Number: COF-C02
Passing Score: 800
Time Limit: 120
File Version: 12.0

Exam Code: COF-C02
Exam Name: SnowPro Core Certification

Exam A

QUESTION 1
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 2
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 3
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 4

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 5
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 6
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 7
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 8
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 9
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 10
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 11
What is true about sharing data in Snowflake? (Choose two.)

A. The Data Consumer pays for data storage as well as for data computing.
B. The shared data is copied into the Data Consumer account, so the Consumer can modify it without impacting the base data of the Provider.
C. A Snowflake account can both provide and consume shared data.
D. The Provider is charged for compute resources used by the Data Consumer to query the shared data.
E. The Data Consumer pays only for compute resources to query the shared data.

Correct Answer: C, E
Section:
Explanation:
In Snowflake's data sharing model, any full Snowflake account can both provide and consume shared data. Additionally, the data consumer pays only for the compute resources used to query the shared data.No actual data is copied or transferred between accounts, and shared data does not take up any storage in a consumer account, so the consumer does not pay for data storage1.
References =Introduction to Secure Data Sharing | Snowflake Documentation

QUESTION 12
The following JSON is stored in a VARIANT column called src of the CAR_SALES table:

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

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

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

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

QUESTION 13
Which of the following significantly improves the performance of selective point lookup queries on a table?

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

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

QUESTION 14
What are the default Time Travel and Fail-safe retention periods for transient tables?

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

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

QUESTION 15
What is a best practice after creating a custom role?

A. Create the custom role using the SYSADMIN role.
B. Assign the custom role to the SYSADMIN role

C. Assign the custom role to the PUBLIC role
D. Add__CUSTOM to all custom role names

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

QUESTION 16
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.
References:Understanding and viewing Fail-safe | Snowflake Documentation

QUESTION 17
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.References:Understanding and viewing Fail-safe | Snowflake Documentation

QUESTION 18
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.References: [COF-C02] SnowPro Core Certification Exam Study Guide

QUESTION 19
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.
References: The information is verified according to the SnowPro Core Certification Study Guide and Snowflake documentation

QUESTION 20
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 21
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 22
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.
References: Preparing your data files | Snowflake Documentation

QUESTION 23
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 24
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 25
What are two ways to create and manage Data Shares in Snowflake? (Choose two.)

A. Via the Snowflake Web Interface (Ul)
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 26
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 27
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 28
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.
References:
SnowPro Core Certification Exam Study Guide (as of 2021)
Snowflake documentation on virtual warehouses and their settings (as of 2021)

QUESTION 29

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.
References:
Continuous Data Protection | Snowflake Documentation1
Data Storage Considerations | Snowflake Documentation2
Snowflake SnowPro Core Certification Study Guide3
Snowflake Data Cloud Glossary
https://docs.snowflake.com/en/user-guide/data-availability.html

QUESTION 30
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 31
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.
References: [COF-C02] SnowPro Core Certification Exam Study Guide
https://docs.snowflake.com/en/user-guide/tables-clustering-micropartitions.html

QUESTION 32
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.
References:
Snowflake Documentation on CREATE MATERIALIZED VIEW1.
Snowflake Documentation on Working with Materialized Views
https://docs.snowflake.com/en/sql-reference/sql/create-materialized-view.html#:~:text=Materialized%20views%20require%20Enterprise%20Edition,upgrading%2C%20please%20contact%20Snowflake%20Support.

QUESTION 33
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.
References:
Snowflake Documentation on Clustering Keys & Clustered Tables1.
Community discussions on how source data's ordering affects a table with a cluster key

QUESTION 34
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.
References: Based on my internal data as of 2021, column-level security is an advanced feature typically reserved for higher-tiered editions like the Enterprise Edition in data warehousing solutions such as Snowflake.
https://docs.snowflake.com/en/user-guide/intro-editions.html

QUESTION 35
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.References:Authentication policies | Snowflake Documentation,Authenticating to the server | Snowflake Documentation,External API authentication and secrets | Snowflake Documentation.

QUESTION 36
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 37

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.
References:
Snowflake Documentation on Clustering Keys & Clustered Tables
Snowflake Documentation on SYSTEM$CLUSTERING_INFORMATION
Stack Overflow discussion on cluster key selection in Snowflake

QUESTION 38
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.
References:
Snowflake Documentation on Pipes
SnowPro Core Certification Study Guide
https://docs.snowflake.com/en/user-guide/data-load-snowpipe-intro.html

QUESTION 39
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.
References:
Snowflake Documentation on Roles
Snowflake Documentation on Materialized Views

QUESTION 40
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 41
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.
References:
Snowflake Documentation on Metadata Management
SnowPro Core Certification Study Guide

QUESTION 42

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.
References:
Snowflake Documentation on Query Results Cache
SnowPro Core Certification Study Guide

QUESTION 43
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.
References:
Snowflake Documentation on Virtual Warehouses
SnowPro Core Certification Study Guide

QUESTION 44
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.

