Snowflake

Exam Questions COF-C02

SnowPro Core Certification Exam (COF-C02)

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

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

Answer: D

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

NEW QUESTION 2
- (Topic 1)
Which account usage views are used to evaluate the details of dynamic data masking? (Select TWO)

A. ROLES
B. POLICY_REFERENCES
C. QUERY_HISTORY
D. RESOURCE_MONIT ORS
E. ACCESS_HISTORY

Answer: BE

Explanation:
To evaluate the details of dynamic data masking,
the POLICY_REFERENCES and ACCESS_HISTORY views in the account_usage schema are used. The POLICY_REFERENCES view provides information about the objects to which a masking policy is applied, and the ACCESS_HISTORY view contains details about access to the masked data, which can be used to audit and verify the application of
dynamic data masking policies.
References:
? [COF-C02] SnowPro Core Certification Exam Study Guide
? Snowflake Documentation on Dynamic Data Masking1

NEW QUESTION 3
- (Topic 1)
Which of the following are valid methods for authenticating users for access into Snowflake? (Select THREE)

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

Answer: BDE

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

NEW QUESTION 4
- (Topic 1)
A company strongly encourages all Snowflake users to self-enroll in Snowflake's default Multi-Factor Authentication (MFA) service to provide increased login security for users connecting to Snowflake.
Which application will the Snowflake users need to install on their devices in order to connect with MFA?

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

Answer: B

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

NEW QUESTION 5
- (Topic 1)
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

Answer: B

Explanation:
The COPY INTO command in Snowflake can be configured to output data into a single file by setting the MAX_FILE_NUMBER option 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

NEW QUESTION 6
- (Topic 1)
Which services does the Snowflake Cloud Services layer manage? (Select TWO).

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

Answer: CE

Explanation:
The Snowflake Cloud Services layer manages a variety of services that are crucial for the operation of the Snowflake platform. Among these services, Authentication and Metadata management are key components. Authentication is essential for controlling access to the Snowflake environment, ensuring that only authorized users can perform actions within the platform. Metadata management involves handling all the metadata related to objects within Snowflake, such as tables, views, and databases, which is vital for the organization and retrieval of data.
References:
? [COF-C02] SnowPro Core Certification Exam Study Guide
? Snowflake Documentation12 https://docs.snowflake.com/en/user-guide/intro-key-concepts.html

NEW QUESTION 7
- (Topic 1)
What is the recommended file sizing for data loading using Snowpipe?

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

Answer: C

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

NEW QUESTION 8
- (Topic 1)
Which of the following describes how multiple Snowflake accounts in a single organization relate to various cloud providers?

A. Each Snowflake account can be hosted in a different cloud vendor and region.
B. Each Snowflake account must be hosted in a different cloud vendor and region
C. All Snowflake accounts must be hosted in the same cloud vendor and region
D. Each Snowflake account can be hosted in a different cloud vendor, but must be in the same region.

Answer: A

Explanation:
Snowflake??s architecture allows for flexibility in account hosting across different cloud vendors and regions. This means that within a single organization, different Snowflake accounts can be set up in various cloud environments, such as AWS, Azure, or
GCP, and in different geographical regions. This allows organizations to leverage the global infrastructure of multiple cloud providers and optimize their data storage and computing needs based on regional requirements, data sovereignty laws, and other considerations.
https://docs.snowflake.com/en/user-guide/intro-regions.html

NEW QUESTION 9
- (Topic 1)

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

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

Answer: CD

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

NEW QUESTION 10
- (Topic 1)
User-level network policies can be created by which of the following roles? (Select TWO).

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

Answer: BD

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

NEW QUESTION 10
- (Topic 1)
What is a best practice after creating a custom role?

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

Answer: B

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

NEW QUESTION 14
- (Topic 1)
Which data types does Snowflake support when querying semi-structured data? (Select TWO)

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

Answer: AB

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

NEW QUESTION 19
- (Topic 1)
Will data cached in a warehouse be lost when the warehouse is resized?

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

Answer: C

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

NEW QUESTION 20
- (Topic 1)
The fail-safe retention period is how many days?

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

Answer: B

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

NEW QUESTION 24
- (Topic 1)
Which of the following is a valid source for an external stage when the Snowflake account is located on Microsoft Azure?

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

Answer: D

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

NEW QUESTION 26
- (Topic 1)
Which feature is only available in the Enterprise or higher editions of Snowflake?

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

Answer: A

Explanation:
Column-level security is a feature that allows fine-grained control over access to specific columns within a table. This is particularly useful for managing sensitive data and ensuring that only authorized users can view or manipulate certain pieces of information. According to my last update, this feature was available in the Enterprise Edition or higher editions of Snowflake.
References: Based on my internal data as of 2021, column-level security is an advanced feature typically reserved for higher-tiered editions like the Enterprise Edition in data warehousing solutions such as Snowflake.
https://docs.snowflake.com/en/user-guide/intro-editions.html

NEW QUESTION 28
- (Topic 1)
When reviewing the load for a warehouse using the load monitoring chart, the chart indicates that a high volume of Queries are always queuing in the warehouse
According to recommended best practice, what should be done to reduce the Queue volume? (Select TWO).

A. Use multi-clustered warehousing to scale out warehouse capacity.

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

Answer: AB

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

NEW QUESTION 30
- (Topic 1)
In which use cases does Snowflake apply egress charges?

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

Answer: C

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

NEW QUESTION 32
- (Topic 1)
Which of the following conditions must be met in order to return results from the results cache? (Select TWO).

A. The user has the appropriate privileges on the objects associated with the query
B. Micro-partitions have been reclustered since the query was last run
C. The new query is run using the same virtual warehouse as the previous query
D. The query includes a User Defined Function (UDF)
E. The query has been run within 24 hours of the previously-run query

Answer: AE

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

NEW QUESTION 37
- (Topic 1)
What happens when a virtual warehouse is resized?

A. When increasing the size of an active warehouse the compute resource for all running and queued queries on the warehouse are affected
B. When reducing the size of a warehouse the compute resources are removed only when they are no longer being used to execute any current statements.
C. The warehouse will be suspended while the new compute resource is provisioned and will resume automatically once provisioning is complete.
D. Users who are trying to use the warehouse will receive an error message until the resizing is complete

Answer: A

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

NEW QUESTION 40
- (Topic 1)
What SQL command would be used to view all roles that were granted to user.1?

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

Answer: A

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

NEW QUESTION 42
- (Topic 1)
True or False: Loading data into Snowflake requires that source data files be no larger than 16MB.

A. True
B. False

Answer: B

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

NEW QUESTION 47
- (Topic 1)
Which semi-structured file formats are supported when unloading data from a table? (Select TWO).

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

Answer: DE

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

NEW QUESTION 50
- (Topic 1)
What is a responsibility of Snowflake's virtual warehouses?

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

Answer: C

Explanation:
The primary responsibility of Snowflake??s virtual warehouses is to execute queries. Virtual warehouses are one of the key components of Snowflake??s architecture, providing the compute power required to perform data processing tasks such as running SQL queries, performing joins, aggregations, and other data manipulations. References:
? [COF-C02] SnowPro Core Certification Exam Study Guide
? Snowflake Documentation on Virtual Warehouses1

NEW QUESTION 53
- (Topic 1)
Which of the following are best practice recommendations that should be considered when loading data into Snowflake? (Select TWO).

