

COF-C02 Dumps

SnowPro Core Certification Exam (COF-C02)

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



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

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

Answer: CE

Explanation:

Snowflake's Continuous Data Protection (CDP) encompasses a set of features that help protect data stored in Snowflake against human error, malicious acts, and software failure. Time Travel allows users to access historical data (i.e., data that has been changed or deleted) for a defined period, enabling querying and restoring of data. Fail-safe is an additional layer of data protection that provides a recovery option in the event of significant data loss or corruption, which can only be performed by Snowflake. References:

? Continuous Data Protection | Snowflake Documentation¹

? Data Storage Considerations | Snowflake Documentation²

? Snowflake SnowPro Core Certification Study Guide³

? Snowflake Data Cloud Glossary

<https://docs.snowflake.com/en/user-guide/data-availability.html>

NEW QUESTION 3

- (Topic 1)

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

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

Answer: BD

Explanation:

Snowflake's architecture is designed to handle high concurrency workloads through several features, two of which are particularly effective:

? B. The use of warehouse auto scaling: This feature allows Snowflake to automatically adjust the compute resources allocated to a virtual warehouse in response to the workload. If there is an increase in concurrent queries, Snowflake can scale up the resources to maintain performance.

? D. Use of multi-clustered warehouses: Multi-clustered warehouses enable Snowflake to run multiple clusters of compute resources simultaneously. This allows for the distribution of queries across clusters, thereby reducing the load on any single cluster and improving the system's ability to handle a high number of concurrent queries.

These features ensure that Snowflake can manage varying levels of demand without manual intervention, providing a seamless experience even during peak usage. References:

? Snowflake Documentation on Virtual Warehouses

? SnowPro® Core Certification Study Guide

NEW QUESTION 4

- (Topic 1)

In which scenarios would a user have to pay Cloud Services costs? (Select TWO).

- A. Compute Credits = 50 Credits Cloud Services = 10
- B. Compute Credits = 80 Credits Cloud Services = 5
- C. Compute Credits = 10 Credits Cloud Services = 9
- D. Compute Credits = 120 Credits Cloud Services = 10
- E. Compute Credits = 200 Credits Cloud Services = 26

Answer: AE

Explanation:

In Snowflake, Cloud Services costs are incurred when the Cloud Services usage exceeds 10% of the compute usage (measured in credits). Therefore, scenarios A and E would result in Cloud Services charges because the Cloud Services usage is more than 10% of the compute credits used.

References:

? [COF-C02] SnowPro Core Certification Exam Study Guide

? Snowflake's official documentation on billing and usage1

NEW QUESTION 5

- (Topic 1)

True or False: When you create a custom role, it is a best practice to immediately grant that role to ACCOUNTADMIN.

A. True

B. False

Answer: B

Explanation:

The ACCOUNTADMIN role is the most powerful role in Snowflake and should be limited to a select number of users within an organization. It is responsible for account-level configurations and should not be used for day-to-day object creation or management. Granting a custom role to ACCOUNTADMIN could inadvertently give broad access to users with this role, which is not a recommended security practice.

Reference: <https://docs.snowflake.com/en/user-guide/security-access-control-considerations.html>

NEW QUESTION 6

- (Topic 1)

Which cache type is used to cache data output from SQL queries?

A. Metadata cache

B. Result cache

C. Remote cache

D. Local file cache

Answer: B

Explanation:

The Result cache is used in Snowflake to cache the data output from SQL queries. This feature is designed to improve performance by storing the results of queries for a period of time. When the same or similar query is executed again, Snowflake can retrieve the result from this cache instead of re-computing the result, which saves time and computational resources.

References:

? Snowflake Documentation on Query Results Cache

? SnowPro® Core Certification Study Guide

NEW QUESTION 7

- (Topic 1)

A developer is granted ownership of a table that has a masking policy. The developer's role is not able to see the masked data. Will the developer be able to modify the table to read the masked data?

A. Yes, because a table owner has full control and can unset masking policies.

B. Yes, because masking policies only apply to cloned tables.

C. No, because masking policies must always reference specific access roles.

D. No, because ownership of a table does not include the ability to change masking policies

Answer: D

Explanation:

Even if a developer is granted ownership of a table with a masking policy, they will not be able to modify the table to read the masked data if their role does not have the necessary permissions. Ownership of a table does not automatically confer the ability to alter masking policies, which are designed to protect sensitive data. Masking policies are applied at the schema level and require specific privileges to modify.

References:

? [COF-C02] SnowPro Core Certification Exam Study Guide

? Snowflake Documentation on Masking Policies

NEW QUESTION 8

- (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 9

- (Topic 1)

Query compilation occurs in which architecture layer of the Snowflake Cloud Data Platform?

- A. Compute layer
- B. Storage layer
- C. Cloud infrastructure layer
- D. Cloud services layer

Answer: D

Explanation:

Query compilation in Snowflake occurs in the Cloud Services layer. This layer is responsible for coordinating and managing all aspects of the Snowflake service, including authentication, infrastructure management, metadata management, query parsing and optimization, and security. By handling these tasks, the Cloud Services layer enables the Compute layer to focus on executing queries, while the Storage layer is dedicated to persistently storing data.

References:

? [COF-C02] SnowPro Core Certification Exam Study Guide

? Snowflake Documentation on Snowflake Architecture1

NEW QUESTION 10

- (Topic 1)

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

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

Answer: D

Explanation:

The size of the virtual warehouse for a task can be configured to handle concurrency automatically using a Multi-cluster warehouse (MCW). This is because tasks are designed to run their body on a schedule, and MCW allows for scaling compute resources to match the task's execution needs without manual intervention.

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

NEW QUESTION 10

- (Topic 1)

What is the purpose of an External Function?

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

Answer: A

Explanation:

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

NEW QUESTION 12

- (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 17

- (Topic 1)

Which of the following indicates that it may be appropriate to use a clustering key for a table? (Select TWO).

- A. The table contains a column that has very low cardinality
- B. DML statements that are being issued against the table are blocked
- C. The table has a small number of micro-partitions
- D. Queries on the table are running slower than expected
- E. The clustering depth for the table is large

Answer: DE

Explanation:

A clustering key in Snowflake is used to co-locate similar data within the same micro-partitions to improve query performance, especially for large tables where data is not naturally ordered or has become fragmented due to extensive DML operations. The appropriate use of a clustering key can lead to improved scan efficiency and better column compression, resulting in faster query execution times.

The indicators that it may be appropriate to use a clustering key for a table include:

? D. Queries on the table are running slower than expected: This can happen when the data in the table is not well-clustered, leading to inefficient scans during query execution.

? E. The clustering depth for the table is large: A large clustering depth indicates that the table's data is spread across many micro-partitions, which can degrade query performance as more data needs to be scanned.

References:

- ? Snowflake Documentation on Clustering Keys & Clustered Tables
- ? Snowflake Documentation on SYSTEM\$CLUSTERING_INFORMATION
- ? Stack Overflow discussion on cluster key selection in Snowflake

NEW QUESTION 18

- (Topic 1)

Which Snowflake object enables loading data from files as soon as they are available in a cloud storage location?

