



Download Snowflake SOL-C01 Exam Dumps Free

Shared by Ruiz on 17-06-2026

For More Free Questions and Preparation Resources

Check the Links on Last Page



Question 1

Question Type: MultipleChoice

Which of the following is true about Snowflake clones?

Options:

- A- Clones consume the same storage as the original object.
- B- Clones initially consume minimal/no storage, with storage only being used for changes made to the clone or the original
- C- Clones can only be created from databases, not tables or schemas.
- D- Clones are automatically synchronized with the original object.

Answer:

B

Explanation:

Snowflake's zero-copy cloning creates a metadata-only copy of an object---database, schema, table, or many other object types. At creation time, the clone shares the same underlying micro-partitions as the source, consuming virtually no additional storage.

Storage consumption occurs only when either the clone or the original object is modified, a behavior known as copy-on-write. This makes cloning ideal for development, testing, QA, troubleshooting, backup creation, or controlled sandboxing.

Incorrect statements:

Clones do not duplicate storage at creation.

Cloning applies to tables, schemas, databases, and other supported objects.

Clones do not synchronize with the original; after creation they diverge independently.

This architecture enables quick, cost-efficient replication of environments and data.

Question 2

Question Type: MultipleChoice

Which of the following are date and time data types in Snowflake? (Select any 3 options)

Options:

- A- TIMEDATE
- B- DATE
- C- Python
- D- TIMESTAMP
- E- DATETIME
- F- JavaScript

Answer:

B, D, E

Explanation:

Snowflake supports several native data types for managing temporal information:

DATE -- Stores dates only (year, month, day) without any time component.

TIMESTAMP -- Represents date and time values, including several variants:

- * TIMESTAMP_NTZ (no time zone)
- * TIMESTAMP_LTZ (local time zone)
- * TIMESTAMP_TZ (explicit time zone)

DATETIME -- An alias for TIMESTAMP_NTZ, storing both date and time.

These data types enable powerful time-based analytics and comparisons.

TIMEDATE is not a valid data type.

Python and JavaScript are programming languages, not Snowflake data types.

Question 3

Question Type: MultipleChoice

What is a Markdown cell in Snowflake Notebooks?

Options:

- A- A cell that contains data in more than one computer languages
- B- A cell that allows the user to format the cell text
- C- A cell that is nested in other cells
- D- A cell that contains older data

Answer:

B

Explanation:

In Snowflake Notebooks, a Markdown cell is a non-executable cell type used for formatted text. It allows the user to format text using Markdown syntax, including headings, lists, tables, emphasis, inline code, and links. This makes it possible to document the analysis, describe steps, and provide commentary alongside SQL and Python cells, improving clarity and collaboration.

Markdown cells do not execute code in "more than one computer language"; they are not code cells at all. Notebook cells are not nested; each cell exists as a separate element in the notebook. There is no concept of "older data" associated specifically with Markdown cells; they simply store text content defined by the user.

=====

Question 4

Question Type: MultipleChoice

What is the primary purpose of Snowflake's Cloning feature?

Options:

- A- To back up data in real-time
- B- To create an exact copy of a database, schema, or table without duplicating the storage
- C- To manage user permissions
- D- To create a backup of the entire Snowflake account

Answer:

B

Explanation:

Snowflake's Cloning feature provides the ability to create a zero-copy clone of tables, schemas, or entire databases. This means that the cloned object references the same underlying micro-partition data as the original, so no duplicate storage costs are incurred upon creation. Instead, Snowflake uses metadata pointers to reference the existing data. If changes occur in the original or cloned objects, Snowflake performs copy-on-write operations to maintain isolation while still optimizing storage. This cloning mechanism is extremely valuable for development, testing environments, analytics validation, and troubleshooting historical states. It is not a real-time backup tool; backup and restore are handled through Time Travel and Fail-safe. Cloning does not grant or manage permissions nor copy the entire Snowflake account; it works solely at the database, schema, and table object levels.

=====

Question 5

Question Type: MultipleChoice

What is a database in Snowflake?

Options:

- A- A logical grouping of schemas.
- B- A single virtual warehouse.
- C- A physical storage location for data files.
- D- A collection of tables and views.