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

Answer: CD

Explanation:

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

NEW QUESTION 55
- (Topic 1)
Which of the following commands cannot be used within a reader account?

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

Answer: A

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

NEW QUESTION 57
- (Topic 1)
Which is the MINIMUM required Snowflake edition that a user must have if they want to use AWS/Azure Privatelink or Google Cloud Private Service Connect?

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

Answer: D

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

NEW QUESTION 58
- (Topic 1)
What Snowflake role must be granted for a user to create and manage accounts?

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

Answer: A

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

NEW QUESTION 61
- (Topic 2)
What occurs when a pipe is recreated using the CREATE OR REPLACE PIPE command?

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

Answer: A

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

NEW QUESTION 64

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

A. Default Time Travel Retention is set to 0 day
B. Maximum Time Travel Retention is 30 day
C. Fail Safe retention time is 1 day.
D. Default Time Travel Retention is set to 1 da
E. Maximum Time Travel Retention is 365 day
F. Fail Safe retention time is 7 days.
G. Default Time Travel Retention is set to 0 day
H. Maximum Time Travel Retention is 90 day
I. Fail Safe retention time is 7 days.
J. Default Time Travel Retention is set to 1 da
K. Maximum Time Travel Retention is 90 day
L. Fail Safe retention time is 7 days.
M. Default Time Travel Retention is set to 7 day
N. Maximum Time Travel Retention is 1 da
O. Fail Safe retention time is 90 days.
P. Default Time Travel Retention is set to 90 day
Q. Maximum Time Travel Retention is 7 day
R. Fail Safe retention time is 356 days.

Answer: D

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

NEW QUESTION 67
- (Topic 2)
Users are responsible for data storage costs until what occurs?

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

Answer: B

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

NEW QUESTION 71
- (Topic 2)
Which of the following statements apply to Snowflake in terms of security? (Choose two.)

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

Answer: AC

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

NEW QUESTION 74
- (Topic 2)
Which command should be used to load data from a file, located in an external stage, into a table in Snowflake?

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

Answer: D

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

NEW QUESTION 79
- (Topic 2)
A Snowflake Administrator needs to ensure that sensitive corporate data in Snowflake tables is not visible to end users, but is partially visible to functional

managers.
How can this requirement be met?

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

Answer: B

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

NEW QUESTION 80
- (Topic 2)
What affects whether the query results cache can be used?

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

Answer: C

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

NEW QUESTION 81
- (Topic 2)
The Snowflake Cloud Data Platform is described as having which of the following architectures?

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

Answer: C

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

NEW QUESTION 84
- (Topic 2)
When cloning a database containing stored procedures and regular views, that have fully qualified table references, which of the following will occur?

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

Answer: A

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

NEW QUESTION 89
- (Topic 2)
A table needs to be loaded. The input data is in JSON format and is a concatenation of multiple JSON documents. The file size is 3 GB. A warehouse size small is being used. The following COPY INTO command was executed:
COPY INTO SAMPLE FROM @~/SAMPLE.JSON (TYPE=JSON)
The load failed with this error:
Max LOB size (16777216) exceeded, actual size of parsed column is 17894470. How can this issue be resolved?

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

Answer: B

Explanation:
The error ??Max LOB size (16777216) exceeded?? indicates that the size of the parsed column exceeds the maximum size allowed for a single column value in

Snowflake, which is 16 MB. To resolve this issue, the file should be split into multiple smaller files that are within the recommended size range of 100 MB to 250 MB. This will ensure that each JSON document within the files is smaller than the maximum LOB size allowed. Compressing the file, using a larger-sized warehouse, or setting STRIP_OUTER_ARRAY=TRUE will not resolve the issue of the column size exceeding the maximum allowed. References: COPY INTO Error during Structured Data Load: ??Max LOB size (16777216) exceeded????

NEW QUESTION 93
- (Topic 2)
What is the purpose of multi-cluster virtual warehouses?

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

Answer: C

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

NEW QUESTION 98
- (Topic 2)
The Snowflake Search Optimization Services supports improved performance of which kind of query?

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

Answer: C

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

NEW QUESTION 102
- (Topic 2)
What actions will prevent leveraging of the ResultSet cache? (Choose two.)

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

Answer: BD

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

NEW QUESTION 104
- (Topic 2)
How should a virtual warehouse be configured if a user wants to ensure that additional multi-clusters are resumed with no delay?

A. Configure the warehouse to a size larger than generally required
B. Set the minimum and maximum clusters to autoscale
C. Use the standard warehouse scaling policy
D. Use the economy warehouse scaling policy

Answer: A

Explanation:
To ensure that additional multi-clusters are resumed with no delay, a virtual warehouse should be configured to a size larger than generally required. This configuration allows for immediate availability of additional resources when needed, without waiting for new clusters to start up

NEW QUESTION 109
- (Topic 2)
Which data types are supported by Snowflake when using semi-structured data? (Choose two.)

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

Answer: AD

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

NEW QUESTION 112
- (Topic 2)
What are best practice recommendations for using the ACCOUNTADMIN system-defined role in Snowflake? (Choose two.)

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

Answer: AD

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

NEW QUESTION 114
- (Topic 2)
Which of the following significantly improves the performance of selective point lookup queries on a table?

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

Answer: D

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

NEW QUESTION 117
- (Topic 2)
Network policies can be set at which Snowflake levels? (Choose two.)

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

Answer: CE

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

NEW QUESTION 120
- (Topic 2)
A user created a new worksheet within the Snowsight Ul and wants to share this with teammates
How can this worksheet be shared?

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

Answer: C

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

NEW QUESTION 125
- (Topic 2)
How can a row access policy be applied to a table or a view? (Choose two.)

A. Within the policy DDL
B. Within the create table or create view DDL

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

Answer: AE

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

NEW QUESTION 129
- (Topic 2)
True or False: Snowpipe via REST API can only reference External Stages as source.

A. True
B. False

Answer: B

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

NEW QUESTION 131
- (Topic 2)
What type of query benefits the MOST from search optimization?

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

Answer: C

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

NEW QUESTION 134
- (Topic 2)
What impacts the credit consumption of maintaining a materialized view? (Choose two.)

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

Answer: CD

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

NEW QUESTION 135
- (Topic 2)
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

Answer: ABD

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.

NEW QUESTION 138
- (Topic 2)
Which minimum Snowflake edition allows for a dedicated metadata store?

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

Answer: B

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

NEW QUESTION 140
- (Topic 2)
A running virtual warehouse is suspended.
What is the MINIMUM amount of time that the warehouse will incur charges for when it is restarted?

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

Answer: B

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

NEW QUESTION 145
- (Topic 2)
What is the following SQL command used for? Select * from table(validate(t1, job_id => '_last'));

A. To validate external table files in table t1 across all sessions
B. To validate task SQL statements against table t1 in the last 14 days
C. To validate a file for errors before it gets executed using a COPY command
D. To return errors from the last executed COPY command into table t1 in the current session

Answer: D

Explanation:
The SQL command Select * from table(validate(t1, job_id => '_last')); is used to return errors from the last executed COPY command into table t1 in the current session. It checks the results of the most recent data load operation and provides details on any errors that occurred during that process1.