References:
Snowflake Documentation on Materialized Views
SnowPro Core Certification Study Guide

QUESTION 45
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 Bl 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.
References:
Snowflake Documentation on Time Travel
SnowPro Core Certification Study Guide

QUESTION 46
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.
References: [COF-C02] SnowPro Core Certification Exam Study Guide

QUESTION 47
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.
References:
Snowflake Documentation on Virtual Warehouses
SnowPro Core Certification Study Guide

QUESTION 48
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.
References:
Snowflake Documentation on Multi-Cluster Warehousing
SnowPro Core Certification best practices

QUESTION 49
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.
References:
Snowflake Documentation on Secure Data Sharing
SnowPro Core Certification Companion: Hands-on Preparation and Practice

QUESTION 50
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.
References:
Snowflake Documentation on Reader Accounts
SnowPro Core Certification Study Guide

QUESTION 51
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.
References:
Snowflake Documentation on Stages
SnowPro Core Certification Study Guide

QUESTION 52
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.
References:
Snowflake Documentation on Security Features
SnowPro Core Certification Exam Study Guide

QUESTION 53
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.
References:
Snowflake Documentation on Unloading Data
Snowflake SnowPro Core: Copy Into Command to Unload Rows to Files in Named Stage

QUESTION 54
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.
References:
Understanding Encryption Key Management in Snowflake

QUESTION 55
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.
References:
Snowflake Documentation on Semi-Structured Data Types
SnowPro Core Certification Study Guide

QUESTION 56
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.
References:
Snowflake Documentation on Snowpipe
SnowPro Core Certification Study Guide

QUESTION 57
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.
References:
Snowflake Documentation on Data Loading Considerations
[COF-C02] SnowPro Core Certification Exam Study Guide

QUESTION 58
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.
References:
Snowflake Documentation on Data Loading
SnowPro Core Certification Study Guide

QUESTION 59
A user has unloaded data from Snowflake to a stage
Which SQL command should be used to validate which data was loaded into the stage?

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

Correct Answer: A
Section:

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

QUESTION 60
What happens when a cloned table is replicated to a secondary database? (Select TWO)

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

Correct Answer: C, E
Section:
Explanation:
When a cloned table is replicated to a secondary database in Snowflake, the following occurs:
C . The physical data is replicated: The actual data of the cloned table is physically replicated to the secondary database.This ensures that the secondary database has its own copy of the data, which can be used for read-only purposes or failover scenarios1.
E . Metadata pointers to cloned tables are replicated: Along with the physical data, the metadata pointers that refer to the cloned tables are also replicated.This metadata includes information about the structure of the table and any associated properties2.
It's important to note that while the physical data and metadata are replicated, the secondary database is typically read-only and cannot be used for write operations. Additionally, while there may be additional storage costs associated with the secondary account, this is not a direct result of the replication process but rather a consequence of storing additional data.
References:
SnowPro Core Exam Prep --- Answers to Snowflake's LEVEL UP: Backup and Recovery
Snowflake SnowPro Core Certification Exam Questions Set 10

QUESTION 61
Which data types does Snowflake support when querying semi-structured data? (Select TWO)

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

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

QUESTION 62
Which of the following Snowflake objects can be shared using a secure share? (Select TWO).

A. Materialized views
B. Sequences
C. Procedures
D. Tables
E. Secure User Defined Functions (UDFs)

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

QUESTION 63
Will data cached in a warehouse be lost when the warehouse is resized?

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

Correct Answer: C
Section:
Explanation:
When a Snowflake virtual warehouse is resized, the data cached in the warehouse is not lost. This is because the cache is maintained independently of the warehouse size. Resizing a warehouse, whether scaling up or down, does not affect the cached data, ensuring that query performance is not impacted by such changes.
References:
[COF-C02] SnowPro Core Certification Exam Study Guide
Snowflake Documentation on Virtual Warehouse Performance1

QUESTION 64
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.
References:
[COF-C02] SnowPro Core Certification Exam Study Guide
Snowflake's official documentation and partner listings

QUESTION 65
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.
References:
[COF-C02] SnowPro Core Certification Exam Study Guide
Snowflake Documentation on Data Loading1

QUESTION 66
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.
References:
[COF-C02] SnowPro Core Certification Exam Study Guide
Snowflake Documentation on Data Unloading

QUESTION 67
Where would a Snowflake user find information about query activity from 90 days ago?

A. account__usage . query history view
B. account__usage.query__history__archive View
C. information__schema . cruery_history view
D. information__schema - query history_by_ses s i on view

Correct Answer: B
Section:
Explanation:
To find information about query activity from 90 days ago, a Snowflake user should use theaccount_usage.query_history_archiveview. This view is designed to provide access to historical query data beyond the default 14-day retention period found in the standardquery_historyview. It allows users to analyze and audit past query activities for up to 365 days after the date of execution, which includes the 90-day period mentioned.
References:
[COF-C02] SnowPro Core Certification Exam Study Guide
Snowflake Documentation on Account Usage Schema1

QUESTION 68
Which Snowflake technique can be used to improve the performance of a query?

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

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

QUESTION 69
User-level network policies can be created by which of the following roles? (Select TWO).

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

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

QUESTION 70
Which command can be used to load data into an internal stage?

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

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

QUESTION 71
What happens when an external or an internal stage is dropped? (Select TWO).

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

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

QUESTION 72
How long is Snowpipe data load history retained?

A. As configured in the create pipe settings
B. Until the pipe is dropped
C. 64 days
D. 14 days

Correct Answer: C
Section:
Explanation:
Snowpipe data load history is retained for 64 days. This retention period allows users to review and audit the data load operations performed by Snowpipe over a significant period of time, which can be crucial for troubleshooting and ensuring data integrity.
References:

[COF-C02] SnowPro Core Certification Exam Study Guide
Snowflake Documentation on Snowpipe1

QUESTION 73
What tasks can be completed using the copy command? (Select TWO)

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

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

QUESTION 74
What feature can be used to reorganize a very large table on one or more columns?