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

Answer: A

Explanation:

In Snowflake, a Pipe is the object designed to enable the continuous, near- real-time loading of data from files as soon as they are available in a cloud storage location. Pipes use Snowflake's COPY command to load data and can be associated with a Stage object to monitor for new files. When new data files appear in the stage, the pipe automatically loads the data into the target table.

References:

- ? Snowflake Documentation on Pipes
- ? SnowPro® Core Certification Study Guide <https://docs.snowflake.com/en/user-guide/data-load-snowpipe-intro.html>

NEW QUESTION 23

- (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 25

- (Topic 1)

What transformations are supported in a CREATE PIPE ... AS COPY ... FROM (....) statement? (Select TWO.)

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

Answer: AD

Explanation:

In a CREATE PIPE ... AS COPY ... FROM (....) statement, the supported transformations include filtering data using an optional WHERE clause and omitting columns. The WHERE clause allows for the specification of conditions to filter the data that is being loaded, ensuring only relevant data is inserted into the table. Omitting columns enables the exclusion of certain columns from the data load, which can be useful when the incoming data contains more columns than are needed for the target table.

References:

- ? [COF-C02] SnowPro Core Certification Exam Study Guide
- ? Simple Transformations During a Load1

NEW QUESTION 28

- (Topic 1)

A user has unloaded data from Snowflake to a stage

Which SQL command should be used to validate which data was loaded into the stage?

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

Answer: A

Explanation:

The list command in Snowflake is used to validate and display the list of files in a specified stage. When a user has unloaded data to a stage, running the list @file stage command will show all the files that have been uploaded to that stage, allowing the user to verify the data that was unloaded.

References:

- ? Snowflake Documentation on Stages
- ? SnowPro® Core Certification Study Guide

NEW QUESTION 33

- (Topic 1)

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

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

Answer: A

Explanation:

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

NEW QUESTION 38

- (Topic 1)

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

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

Answer: B

Explanation:

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

? Snowflake Documentation on CREATE MATERIALIZED VIEW1.

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

NEW QUESTION 42

- (Topic 1)

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

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

Answer: D

Explanation:

To generate a report listing all Snowflake logins within the last 90 days, the ACCOUNT_USAGE.LOGIN_HISTORY view should be used. This view provides information about login attempts, including successful and unsuccessful logins, and is suitable for security audits.

NEW QUESTION 43

- (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 47

- (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. DESCRIBE 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 50

- (Topic 1)

Which command can be used to stage local files from which Snowflake interface?

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

Answer: A

Explanation:

SnowSQL is the command-line client for Snowflake that allows users to execute SQL queries and perform all DDL and DML operations, including staging files for bulk data loading. It is specifically designed for scripting and automating tasks. References:

? SnowPro Core Certification Exam Study Guide

? Snowflake Documentation on SnowSQL <https://docs.snowflake.com/en/user-guide/snowsql-use.html>

NEW QUESTION 52

- (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 53

- (Topic 2)

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

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

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

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

Answer: B

Explanation:

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

NEW QUESTION 54

- (Topic 2)

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

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

Answer: AD

Explanation:

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

NEW QUESTION 58

- (Topic 2)

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

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

Answer: BC

Explanation:

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

NEW QUESTION 60

- (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 63

- (Topic 2)

In a Snowflake role hierarchy, what is the top-level role?

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

Answer: C

Explanation:

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

NEW QUESTION 65

- (Topic 2)

Which of the following describes a Snowflake stored procedure?

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

Answer: D

Explanation:

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

NEW QUESTION 69

- (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 74

- (Topic 2)

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

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

Answer: AB

Explanation:

When cloning a database in Snowflake, the clone includes all privileges on the database as well as existing child objects within the database, such as schemas, tables, views, etc. However, it does not include future child objects or privileges on schemas within the database2.

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

NEW QUESTION 78

- (Topic 2)

How are serverless features billed?

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

Answer: B

Explanation:

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

NEW QUESTION 82

- (Topic 2)

A single user of a virtual warehouse has set the warehouse to auto-resume and auto-suspend after 10 minutes. The warehouse is currently suspended and the user performs the following actions:

- * 1. Runs a query that takes 3 minutes to complete
 - * 2. Leaves for 15 minutes
 - * 3. Returns and runs a query that takes 10 seconds to complete
 - * 4. Manually suspends the warehouse as soon as the last query was completed
- When the user returns, how much billable compute time will have been consumed?

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

Answer: C

Explanation:

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

NEW QUESTION 85

- (Topic 2)

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

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

Answer: B

Explanation:

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

NEW QUESTION 86

- (Topic 2)

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

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

Answer: C

Explanation:

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

NEW QUESTION 91

- (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 95

- (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 100

- (Topic 2)

What is the MINIMUM edition of Snowflake that is required to use a SCIM security integration?

- A. Business Critical Edition
- B. Standard Edition
- C. Virtual Private Snowflake (VPS)
- D. Enterprise Edition

Answer: D

Explanation:

The minimum edition of Snowflake required to use a SCIM security integration is the Enterprise Edition. SCIM integrations are used for automated management of user identities and groups, and this feature is available starting from the Enterprise Edition of Snowflake. References: [COF-C02] SnowPro Core Certification Exam Study Guide

NEW QUESTION 101

- (Topic 2)

A user created a new worksheet within the Snowsight UI and wants to share this with teammates
How can this worksheet be shared?

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

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 106

- (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 109

- (Topic 2)

Which statement is true about running tasks in Snowflake?

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

Answer: B

Explanation:

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

NEW QUESTION 113

- (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 117

- (Topic 2)

Which command should be used to download files from a Snowflake stage to a local folder on a client's machine?

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

Answer: B

Explanation:

The GET command is used to download files from a Snowflake stage to a local folder on a client's machine².

Reference: <https://docs.snowflake.com/en/sql-reference/sql/get.html>

NEW QUESTION 118

- (Topic 2)

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

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

Answer: AD

Explanation:

In Snowflake, both materialized views and transient tables will incur storage charges because they store data. They will also incur compute charges when queries are run against them, as compute resources are used to process the queries. References:

[COF-C02] SnowPro Core Certification Exam Study Guide

NEW QUESTION 119

- (Topic 2)

Which of the following accurately describes shares?

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

Answer: A

Explanation:

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

NEW QUESTION 121

- (Topic 2)

A user created a transient table and made several changes to it over the course of several days. Three days after the table was created, the user would like to go

back to the first version of the table.
How can this be accomplished?

- A. Use Time Travel, as long as DATA_RETENTION_TIME_IN_DAYS was set to at least 3 days.
- B. The transient table version cannot be retrieved after 24 hours.
- C. Contact Snowflake Support to have the data retrieved from Fail-safe storage.
- D. Use the FAIL_SAFE parameter for Time Travel to retrieve the data from Fail-safe storage.

Answer: A

Explanation:

To go back to the first version of a transient table created three days prior, one can use Time Travel if the DATA_RETENTION_TIME_IN_DAYS was set to at least 3 days. This allows the user to access historical data within the specified retention period. References: [COF-C02] SnowPro Core Certification Exam Study Guide

NEW QUESTION 125

- (Topic 2)

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

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

Answer: ADE

Explanation:

Snowflake supports the use of external stages with Amazon Web Services (AWS), Microsoft Azure Cloud, and Google Cloud Platform (GCP). These platforms allow users to stage data externally and integrate with Snowflake for data loading operations

NEW QUESTION 127

- (Topic 2)

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

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

Answer: D

Explanation:

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

NEW QUESTION 128

- (Topic 2)

Which of the following objects can be directly restored using the UNDROP command? (Choose two.)