NEW QUESTION 148
- (Topic 2)
A user has unloaded data from a Snowflake table to an external stage.
Which command can be used to verify if data has been uploaded to the external stage named my_stage?

A. view @my_stage
B. list @my_stage
C. show @my_stage
D. display @my_stage

Answer: B

Explanation:
The list @my_stage command in Snowflake can be used to verify if data has been uploaded to an external stage named my_stage. This command provides a list of files that are present in the specified stage2.

NEW QUESTION 153
- (Topic 2)
Which statements are correct concerning the leveraging of third-party data from the Snowflake Data Marketplace? (Choose two.)

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

Answer: AD

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

NEW QUESTION 156
- (Topic 3)
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.

Answer: A

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.

NEW QUESTION 161
- (Topic 3)
Which parameter prevents streams on tables from becoming stale?

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

Answer: A

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

NEW QUESTION 165
- (Topic 3)
Using variables in Snowflake is denoted by using which SQL character?

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

Answer: C

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

NEW QUESTION 170
- (Topic 3)
Which of the following can be used when unloading data from Snowflake? (Choose two.)

A. When unloading semi-structured data, it is recommended that the STRIP_OUTER_ARRAY option be used.
B. Use the ENCODING file format option to change the encoding from the default UTF-8.
C. The OBJECT_CONSTRUCT function can be used to convert relational data to semi- structured data.
D. By using the SINGLE = TRUE parameter, a single file up to 5 GB in size can be exported to the storage layer.
E. Use the PARSE_JSON function to ensure structured data will be unloaded into the VARIANT data type.

Answer: CD

Explanation:
The OBJECT_CONSTRUCT function is used in Snowflake to create a JSON object from relational data, which is useful when unloading semi-structured data. The SINGLE = TRUE parameter is used when unloading data to ensure that the data is exported as a single file, which can be up to 5 GB in size. References: [COF-C02] SnowPro Core Certification Exam Study Guide

NEW QUESTION 172
- (Topic 3)
Which SQL command can be used to see the CREATE definition of a masking policy?

A. SHOW MASKING POLICIES
B. DESCRIBE MASKING POLICY
C. GET_DDL
D. LIST MASKING POLICIES

Answer: C

Explanation:
The SQL command GET_DDL can be used to retrieve the CREATE definition of a masking policy in Snowflake. This command generates the DDL statement required to recreate the masking policy

NEW QUESTION 175
- (Topic 3)
What internal stages are available in Snowflake? (Choose three.)

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

Answer: BCE

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

NEW QUESTION 176
- (Topic 3)
Which TABLE function helps to convert semi-structured data to a relational representation?

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

Answer: C

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

NEW QUESTION 181
- (Topic 3)
Which role has the ability to create and manage users and roles?

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

Answer: B

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

NEW QUESTION 186
- (Topic 3)
In which Snowflake layer does Snowflake reorganize data into its internal optimized, compressed, columnar format?

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

Answer: B

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

NEW QUESTION 188
- (Topic 3)
Which of the following is the Snowflake Account_Usage.Metering_History view used for?

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

Answer: A

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

NEW QUESTION 190
- (Topic 3)
What effect does WAIT_FOR_COMPLETION = TRUE have when running an ALTER WAREHOUSE command and changing the warehouse size?

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

Answer: D

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

NEW QUESTION 195
- (Topic 3)
A user needs to create a materialized view in the schema MYDB.MYSCHEMA. Which statements will provide this access?

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

Answer: A

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

NEW QUESTION 199
- (Topic 3)
What is the purpose of using the OBJECT_CONSTRUCT function with me COPY INTO command?

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

Answer: B

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

NEW QUESTION 201
- (Topic 3)
What MINIMUM privilege is required on the external stage for any role in the GET REST API to access unstructured data files using a file URL?

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

Answer: A

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

NEW QUESTION 204
- (Topic 3)
How can a Snowflake user optimize query performance in Snowflake? (Select TWO).

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

Answer: BC

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

NEW QUESTION 205
- (Topic 3)
Which file format will keep floating-point numbers from being truncated when data is unloaded?

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

Answer: D

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

NEW QUESTION 209
- (Topic 3)
If a virtual warehouse runs for 61 seconds, shuts down, and then restarts and runs for 30 seconds, for how many seconds is it billed?

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

Answer: D

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

NEW QUESTION 211
- (Topic 3)
Which features could be used to improve the performance of queries that return a small subset of rows from a large table? (Select TWO).

A. Search optimization service
B. Automatic clustering
C. Row access policies
D. Multi-cluster virtual warehouses
E. Secure views

Answer: AB

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

NEW QUESTION 213
- (Topic 3)
What service is provided as an integrated Snowflake feature to enhance Multi-Factor Authentication (MFA) support?

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

Answer: A

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

NEW QUESTION 217
- (Topic 3)
Which operations are handled in the Cloud Services layer of Snowflake? (Select TWO).

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

Answer: AE

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

NEW QUESTION 218
- (Topic 3)
What file formats does Snowflake support for loading semi-structured data? (Choose three.)

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

Answer: BDE

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

NEW QUESTION 220
- (Topic 3)
Which stages are used with the Snowflake PUT command to upload files from a local file system? (Choose three.)

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

Answer: BDF

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

NEW QUESTION 222
- (Topic 3)
A company needs to read multiple terabytes of data for an initial load as part of a Snowflake migration. The company can control the number and size of CSV extract files.
How does Snowflake recommend maximizing the load performance?

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

Answer: B

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

NEW QUESTION 226
- (Topic 3)
Which feature allows a user the ability to control the organization of data in a micro- partition?

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

Answer: C

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

NEW QUESTION 228
- (Topic 3)
Network policies can be applied to which of the following Snowflake objects? (Choose two.)

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

Answer: DE

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

NEW QUESTION 229
- (Topic 3)
If file format options are specified in multiple locations, the load operation selects which option FIRST to apply in order of precedence?

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

Answer: D

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

NEW QUESTION 234
- (Topic 3)
How can a Snowflake user access a JSON object, given the following table? (Select TWO).

A. src:salesperson.name
B. src:sa1esPerso
C. name
D. src:salesperson.Name
E. SRC:salesperson.name
F. SRC:salesperson.Name

Answer: AC

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

NEW QUESTION 237
- (Topic 3)
If queries start to queue in a multi-cluster virtual warehouse, an additional compute cluster starts immediately under what setting?

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

Answer: A

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

NEW QUESTION 240
- (Topic 3)
Which features make up Snowflake's column level security? (Select TWO).

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

Answer: BC

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

NEW QUESTION 243
- (Topic 3)
Which URL type allows users to access unstructured data without authenticating into Snowflake or passing an authorization token?

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

Answer: A

Explanation:
Pre-signed URLs in Snowflake allow users to access unstructured data without the need for authentication into Snowflake or passing an authorization token. These URLs are open and can be directly accessed or downloaded by any user or application, making them ideal for business intelligence applications or reporting tools that need to display unstructured file contents

NEW QUESTION 244
- (Topic 3)
The bulk data load history that is available upon completion of the COPY statement is stored where and for how long?

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

Answer: D

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

NEW QUESTION 245
- (Topic 3)
What can a Snowflake user do in the Activity section in Snowsight?

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

Answer: D

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

NEW QUESTION 250
- (Topic 3)
What action can a user take to address query concurrency issues?

