

## Give Push to your Success with SnowPro Core SnowPro-Core Exam Questions [Q274-Q297]



Give Push to your Success with SnowPro Core SnowPro-Core Exam Questions  
SnowPro-Core 100% Guarantee Download SnowPro-Core Exam PDF Q&A

The SnowPro-Core Certification Exam consists of multiple-choice questions, covering various aspects of Snowflake, including architecture, security, data loading, querying, and optimization. SnowPro-Core exam is conducted online, and candidates have two hours to complete it. SnowPro-Core exam is available in English and Japanese languages.

SnowPro-Core certification is suitable for professionals who work with Snowflake's data warehousing platform. This includes database administrators, data engineers, data architects, and data analysts. SnowPro-Core exam is also beneficial for professionals who are looking to enhance their skills and knowledge of Snowflake's platform. Additionally, businesses that use Snowflake can encourage their employees to take the exam to ensure that they have the necessary skills to work with the platform.

**NO.274** In which use cases does Snowflake apply egress charges?

\* Data sharing within a specific region

- \* Query result retrieval
- \* Database replication
- \* Loading data into Snowflake

Explanation

Cloud providers apply data egress

- \* Data is transferred from one region to another within the same cloud platform.
- \* Data is transferred out of the cloud platform.

<https://docs.snowflake.com/en/user-guide/billing-data-transfer.html#:~:text=Cloud%20providers%20apply%20d>

**NO.275** What service is provided as an integrated Snowflake feature to enhance Multi-Factor Authentication (MFA) support?

- \* Duo Security
- \* OAuth
- \* Okta
- \* Single Sign-On (SSO)

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

**NO.276** What is the minimum Snowflake edition needed for database failover and fail-back between Snowflake accounts for business continuity and disaster recovery?

- \* Standard
- \* Enterprise
- \* Business Critical
- \* Virtual Private Snowflake

The minimum Snowflake edition required for database failover and fail-back between Snowflake accounts for business continuity and disaster recovery is the Business Critical edition. References: Snowflake Documentation3.

**NO.277** Which Snowflake feature records changes made to a table so actions can be taken using that change data capture?

- \* Materialized View
- \* Pipe
- \* Stream
- \* Task

Snowflake's Streams feature is specifically designed for change data capture (CDC). A stream records insert, update, and delete operations performed on a table, and allows users to query these changes. This enables actions to be taken on the changed data, facilitating processes like incremental data loads and real-time analytics. Streams provide a powerful mechanism for applications to respond to data changes in Snowflake tables efficiently. References: Snowflake Documentation on Streams

**NO.278** True or False: A Snowflake account is charged for data stored in both Internal and External Stages.

- \* True
- \* False

**NO.279** Which privilege is required for a role to be able to resume a suspended warehouse if auto-resume is not enabled?

- \* USAGE
- \* OPERATE
- \* MONITOR
- \* MODIFY

Reference:

<https://community.snowflake.com/s/question/0D50Z00008yHYdqSAG/auto-resume-operate-warehouse-privileg>

**NO.280** What is the purpose of a Query Profile?

- \* To profile how many times a particular query was executed and analyze its usage statistics over time.
- \* To profile a particular query to understand the mechanics of the query, its behavior, and performance.
- \* To profile the user and/or executing role of a query and all privileges and policies applied on the objects within the query.
- \* To profile which queries are running in each warehouse and identify proper warehouse utilization and sizing for better performance and cost balancing.

The purpose of a Query Profile is to provide a detailed analysis of a particular query's execution plan, including the mechanics, behavior, and performance. It helps in identifying potential performance bottlenecks and areas for optimization

**NO.281** What are the correct parameters for time travel and fail-safe in the Snowflake Enterprise Edition?

- \* Default Time Travel Retention is set to 0 days.

Maximum Time Travel Retention is 30 days.

Fail Safe retention time is 1 day.

- \* Default Time Travel Retention is set to 1 day.

Maximum Time Travel Retention is 365 days.

Fail Safe retention time is 7 days.

- \* Default Time Travel Retention is set to 0 days.

Maximum Time Travel Retention is 90 days.

Fail Safe retention time is 7 days.

- \* Default Time Travel Retention is set to 1 day.

Maximum Time Travel Retention is 90 days.

Fail Safe retention time is 7 days.

- \* Default Time Travel Retention is set to 7 days.

Maximum Time Travel Retention is 1 day.

Fail Safe retention time is 90 days.

- \* Default Time Travel Retention is set to 90 days.

Maximum Time Travel Retention is 7 days.

Fail Safe retention time is 356 days.