- A. Schema
- B. View
- C. Internal stage
- D. Table
- E. User
- F. Role

Answer: BD

Explanation:

The UNDROP command in Snowflake can be used to directly restore Views and Tables. These objects, when dropped, are moved to a ??Recycle Bin?? where they can be restored within a time limit before they are permanently deleted. References: [COF-C02] SnowPro Core Certification Exam Study Guide

NEW QUESTION 129

- (Topic 2)

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

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

Answer: B

Explanation:

The minimum Snowflake edition required to use Dynamic Data Masking is the Enterprise edition. This feature is not available in the Standard edition.

NEW QUESTION 131

- (Topic 2)

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

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

Answer: AC

Explanation:

The SQL commands that consume a stream and advance the stream offset are those that result in changes to the data, such as UPDATE and INSERT operations. Specifically, `UPDATE TABLE FROM STREAM` and `INSERT INTO TABLE SELECT FROM STREAM` will consume the stream and move the offset forward, reflecting the changes made to the data.

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

NEW QUESTION 135

- (Topic 3)

What is a responsibility of Snowflake's virtual warehouses?

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

Answer: C

Explanation:

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

NEW QUESTION 137

- (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 141

- (Topic 3)

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

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

Answer: D

Explanation:

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

NEW QUESTION 144

- (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 days.

NEW QUESTION 146

- (Topic 3)

Which of the following practices are recommended when creating a user in Snowflake? (Choose two.)

- A. Configure the user to be initially disabled.
- B. Force an immediate password change.
- C. Set a default role for the user.
- D. Set the number of minutes to unlock to 15 minutes.
- E. Set the user's access to expire within a specified timeframe.

Answer: BC

NEW QUESTION 148

- (Topic 3)

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

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

Answer: B

Explanation:

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

NEW QUESTION 151

- (Topic 3)

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

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

Answer: BC

Explanation:

In Snowflake, a namespace is comprised of a database and a schema. The combination of a database and schema uniquely identifies database objects within an account

NEW QUESTION 155

- (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 159

- (Topic 3)

How can a user change which columns are referenced in a view?

- A. Modify the columns in the underlying table
- B. Use the ALTER VIEW command to update the view
- C. Recreate the view with the required changes
- D. Materialize the view to perform the changes

Answer: C

Explanation:

In Snowflake, to change the columns referenced in a view, the view must be recreated with the required changes. The ALTER VIEW command does not allow changing the definition of a view; it can only be used to rename a view, convert it to or from a secure view, or add, overwrite, or remove a comment for a view. Therefore, the correct approach is to drop the existing view and create a new one with the desired column references.

NEW QUESTION 160

- (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 seconds.

NEW QUESTION 164

- (Topic 3)

Which Snowflake feature will allow small volumes of data to continuously load into Snowflake and will incrementally make the data available for analysis?

- A. COPY INTO
- B. CREATE PIPE
- C. INSERT INTO
- D. TABLE STREAM

Answer: B

Explanation:

The Snowflake feature that allows for small volumes of data to be continuously loaded into Snowflake and incrementally made available for analysis is Snowpipe. Snowpipe is designed for near-real-time data loading, enabling data to be loaded as soon as it's available in the storage layer.

NEW QUESTION 165

- (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 Duo.

NEW QUESTION 169

- (Topic 3)

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

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

Answer: A

Explanation:

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

NEW QUESTION 173

- (Topic 3)

What is the MINIMUM Snowflake edition required to use the periodic rekeying of micro-partitions?

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

Answer: A

Explanation:

Periodic rekeying of micro-partitions is a feature that requires the Enterprise Edition of Snowflake or higher. This feature is part of Snowflake's comprehensive approach to encryption key management, ensuring data security through regular rekeying. References: [COF-C02] SnowPro Core Certification Exam Study Guide

NEW QUESTION 175

- (Topic 3)

A materialized view should be created when which of the following occurs? (Choose two.)

- A. There is minimal cost associated with running the query.
- B. The query consumes many compute resources every time it runs.
- C. The base table gets updated frequently.
- D. The query is highly optimized and does not consume many compute resources.
- E. The results of the query do not change often and are used frequently.

Answer: BE

Explanation:

A materialized view is beneficial when the query consumes many compute resources every time it runs (B), and when the results of the query do not change often and are used frequently (E). This is because materialized views store pre-computed data, which can speed up query performance for workloads that are run frequently or are complex

NEW QUESTION 178

- (Topic 3)

Which clients does Snowflake support Multi-Factor Authentication (MFA) token caching for? (Select TWO).

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

Answer: CD

Explanation:

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

NEW QUESTION 182

- (Topic 3)

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

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

Answer: C

Explanation:

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

NEW QUESTION 184

- (Topic 3)

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

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

Answer: A

Explanation:

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

NEW QUESTION 185

- (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 187

- (Topic 3)

Which database objects can be shared with the Snowflake secure data sharing feature? (Choose two.)

- A. Files
- B. External tables
- C. Secure User-Defined Functions (UDFs)
- D. Sequences
- E. Streams

Answer: BC

Explanation:

Snowflake's secure data sharing feature allows sharing of certain database objects with other Snowflake accounts. Among the options provided, external tables and secure UDFs can be shared

NEW QUESTION 189

- (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 192

- (Topic 3)

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

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

Answer: A

Explanation:

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

NEW QUESTION 197

- (Topic 3)

How would a user execute a series of SQL statements using a task?

- A. Include the SQL statements in the body of the task `CREATE TASK mytask .. AS INSERT INTO target1 SELECT .. FROM stream_s1 WHERE .. INSERT INTO target2 SELECT .. FROM stream_s1 WHERE ..`
- B. A stored procedure can have only one DML statement per stored procedure invocation and therefore the user should sequence stored procedure calls in the task definition `CREATE TASK mytask AS call stored_proc1(); call stored_proc2();`
- C. Use a stored procedure executing multiple SQL statements and invoke the stored procedure from the task
- D. `CREATE TASK mytask AS call stored_proc_multiple_statements_inside();`
- E. Create a task for each SQL statement (e.
- F. resulting in task1, task2, etc.) and string the series of SQL statements by having a control task calling task1, task2, et
- G. sequentially.

Answer: C

Explanation:

To execute a series of SQL statements using a task, a user would use a stored procedure that contains multiple SQL statements and invoke this stored procedure from the task. References: Snowflake Documentation².

NEW QUESTION 201

- (Topic 3)

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

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

D. Standard

Answer: B

Explanation:

The stream type that can be used for tracking the records in external tables is `EXTERNAL`. This type of stream is specifically designed to track changes in external tables

NEW QUESTION 202

- (Topic 3)

For non-materialized views, what column in Information Schema and Account Usage identifies whether a view is secure or not?

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

Answer: B

Explanation:

In the Information Schema and Account Usage, the column that identifies whether a view is secure or not is `IS_SECURE`.