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

Answer: C

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

NEW QUESTION 255
- (Topic 3)
When unloading data to an external stage, what is the MAXIMUM file size supported?

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

Answer: B

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

NEW QUESTION 259
- (Topic 3)
At what levels can a resource monitor be configured? (Select TWO).

A. Account
B. Database

C. Organization
D. Schema
E. Virtual warehouse

Answer: AE

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

NEW QUESTION 261
- (Topic 3)
A data provider wants to share data with a consumer who does not have a Snowflake account. The provider creates a reader account for the consumer following these steps:
* 1. Created a user called "CONSUMER"
* 2. Created a database to hold the share and an extra-small warehouse to query the data
* 3. Granted the role PUBLIC the following privileges: Usage on the warehouse, database, and schema, and SELECT on all the objects in the share
Based on this configuration what is true of the reader account?

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

Answer: B

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

NEW QUESTION 265
- (Topic 3)
What role is required to use Partner Connect?

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

Answer: A

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

NEW QUESTION 268
- (Topic 3)
Which Snowflake URL type is used by directory tables?

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

Answer: C

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

NEW QUESTION 272
- (Topic 3)
Which pages are included in the Activity area of Snowsight? (Select TWO).

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

Answer: DE

Explanation:
The Activity area of Snowsight includes the Query History page, which allows users to monitor and view details about queries executed in their account, including

performance data1. It also includes the Automatic Clustering History, which provides insights into the automatic clustering operations performed on tables2.

NEW QUESTION 274
- (Topic 3)
A user has a standard multi-cluster warehouse auto-scaling policy in place.
Which condition will trigger a cluster to shut-down?

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

Answer: D

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

NEW QUESTION 279
- (Topic 3)
Credit charges for Snowflake virtual warehouses are calculated based on which of the following considerations? (Choose two.)

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

Answer: CD

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

NEW QUESTION 282
- (Topic 3)
What is cached during a query on a virtual warehouse?

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

Answer: C

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

NEW QUESTION 285
- (Topic 3)
A Snowflake user executed a query and received the results. Another user executed the same query 4 hours later. The data had not changed.
What will occur?

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

Answer: A

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

NEW QUESTION 287
- (Topic 3)
Which feature is integrated to support Multi-Factor Authentication (MFA) at Snowflake?

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

Answer: B

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

NEW QUESTION 290
- (Topic 3)
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.

Answer: B

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

NEW QUESTION 294
- (Topic 4)
What is the purpose of the Snowflake SPLIT TO_TABLE function?

A. To count the number of characters in a string
B. To split a string into an array of sub-strings
C. To split a string and flatten the results into rows
D. To split a string and flatten the results into columns

Answer: C

Explanation:
The purpose of the Snowflake SPLIT_TO_TABLE function is to split a string based on a specified delimiter and flatten the results into rows. This table function is useful for transforming a delimited string into a set of rows that can be further processed or queried5.

NEW QUESTION 298
- (Topic 4)
Which statistics can be used to identify queries that have inefficient pruning? (Select TWO).

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

Answer: CD

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

NEW QUESTION 302
- (Topic 4)
A permanent table and temporary table have the same name, TBL1, in a schema. What will happen if a user executes select * from TBL1 ;?

A. The temporary table will take precedence over the permanent table.
B. The permanent table will take precedence over the temporary table.
C. An error will say there cannot be two tables with the same name in a schema.
D. The table that was created most recently will take precedence over the older table.

Answer: A

Explanation:
In Snowflake, if a temporary table and a permanent table have the same name within the same schema, the temporary table takes precedence over the permanent table within the session where the temporary table was created4.

NEW QUESTION 304
- (Topic 4)
A user wants to access files stored in a stage without authenticating into Snowflake. Which type of URL should be used?

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

Answer: D

Explanation:

A Pre-signed URL should be used to access files stored in a Snowflake stage without requiring authentication into Snowflake. Pre-signed URLs are simple HTTPS URLs that provide temporary access to a file via a web browser, using a pre-signed access token. The expiration time for the access token is configurable, and this type of URL allows users or applications to directly access or download the files without needing to authenticate into Snowflake5.
References: [COF-C02] SnowPro Core Certification Exam Study Guide

NEW QUESTION 306
- (Topic 4)
What will prevent unauthorized access to a Snowflake account from an unknown source?

A. Network policy
B. End-to-end encryption
C. Multi-Factor Authentication (MFA)
D. Role-Based Access Control (RBAC)

Answer: A

Explanation:
A network policy in Snowflake is used to restrict access to the Snowflake account from unauthorized or unknown sources. It allows administrators to specify allowed IP address ranges, thus preventing access from any IP addresses not listed in the policy1.

NEW QUESTION 311
- (Topic 4)
What type of query will benefit from the query acceleration service?

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

Answer: B

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

NEW QUESTION 316
- (Topic 4)
When enabling access to unstructured data, which URL permits temporary access to a staged file without the need to grant privileges to the stage or to issue access tokens?

A. File URL
B. Scoped URL
C. Relative URL
D. Pre-Signed URL

Answer: B

Explanation:
A Scoped URL permits temporary access to a staged file without the need to grant privileges to the stage or to issue access tokens. It provides a secure way to share access to files stored in Snowflake

NEW QUESTION 320
- (Topic 4)
Which type of loop requires a BREAK statement to stop executing?

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

Answer: B

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

NEW QUESTION 323
- (Topic 4)
Which object can be used with Secure Data Sharing?

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

Answer: A

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

NEW QUESTION 328
- (Topic 4)
What does a masking policy consist of in Snowflake?

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

Answer: A

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

NEW QUESTION 332
- (Topic 4)
What is the primary purpose of a directory table in Snowflake?

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

Answer: D

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

NEW QUESTION 335
- (Topic 4)
Which ACCOUNT_USAGE schema database role provides visibility into policy-related information?

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

Answer: B

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

NEW QUESTION 340
- (Topic 4)
Which parameter can be set at the account level to set the minimum number of days for which Snowflake retains historical data in Time Travel?

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

Answer: A

Explanation:
The parameter DATA_RETENTION_TIME_IN_DAYS can be set at the account level to define the minimum number of days Snowflake retains historical data for Time Travel1.

NEW QUESTION 341
- (Topic 4)
What is a directory table in Snowflake?

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

Answer: B

Explanation:
A directory table in Snowflake is an object layered on a stage that is used to store file-level metadata. It is not a separate database object but is conceptually

similar to an external table because it stores metadata about the data files in the stage5.

NEW QUESTION 343
- (Topic 4)
What does the LATERAL modifier for the FLATTEN function do?

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

Answer: C

Explanation:
The LATERAL modifier for the FLATTEN function allows joining information outside the object (such as other columns in the source table) with the flattened data, creating a lateral view that correlates with the preceding tables in the FROM clause2345. References: [COF-C02] SnowPro Core Certification Exam Study Guide

NEW QUESTION 345
- (Topic 4)
Which Snowflake object does not consume any storage costs?

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

Answer: C

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

NEW QUESTION 346
- (Topic 4)
Who can activate and enforce a network policy for all users in a Snowflake account? (Select TWO).