A. Micro-partitions
B. Clustering keys
C. Key partitions
D. Clustered partitions

Correct Answer: B
Section:
Explanation:
Clustering keys in Snowflake are used to reorganize large tables based on one or more columns. This feature optimizes the arrangement of data within micro-partitions to improve query performance, especially for large tables where efficient data retrieval is crucial. References: [COF-C02] SnowPro Core Certification Exam Study Guide
https://docs.snowflake.com/en/user-guide/tables-clustering-keys.html

QUESTION 75
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. References:SHOW GRANTS | Snowflake Documentation

QUESTION 76

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. References:Snowpipe | Snowflake Documentation

QUESTION 77
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 78
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 79
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 80
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 81
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 82
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 83
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. References: [COF-C02] SnowPro Core Certification Exam Study Guide

QUESTION 84
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. References: [COF-C02] SnowPro Core Certification Exam Study Guide

QUESTION 85
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 86
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 87

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 88
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 89
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.
References = [COF-C02] SnowPro Core Certification Exam Study Guide, Snowflake Documentation1

QUESTION 90
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 91
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 92
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.
References = [COF-C02] SnowPro Core Certification Exam Study Guide, Snowflake Documentation1

QUESTION 93
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.
References = [COF-C02] SnowPro Core Certification Exam Study Guide, Snowflake Documentation

QUESTION 94
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.
References = [COF-C02] SnowPro Core Certification Exam Study Guide, Snowflake Documentation1

QUESTION 95
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.
References = [COF-C02] SnowPro Core Certification Exam Study Guide, Snowflake Documentation2

QUESTION 96
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.

References = [COF-C02] SnowPro Core Certification Exam Study Guide, Snowflake Documentation

QUESTION 97
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.
References = [COF-C02] SnowPro Core Certification Exam Study Guide, Snowflake Documentation

QUESTION 98
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.
References = [COF-C02] SnowPro Core Certification Exam Study Guide, Snowflake Documentation

QUESTION 99
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.
References = [COF-C02] SnowPro Core Certification Exam Study Guide, Snowflake Documentation

QUESTION 100

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.
References = [COF-C02] SnowPro Core Certification Exam Study Guide, Snowflake Documentation

QUESTION 101
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 102
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 103
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 104
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 105
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 106
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 107
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.

