

Snowflake COF-C02

SnowPro Core Certification Exam

For More Information – Visit link below:

<https://www.examsempire.com/>

Product Version

- 1. Up to Date products, reliable and verified.**
- 2. Questions and Answers in PDF Format.**



<https://examsempire.com/>

Visit us at: <https://www.examsempire.com/cof-c02>

Latest Version: 19.0

Question: 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.

Question: 2

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

- A. True
- B. False

Answer: A

Explanation:

Provisioning time can vary based on the size of the warehouse. A 4X-Large Warehouse typically has more resources and may take longer to provision compared to a X-Small Warehouse, which has fewer resources and can generally be provisioned more quickly. Understanding and viewing Fail-safe | Snowflake Documentation

Question: 3

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

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

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

Question: 4

The Information Schema and Account Usage Share provide storage information for which of the following objects? (Choose three.)

- A. Users
- B. Tables
- C. Databases
- D. Internal Stages

Answer: B, C, D

Explanation:

The Information Schema and Account Usage Share in Snowflake provide metadata and historical usage data for various objects within a Snowflake account. Specifically, they offer storage information for Tables, Databases, and Internal Stages. These schemas contain views and table functions that allow users to query object metadata and usage metrics, such as the amount of data stored and historical activity.

Tables: The storage information includes data on the daily average amount of data in database tables.

Databases: For databases, the storage usage is calculated based on all the data contained within the database, including tables and stages.

Internal Stages: Internal stages are locations within Snowflake for temporarily storing data, and their storage usage is also tracked.

Question: 5

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.

Question: 6

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

- A. True
- B. False

Answer: B

Explanation:

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

Question: 7

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

- A. True
- B. False

Answer: B

Explanation:

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

Question: 8

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

- A. True
- B. False

Answer: A

Explanation:

Question: 9

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.

Question: 10

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

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

Answer: A, C

Explanation:

In Snowflake, Data Shares can be created and managed in two primary ways:

Via the Snowflake Web Interface (UI): Users can create and manage shares through the graphical interface provided by Snowflake, which allows for a user-friendly experience.

Via SQL commands: Snowflake also allows the creation and management of shares using SQL commands. This method is more suited for users who prefer scripting or need to automate the process.

Thank You for Trying Our Product

Special 16 USD Discount Coupon: NSZUBG3X

Email: support@examsempire.com

**Check our Customer Testimonials and ratings
available on every product page.**

Visit our website.

<https://examsempire.com/>