A. A user with an USERADMIN or higher role
B. A user with a SECURITYADMIN or higher role
C. A role that has been granted the ATTACH POLICY privilege
D. A role that has the NETWORK_POLICY account parameter set
E. A role that has the ownership of the network policy

Answer: BE

Explanation:
In Snowflake, a user with the SECURITYADMIN role or higher can activate and enforce a network policy for all users in an account. Additionally, a role that has ownership of the network policy can also activate and enforce it

NEW QUESTION 350
- (Topic 4)
While working with unstructured data, which file function generates a Snowflake-hosted file URL to a staged file using the stage name and relative file path as inputs?

A. GET_PRESIGNED_URL
B. GET_ABSOLUTE_PATH
C. BUILD_STAGE_FILE_URL
D. BUILD SCOPED FILE URL

Answer: C

Explanation:
The BUILD_STAGE_FILE_URL function generates a Snowflake-hosted file URL to a staged file using the stage name and relative file path as inputs2.

NEW QUESTION 355
- (Topic 4)
What happens to the objects in a reader account when the DROP MANAGED ACCOUNT command is executed?

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

Answer: A

Explanation:
When the DROP MANAGED ACCOUNT command is executed in Snowflake, it removes the managed account, including all objects created within the account, and access to the account is immediately restricted2.

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

NEW QUESTION 357
- (Topic 4)
Which Snowflake role can manage any object grant globally, including modifying and revoking grants?

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

Answer: D

Explanation:
The SECURITYADMIN role in Snowflake can manage any object grant globally, including modifying and revoking grants. This role has the necessary privileges to oversee and control access to all securable objects within the Snowflake environment4.

NEW QUESTION 362
- (Topic 4)
Which metadata table will store the storage utilization information even for dropped tables?

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

Answer: B

Explanation:
The TABLE_STORAGE_METRICS metadata table stores the storage utilization information, including for tables that have been dropped but are still incurring storage costs2.

NEW QUESTION 366
- (Topic 5)
What is the Fail-safe period for a transient table in the Snowflake Enterprise edition and higher?

A. 0 days
B. 1 day
C. 7 days
D. 14 days

Answer: A

Explanation:
The Fail-safe period for a transient table in Snowflake, regardless of the edition (including Enterprise edition and higher), is 0 days. Fail-safe is a data protection feature that provides additional retention beyond the Time Travel period for recovering data in case of accidental deletion or corruption. However, transient tables are designed for temporary or short-term use and do not benefit from the Fail-safe feature, meaning that once their Time Travel period expires, data cannot be recovered.
References:
? Snowflake Documentation: Understanding Fail-safe

NEW QUESTION 368
- (Topic 5)
Which Snowflow object does not consume and storage costs?

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

Answer: C

Explanation:
Temporary tables in Snowflake do not consume storage costs. They are designed for transient data that is needed only for the duration of a session. Data stored in temporary tables is held in the virtual warehouse's cache and does not persist beyond the session's lifetime, thereby not incurring any storage charges.
References:
? Snowflake Documentation: Temporary Tables

NEW QUESTION 371
- (Topic 5)
The effects of query pruning can be observed by evaluating which statistics? (Select TWO).

A. Partitions scanned
B. Partitions total
C. Bytes scanned
D. Bytes read from result
E. Bytes written

Answer: AC

Explanation:
Query pruning in Snowflake refers to the optimization technique where the system reduces the amount of data scanned by a query based on the query conditions. This typically involves skipping unnecessary data partitions that do not contribute to the query result. The effectiveness of this technique can be observed through:
? Option A: Partitions scanned. This statistic indicates how many data partitions were actually scanned as a result of query pruning, showing the optimization in action.
? Option C: Bytes scanned. This measures the volume of data physically read during query execution, and a reduction in this number indicates effective query pruning, as fewer bytes are read when unnecessary partitions are skipped.
Options B, D, and E do not directly relate to observing the effects of query pruning. "Partitions total" shows the total available, not the impact of pruning, while "Bytes read from result" and "Bytes written" relate to output rather than the efficiency of data scanning. References:Snowflake documentation on performance tuning and query optimization techniques, specifically how query pruning affects data access.

NEW QUESTION 373
- (Topic 5)
Which activities are included in the Cloud Services layer? {Select TWO).

A. Data storage
B. Dynamic data masking
C. Partition scanning
D. User authentication
E. Infrastructure management

Answer: DE

Explanation:
The Cloud Services layer in Snowflake is responsible for a wide range of services that facilitate the management and use of Snowflake, including:
? D. User authentication: This service handles identity and access management, ensuring that only authorized users can access Snowflake resources.
? E. Infrastructure management: This service manages the allocation and scaling of resources to meet user demands, including the management of virtual warehouses, storage, and the orchestration of query execution.
These services are part of Snowflake's fully managed, cloud-based architecture, which abstracts and automates many of the complexities associated with data warehousing. References:
? Snowflake Documentation: Overview of Snowflake Cloud Services

NEW QUESTION 374
- (Topic 5)
How does Snowflake describe its unique architecture?

A. A single-cluster shared data architecture using a central data repository and massively parallel processing (MPP)
B. A multi-duster shared nothing architecture using a soloed data repository and massively parallel processing (MPP)
C. A single-cluster shared nothing architecture using a sliced data repository and symmetric multiprocessing (SMP)
D. A multi-cluster shared nothing architecture using a siloed data repository and symmetric multiprocessing (SMP)

Answer: A

Explanation:
Snowflake's unique architecture is described as a multi-cluster, shared data architecture that leverages massively parallel processing (MPP). This architecture separates compute and storage resources, enabling Snowflake to scale them independently. It does not use a single cluster or rely solely on symmetric multiprocessing (SMP); rather, it uses acombination of shared-nothing architecture for compute clusters (virtual warehouses) and a centralized storage layer for data, optimizing for both performance and scalability. References:
? Snowflake Documentation: Snowflake Architecture Overview

NEW QUESTION 379
- (Topic 5)
Which data formats are supported by Snowflake when unloading semi-structured data? (Select TWO).

A. Binary file in Avro
B. Binary file in Parquet
C. Comma-separated JSON
D. Newline Delimited JSON
E. Plain text file containing XML elements

Answer: BD

Explanation:
Snowflake supports a variety of file formats for unloading semi-structured data, among which Parquet and Newline Delimited JSON (NDJSON) are two widely used formats.
? B. Binary file in Parquet: Parquet is a columnar storage file format optimized for large-scale data processing and analysis. It is especially suited for complex nested data structures.
? D. Newline Delimited JSON (NDJSON): This format represents JSON records separated by newline characters, facilitating the storage and processing of multiple, separate JSON objects in a single file.
These formats are chosen for their efficiency and compatibility with data analytics tools and ecosystems, enabling seamless integration and processing of exported data.
References:
? Snowflake Documentation: Data Unloading

NEW QUESTION 384
- (Topic 5)
Which statement accurately describes Snowflake's architecture?

A. It uses a local data repository for all compute nodes in the platform.
B. It is a blend of shared-disk and shared-everything database architectures.
C. It is a hybrid of traditional shared-disk and shared-nothing database architectures.
D. It reorganizes loaded data into internal optimized, compressed, and row-based format.

Answer: C

Explanation:
Snowflake's architecture is unique in that it combines elements of both traditional shared-disk and shared-nothing database architectures. This hybrid approach allows Snowflake to offer the scalability and performance benefits of a shared-nothing architecture (with compute and storage separated) while maintaining the simplicity and flexibility of a shared-disk architecture in managing data across all nodes in the system. This results in an architecture that provides on-demand scalability, both vertically and horizontally, without sacrificing performance or data cohesion.
References:
? Snowflake Documentation: Snowflake Architecture

NEW QUESTION 387
- (Topic 5)
What happens when a network policy includes values that appear in both the allowed and blocked IP address list?

