



Snowflake SOL-C01 Practice Questions

Shared by Solomon on 19-08-2026

For More Free Questions and Preparation Resources

Check the Links on Last Page



Question 1

Question Type: MultipleChoice

When using the TRANSLATE function, what type of input is expected for the source and target languages?

Options:

- A- Full language names (e.g., 'English', 'French')
- B- Standard language codes (e.g., 'en', 'fr')
- C- File paths to language dictionaries
- D- Numerical language identifiers

Answer:

B

Explanation:

The TRANSLATE function in Snowflake Cortex expects ISO-standard language codes such as 'en' for English or 'fr' for French. These are two-letter or multi-letter codes defined by international standards. Snowflake uses these codes to instruct its translation model which languages to interpret and produce. If the source language is unknown, users can pass an empty string to trigger automatic language detection. Full-language names, dictionary file paths, or numeric identifiers are not required or supported. TRANSLATE operates entirely with text input and language codes, ensuring easy integration into SQL workflows for multi-language data pipelines, localization tasks, and multilingual analytics.

=====

Question 2

Question Type: MultipleChoice

Which of the following are valid Snowflake Virtual Warehouse sizes? (Choose any 3 options)

Options:

- A- Small

- B- Tiny
- C- Medium
- D- Xsmall

Answer:

A, C, D

Explanation:

Snowflake warehouse sizes follow T-shirt sizing, each representing a doubling of compute resources relative to the previous size. Valid Snowflake warehouse sizes include X-Small, Small, Medium, Large, and up to 6X-Large.

X-Small is the smallest supported warehouse size and is commonly used for light data processing, development tasks, and low-volume operations.

Small doubles the compute capacity of X-Small, providing better performance for moderate workloads.

Medium continues scaling upward, offering increased CPU, memory, and parallelism for heavier analytics or transformation workloads.

Tiny is not a Snowflake warehouse size. Snowflake never exposes a warehouse smaller than X-Small, ensuring consistent minimum compute capabilities.

These standardized size tiers allow predictable resource scaling and cost forecasting.

Question 3

Question Type: MultipleChoice

Which of the following components are included in a fully qualified name in Snowflake? (Select any 3 options)

Options:

- A- Database name
- B- Role name
- C- Schema name
- D- Table or view name

Answer:

A, C, D

Explanation:

A fully qualified name in Snowflake specifies the exact location of an object and follows this structure:

`<database_name>.<schema_name>.<object_name>`

This enables precise referencing, especially in environments with multiple databases or schemas.

Components:

Database name-- the top-level container.

Schema name-- the logical grouping of related objects.

Table or view name-- the specific object being referenced.

Example:

`SALES_DB.PUBLIC.ORDERS`

Roles are not part of object names. Roles control authorization, while fully qualified names handle namespace resolution. Using fully qualified names helps avoid ambiguity in pipelines, stored procedures, views, and cross-database queries.

Snowflake also supports partially qualified names depending on active context (current database and schema).

Question 4

Question Type: MultipleChoice

What does "warehouse scaling up/down" refer to in Snowflake?

Options:

- A- Moving data between different storage locations
- B- Changing the region of the warehouse
- C- Changing the size of the warehouse (e.g., from Small to Medium or Vice Versa).
- D- Adjusting the number of clusters in a multi-cluster warehouse.

Answer:

C

Explanation:

Scaling up or down refers to vertical scaling, meaning the warehouse's compute size is increased or decreased. For example, moving from Small Medium Large increases CPU, memory, and I/O capacity, enabling faster processing for compute-intensive workloads.

Vertical scaling improves single-query performance, large ETL jobs, complex joins, or transformations. It does not improve concurrency unless multi-cluster mode is also used.

Horizontal scaling (scaling out/in), by contrast, adjusts the number of clusters and is used for concurrency.

Region selection is fixed at account creation and cannot be changed by resizing a warehouse. Storage movement is unrelated to compute rescaling.

Question 5

Question Type: MultipleChoice

How can you search for listings on Snowflake Marketplace? (Select any 2 options)

Options:

- A- Browsing categories and tags
- B- Contact Snowflake Support
- C- User Universal Search
- D- Using the search bar at the top of the marketplace

Answer:

A, C, D

Explanation:

Snowflake Marketplace provides multiple intuitive discovery mechanisms to help users explore available listings efficiently. The search bar at the top of the Marketplace is a primary entry

point, enabling keyword-based discovery across providers, industries, and dataset types. This facilitates targeted searching for specific topics or data categories.

Universal Search, available in Snowsight, extends search beyond databases and worksheets to include Marketplace listings. This unified search experience allows users to locate datasets without navigating through menus, improving productivity and discoverability.

Browsing categories and tags supports exploratory search by grouping listings under predefined themes such as financial data, geospatial data, demographics, and more. These tags promote structured navigation, especially useful when users are exploring rather than seeking something specific.

Contacting Snowflake Support is not required for accessing or discovering listings and is irrelevant to Marketplace search workflows.

Question 6

Question Type: MultipleChoice

What is the name for a Snowflake account holder who shares data with other Snowflake accounts?

Options:

- A- A data provider
- B- A data consumer
- C- A data reader
- D- A data administrator

Answer:

A

Explanation:

In Snowflake's Secure Data Sharing model, the account sharing data is known as the data provider. Providers create shares that include selected objects (tables, views, secure views, etc.) and grant access to consumer accounts.

A data consumer receives the shared data. Data reader and data administrator are not defined terms in Snowflake's sharing architecture.

=====

Question 7

Question Type: MultipleChoice

What is the maximum size limit for a single VARIANT value in Snowflake?

Options:

- A- 1MB
- B- 8MB
- C- 32MB
- D- 16MB

Answer:

D

Explanation:

A single VARIANT value in Snowflake has a maximum size limit of 16 MB (uncompressed). This limit applies to semi-structured data stored in VARIANT, such as JSON, XML, Avro, and Parquet objects.

This limit ensures efficient storage, query performance, and micro-partitioning. Exceeding the size results in load or processing errors. For ingestion, when loading large JSON files, it's recommended to break them into multiple objects or use JSON Lines formatting.

Snowflake has introduced an optional behavior change bundle (2025_03), increasing the limit to 128 MB for accounts that enable it, but the standard default remains 16 MB.

Thus, under default behavior, 16 MB is the correct maximum value size.

Question 8

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 9

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 10

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 11

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 12

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 13

Question Type: MultipleChoice

Which option best 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.



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