signed URLs are simple HTTPS URLs that provide temporary access to a file via a web browser, using a pre-signed access token.The expiration time for the access token is configurable, and this type of URL allows users or applications to directly access or download the files without needing to authenticate into Snowflake5.

QUESTION 170
What does a Notify & Suspend action for a resource monitor do?

A. Send an alert notification to all account users who have notifications enabled.
B. Send an alert notification to all virtual warehouse users when thresholds over 100% have been met.
C. Send a notification to all account administrators who have notifications enabled, and suspend all assigned warehouses after all statements being executed by the warehouses have completed.
D. Send a notification to all account administrators who have notifications enabled, and suspend all assigned warehouses immediately, canceling any statements being executed by the warehouses.

Correct Answer: C
Section:
Explanation:
The Notify & Suspend action for a resource monitor in Snowflake sends a notification to all account administrators who have notifications enabled and suspends all assigned warehouses.However, the suspension only occurs after all currently running statements in the warehouses have been completed1.Reference:[COF-C02] SnowPro Core Certification Exam Study Guide

QUESTION 171
What type of query will benefit from the query acceleration service?

A. Queries without filters or aggregation
B. Queries with large scans and selective filters
C. Queries where the GROUP BY has high cardinality
D. Queries of tables that have search optimization service enabled

Correct Answer: B
Section:
Explanation:
The query acceleration service in Snowflake is designed to benefit queries that involve large scans and selective filters.This service can offload portions of the query processing work to shared compute resources, which can handle these types of workloads more efficiently by performing more work in parallel and reducing the wall-clock time spent in scanning and filtering2.Reference:[COF-C02] SnowPro Core Certification Exam Study Guide

QUESTION 172
What Snowflake feature provides a data hub for secure data collaboration, with a selected group of invited members?

A. Data Replication
B. Secure Data Sharing
C. Data Exchange
D. Snowflake Marketplace

Correct Answer: C
Section:
Explanation:
Snowflake's Data Exchange feature provides a data hub for secure data collaboration.It allows providers to publish data that can be discovered and accessed by a selected group of invited members, facilitating secure and controlled data sharing within a collaborative environment3.Reference:[COF-C02] SnowPro Core Certification Exam Study Guide

QUESTION 173
What is a characteristic of materialized views in Snowflake?

A. Materialized views do not allow joins.
B. Clones of materialized views can be created directly by the user.
C. Multiple tables can be joined in the underlying query of a materialized view.
D. Aggregate functions can be used as window functions in materialized views.

Correct Answer: C
Section:
Explanation:
One of the characteristics of materialized views in Snowflake is that they allow multiple tables to be joined in the underlying query.This enables the pre-computation of complex queries involving joins, which can significantly improve the performance of subsequent queries that access the materialized view4.Reference:[COF-C02] SnowPro Core Certification Exam Study Guide

QUESTION 174
Which command is used to start configuring Snowflake for Single Sign-On (SSO)?

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

Correct Answer: C
Section:
Explanation:
To start configuring Snowflake for Single Sign-On (SSO), theCREATE SECURITY INTEGRATIONcommand is used.This command sets up a security integration object in Snowflake, which is necessary for enabling SSO with external identity providers using SAML 2.01.

QUESTION 175
What happens to the objects in a reader account when the DROP MANAGED ACCOUNT command is executed?

A. The objects are dropped.
B. The objects enter the Fail-safe period.
C. The objects enter the Time Travel period.
D. The objects are immediately moved to the provider account.

Correct Answer: A
Section:
Explanation:
When theDROP MANAGED ACCOUNTcommand is executed in Snowflake, it removes the managed account, including all objects created within the account, and access to the account is immediately restricted2.

QUESTION 176
Which operation can be performed on Snowflake external tables?

A. INSERT
B. JOIN
C. RENAME
D. ALTER

Correct Answer: B
Section:
Explanation:
Snowflake external tables are read-only, which means data manipulation language (DML) operations like INSERT, RENAME, or ALTER cannot be performed on them.However, external tables can be used for query and join operations3.

QUESTION 177
How should clustering be used to optimize the performance of queries that run on a very large table?

A. Manually re-cluster the table regularly.
B. Choose one high cardinality column as the clustering key.
C. Use the column that is most-frequently used in query select clauses as the clustering key.
D. Assess the average table depth to identify how clustering is impacting the query.