NEW QUESTION 206

- (Topic 3)

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 includes activities such as user authentication and infrastructure management. This layer coordinates activities across Snowflake, including security enforcement, query compilation and optimization, and more

NEW QUESTION 209

- (Topic 3)

The first user assigned to a new account, `ACCOUNTADMIN`, should create at least one additional user with which administrative privilege?

- A. `USERADMIN`
- B. `PUBLIC`
- C. `ORGADMIN`
- D. `SYSADMIN`

Answer: A

Explanation:

The first user assigned to a new Snowflake account, typically with the `ACCOUNTADMIN` role, should create at least one additional user with the `USERADMIN` administrative privilege. This role is responsible for creating and managing users and roles within the Snowflake account. References: Access control considerations | Snowflake Documentation

NEW QUESTION 214

- (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 runtime.

NEW QUESTION 217

- (Topic 3)

Which statement describes how Snowflake supports reader accounts?

- A. A reader account can consume data from the provider account that created it and combine it with its own data.
- B. A consumer needs to become a licensed Snowflake customer as data sharing is only supported between Snowflake accounts.

- C. The users in a reader account can query data that has been shared with the reader account and can perform DML tasks.
- D. The SHOW MANAGED ACCOUNTS command will view all the reader accounts that have been created for an account.

Answer: B

Explanation:

Snowflake supports reader accounts, which are a type of account that allows data providers to share data with consumers who are not Snowflake customers. However, for data sharing to occur, the consumer needs to become a licensed Snowflake customer because data sharing is only supported between Snowflake accounts. References: Introduction to Secure Data Sharing | Snowflake Documentation².

NEW QUESTION 218

- (Topic 3)

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

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

Answer: A

Explanation:

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

NEW QUESTION 219

- (Topic 3)

Which statements reflect key functionalities of a Snowflake Data Exchange? (Choose two.)

- A. If an account is enrolled with a Data Exchange, it will lose its access to the Snowflake Marketplace.
- B. A Data Exchange allows groups of accounts to share data privately among the accounts.
- C. A Data Exchange allows accounts to share data with third, non-Snowflake parties.
- D. Data Exchange functionality is available by default in accounts using the Enterprise edition or higher.
- E. The sharing of data in a Data Exchange is bidirectional.
- F. An account can be a provider for some datasets and a consumer for others.

Answer: BE

Explanation:

A Snowflake Data Exchange allows groups of accounts to share data privately among the accounts (B), and it supports bidirectional sharing, meaning an account can be both a provider and a consumer of data (E). This facilitates secure and governed data collaboration within a selected group³.

NEW QUESTION 222

- (Topic 3)

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

- A. The data from the permanent table.
- B. The data from the temporary table.
- C. An error stating that the view could not be compiled.
- D. An error stating that the referenced object could not be uniquely identified.

Answer: A

Explanation:

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

NEW QUESTION 226

- (Topic 3)

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

- A. history
- B. config
- C. snowsql.cnf
- D. snowsql.pubkey

Answer: B

Explanation:

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

NEW QUESTION 227

- (Topic 3)

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

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

Answer: CD

Explanation:

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

NEW QUESTION 232

- (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 234

- (Topic 3)

What happens to the shared objects for users in a consumer account from a share, once a database has been created in that account?

- A. The shared objects are transferred.
- B. The shared objects are copied.
- C. The shared objects become accessible.
- D. The shared objects can be re-shared.

Answer: C

Explanation:

Once a database has been created in a consumer account from a share, the shared objects become accessible to users in that account. The shared objects are not transferred or copied; they remain in the provider's account and are accessible to the consumer account

NEW QUESTION 236

- (Topic 3)

How does Snowflake recommend handling the bulk loading of data batches from files already available in cloud storage?

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

Answer: D

Explanation:

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

NEW QUESTION 237

- (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 queried.

NEW QUESTION 239

- (Topic 4)

Which data types can be used in Snowflake to store semi-structured data? (Select TWO)

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

Answer: AE

Explanation:

Snowflake supports the storage of semi-structured data using the ARRAY and VARIANT data types. The ARRAY data type can directly contain VARIANT, and thus indirectly contain any other data type, including itself. The VARIANT data type can store a value of any other type, including OBJECT and ARRAY, and is often used to represent semi-structured data formats like JSON, Avro, ORC, Parquet, or XML34.

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

NEW QUESTION 240

- (Topic 4)

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

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

Answer: C

Explanation:

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

NEW QUESTION 241

- (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 246

- (Topic 4)

Which views are included in the DATA SHARING USAGE schema? (Select TWO).

- A. ACCESS_HISTORY
- B. DATA_TRANSFER_HISTORY
- C. WAREHOUSE_METERING_HISTORY
- D. MONETIZED_USAGE_DAILY
- E. LISTING_TELEMETRY_DAILY

Answer: DE

Explanation:

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

NEW QUESTION 248

- (Topic 4)

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

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

Answer: BC

Explanation:

Running a virtual warehouse is required when loading data into a table and when unloading data from a table because these operations require compute resources that are provided by the virtual warehouse23.

NEW QUESTION 250

- (Topic 4)

Which function unloads data from a relational table to JSON?

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

Answer: B

Explanation:

The TO_JSON function is used to convert a VARIANT value into a string containing the JSON representation of the value. This function is suitable for unloading data from a relational table to JSON format. References: [COF-C02] SnowPro Core Certification Exam Study Guide

NEW QUESTION 253

- (Topic 4)

What metadata does Snowflake store for rows in micro-partitions? (Select TWO).

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

Answer: AB

Explanation:

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

NEW QUESTION 254

- (Topic 4)

What factors impact storage costs in Snowflake? (Select TWO).

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

Answer: AC

Explanation:

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

NEW QUESTION 256

- (Topic 4)

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

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

Answer: BC

Explanation:

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

NEW QUESTION 257

- (Topic 4)

How can a Snowflake user traverse semi-structured data?

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

Answer: A

Explanation:

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

NEW QUESTION 260

- (Topic 4)

How can a Snowflake user validate data that is unloaded using the COPY INTO <location> command?

- A. Load the data into a CSV file.
- B. Load the data into a relational table.
- C. Use the VALIDATION_MODE - SQL statement.
- D. Use the validation mode = return rows statement.

Answer: C

Explanation:

To validate data unloaded using the COPY INTO <location> command, a Snowflake user can use the VALIDATION_MODE parameter within the SQL statement to test the files for errors without loading them³.

NEW QUESTION 262

- (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 264

- (Topic 4)

What feature of Snowflake Continuous Data Protection can be used for maintenance of historical data?

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

Answer: D

Explanation:

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

NEW QUESTION 266

- (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 270

- (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 273

- (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 277

- (Topic 4)

What step can reduce data spilling in Snowflake?

- A. Using a larger virtual warehouse
- B. Increasing the virtual warehouse maximum timeout limit
- C. Increasing the amount of remote storage for the virtual warehouse
- D. Using a common table expression (CTE) instead of a temporary table

Answer: A

Explanation:

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

NEW QUESTION 282

- (Topic 4)

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

- A. INTEGER
- B. OBJECT
- C. VARIANT
- D. VARCHAR
- E. LIST

Answer: BC

Explanation:

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

NEW QUESTION 287

- (Topic 4)

Which operation can be performed on Snowflake external tables?