Answer:

A

Explanation:

A Snowflake database is a logical container grouping one or more schemas. It does not store data physically in the traditional sense---Snowflake abstracts physical storage into its cloud-native micro-partitioned storage layer. Because schemas contain tables, views, and other objects, a database organizes schemas logically, enabling governance, naming conventions, and security boundaries.

A virtual warehouse is compute, not storage. A "collection of tables and views" better describes a

schema, not a database. Snowflake separates logical structure from physical storage, so a database is not a physical file repository.

Thus, a database is best defined as a logical grouping of schemas.

Question 6

Question Type: MultipleChoice

Which of the following are valid components of Snowflake's core architecture? (Select THREE)

Options:

- A- On-premises Data Centers
- B- User Interface Layer
- C- Third-Party ETL Tools
- D- Database Storage Layer
- E- Query Processing Layer
- F- Cloud Services Layer

Answer:

D, E, F

Explanation:

Snowflake's architecture consists of:

- * Database Storage Layer -- Stores optimized columnar data.
- * Query Processing Layer -- Executes workloads via virtual warehouses.
- * Cloud Services Layer -- Handles metadata, optimization, access control.

ETL tools and UI are external integrations. Snowflake is cloud-native and has no on-premises components.

Question 7

Question Type: MultipleChoice

What is a "provider" Given the Snowflake Marketplace or Data Exchange?

Options:

- A- A user who only consumes data
- B- A user who manages Snowflake infrastructure
- C- A user or organization that shares data products or data listings
- D- A user who is a Snowflake employee

Answer:

C

Explanation:

In Snowflake Marketplace and Data Exchange, a provider is any Snowflake user or organization that publishes and shares datasets, data products, or applications with other Snowflake consumers. Providers control access to their listings, configure data sharing models, and can offer either free or paid data products. They manage listing descriptions, update schedules, and refresh logic, ensuring customers or subscribers receive up-to-date and reliable information. Providers can be enterprises, analytics vendors, public data contributors, or research institutions. They are distinct from consumers, who access or use the shared data. Providers also do not handle Snowflake platform infrastructure---that is Snowflake's responsibility. Snowflake employees are not categorized as providers unless they publish official Snowflake datasets. Thus, the primary purpose of a provider is to share curated data offerings on the Marketplace or within private exchanges.

=====

Question 8

Question Type: MultipleChoice

What filter can be used when searching for data products on the Snowflake Marketplace?

Options:

- A- Size
- B- Compatibility

- C- Price
- D- Performance

Answer:

C

Explanation:

On the Snowflake Marketplace, users can apply a Price filter when searching for data products. This filter helps distinguish between free offerings and paid listings, making it easier to align data acquisition with budget and commercial terms. Marketplace metadata includes pricing information that this filter uses to refine search results.

Dataset "Size" is not exposed as a primary Marketplace search filter. "Compatibility" is implied, as all Marketplace products are consumable within Snowflake. "Performance" is not a filterable attribute, because performance is influenced mainly by the consumer's compute configuration and query design, rather than a fixed property of the Marketplace listing.

=====

Question 9

Question Type: MultipleChoice

What is the purpose of assigning roles to users in Snowflake?

Options:

- A- To manage and enforce permissions on objects
- B- To define user tasks
- C- To optimize query execution plans
- D- To control which data types users can query

Answer:

A

Explanation:

Snowflake uses a Role-Based Access Control (RBAC) model, where roles are the containers of

privileges. Assigning roles to users ensures that permissions on database objects (such as tables, schemas, warehouses, and functions) are enforced consistently and securely. Users do not receive privileges directly; instead, privileges are granted to roles, and roles are assigned to users.

This enables scalable, auditable, and manageable access control.

Roles do not determine tasks, do not affect query optimization, and do not govern which data types a user may query---permissions are object-based, not datatype-based.

=====



To Get Premium Files for SOL-C01 Visit

<https://www.p2pexams.com/products/sol-c01>

For More Free Questions Visit

<https://www.p2pexams.com/snowflake/pdf/sol-c01>

20%
DISCOUNT

P2P
exams