A. Those IP addresses are allowed access to the Snowflake account as Snowflake applies the allowed IP address list first.
B. Those IP addresses are denied access lei the Snowflake account as Snowflake applies the blocked IP address list first.
C. Snowflake issues an alert message and adds the duplicate IP address values lo both 'he allowed and blocked IP address lists.
D. Snowflake issues an error message and adds the duplicate IP address values to both the allowed and blocked IP address list

Answer: B

Explanation:
In Snowflake, when setting up a network policy that specifies both allowed and blocked IP address lists, if an IP address appears in both lists, access from that IP address will be denied. The reason is that Snowflake prioritizes security, and the presence of an IP address in the blocked list indicates it should not be allowed regardless of its presence in the allowed list. This ensures that access controls remain stringent and that any potentially unsafe IP addresses are not inadvertently permitted access.
References:
? Snowflake Documentation: Network Policies

NEW QUESTION 388
- (Topic 5)
Which role has the ability to create a share from a shared database by default?

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

Answer: A

Explanation:
By default, the ACCOUNTADMIN role in Snowflake has the ability to create a share from a shared database. This role has the highest level of access within a Snowflake account, including the management of all aspects of the account, such as users, roles, warehouses, and databases, as well as the creation and management of shares for secure data sharing with other Snowflake accounts.
References:
? Snowflake Documentation: Roles

NEW QUESTION 393
- (Topic 5)
What are characteristics of reader accounts in Snowflake? (Select TWO).

A. Reader account users cannot add new data to the account.
B. Reader account users can share data to other reader accounts.
C. A single reader account can consume data from multiple provider accounts.
D. Data consumers are responsible for reader account setup and data usage costs.
E. Reader accounts enable data consumers to access and query data shared by the provider.

Answer: AE

Explanation:
Characteristics of reader accounts in Snowflake include:
? A. Reader account users cannot add new data to the account: Reader accounts are intended for data consumption only. Users of these accounts can query and analyze the data shared with them but cannot upload or add new data to the account.
? E. Reader accounts enable data consumers to access and query data shared by the provider: One of the primary purposes of reader accounts is to allow data consumers to access and perform queries on the data shared by another Snowflake account, facilitating secure and controlled data sharing.
References:
? Snowflake Documentation: Reader Accounts

NEW QUESTION 396
- (Topic 5)
How are network policies defined in Snowflake?

A. They are a set of rules that define the network routes within Snowflake.
B. They are a set of rules that dictate how Snowflake accounts can be used between multiple users.

C. They are a set of rules that define how data can be transferred between different Snowflake accounts within an organization.
D. They are a set of rules that control access to Snowflake accounts by specifying the IP addresses or ranges of IP addresses that are allowed to connectto Snowflake.

Answer: D

Explanation:
Network policies in Snowflake are defined as a set of rules that manage the network-level access to Snowflake accounts. These rules specify which IP addresses or IP ranges are permitted to connect to Snowflake, enhancing the security of Snowflake accounts by preventing unauthorized access. Network policies are an essential aspect of Snowflake's security model, allowing administrators to enforce access controls based on network locations.
References:
? Snowflake Documentation: Network Policies

NEW QUESTION 398
- (Topic 5)
Why would a Snowflake user decide to use a materialized view instead of a regular view?

A. The base tables do not change frequently.
B. The results of the view change often.
C. The query is not resource intensive.
D. The query results are not used frequently.

Answer: A

Explanation:
A Snowflake user would decide to use a materialized view instead of a regular view primarily when the base tables do not change frequently. Materialized views store the result of the view query and update it as the underlying data changes, making them ideal for situations where the data is relatively static and query performance is critical. By precomputing and storing the query results, materialized views can significantly reduce query execution times for complex aggregations, joins, and calculations.
References:
? Snowflake Documentation: Materialized Views

NEW QUESTION 400
- (Topic 5)
Which function should be used to insert JSON format string data inot a VARIANT field?

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

Answer: C

Explanation:
To insert JSON formatted string data into a VARIANT field in Snowflake, the correct function to use isPARSE_JSON. ThePARSE_JSONfunction is specifically designed to interpret a JSON formatted string and convert it into a VARIANT type, which is Snowflake's flexible format for handling semi-structured data like JSON, XML, and Avro. This function is essential for loading and querying JSON data within Snowflake, allowing users to store and manage JSON data efficiently while preserving its structure for querying purposes. This function's usage and capabilities are detailed in the Snowflake documentation, providing users with guidance on how to handle semi-structured data effectively within their Snowflake environments.
References:
? Snowflake Documentation: PARSE_JSON

NEW QUESTION 401
- (Topic 5)
What will happen if a Snowflake user increases the size of a suspended virtual warehouse?

A. The provisioning of new compute resources for the warehouse will begin immediately.
B. The warehouse will remain suspended but new resources will be added to the query acceleration service.
C. The provisioning of additional compute resources will be in effect when the warehouse is next resumed.
D. The warehouse will resume immediately and start to share the compute load with other running virtual warehouses.

Answer: C

Explanation:
When a Snowflake user increases the size of a suspended virtual warehouse, the changes to compute resources are queued but do not take immediate effect. The provisioning of additional compute resources occurs only when the warehouse is resumed. This ensures that resources are allocated efficiently, aligning with Snowflake's commitment to cost-effective and on-demand scalability.
References:
? Snowflake Documentation: Virtual Warehouses

NEW QUESTION 405
- (Topic 5)
Which Snowflake mechanism is used to limit the number of micro-partitions scanned by a query?

A. Caching
B. Cluster depth
C. Query pruning
D. Retrieval optimization

Answer: C

Explanation:
Query pruning in Snowflake is the mechanism used to limit the number of micro-partitions scanned by a query. By analyzing the filters and conditions applied in a query, Snowflake can skip over micro-partitions that do not contain relevant data, thereby reducing the amount of data processed and improving query performance. This technique is particularly effective for large datasets and is a key component of Snowflake's performance optimization features.
References:
? Snowflake Documentation: Query Performance Optimization

NEW QUESTION 410
- (Topic 5)
How can a user get the MOST detailed information about individual table storage details in Snowflake?

A. SHOW TABLES command
B. SHOW EXTERNAL TABLES command
C. TABLES view
D. TABLE STORAGE METRICS view

Answer: D

Explanation:
To obtain the most detailed information about individual table storage details in Snowflake, the TABLE STORAGE METRICS view is the recommended option. This view provides comprehensive metrics on storage usage, including data size, time travel size, fail-safe size, and other relevant storage metrics for each table. This level of detail is invaluable for monitoring, managing, and optimizing storage costs and performance. References:
? Snowflake Documentation: Information Schema

NEW QUESTION 414
- (Topic 5)
What is the only supported character set for loading and unloading data from all supported file formats?