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

Answer: B

Explanation:

Snowflake external tables are read-only, which means data manipulation language (DML) operations like INSERT, RENAME, or ALTER cannot be performed on them. However, external tables can be used for query and join operations³. References: [COF-C02] SnowPro Core Certification Exam Study Guide

NEW QUESTION 292

- (Topic 4)

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

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

Answer: B

Explanation:

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

NEW QUESTION 297

- (Topic 4)

Which Snowflake view is used to support compliance auditing?

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

Answer: A

Explanation:

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

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

NEW QUESTION 300

- (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 303

- (Topic 4)

How can a dropped internal stage be restored?

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

Answer: D

Explanation:

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

NEW QUESTION 306

- (Topic 4)

What transformations are supported when loading data into a table using the COPY INTO <table> command? (Select TWO).

- A. Column reordering
- B. Column omission
- C. JOIN function
- D. FLATTEN function
- E. GROUP BY function

Answer: AB

Explanation:

The COPY INTO <table> command in Snowflake supports column reordering and column omission as part of its data transformation capabilities during the load process⁴⁵.

NEW QUESTION 310

- (Topic 4)

Which Snowflake function is maintained separately from the data and helps to support features such as Time Travel, Secure Data Sharing, and pruning?

- A. Column compression
- B. Data clustering
- C. Micro-partitioning
- D. Metadata management

Answer: C

Explanation:

Micro-partitioning is a Snowflake function that is maintained separately from the data and supports features such as Time Travel, Secure Data Sharing, and

pruning. It allows Snowflake to efficiently manage and query large datasets by organizing them into micro-partitions1.

NEW QUESTION 314

- (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 315

- (Topic 5)

The VALIDATE table function has which parameter as an input argument for a Snowflake user?

- A. Last_QUERY_ID
- B. CURRENT_STATEMENT
- C. UUID_STRING
- D. JOB_ID

Answer: C

Explanation:

The VALIDATE table function in Snowflake would typically use a unique identifier, such as a UUID_STRING, as an input argument. This function is designed to validate the data within a table against a set of constraints or conditions, often requiring a specific identifier to reference the particular data or job being validated.

References:

? There is no direct reference to a VALIDATE table function with these specific parameters in Snowflake documentation. It seems like a theoretical example for understanding function arguments. Snowflake documentation on UDFs and system functions can provide guidance on how to create and use custom functions for similar purposes.

NEW QUESTION 319

- (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 321

- (Topic 5)

Which URL provides access to files in Snowflake without authorization?

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

Answer: C

Explanation:

A Pre-signed URL provides access to files stored in Snowflake without requiring authorization at the time of access. This feature allows users to generate a URL with a limited validity period that grants temporary access to a file in a secure manner. It's particularly useful for sharing data with external parties or applications without the need for them to authenticate directly with Snowflake.

References:

? Snowflake Documentation: Using Pre-signed URLs

NEW QUESTION 322

- (Topic 5)

Which function will provide the proxy information needed to protect Snowsight?

- A. SYSTEMADMIN_TAG
- B. SYSTEM\$GET_PRIVATELINK
- C. SYSTEMSALLONTLIST
- D. SYSTEMAUTHORIZE

Answer: B

Explanation:

The SYSTEM\$GET_PRIVATELINK function in Snowflake provides proxy information necessary for configuring PrivateLink connections, which can protect Snowsight as well as other Snowflake services. PrivateLink enhances security by allowing Snowflake to be accessed via a private connection within a cloud provider's network, reducing exposure to the public internet.

References:

? Snowflake Documentation: PrivateLink Setup

NEW QUESTION 326

- (Topic 5)

Which common query problems are identified by the Query Profile? (Select TWO.)

- A. Syntax error
- B. Inefficient pruning
- C. Ambiguous column names
- D. Queries too large to fit in memory
- E. Object does not exist or not authorized

Answer: BD

Explanation:

The Query Profile in Snowflake can identify common query problems, including:

? B. Inefficient pruning: This refers to the inability of a query to effectively limit the amount of data being scanned, potentially leading to suboptimal performance.

? D. Queries too large to fit in memory: This indicates that a query requires more memory than is available in the virtual warehouse, which can lead to spilling to disk and degraded performance.

The Query Profile helps diagnose these issues by providing detailed execution statistics and visualizations, aiding in query optimization and troubleshooting.

References:

? Snowflake Documentation: Query Profile Top of Form

NEW QUESTION 329

- (Topic 5)

What are valid sub-clauses to the OVER clause for a window function? (Select TWO).

- A. GROUP BY
- B. LIMIT
- C. ORDER BY
- D. PARTITION BY
- E. UNION ALL

Answer: CD

Explanation:

Valid sub-clauses to the OVER clause for a window function in SQL are:

? C. ORDER BY: This clause specifies the order in which the rows in a partition are processed by the window function. It is essential for functions that depend on the row order, such as ranking functions.

? D. PARTITION BY: This clause divides the result set into partitions to which the window function is applied. Each partition is processed independently of other partitions, making it crucial for functions that compute values across sets of rows that share common characteristics.

These clauses are fundamental to defining the scope and order of data over which the window function operates, enabling complex analytical computations within SQL queries. References:

? Snowflake Documentation: Window Functions

NEW QUESTION 330

- (Topic 5)

How does Snowflake reorganize data when it is loaded? (Select TWO).

- A. Binary format
- B. Columnar format
- C. Compressed format
- D. Raw format
- E. Zipped format

Answer: BC

Explanation:

When data is loaded into Snowflake, it undergoes a reorganization process where the data is stored in a columnar format and compressed. The columnar storage format enables efficient querying and data retrieval, as it allows for reading only the necessary columns for a query, thereby reducing IO operations. Additionally, Snowflake uses advanced compression techniques to minimize storage costs and improve performance. This combination of columnar storage and compression is key to Snowflake's data warehousing capabilities.

References:

? Snowflake Documentation: Data Storage and Organization

NEW QUESTION 333

- (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 335

- (Topic 5)

There are two Snowflake accounts in the same cloud provider region: one is production and the other is non-production. How can data be easily transferred from the production account to the non-production account?

- A. Clone the data from the production account to the non-production account.
- B. Create a data share from the production account to the non-production account.
- C. Create a subscription in the production account and have it publish to the non-production account.
- D. Create a reader account using the production account and link the reader account to the non-production account.

Answer: B

Explanation:

To easily transfer data from a production account to a non-production account in Snowflake within the same cloud provider region, creating a data share is the most efficient approach. Data sharing allows for live, read-only access to selected data objects from the production account to the non-production account without the need to duplicate or move the actual data. This method facilitates seamless access to the data for development, testing, or analytics purposes in the non-production environment. References:

? Snowflake Documentation: Data Sharing

NEW QUESTION 337

- (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 to 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 to both the 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 338

- (Topic 5)

What does the worksheet and database explorer feature in Snowsight allow users to do?

- A. Add or remove users from a worksheet.
- B. Move a worksheet to a folder or a dashboard.
- C. Combine multiple worksheets into a single worksheet.
- D. Tag frequently accessed worksheets for ease of access.