Correct Answer: B
Section:
Explanation:
For optimizing the performance of queries that run on a very large table, it is recommended to choose one high cardinality column as the clustering key.This helps to co-locate similar rows in the same micro-partitions, improving scan efficiency in queries by skipping data that does not match filtering predicates4.

QUESTION 178
What step can reduce data spilling in Snowflake?

A. Using a larger virtual warehouse
B. Increasing the virtual warehouse maximum timeout limit
C. Increasing the amount of remote storage for the virtual warehouse

D. Using a common table expression (CTE) instead of a temporary table

Correct Answer: A
Section:
Explanation:
To reduce data spilling in Snowflake, using a larger virtual warehouse is effective because it provides more memory and local disk space, which can accommodate larger data operations and minimize the need to spill data to disk or remote storage1.
Reference: [COF-C02] SnowPro Core Certification Exam Study Guide

QUESTION 179
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

Correct Answer: C
Section:
Explanation:
TheRETURN_ERRORSvalue in theVALIDATION_MODEoption 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.
Reference: [COF-C02] SnowPro Core Certification Exam Study Guide

QUESTION 180
Which solution improves the performance of point lookup queries that return a small number of rows from large tables using highly selective filters?

A. Automatic clustering
B. Materialized views
C. Query acceleration service
D. Search optimization service

Correct Answer: D
Section:
Explanation:
The search optimization service improves the performance of point lookup queries on large tables by using selective filters to quickly return a small number of rows.It creates an optimized data structure that helps in pruning the micro-partitions that do not contain the queried values3.
Reference: [COF-C02] SnowPro Core Certification Exam Study Guide

QUESTION 181
Which Snowflake object does not consume any storage costs?

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

Correct Answer: C

Section:
Explanation:
Temporary tables do not consume any storage costs in Snowflake because they only exist for the duration of the session that created them and are automatically dropped when the session ends, thus incurring no long-term storage charges4.
Reference: [COF-C02] SnowPro Core Certification Exam Study Guide

QUESTION 182
Which function unloads data from a relational table to JSON?

A. TO_OBJECT
B. TO_JSON
C. TO_VARIANT
D. OBJECT CONSTRUCT

Correct Answer: B
Section:
Explanation:
TheTO_JSONfunction is used to convert a VARIANT value into a string containing the JSON representation of the value. This function is suitable for unloading data from a relational table to JSON format.
Reference: [COF-C02] SnowPro Core Certification Exam Study Guide

QUESTION 183
Which commands can only be executed using SnowSQL? (Select TWO).

A. COPY INTO
B. GET
C. LIST
D. PUT
E. REMOVE

Correct Answer: C, D
Section:
Explanation:
TheLISTandPUTcommands are specific to SnowSQL and cannot be executed in the web interface or other SQL clients.LISTis used to display the contents of a stage, andPUTis used to upload files to a stage.
Reference: [COF-C02] SnowPro Core Certification Exam Study Guide

QUESTION 184
What is the relationship between a Query Profile and a virtual warehouse?

A. A Query Profile can help users right-size virtual warehouses.
B. A Query Profile defines the hardware specifications of the virtual warehouse.
C. A Query Profile can help determine the number of virtual warehouses available.
D. A Query Profile automatically scales the virtual warehouse based on the query complexity.

Correct Answer: A
Section:
Explanation:
A Query Profile provides detailed execution information for a query, which can be used to analyze the performance and behavior of queries. This information can help users optimize and right-size their virtual warehouses for better efficiency.

Reference: [COF-C02] SnowPro Core Certification Exam Study Guide

QUESTION 185
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

Correct Answer: D
Section:
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.
Reference: [COF-C02] SnowPro Core Certification Exam Study Guide

QUESTION 186
Which statements describe benefits of Snowflake's separation of compute and storage? (Select TWO).

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

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

QUESTION 187
Which commands are restricted in owner's rights stored procedures? (Select TWO).

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

Correct Answer: A, E
Section:
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.

QUESTION 188

Which parameter can be set at the account level to set the minimum number of days for which Snowflake retains historical data in Time Travel?

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

Correct Answer: A
Section:
Explanation:
The parameterDATA_RETENTION_TIME_IN_DAYScan be set at the account level to define the minimum number of days Snowflake retains historical data for Time Travel1.