A. UTF-8
B. UTF-16
C. ISO-8859-1
D. WINDOWS-1253

Answer: A

Explanation:
UTF-8 is the only supported character set for loading and unloading data from all supported file formats in Snowflake. UTF-8 is a widely used encoding that supports a large range of characters from various languages, making it suitable forinternationalization and ensuring data compatibility across different systems and platforms.
References:
? Snowflake Documentation: Data Loading and Unloading

NEW QUESTION 418
- (Topic 6)
Which Snowflake function and command combination should be used to convert rows in a relational table to a single VARIANT column, and unload the rows Into a file in JSON format? (Select TWO).

A. PUT
B. GET
C. COPY
D. EXPORT
E. OBJECT CONSTRUCT

Answer: CE

Explanation:
To convert rows in a relational table to a single VARIANT column and unload the rows into a file in JSON format, you can use the COPY command in combination with the OBJECT_CONSTRUCT function. The OBJECT_CONSTRUCT function converts the row into a JSON object stored in a VARIANT column, and the COPY command can then be used to unload this data into a JSON file.
References:
? Snowflake Documentation: OBJECT_CONSTRUCT
? Snowflake Documentation: COPY INTO <location> Top of Form
Bottom of Form

NEW QUESTION 419
- (Topic 6)
Who can create network policies within Snowflake? (Select TWO).

A. SYSADMIN only
B. ORCADMIN only
C. SECURITYADMIN or higher roles
D. A role with the CREATE NETWORK POLICY privilege
E. A role with the CREATE SECURITY INTEGRATION privilege

Answer: CD

Explanation:

In Snowflake, network policies define the allowed IP address ranges from which users can connect to Snowflake, enhancing security by restricting access based on network location. The creation and management of network policies require sufficient privileges. Specifically, a user with theSECURITYADMINrole or any role with higher privileges, such asACCOUNTADMIN, can create network policies. Additionally, a custom role can be granted theCREATE NETWORK POLICYprivilege, enabling users assigned to that role to also create network policies. This approach allows for flexible and secure management of network access to Snowflake.References:Snowflake Documentation on Network Policies

NEW QUESTION 424
- (Topic 6)
What Snowflake recommendation is designed to ensure that staged data is only loaded once"?

A. Partitioning staged data files
B. Loading only the most recently-staged data files
C. Removing data files after loading
D. Identifying and removing duplicates after each data load

Answer: C

Explanation:
Snowflake recommends removing data files from the staging area after they have been loaded into the target table. This practice ensures that the data is only loaded once and prevents accidental reloading of the same data. By removing the files, you eliminate the risk of duplicate data loads.
? Stage the Data:Upload the data files to a Snowflake stage (internal or external).
? Load the Data:Use theCOPY INTOcommand to load the data from the stage into the Snowflake table.
? Remove the Data Files:After successfully loading the data, remove the data files from the stage using theREMOVEcommand.
References:
? Snowflake Documentation: Loading Data into Snowflake
? Snowflake Documentation: Staging Data Files
? Snowflake Documentation: COPY INTO Command

NEW QUESTION 426
- (Topic 6)
Which MINIMUM set of privileges is required to temporarily bypass an active network policy by configuring the user object property MINS_TO_BYPASS_NETWORK_POLICY?

A. Only while in the ACCOUNTADMIH role
B. Only while in the securityadmin role
C. Only the role with the ownership privilege on the network policy
D. Only Snowflake Support can set the value for this object property

Answer: A

Explanation:
To temporarily bypass an active network policy by configuring the user object propertyMINS_TO_BYPASS_NETWORK_POLICY, the minimum set of privileges required is having the ACCOUNTADMIN role. This role has the necessary privileges to make such changes, including modifying user properties that affect network policies.
References:
? Snowflake Documentation: Network Policy Management

NEW QUESTION 429
- (Topic 6)
Which security models are used in Snowflake to manage access control? (Select TWO).

A. Discretionary Access Control (DAC)
B. Identity Access Management (1AM)
C. Mandatory Access Control (MAC)
D. Role-Based Access Control (RBAC)
E. Security Assertion Markup Language (SAML)

Answer: AD

Explanation:
Snowflake uses both Discretionary Access Control (DAC) and Role-Based Access Control (RBAC) to manage access control. DAC allows object owners to grant access privileges to other users. RBAC assigns permissions to roles, and roles are then granted to users, making it easier to manage permissions based on user roles within the organization.
References:
? Snowflake Documentation: Access Control in Snowflake

NEW QUESTION 432
- (Topic 6)
Which key access control concept does Snowflake descibe as a defined level of access to an object?

A. Grant
B. Privilege
C. Role
D. Session

Answer: B

Explanation:
In Snowflake, the term "privilege" refers to a defined level of access to an object. Privileges are specific actions that roles can perform on securable objects in

Snowflake, such as tables, views, warehouses, databases, and schemas. These privileges are granted to roles and can be further granted to users through their roles, forming the basis of Snowflake??s access control framework.References:Snowflake Documentation on Access Control Privileges

NEW QUESTION 434
- (Topic 6)
Which function should be used to find the query ID of the second query executed in a current session?

A. Select LAST_QUERY_ID(-2)
B. Select LAST_QUERY_ID(2)
C. Select LAST_QUERY_ID(1)
D. Select LAST_QUERY_ID(2)

Answer: A

Explanation:
The correct function to find the query ID of the second query executed in the current session isSELECT LAST_QUERY_ID(-2). TheLAST_QUERY_IDfunction returns the query ID for the most recent query executed in the session when called with no arguments. When used with an argument, it can retrieve the ID of previous queries within the same session, where-2would reference the second most recent query executed.
References:
? There's a clarification needed here; Snowflake's documentation indicatesLAST_QUERY_ID()function does not accept arguments. It returns the ID of the last query executed in the session. To find the query ID of the second last executed query, users typically need to track query IDs manually or use session history views.

NEW QUESTION 438
- (Topic 6)
How can staged files be removed during data loading once the files have loaded successfully?

A. Use the DROP command
B. Use the purge copy option.
C. Use the FORCE = TRUE parameter
D. Use the LOAD UNCERTAIN FILES copy option.

Answer: B

Explanation:
To remove staged files during data loading after they have been successfully loaded, thePURGEcopy option is used in Snowflake.
? PURGE Option: This option automatically deletes files from the stage after they
have been successfully copied into the target table.
? Usage: FROM @my_stage
FILE_FORMAT = (type = 'csv') PURGE = TRUE;
References:
? Snowflake Documentation on COPY INTO

NEW QUESTION 440
- (Topic 6)
What can be used to process unstructured data?

A. External tables
B. The copy into <table> command
C. External functions
D. Snowpipe

Answer: C

Explanation:
To process unstructured data in Snowflake, external functions can be used.
? External Functions: These allow you to call external services and processing engines from within Snowflake SQL. External functions can be used to handle complex processing tasks that are not natively supported by Snowflake, including those involving unstructured data.
? Implementation: You define an external function in Snowflake that points to an external processing service (e.g., AWS Lambda, Google Cloud Functions).
References:
? Snowflake Documentation on External Functions

NEW QUESTION 443
- (Topic 6)
What are characteristics of Snowflake network policies? (Select TWO).

A. They can be set for any Snowflake Edition.
B. They can be applied to roles.
C. They restrict or enable access to specific IP addresses.
D. They are activated using ALTER DATABASE SQL commands.
E. They can only be managed using the ORGADMIN role.