Answer: D

Explanation:

The worksheet and database explorer feature in Snowsight allows users to tag frequently accessed worksheets for ease of access. This functionality helps users organize and quickly navigate to the worksheets they use most often, enhancing productivity and streamlining the data exploration and analysis process within Snowsight, Snowflake's web-based query and visualization interface.

References:

? Snowflake Documentation: Snowsight (UI for Snowflake)

NEW QUESTION 342

- (Topic 5)

For Directory tables, what stage allows for automatic refreshing of metadata?

- A. User stage
- B. Table stage
- C. Named internal stage
- D. Named external stage

Answer: D

Explanation:

For directory tables, a named external stage allows for the automatic refreshing of metadata. This capability is particularly useful when dealing with files stored on external storage services (like Amazon S3, Google Cloud Storage, or Azure Blob Storage) and accessed through Snowflake. The external stage references these files, and the directory table's metadata can be automatically updated to reflect changes in the underlying files. References:

? Snowflake Documentation: External Stages

NEW QUESTION 343

- (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 get the most detailed information about individual table storage details in Snowflake, the TABLE STORAGE METRICS view should be used. This Information Schema view provides granular storage metrics for tables within Snowflake, including data related to the size of the table, the amount of data stored, and storage usage over time. It's an essential tool for administrators and users looking to monitor and optimize storage consumption and costs.

References:

? Snowflake Documentation: Information Schema - TABLE STORAGE METRICS View

NEW QUESTION 348

- (Topic 5)

Which function is used to convert rows in a relational table to a single VARIANT column?

- A. ARRAY_AGG
- B. OBJECT_AGG
- C. ARRAY_CONSTRUCT
- D. OBJECT_CONSTRUCT

Answer: D

Explanation:

The OBJECT_CONSTRUCT function in Snowflake is used to convert rows in a relational table into a single VARIANT column that represents each row as a JSON object. This function dynamically creates a JSON object from a list of key-value pairs, where each key is a column name and each value is the corresponding column value for a row. This is particularly useful for aggregating and transforming structured data into semi-structured JSON format for further processing or analysis.

References:

? Snowflake Documentation: Semi-structured Data Functions

NEW QUESTION 351

- (Topic 5)

What are characteristics of transient tables in Snowflake? (Select TWO).

- A. Transient tables have a Fail-safe period of 7 days.
- B. Transient tables can be cloned to permanent tables.
- C. Transient tables persist until they are explicitly dropped.
- D. Transient tables can be altered to make them permanent tables.
- E. Transient tables have Time Travel retention periods of 0 or 1 day.

Answer: BC

Explanation:

Transient tables in Snowflake are designed for temporary or intermediate workloads with the following characteristics:

? B. Transient tables can be cloned to permanent tables: This feature allows users to create copies of transient tables for permanent use, providing flexibility in managing data lifecycles.

? C. Transient tables persist until they are explicitly dropped: Unlike temporary tables that exist for the duration of a session, transient tables remain in the database until explicitly removed by a user, offering more durability for short-term data storage needs.

References:

? Snowflake Documentation: Transient Tables

NEW QUESTION 355

- (Topic 5)

What is the MAXIMUM number of clusters that can be provisioned with a multi-cluster virtual warehouse?

- A. 1
- B. 5
- C. 10
- D. 100

Answer: C

Explanation:

In Snowflake, the maximum number of clusters that can be provisioned within a multi-cluster virtual warehouse is 10. This allows for significant scalability and performance management by enabling Snowflake to handle varying levels of query load by adjusting the number of active clusters within the warehouse. References: Snowflake documentation on virtual warehouses, particularly the scalability options available in multi-cluster configurations.

NEW QUESTION 358

- (Topic 5)

What is the Fail-safe retention period for transient and temporary tables?

- A. 0 days
- B. 1 day
- C. 7 days
- D. 90 days

Answer: A

Explanation:

The Fail-safe retention period for transient and temporary tables in Snowflake is 0 days. Fail-safe is a feature designed to protect data against accidental loss or deletion by retaining historical data for a period after its Time Travel retention period expires. However, transient and temporary tables, which are designed for temporary or short-term storage and operations, do not have a Fail-safe period. Once the data is deleted or the table is dropped, it cannot be recovered.

References:

? Snowflake Documentation: Understanding Fail-safe

NEW QUESTION 359

- (Topic 5)

User1, who has the SYSADMIN role, executed a query on Snowsight. User2, who is in the same Snowflake account, wants to view the result set of the query executed by User1 using the Snowsight query history.

What will happen if User2 tries to access the query history?

- A. If User2 has the sysadmin role they will be able to see the results.
- B. If User2 has the securityadmin role they will be able to see the results.
- C. If User2 has the ACCOUNTADMIN role they will be able to see the results.
- D. User2 will be unable to view the result set of the query executed by User1.

Answer: C

Explanation:

In Snowflake, the query history and the results of queries executed by a user are accessible based on the roles and permissions. If User1 executed a query with the SYSADMIN role, User2 would be able to view the result set of that query executed by User1 only if User2 has the ACCOUNTADMIN role. The ACCOUNTADMIN role has the broadest set of privileges, including the ability to access all aspects of the account's operation, data, and query history, thus enabling User2 to view the results of queries executed by other users.

References:

? Snowflake Documentation: Understanding Snowflake Roles

NEW QUESTION 360

- (Topic 5)

What Snowflake database object is derived from a query specification, stored for later use, and can speed up expensive aggregation on large data sets?

- A. Temporary table
- B. External table
- C. Secure view
- D. Materialized view

Answer: D

Explanation:

A materialized view in Snowflake is a database object derived from a query specification, stored for later use, and can significantly speed up expensive aggregations on large data sets. Materialized views store the result of their underlying query, reducing the need to recompute the result each time the view is accessed. This makes them ideal for improving the performance of read-heavy, aggregate-intensive queries.

References:

? Snowflake Documentation: Using Materialized Views

NEW QUESTION 363

- (Topic 5)

What should be used when creating a CSV file format where the columns are wrapped by single quotes or double quotes?

- A. BINARY_FORMAT
- B. ESCAPE_UNENCLOSED_FIELD
- C. FIELD_OPTIONALLY_ENCLOSED_BY
- D. SKIP_BYTE_ORDER_MARK

Answer: C

Explanation:

When creating a CSV file format in Snowflake and the requirement is to wrap columns by single quotes or double quotes, the FIELD_OPTIONALLY_ENCLOSED_BY parameter should be used in the file format specification. This parameter allows you to define a character (either a

single quote or a double quote) that can optionally enclose each field in the CSV file, providing flexibility in handling fields that contain special characters or delimiters as part of their data.

References:

? Snowflake Documentation: CSV File Format

NEW QUESTION 367

- (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 368

- (Topic 6)

Which object type is granted permissions for reading a table?

- A. User
- B. Role
- C. Attribute
- D. Schema

Answer: B

Explanation:

In Snowflake, permissions for accessing database objects, including tables, are not granted directly to users but rather to roles. A role encapsulates a collection of privileges on various Snowflake objects. Users are then granted roles, and through those roles, they inherit the permissions necessary to read a table or perform other actions. This approach adheres to the principle of least privilege, allowing for granular control over database access and simplifying the management of user permissions.