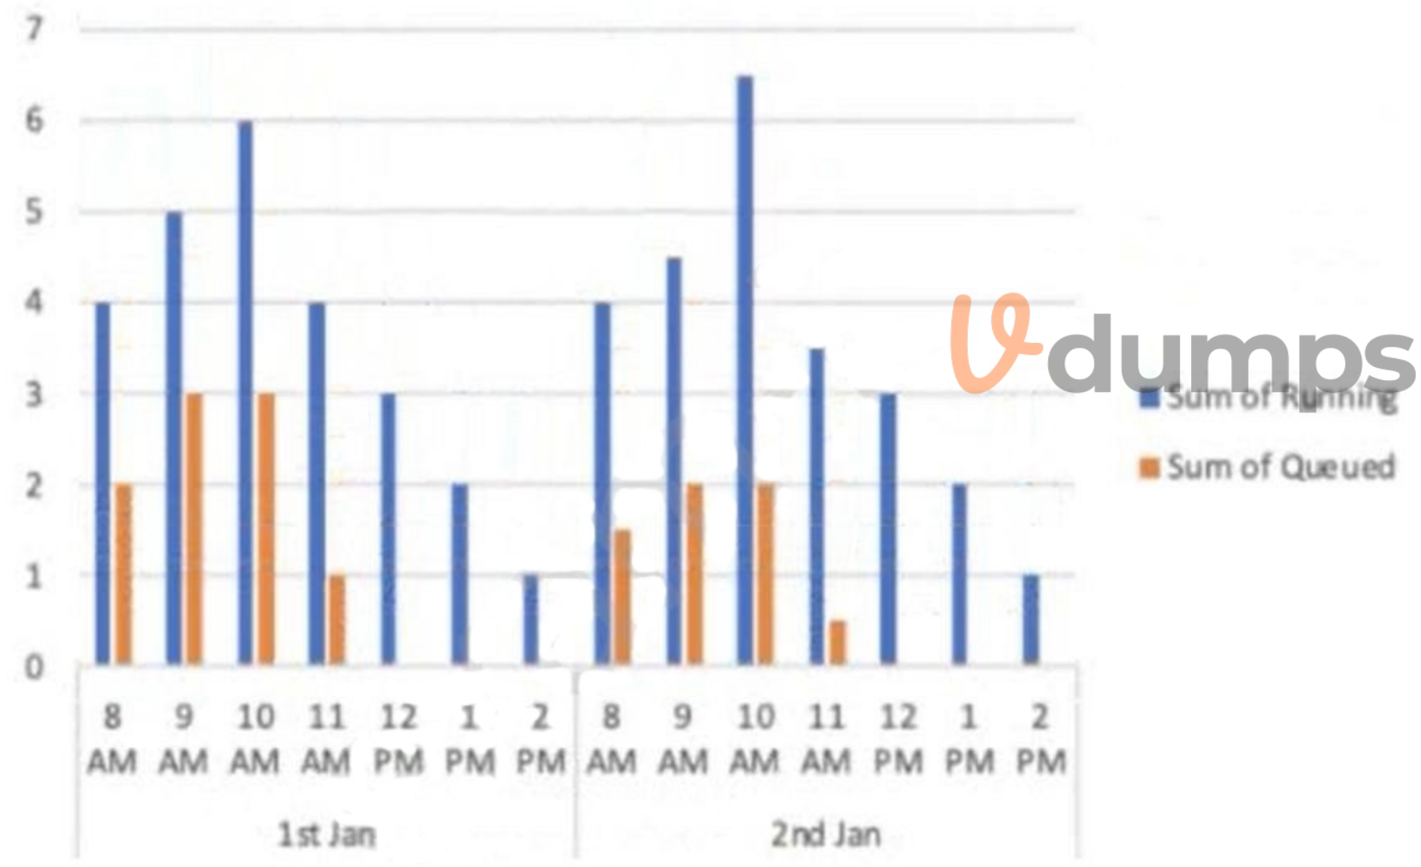

What action should be taken to address this situation?

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.
References = Snowflake Documentation on Multi-cluster Warehouses1

QUESTION 108
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 109
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 110
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 111
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. References:SHOW OBJECTS,Database, Schema, & Share DDL

QUESTION 112
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. References:Caching in the Snowflake Cloud Data Platform,Optimizing the warehouse cache

QUESTION 113
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. References:COPY INTO Error during Structured Data Load: ''Max LOB size (16777216) exceeded...''

QUESTION 114
Which of the following describes a Snowflake stored procedure?

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

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

QUESTION 115
Which columns are part of the result set of the Snowflake LATERAL FLATTEN command? (Choose two.)

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

Correct Answer: B, D
Section:
Explanation:
TheLATERAL FLATTENcommand in Snowflake produces a result set that includes several columns, among whichPATHandINDEXare includedPATHindicates the path to the element within a data structure that needs to be flattened, andINDEXrepresents the index of the element if it is an array2.

QUESTION 116
Which Snowflake function will interpret an input string as a JSON document, and produce a VARIANT value?

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

Correct Answer: A
Section:
Explanation:

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

QUESTION 117
How are serverless features billed?

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

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

QUESTION 118
Which Snowflake architectural layer is responsible for a query execution plan?

A. Compute
B. Data storage
C. Cloud services
D. Cloud provider

Correct Answer: C
Section:
Explanation:
In Snowflake's architecture, the Cloud Services layer is responsible for generating the query execution plan. This layer handles all the coordination, optimization, and management tasks, including query parsing, optimization, and compilation into an execution plan that can be processed by the Compute layer.

QUESTION 119
Which SQL commands, when committed, will consume a stream and advance the stream offset? (Choose two.)

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

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

QUESTION 120

Which methods can be used to delete staged files from a Snowflake stage? (Choose two.)

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

Correct Answer: C, D
Section:
Explanation:
To delete staged files from a Snowflake stage, you can specify thePURGEoption in theCOPY INTO <table>command, which will automatically delete the files after they have been successfully loaded.Additionally, you can use theREMOVEcommand after the load completes to manually delete the files from the stage12.
References =DROP STAGE,REMOVE

QUESTION 121
Assume there is a table consisting of five micro-partitions with values ranging from A to Z.
Which diagram indicates a well-clustered table?

A.

B.

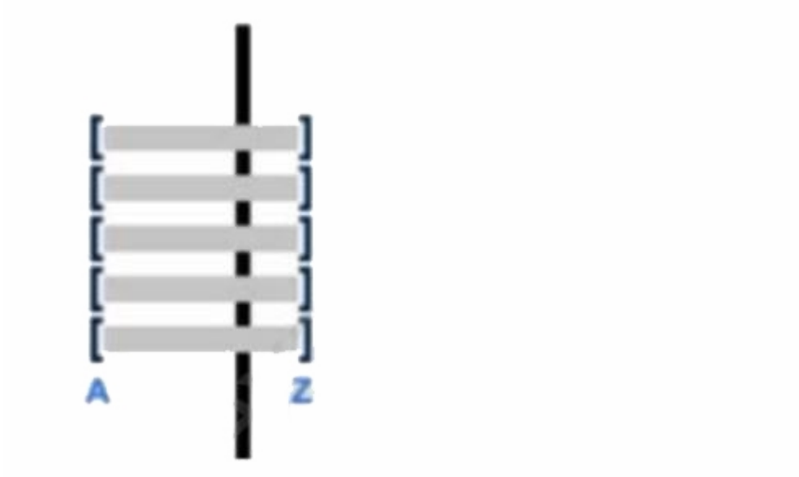

C.

D.

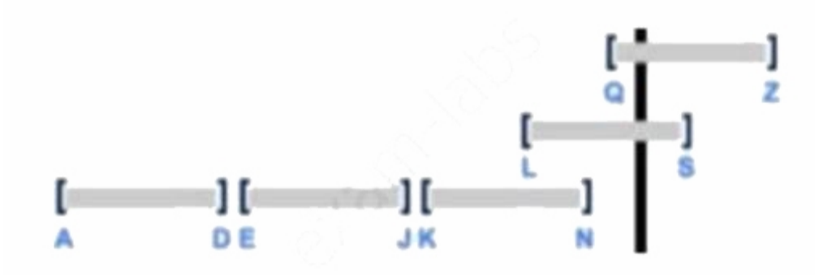

E. Option A
F. Option B
G. Option C
H. Option D

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

QUESTION 122
What is an advantage of using an explain plan instead of the query profiler to evaluate the performance of a query?