Answer: AC

Explanation:
Snowflake network policies are a security feature that allows administrators to control access to Snowflake by specifying allowed and blocked IP address ranges. These policies apply to all editions of Snowflake, making them widely applicable across different Snowflake environments. They are specifically designed to restrict or enable accessbased on the originating IP addresses of client requests, adding an extra layer of security. Network policies are not applied to roles but are set at the account or user level. They are not activated usingALTER DATABASESQL commands but are managed throughALTER ACCOUNTorALTER NETWORK

POLICYcommands. The management of network policies does not exclusively require theORGADMINrole; instead, they can be managed by users with the necessary privileges on the account.
Reference: Snowflake Documentation on Network Policies:
https://docs.snowflake.com/en/user-guide/network-policies.html

NEW QUESTION 448
- (Topic 6)
Which roles can make grant decisions to objects within a managed access schema? (Select TWO)

A. ACCOUNTADMIN
B. SECURITYADMIN
C. SYSTEMADMIN
D. ORGADMIN
E. USERADMIN

Answer: AB

Explanation:
? Managed Access Schemas:These are a special type of schema designed for fine- grained access control in Snowflake.
? Roles with Grant Authority:
? Important Note:The ORGADMIN role focuses on organization-level management, not object access control.

NEW QUESTION 449
- (Topic 6)
Which Snowflake database object can be shared with other accounts?

A. Tasks
B. Pipes
C. Secure User-Defined Functions (UDFs)
D. Stored Procedures

Answer: C

Explanation:
In Snowflake, Secure User-Defined Functions (UDFs) can be shared with other accounts using Snowflake's data sharing feature. This allows different Snowflake accounts to securely execute the UDFs without having direct access to the underlying data the functions operate on, ensuring privacy and security. The sharing is facilitated through shares created in Snowflake, which can contain Secure UDFs along with other database objects like tables and views.References:Snowflake Documentation on Data Sharing and Secure UDFs

NEW QUESTION 451
- (Topic 6)
Which data type can be used for floating-point numbers without losing precision?

A. BINARY
B. VARIANT
C. INTEGER
D. DOUBLE

Answer: D

Explanation:
In Snowflake, theDOUBLEdata type is used for floating-point numbers and can represent a wide range of values without losing precision. This data type is ideal for storing numerical values that require decimal precision.
? Define the Column:When creating a table, specify the column with theDOUBLEdata type to store floating-point numbers. CREATE TABLE example_table (
id INTEGER,
value DOUBLE
);
? Insert Data:Insert floating-point numbers into theDOUBLEcolumn. INSERT INTO example_table (id, value) VALUES (1, 123.456); References:
? Snowflake Documentation: Data Types
? Snowflake Documentation: Numeric Data Types

NEW QUESTION 454
- (Topic 6)
When should a stored procedure be created with caller's rights?

A. When the caller needs to be prevented from viewing the source code of the stored procedure
B. When the caller needs to run a statement that could not execute outside of the stored procedure
C. When the stored procedure needs to run with the privileges of the role that called the stored procedure
D. When the stored procedure needs to operate on objects that the caller does not have privileges on

Answer: C

Explanation:
Stored procedures in Snowflake can be created with either 'owner's rights' or 'caller's rights'. A stored procedure created with caller's rights executes with the privileges of the role that calls the procedure, not the privileges of the role that owns the procedure. This is particularly useful in scenarios where the procedure needs to perform operations that depend on the caller's access permissions, ensuring that the procedure can only access objects that the caller is authorized to access.
Reference: Snowflake's official documentation on stored procedures, specifically the section on execution context (caller's rights vs. owner's rights), provides detailed guidance on when to use caller's rights: https://docs.snowflake.com/en/sql-reference/stored-procedures-usage.html#caller-s-rights-vs-owner-s-rights

NEW QUESTION 456
- (Topic 6)
While unloading data into a stage, how can the user ensure that the output will be a single file?

A. Use the copy option files=single.
B. Use the COPY Option SINGLE=TRUE .
C. Use the get option SINGLE-TRUE.
D. Use the GET option FILES-SINGLE.

Answer: B

Explanation:
To ensure that the output will be a single file when unloading data into a stage, you should use the COPY option SINGLE=TRUE. This option specifies that the result of the COPY INTO command should be written to a single file, rather than multiple files.
References:
? Snowflake Documentation: COPY INTO <location>

NEW QUESTION 459
- (Topic 6)
Masking policies are created at what level in Snowflake?

A. Table
B. Column
C. Schema
D. Database

Answer: B

Explanation:
Masking policies in Snowflake are created and applied at the column level. These policies are used to obfuscate sensitive data by masking the values in a specific column. Masking policies can enforce data privacy and protection measures by dynamically masking the data based on the role of the user querying the data.
References:
? Snowflake Documentation: Dynamic Data Masking

NEW QUESTION 461
- (Topic 6)
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.

Answer: DE

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:
? Application of Masking Policies:
References:
? Snowflake Documentation: Column-Level Security
? Snowflake Documentation: Dynamic Data Masking

NEW QUESTION 464
- (Topic 6)
What information is stored in the ACCESS_HISTORY view?

A. History of the files that have been loaded into Snowflake
B. Names and owners of the roles that are currently enabled in the session
C. Query details such as the objects included and the user who executed the query
D. Details around the privileges that have been granted for all objects in an account

Answer: D

Explanation:
Details around the privileges that have been granted for all objects in an account. The ACCESS_HISTORYview in Snowflake provides a comprehensive log of access control changes, including grants and revocations of privileges on all securable objects within the account. This information is crucial for auditing and monitoring the security posture of your Snowflake environment.
Here's how to understand and use theACCESS_HISTORYview:
? Purpose of ACCESS_HISTORY View:It is designed to track changes in access controls, such as when a user or role is granted or revoked privileges on various Snowflake objects. This includes tables, schemas, databases, and more.
? Querying ACCESS_HISTORY:To access this view, you can use the following SQL query pattern:
SELECT*FROMSNOWFLAKE.ACCOUNT_USAGE.ACCESS_HISTORYWHEREEVENT_T YPE ='GRANT'OREVENT_TYPE ='REVOKE';
? Interpreting the Results:The results from theACCESS_HISTORYview include the
object type, the specific privilege granted or revoked, the grantee (who received or lost the privilege), and the timestamp of the event. This data is invaluable for audits and compliance checks.
Reference:For detailed information on theACCESS_HISTORYview and how to interpret its data, please visit the official Snowflake documentation:
https://docs.snowflake.com/en/sql- reference/account-usage/access_history.html

NEW QUESTION 465
- (Topic 6)
Which Snowflake table supports unstructured data?

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

Answer: D

Explanation:
While Snowflake primarily deals with structured and semi-structured data, it also has the capability to handle unstructured data. Unstructured data can be stored in Snowflake using variants of SQL data types in tables, which can be permanent tables. These permanent tables, while traditionally used for structured or semi-structured data (like JSON, Avro, or Parquet), can also accommodate unstructured data in the form of binary formats or strings, offering flexibility in data storage and analysis. However, the management and querying of unstructured data in Snowflake may require additional considerations compared to structured data.References:Snowflake Documentation on Data Types

NEW QUESTION 466
.....

Thank You for Trying Our Product

We offer two products:

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

2nd - Questons and Answers in PDF Format

COF-C02 Practice Exam Features:

* COF-C02 Questions and Answers Updated Frequently

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

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

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