Reference: Snowflake's official documentation on access control introduces the concept of

roles and how they are used to manage permissions: <https://docs.snowflake.com/en/user-guide/security-access-control-overview.html#roles>

NEW QUESTION 369

- (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 the SECURITYADMIN role or any role with higher privileges, such as ACCOUNTADMIN, can create network policies. Additionally, a custom role can be granted the CREATE NETWORK POLICY privilege, 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 371

- (Topic 6)

Which actions can be performed using a resource monitor in Snowflake? (Select TWO).

- A. Monitor the performance of individual queries in real-time
- B. Automatically allocate more storage space to a virtual warehouse
- C. Modify the queries being executed within a virtual warehouse.
- D. Suspend a virtual warehouse when its credit usage reaches a defined limit.
- E. Trigger a notification to account administrators when credit usage reaches a specified threshold

Answer: DE

Explanation:

Resource monitors in Snowflake can perform actions such as suspending a virtual warehouse when its credit usage reaches a defined limit and triggering a

notification to account administrators when credit usage reaches a specified threshold. These actions help manage and control resource usage and costs within Snowflake.

References:

? Snowflake Documentation: Resource Monitors

NEW QUESTION 376

- (Topic 6)

Which categories are included in the execution time summary in a Query Profile? (Select TWO).

- A. Pruning
- B. Spilling
- C. Initialization
- D. Local Disk I/O
- E. Percentage of data read from cache

Answer: AC

Explanation:

In the execution time summary of a Query Profile in Snowflake, the categories included provide insights into various aspects of query execution. "Pruning" refers to the process by which Snowflake reduces the amount of data scanned by eliminating partitions of data that are not relevant to the query, thus improving performance. "Initialization" represents the time taken for query planning and setup before actual execution begins. These metrics are crucial for understanding and optimizing query performance.

Reference: Snowflake Documentation on the Query Profile, which outlines the different metrics and categories included in the execution summary:

<https://docs.snowflake.com/en/user-guide/ui-query-profile.html#execution-summary>

NEW QUESTION 377

- (Topic 6)

What are the possible values within a METADATASACTION column in a Snowflake stream? (Select TWO).

- A. INSERT
- B. UPDATE
- C. DELETE
- D. TRUNCATE
- E. UPSERT

Answer: AC

Explanation:

In Snowflake streams, the METADATASACTION column indicates the type of data manipulation operation that has occurred. The possible values include INSERT and DELETE.

? INSERT: Indicates that a new row has been inserted into the table.

? DELETE: Indicates that a row has been deleted from the table.

References:

? Snowflake Documentation: Change Data Capture (CDC) with Streams

? Snowflake Documentation: Stream Data Capture

NEW QUESTION 380

- (Topic 6)

When an object is created in Snowflake, who owns the object?

- A. The public role
- B. The user's default role
- C. The current active primary role
- D. The owner of the parent schema

Answer: C

Explanation:

In Snowflake, when an object is created, it is owned by the role that is currently active. This active role is the one that is being used to execute the creation command. Ownership implies full control over the object, including the ability to grant and revoke access privileges. This is specified in Snowflake's documentation under the topic of Access Control, which states that "the role in use at the time of object creation becomes the owner of the object."

References:

? Snowflake Documentation: Object Ownership

NEW QUESTION 382

- (Topic 6)

What optional properties can a Snowflake user set when creating a virtual warehouse? (Select TWO).

- A. Auto-suspend
- B. Cache size
- C. Default role
- D. Resource monitor
- E. Storage size

Answer: AD

Explanation:

When creating a virtual warehouse in Snowflake, users have the option to set several properties to manage its behavior and resource usage. Two of these optional

properties are Auto-suspend and Resource monitor.

? Auto-suspend: This property defines the period of inactivity after which the warehouse will automatically suspend. This helps in managing costs by stopping the warehouse when it is not in use.

```
CREATE WAREHOUSE my_warehouse WITH WAREHOUSE_SIZE = 'XSMALL'
```

```
AUTO_SUSPEND = 300; -- Auto-suspend after 5 minutes of inactivity
```

? Resource monitor: Users can assign a resource monitor to a warehouse to control and limit the amount of credit usage. Resource monitors help in setting quotas and alerts for warehouse usage.

```
CREATE WAREHOUSE my_warehouse WITH WAREHOUSE_SIZE = 'XSMALL'
```

```
RESOURCE_MONITOR = 'my_resource_monitor';
```

References:

? Snowflake Documentation: Creating Warehouses

? Snowflake Documentation: Resource Monitors

NEW QUESTION 383

- (Topic 6)

If a virtual warehouse is suspended, what happens to the warehouse cache?

- A. The cache is dropped when the warehouse is suspended and is no longer available upon restart.
- B. The warehouse cache persists for as long the warehouse exists, regardless of its suspension status.
- C. The cache is maintained for up to two hours and can be restored if the warehouse is restarted within this limit.
- D. The cache is maintained for the auto suspend duration and can be restored if the warehouse is restarted within this limit.

Answer: A

Explanation:

When a virtual warehouse in Snowflake is suspended, the cache is dropped and is no longer available upon restart. This means that all cached data, including results and temporary data, are cleared from memory. The purpose of this behavior is to conserve resources while the warehouse is not active. Upon restarting the warehouse, it will need to reload any data required for queries from storage, which may result in a slower initial performance until the cache is repopulated. This is a critical consideration for managing performance and cost in Snowflake.

NEW QUESTION 387

- (Topic 6)

Which service or tool is a Command Line Interface (CLI) client used for connecting to Snowflake to execute SQL queries?

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

Answer: D

Explanation:

SnowSQL is the Command Line Interface (CLI) client provided by Snowflake for executing SQL queries and performing various tasks. It allows users to connect to their Snowflake accounts and interact with the Snowflake data warehouse.

? Installation: SnowSQL can be downloaded and installed on various operating systems.

? Configuration: Users need to configure SnowSQL with their Snowflake account credentials.

? Usage: Once configured, users can run SQL queries, manage data, and perform administrative tasks through the CLI.

References:

? Snowflake Documentation: SnowSQL

? Snowflake Documentation: Installing SnowSQL

NEW QUESTION 391

- (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 392

- (Topic 6)

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

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

Answer: A

Explanation:

Granting a user the MODIFY, MONITOR, and OPERATE privileges on all virtual warehouses in Snowflake is equivalent to granting the global MANAGE WAREHOUSES privilege. These privileges collectively provide comprehensive control over virtual warehouses.

? MODIFY Privilege:

? MONITOR Privilege:

? OPERATE Privilege:

References:

? Snowflake Documentation: Warehouse Privileges

NEW QUESTION 395

- (Topic 6)

Which Snowflake object contains all the information required to share a database?

- A. Private listing
- B. Secure view
- C. Sequence
- D. Share

Answer: D

Explanation:

In Snowflake, a Share is the object that contains all the information required to share a database with other Snowflake accounts. Shares are used to securely share data stored in Snowflake tables and views, enabling data providers to grant data consumers access to their datasets without duplicating data. When a database is shared, it can include one or more schemas, and each schema can contain tables, views, or both.

References:

? Snowflake Documentation on Shares: Shares

NEW QUESTION 398

- (Topic 6)

In the Data Exchange, who can get or request data from the listings? (Select TWO).