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

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

QUESTION 123
Which data types are supported by Snowflake when using semi-structured data? (Choose two.)

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

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

QUESTION 124
Why does Snowflake recommend file sizes of 100-250 MB compressed when loading data?

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

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

QUESTION 125
Which of the following features are available with the Snowflake Enterprise edition? (Choose two.)

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

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

QUESTION 126
What is the default file size when unloading data from Snowflake using the COPY command?

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

D. 32 MB

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

QUESTION 127
What features that are part of the Continuous Data Protection (CDP) feature set in Snowflake do not require additional configuration? (Choose two.)

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

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

QUESTION 128
Which Snowflake layer is always leveraged when accessing a query from the result cache?

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

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

QUESTION 129
A Snowflake Administrator needs to ensure that sensitive corporate data in Snowflake tables is not visible to end users, but is partially visible to functional managers.
How can this requirement be met?

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

Correct Answer: B
Section:

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

QUESTION 130
Users are responsible for data storage costs until what occurs?

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

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

QUESTION 131
What affects whether the query results cache can be used?

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

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

QUESTION 132
Which of the following is an example of an operation that can be completed without requiring compute, assuming no queries have been executed previously?

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

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

QUESTION 133
How many days is load history for Snowpipe retained?

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

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

QUESTION 134
How can a row access policy be applied to a table or a view? (Choose two.)

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

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

QUESTION 135
Which command can be used to load data files into a Snowflake stage?

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

Correct Answer: C
Section:
Explanation:
ThePUTcommand is used to load data files into a Snowflake stage.This command uploads data files from a local file system to a specified stage in Snowflake

QUESTION 136
What types of data listings are available in the Snowflake Data Marketplace? (Choose two.)

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

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

QUESTION 137
What is the maximum Time Travel retention period for a temporary Snowflake table?

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

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

QUESTION 138
When should a multi-cluster warehouse be used in auto-scaling mode?

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

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

QUESTION 139
Snowflake supports the use of external stages with which cloud platforms? (Choose three.)

A. Amazon Web Services
B. Docker
C. IBM Cloud
D. Microsoft Azure Cloud
E. Google Cloud Platform
F. Oracle Cloud

Correct Answer: A, D, E
Section:
Explanation:
Snowflake supports the use of external stages with Amazon Web Services (AWS), Microsoft Azure Cloud, and Google Cloud Platform (GCP).These platforms allow users to stage data externally and integrate with Snowflake for

data loading operations

QUESTION 140
In the Snowflake access control model, which entity owns an object by default?

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

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

QUESTION 141
What is the minimum Snowflake edition required to use Dynamic Data Masking?

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

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

QUESTION 142
Which services does the Snowflake Cloud Services layer manage? (Choose two.)

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

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

QUESTION 143
A company needs to allow some users to see Personally Identifiable Information (PII) while limiting other users from seeing the full value of the PII.
Which Snowflake feature will support this?

A. Row access policies
B. Data masking policies
C. Data encryption
D. Role based access control

Correct Answer: B
Section:
Explanation:
Data masking policies in Snowflake allow for the obfuscation of specific data within a field, enabling some users to see the full data while limiting others.This feature is particularly useful for handling PII, ensuring that sensitive information is only visible to authorized users1.

QUESTION 144
Which of the following accurately describes shares?

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

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

QUESTION 145
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.References:Based on general data exchange practices in cloud services as of 2021.

QUESTION 146
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 147
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 148
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 149
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 150
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 151
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 152
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 153
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 154
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 155
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.References:[COF-C02] SnowPro Core Certification Exam Study Guide

QUESTION 156
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.References:[COF-C02] SnowPro Core Certification Exam Study Guide

QUESTION 157
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.References:[COF-C02] SnowPro Core Certification Exam Study Guide

QUESTION 158
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.References:Based on my internal knowledge as of 2021.

QUESTION 159
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.References:Based on my internal knowledge as of 2021.

QUESTION 160
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.References:Based on my internal knowledge as of 2021.

QUESTION 161
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 162
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 163
Which feature is integrated to support Multi-Factor Authentication (MFA) at Snowflake?

A. Authy
B. Duo Security

C. One Login
D. RSA SecurlD 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 164
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 165
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 166
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 167
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. References: Based on general database optimization principles.

QUESTION 168
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. References: Based on general cloud database architecture knowledge.

QUESTION 169
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. References: Based on best practices for database object ownership and role management.

QUESTION 170
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. References: [COF-C02] SnowPro Core Certification Exam Study Guide

QUESTION 171
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. References: [COF-C02] SnowPro Core Certification Exam Study Guide

QUESTION 172
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 173
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 174
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 175
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 176
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 177
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 178
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 179
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 180
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 181
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 182
Which statement accurately describes a characteristic of a materialized view?

A. A materialized view can query only a single table.
B. Data accessed through materialized views can be stale.
C. Materialized view refreshes need to be maintained by the user.
D. Querying a materialized view is slower than executing a query against the base table of the view.