QUESTION 189
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;

Correct Answer: C
Section:
Explanation:
When a tag object is assigned to a table, using the statementCREATE TABLE <table_name> AS SELECT * FROM TABLE_Awill automatically assign the TABLE_A tag to the newly created table2.

QUESTION 190
What objects in Snowflake are supported by Dynamic Data Masking? (Select TWO).'

A. Views
B. Materialized views
C. Tables
D. External tables
E. Future grants

Correct Answer: A, C
Section:
Explanation:
Dynamic Data Masking in Snowflake supportstables and views.These objects can have masking policies applied to their columns to dynamically mask data at query time3.

QUESTION 191
Which Snowflake feature allows administrators to identify unused data that may be archived or deleted?

A. Access history
B. Data classification
C. Dynamic Data Masking
D. Object tagging

Correct Answer: A
Section:
Explanation:
The Access History feature in Snowflake allows administrators to track data access patterns and identify unused data. This information can be used to make decisions about archiving or deleting data to optimize storage and reduce costs.

QUESTION 192
Which privilege must be granted by one role to another role, and cannot be revoked?

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

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

QUESTION 193
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

Correct Answer: C
Section:
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.
Reference: [COF-C02] SnowPro Core Certification Exam Study Guide

QUESTION 194
What does the LATERAL modifier for the FLATTEN function do?

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

Correct Answer: C
Section:
Explanation:
The LATERAL modifier for the FLATTEN function allows joining information outside the object (such as other columns in the source table) with the flattened data, creating a lateral view that correlates with the preceding tables in the FROM clause2345.

Reference: [COF-C02] SnowPro Core Certification Exam Study Guide

QUESTION 195
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

Correct Answer: A
Section:
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.
Reference: [COF-C02] SnowPro Core Certification Exam Study Guide

QUESTION 196
What are key characteristics of virtual warehouses in Snowflake? (Select TWO).

A. Warehouses that are multi-cluster can have nodes of different sizes.
B. Warehouses can be started and stopped at any time.
C. Warehouses can be resized at any time, even while running.
D. Warehouses are billed on a per-minute usage basis.
E. Warehouses can only be used for querying and cannot be used for data loading.

Correct Answer: B, C
Section:
Explanation:
Virtual warehouses in Snowflake can be started and stopped at any time, providing flexibility in managing compute resources.They can also be resized at any time, even while running, to accommodate varying workloads910.
Reference: [COF-C02] SnowPro Core Certification Exam Study Guide

QUESTION 197
What happens when a Snowflake user changes the data retention period at the schema level?

A. All child objects will retain data for the new retention period.
B. All child objects that do not have an explicit retention period will automatically inherit the new retention period.
C. All child objects with an explicit retention period will be overridden with the new retention period.
D. All explicit child object retention periods will remain unchanged.

Correct Answer: B
Section:
Explanation:
When the data retention period is changed at the schema level, all child objects that do not have an explicit retention period set will inherit the new retention period from the schema4.

QUESTION 198
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

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

QUESTION 199
What is a directory table in Snowflake?

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

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

QUESTION 200
Which Snowflake command can be used to unload the result of a query to a single file?

A. Use COPY INTO <external stage> followed by a GET command to download the file.
B. Use COPY INTO <internal stage> followed by a put command to download the file.
C. Use COPY INTO <internal stage> with SINGLE = TRUE followed by a GET command to download the file.
D. Use COPY INTO <external stage> with SINGLE = TRUE followed by a PUT command to download the file.

Correct Answer: C
Section:
Explanation:
The Snowflake command to unload the result of a query to a single file isCOPY INTO <internal stage> with SINGLE = TRUEfollowed by aGETcommand to download the file.This command unloads the query result into a single file in the specified internal stage
Explanation:

QUESTION 201
When working with a managed access schema, who has the OWNERSHIP privilege of any tables added to the schema?

A. The database owner
B. The object owner
C. The schema owner
D. The Snowflake user's role

Correct Answer: C
Section:
Explanation:
In a managed access schema, the schema owner retains the OWNERSHIP privilege of any tables added to the schema.This means that while object owners have certain privileges over the objects they create, only the schema owner can manage privilege grants on these objects1.