- A. Users with ACCOUNTADMIN role
- B. Users with sysadmin role
- C. Users with ORGADMIN role
- D. Users with import share privilege
- E. Users with manage grants privilege

Answer: AD

Explanation:

In the Snowflake Data Exchange, the ability to get or request data from listings is generally controlled by specific roles and privileges:

? A. Users with ACCOUNTADMIN role: This role typically has the highest level of

access within a Snowflake account, including the ability to manage and access all features and functions. Users with this role can access data listings within the Data Exchange.

? D. Users with import share privilege: This specific privilege is necessary for users

who need to import shared data from the Data Exchange. This privilege allows them to request and access data listings explicitly shared with them.

NEW QUESTION 402

- (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 404

- (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 406

- (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 409

- (Topic 6)

What Snowflake objects can contain custom application logic written in JavaScript? (Select TWO)

- A. Stored procedures
- B. Stages
- C. Tasks
- D. Views
- E. User-Defined Functions (UDFs)

Answer: AE

Explanation:

Snowflake allows users to write custom application logic in JavaScript for two types of objects: Stored Procedures and User-Defined Functions (UDFs).

? Stored Procedures: Snowflake stored procedures can be written in JavaScript to encapsulate complex business logic and procedural operations. CREATE OR

REPLACE PROCEDURE my_procedure()

RETURNS STRING LANGUAGE JAVASCRIPT EXECUTE AS CALLER AS

\$\$

// JavaScript logic here

\$\$;

? User-Defined Functions (UDFs): Snowflake UDFs can be written in JavaScript to perform custom calculations or operations on data.

CREATE OR REPLACE FUNCTION my_function(x FLOAT) RETURNS FLOAT

LANGUAGE JAVASCRIPT AS

\$\$

return x * 2;

\$\$;

References:

? Snowflake Documentation: Stored Procedures

? Snowflake Documentation: User-Defined Functions (UDFs)

NEW QUESTION 413

- (Topic 6)

How can the Query Profile be used to troubleshoot a problematic query?

- A. It will indicate if a virtual warehouse memory is too small to run the query
- B. It will indicate if a user lacks the privileges needed to run the query.
- C. It will indicate if a virtual warehouse is in auto-scale mode
- D. It will indicate if the user has enough Snowflake credits to run the query

Answer: A

Explanation:

The Query Profile in Snowflake provides detailed insights into the execution of a query. It helps in troubleshooting performance issues by showing the steps of the query execution and the resources consumed. One of the key aspects it reveals is whether the virtual warehouse memory was sufficient for the query.

? Access Query Profile: Navigate to the Query History page and select the query you want to analyze.

? Examine Query Execution Steps: The Query Profile displays the different stages of the query execution, including the time taken and resources used at each step.

? Identify Memory Issues: Look for indicators of memory issues, such as spilling to disk or memory errors, which suggest that the virtual warehouse memory might be too small.

References:

? Snowflake Documentation: Query Profile

? Snowflake Documentation: Optimizing Queries

NEW QUESTION 417

- (Topic 6)

How can an administrator check for updates (for example, SCIM API requests) sent to Snowflake by the identity provider?

- A. ACCESS_HISTORY
- B. LOAD_HISTORY
- C. QUERY_HISTORY
- D. REST EVENT HISTORY

Answer: D

Explanation:

To monitor updates, such as SCIM API requests sent to Snowflake by the identity provider, an administrator can use the REST EVENT HISTORY feature. This feature allows administrators to query historical data about REST API calls made to Snowflake, including those related to user and role management through SCIM (System for Cross-domain Identity Management).

The REST EVENT HISTORY table function returns information about REST API calls made over a specified period. It is particularly useful for auditing and monitoring purposes, especially when integrating Snowflake with third-party identity providers that use SCIM for automated user provisioning and deprovisioning. An example query to check for SCIM API requests might look like this: `SELECT * FROM TABLE (information_schema.rest_event_history (date_range_start => datediff ('hours', -1, current_timestamp ()))) WHERE request_type = 'SCIM';`

This query returns details on SCIM API requests made in the last hour, including the request type, the identity provider's details, and the outcome of each request. Reference: Snowflake Documentation on REST EVENT HISTORY (https://docs.snowflake.com/en/sql-reference/functions/rest_event_history.html)

NEW QUESTION 420

- (Topic 6)

Which statements reflect valid commands when using secondary roles? (Select TWO).

- A. Use SECONDARY ROLES RESUME
- B. USE SECONDARY ROLES SUSPEND
- C. USE SECONDARY RLES ALL
- D. USE SECONDARY ROLES ADD <Role Name>
- E. Use SECONDARY ROLES NONE

Answer: CE

Explanation:

? Incorrect Commands: The options referencing "RESUME", "SUSPEND", and "ADD" are not valid commands in the context of secondary roles.

References:

? Snowflake Documentation (USE SECONDARY ROLES): <https://docs.snowflake.com/en/sql-reference/sql/use-secondary-roles.html>

NEW QUESTION 422

- (Topic 6)

Snowflake users can create a resource monitor at which levels? (Select TWO).

- A. User level
- B. Pipe level
- C. Account level
- D. Cloud services level
- E. Virtual warehouse level

Answer: CE

Explanation:

Resource monitors in Snowflake are tools used to track and control the consumption of compute resources, ensuring that usage stays within defined limits. These monitors can be created at the account level, allowing administrators to set overall resource consumption limits for the entire Snowflake account. Additionally, resource monitors can be set at the virtual warehouse level, enabling more granular control over the resources consumed by individual warehouses. This dual-level capability allows organizations to manage their Snowflake usage efficiently, preventing unexpected costs and optimizing performance. References: Snowflake Documentation on Resource Monitors Top of Form

NEW QUESTION 423

- (Topic 6)

What command is used to export or unload data from Snowflake?

- A. PUT @mystage
- B. GET @mystage
- C. COPY INTO @mystage
- D. INSERT @mystage

Answer: A

Explanation:

The command used to export or unload data from Snowflake to a stage (such as a file in an S3 bucket, Azure Blob Storage, or Google Cloud Storage) is the PUT command. The PUT command is designed to upload data files from a local file system (in the case of SnowSQL or other client) or a virtual warehouse to a specified stage. This functionality is critical for scenarios where data needs to be extracted from Snowflake for use in external systems, backups, or further processing.

The syntax for the PUT command follows the structure: `PUT file://<local_file_path>`

`@<stage_name>`, where `<local_file_path>` specifies the path to the file(s) on the local file system that you wish to upload, and `<stage_name>` specifies the destination stage in Snowflake.

It's important to distinguish that the PUT command is used for exporting data out of Snowflake, whereas the COPY INTO <table> command is used for importing data into Snowflake from a stage. The GET command, on the other hand, is used to download files from a stage to the local file system, essentially the inverse operation of the PUT command.

References:

? Snowflake Documentation on Loading and Unloading Data: [Loading and Unloading Data] (<https://docs.snowflake.com/en/user-guide/data-load>)

NEW QUESTION 425

.....

Thank You for Trying Our Product

* 100% Pass or Money Back

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

* One year free update

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

* Trusted by Millions

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

* Shop Securely

All transactions are protected by VeriSign!

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

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