Correct Answer: B
Section:
Explanation:
A characteristic of a materialized view is that the data accessed through it can be stale.This is because the data in a materialized view may not reflect the latest changes in the base tables until the view is refreshed

QUESTION 183
Which stream type can be used for tracking the records in external tables?

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

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

QUESTION 184
What can a Snowflake user do in the Activity section in Snowsight?

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

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

QUESTION 185
Which Snowflake URL type is used by directory tables?

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

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

QUESTION 186
For the ALLOWED VALUES tag property, what is the MAXIMUM number of possible string values for a single tag?

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

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

QUESTION 187
What does Snowflake's search optimization service support?

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

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

QUESTION 188
What action can a user take to address query concurrency issues?

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

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

QUESTION 189
How long can a data consumer who has a pre-signed URL access data files using Snowflake?

A. Indefinitely
B. Until the result_cache expires
C. Until the retention_time is met
D. Until the expiration time is exceeded

Correct Answer: D
Section:
Explanation:
A data consumer who has a pre-signed URL can access data files using Snowflake until the expiration time is exceeded.The expiration time is set when the pre-signed URL is generated and determines how long the URL remains valid3.

QUESTION 190
Which pages are included in the Activity area of Snowsight? (Select TWO).

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

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

QUESTION 191
Which kind of Snowflake table stores file-level metadata for each file in a stage?

A. Directory
B. External
C. Temporary
D. Transient

Correct Answer: A
Section:
Explanation:
The kind of Snowflake table that stores file-level metadata for each file in a stage is a directory table.A directory table is an implicit object layered on a stage and stores file-level metadata about the data files in the stage3.

QUESTION 192
Which parameter prevents streams on tables from becoming stale?

A. MAXDATAEXTENSIONTIMEINDAYS
B. MTN_DATA_RETENTION_TTME_TN_DAYS
C. LOCK_TIMEOUT
D. STALE_AFTER

Correct Answer: A
Section:
Explanation:
The parameter that prevents streams on tables from becoming stale is MAXDATAEXTENSIONTIMEINDAYS.This parameter specifies the maximum number of days for which Snowflake can extend the data retention period for the table to prevent streams on the table from becoming stale4.

QUESTION 193
A user needs to create a materialized view in the schema MYDB.MYSCHEM

A. Which statements will provide this access?
B. GRANT ROLE MYROLE TO USER USER1; GRANT CREATE MATERIALIZED VIEW ON SCHEMA MYDB.MYSCHEMA TO ROLE MYROLE;
C. GRANT ROLE MYROLE TO USER USER1; GRANT CREATE MATERIALIZED VIEW ON SCHEMA MYDB.MYSCHEMA TO USER USER1;
D. GRANT ROLE MYROLE TO USER USER1; GRANT CREATE MATERIALIZED VIEW ON SCHEMA MYDB. K'-'SCHEMA TO USER! ;
E. GRANT ROLE MYROLE TO USER USER1; GRANT CREATE MATERIALIZED VIEW ON SCHEMA MYDB.MYSCHEMA TO MYROLE;

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

QUESTION 194
Which objects together comprise a namespace in Snowflake? (Select TWO).

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

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

QUESTION 195
Two users share a virtual warehouse named wh dev 01. When one of the users loads data, the other one experiences performance issues while querying data.
How does Snowflake recommend resolving this issue?

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

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

QUESTION 196
How long does Snowflake retain information in the ACCESS HISTORY view?

A. 7 days
B. 14 days
C. 28 days
D. 365 days

Correct Answer: D
Section:
Explanation:
Snowflake retains information in the ACCESS HISTORY view for 365 days.This allows users to query the access history of Snowflake objects within the last year1.

QUESTION 197
A view is defined on a permanent table. A temporary table with the same name is created in the same schema as the referenced table. What will the query from the view return?

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

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

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

QUESTION 198
What is used to diagnose and troubleshoot network connections to Snowflake?

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

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

QUESTION 199
Which REST API can be used with unstructured data?

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

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

QUESTION 200
What column type does a Kafka connector store formatted information in a single column?

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

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

QUESTION 201
When unloading data to an external stage, what is the MAXIMUM file size supported?

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

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

QUESTION 202
By definition, a secure view is exposed only to users with what privilege?

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

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

QUESTION 203
What is a characteristic of the Snowflake Query Profile?

A. It can provide statistics on a maximum number of 100 queries per week.
B. It provides a graphic representation of the main components of the query processing.
C. It provides detailed statistics about which queries are using the greatest number of compute resources.
D. It can be used by third-party software using the Query Profile API.

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

QUESTION 204
What statistical information in a Query Profile indicates that the query is too large to fit in memory? (Select TWO).

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

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

QUESTION 205
What is the MAXIMUM Time Travel retention period for a transient table?

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

Correct Answer: B
Section:
Explanation:
The maximum Time Travel retention period for a transient table in Snowflake is 1 day.This is the default and maximum duration for which Snowflake maintains the historical data for transient tables, allowing users to query data as it appeared at any point within the past 24 hours2.

QUESTION 206
Which query contains a Snowflake hosted file URL in a directory table for a stage named bronzestage?

A. list @bronzestage;
B. select * from directory(@bronzestage);
C. select metadata$filename from @bronzestage;
D. select * from table(information_schema.stage_directory_file_registration_history(stage name=>'bronzestage1));

Correct Answer: B
Section:
Explanation:
The query that contains a Snowflake hosted file URL in a directory table for a stage named bronzestage isselect * from directory(@bronzestage).This query retrieves a list of all files on the stage along with metadata, including the Snowflake file URL for each file3.

QUESTION 207
Which Snowflake edition enables data sharing only through Snowflake Support?

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

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

QUESTION 208
When would Snowsight automatically detect if a target account is in a different region and enable cross-cloud auto-fulfillment?

A. When using a paid listing on the Snowflake Marketplace
B. When using a private listing on the Snowflake Marketplace
C. When using a personalized listing on the Snowflake Marketplace
D. When using a Direct Share with another account

Correct Answer: A
Section:
Explanation:
Snowsight automatically detects if a target account is in a different region and enables cross-cloud auto-fulfillment when using a paid listing on the Snowflake Marketplace.This feature allows Snowflake to manage the replication of data products to consumer regions as needed, without manual intervention1.

QUESTION 209
Which languages requite that User-Defined Function (UDF) handlers be written inline? (Select TWO).

A. Java
B. Javascript
C. Scala
D. Python
E. SQL

Correct Answer: B, E
Section:
Explanation:
User-Defined Function (UDF) handlers must be written inline for Javascript and SQL.These languages allow the UDF logic to be included directly within the SQL statement that creates the UDF2.

QUESTION 210
Which semi-structured data function interprets an input string as a JSON document that produces a VARIANT value?

A. PARSE_JSON
B. CHECK_JSON
C. JSON_EXTRACT_PATH_TEXT
D. PARSE_XML

Correct Answer: A
Section:
Explanation:
The semi-structured data function that interprets an input string as a JSON document and produces a VARIANT value is PARSE_JSON.This function is used to parse a JSON formatted string and return it as a VARIANT data type, which can then be used for further processing within Snowflake3.

QUESTION 211
Which items are considered schema objects in Snowflake? (Select TWO).

A. Pipe
B. File format

C. Resource monitor
D. Storage integration
E. Virtual warehouse

Correct Answer: A, B
Section:
Explanation:
In Snowflake, schema objects include Pipes and File formats.Pipes are used for continuous data loading, and File formats specify the format of data files used in loading and unloading operations within Snowflake

QUESTION 212
What are benefits of using Snowpark with Snowflake? (Select TWO).

A. Snowpark uses a Spark engine to generate optimized SQL query plans.
B. Snowpark automatically sets up Spark within Snowflake virtual warehouses.
C. Snowpark does not require that a separate cluster be running outside of Snowflake.
D. Snowpark allows users to run existing Spark code on virtual warehouses without the need to reconfigure the code.
E. Snowpark executes as much work as possible in the source databases for all operations including User-Defined Functions (UDFs).

Correct Answer: C, D
Section:
Explanation:
Snowpark is designed to bring the data programmability to Snowflake, enabling developers to write code in familiar languages like Scala, Java, and Python. It allows for the execution of these codes directly within Snowflake's virtual warehouses, eliminating the need for a separate cluster.Additionally, Snowpark's compatibility with Spark allows users to leverage their existing Spark code with minimal changes1.

QUESTION 213
If a virtual warehouse runs for 61 seconds, shuts down, and then restarts and runs for 30 seconds, for how many seconds is it billed?

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

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

QUESTION 214
Which transformation is supported by a COPY INTO <table> command?

A. Filter using a where clause
B. Filter using a limit keyword
C. Cast using a SELECT statement
D. Order using an ORDER BY clause

Correct Answer: C

Section:
Explanation:
The COPY INTO <table> command in Snowflake supports transformations such as casting using a SELECT statement.This allows for the transformation of data types as the data is being loaded into the table, which can be particularly useful when the data types in the source files do not match the data types in the target table

QUESTION 215
Which task privilege does a Snowflake role need in order to suspend or resume a task?

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

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

QUESTION 216
When a Snowflake user loads CSV data from a stage, which copy into <table> command guideline should they follow?

A. The CSV field delimiter must be a comma character (*,'),
B. The number of columns in each row should be consistent.
C. The data file in the stage must be in a compressed format.
D. The data file must have the same number of columns as the target table.

Correct Answer: B
Section:
Explanation:
When using the COPY INTO <table> command to load CSV data from a stage into a Snowflake table, one of the crucial guidelines to follow is that the number of columns in each row within the CSV file should be consistent. This ensures data integrity and allows for a smooth data loading process, as Snowflake expects each row in the CSV file to map directly to a row in the target table based on the number of columns. If there's a mismatch in the number of columns between any row in the file and the target table, Snowflake might return an error or produce unexpected results during the load operation.
It's important to note that while the CSV field delimiter can be specified to something other than a comma, ensuring the consistency in the number of columns across all rows is fundamental to successfully loading data.
References:
Snowflake Documentation on Loading Data: Loading CSV Data

QUESTION 217
In addition to performing all the standard steps to share data, which privilege must be granted on each database referenced by a secure view in order to be shared?

A. READ
B. REFERENCES
C. REFERENCEJJSAGE
D. USAGE

Correct Answer: D
Section:
Explanation:

In addition to performing all the standard steps to share data, the USAGE privilege must be granted on each database referenced by a secure view in order to be shared. When sharing a database or specific objects like secure views, the receiving account needs to have the USAGE privilege on the database and schema to access the shared data. This privilege enables the receiving account to access the database and its schemas but does not allow for any DML operations. It's a prerequisite for accessing any objects within the database.
For a secure view to be part of a share, not only does the view itself need to be shared, but the underlying database (and schema, if applicable) must also be accessible to the recipients. Granting USAGE privilege on the database ensures that the receiving account can access the database in a read-only mode to utilize the shared view.
References:
Snowflake Documentation on Shares: Creating and Managing Shares

QUESTION 218
Which Snowsight feature can be used to perform data manipulations and transformations using a programming language?

A. SnowSQL
B. Dashboards
C. Python worksheets
D. Provider Studio

Correct Answer: C
Section:
Explanation:
Python worksheets in Snowsight enable users to perform data manipulations and transformations using the Python programming language directly within the Snowflake environment. This feature integrates the power of Python with Snowflake's data warehousing capabilities, allowing for sophisticated data analysis and manipulation.
Introduction to Python Worksheets:
Python worksheets provide an interactive environment to write and execute Python code.
They are designed to facilitate data science and data engineering tasks.
Functionality:
Users can run Python scripts to manipulate data stored in Snowflake.
It allows for leveraging Python's extensive libraries for data analysis, machine learning, and more.
Integration with Snowflake:
Python worksheets run on Snowflake's compute infrastructure, ensuring scalability and performance.
They can access and manipulate Snowflake tables directly, making them a powerful tool for data transformation.
References:
Snowflake Documentation: Snowsight Python Worksheets

QUESTION 219
Who can activate a network policy for users in a Snowflake account? (Select TWO)

A. ACCOUNTADMIN
B. USERADMIN
C. PUBLIC
D. SYSADMIN
E. Any role that has the global ATTACH POLICY privilege

Correct Answer: A, E
Section:
Explanation:
Network policies in Snowflake are used to control access to Snowflake accounts based on IP address ranges. These policies can be activated by specific roles that have the necessary privileges.
Role: ACCOUNTADMIN:
The ACCOUNTADMIN role has full administrative rights across the Snowflake account.

This role can manage all aspects of the Snowflake environment, including network policies.
Role with Global ATTACH POLICY Privilege:
Any role that has been granted the global ATTACH POLICY privilege can activate network policies.
This privilege allows the role to attach policies that control network access to the account.
References:
Snowflake Documentation: Network Policies

QUESTION 220
What is the MINIMUM size requirement when creating a Snowpark-optimized virtual warehouse?

A. X-Small
B. Small
C. Medium
D. Large

Correct Answer: B
Section:
Explanation:
When creating a Snowpark-optimized virtual warehouse in Snowflake, the minimum size requirement is Small. Snowpark is designed to handle data processing workloads efficiently, and the Small size ensures adequate resources for such tasks.
Virtual Warehouse Sizes:
Snowflake offers different sizes for virtual warehouses, ranging from X-Small to 6X-Large.
Each size corresponds to a specific level of compute resources.
Minimum Size Requirement for Snowpark:
A Small virtual warehouse is the minimum size required to optimize performance and resource allocation for Snowpark workloads.
This ensures that the warehouse has sufficient capacity to handle data processing and transformation tasks efficiently.
References:
Snowflake Documentation: Virtual Warehouse Sizes

QUESTION 221
Awarding a user which privileges on all virtual warehouses is equivalent to granting the user the global MANAGE WAREHOUSES privilege?

A. MODIFY, MONITOR and OPERATE privileges
B. ownership and usage privileges
C. APPLYBUDGET and audit privileges
D. MANAGE LISTING ADTOTOLFillment and resolve all privileges

Correct Answer: A
Section:
Explanation:
Granting a user the MODIFY, MONITOR, and OPERATE privileges on all virtual warehouses in Snowflake is equivalent to granting the global MANAGE WAREHOUSES privilege. These privileges collectively provide comprehensive control over virtual warehouses.
MODIFY Privilege:
Allows users to change the configuration of the virtual warehouse.
Includes resizing, suspending, and resuming the warehouse.
MONITOR Privilege:
Allows users to view the status and usage metrics of the virtual warehouse.
Enables monitoring of performance and workload.
OPERATE Privilege:

Grants the ability to start and stop the virtual warehouse.
Includes pausing and resuming operations as needed.
References:
Snowflake Documentation: Warehouse Privileges

QUESTION 222
In Snowflake's data security framework, how does column-level security contribute to the protection of sensitive information? (Select TWO).

A. Implementation of column-level security will optimize query performance.
B. Column-level security supports encryption of the entire database.
C. Column-level security ensures that only the table owner can access the data.
D. Column-level security limits access to specific columns within a table based on user privileges
E. Column-level security allows the application of a masking policy to a column within a table or view.

Correct Answer: D, E
Section:
Explanation:
Column-level security in Snowflake enhances data protection by restricting access and applying masking policies to sensitive data at the column level.
Limiting Access Based on User Privileges:
Column-level security allows administrators to define which users or roles have access to specific columns within a table.
This ensures that sensitive data is only accessible to authorized personnel, thereby reducing the risk of data breaches.
Application of Masking Policies:
Masking policies can be applied to columns to obfuscate sensitive data.
For example, credit card numbers can be masked to show only the last four digits, protecting the full number from being exposed.
References:
Snowflake Documentation: Column-Level Security
Snowflake Documentation: Dynamic Data Masking