**NO.282** 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?

- \* select src:dealership from car\_sales;
- \* select src.dealership from car\_sales;
- \* select src:Dealership from car\_sales;
- \* select dealership from car\_sales;

**NO.283** True or False: Snowpipe via RFST API can only reference External Stages as source.

- \* True
- \* False

**NO.284** Which of the following statements are true of VALIDATION\_MODE in Snowflake? (Choose two.)

- \* The validation\_mode option is used when creating an Internal Stage
- \* validation\_mode=return\_all\_errors is a parameter of the copy command
- \* The validation\_mode option will validate data to be loaded by the copy statement while completing the load and will return the rows that could not be loaded without error
- \* The validation\_mode option will validate data to be loaded by the copy statement without completing the load and will return possible errors

**NO.285** Two users share a virtual warehouse named wh dev 01. When one of the users loads data, the other one experiences performance issues while querying data.

How does Snowflake recommend resolving this issue?

- \* Scale up the existing warehouse.
- \* Create separate warehouses for each user.
- \* Create separate warehouses for each workload.
- \* Stop loading and querying data at the same time.

**NO.286** Which Snowflake table objects can be shared with other accounts? (Select TWO).

- \* Temporary tables
- \* Permanent tables
- \* Transient tables
- \* External tables
- \* User-Defined Table Functions (UDTFs)

In Snowflake, permanent tables and external tables can be shared with other accounts using Secure Data Sharing. Temporary tables, transient tables, and UDTFs are not shareable objects

**NO.287** Which SQL commands, when committed, will consume a stream and advance the stream offset? (Choose two.)

- \* UPDATE TABLE FROM STREAM
- \* SELECT FROM STREAM
- \* INSERT INTO TABLE SELECT FROM STREAM
- \* ALTER TABLE AS SELECT FROM STREAM
- \* BEGIN COMMIT

**NO.288** A virtual warehouse's auto-suspend and auto-resume settings apply to which of the following?

- \* The primary cluster in the virtual warehouse
- \* The entire virtual warehouse
- \* The database in which the virtual warehouse resides
- \* The Queries currently being run on the virtual warehouse

The auto-suspend and auto-resume settings in Snowflake apply to the entire virtual warehouse. These settings allow the warehouse to automatically suspend when it's not in use, helping to save on compute costs. When queries or tasks are submitted to the warehouse, it can automatically resume operation. This functionality is designed to optimize resource usage and cost-efficiency.

References:

- \* SnowPro Core Certification Exam Study Guide (as of 2021)
- \* Snowflake documentation on virtual warehouses and their settings (as of 2021)

**NO.289** The PUT command: (Choose two.)

- \* Automatically creates a File Format object
- \* Automatically uses the last Stage created
- \* Automatically compresses files using Gzip
- \* Automatically encrypts files

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

**NO.290** A user has an application that writes a new Tile to a cloud storage location every 5 minutes.

What would be the MOST efficient way to get the files into Snowflake?

- \* Create a task that runs a copy into operation from an external stage every 5 minutes
- \* Create a task that puts the files in an internal stage and automate the data loading wizard
- \* Create a task that runs a GET operation to intermittently check for new files
- \* Set up cloud provider notifications on the Tile location and use Snowpipe with auto-ingest

Explanation

<https://docs.snowflake.com/en/user-guide/data-load-snowpipe-intro.html>

**NO.291** If auto-suspend is enable for a Virtual Warehouse, he Warehouse is automatically suspended when:

- \* All Snowflake sessions using the warehouse are terminated.

- \* The last query using the warehouse completes.
- \* There are no users loaned into Snowflake.
- \* The Warehouse is inactive for a specified period of time.

**NO.292** Which interfaces can be used to create and/or manage Virtual Warehouses?

- \* The Snowflake Web Interface (UI)
- \* SQL commands
- \* Data integration tools
- \* All of the above

Reference: <https://docs.snowflake.com/en/user-guide/warehouses.html>

**NO.293** Which services does the Snowflake Cloud Services layer manage? (Select TWO).

- \* Compute resources
- \* Query execution
- \* Authentication
- \* Data storage
- \* Metadata

The Snowflake Cloud Services layer manages a variety of services that are crucial for the operation of the Snowflake platform. Among these services, Authentication and Metadata management are key components.

Authentication is essential for controlling access to the Snowflake environment, ensuring that only authorized users can perform actions within the platform. Metadata management involves handling all the metadata related to objects within Snowflake, such as tables, views, and databases, which is vital for the organization and retrieval of data.

References:

- \* [COF-C02] SnowPro Core Certification Exam Study Guide
- \* Snowflake Documentation12

<https://docs.snowflake.com/en/user-guide/intro-key-concepts.html>

**NO.294** What computer language can be selected when creating User-Defined Functions (UDFs) using the Snowpark API?

- \* Swift
- \* JavaScript
- \* Python
- \* SQL

**NO.295** True or False: Fail-safe can be disabled within a Snowflake account.

- \* True
- \* False

**NO.296** Which of the following statements is true of Snowflake micro-partitioning?

- \* Micro-partitioning has been known to introduce data skew
- \* Micro-partitioning: requires a partitioning schema to be defined up front
- \* Micro-partitioning is transparently completed using the ordering that occurs when the data is inserted/loaded
- \* Micro-partitioning can be disabled within a Snowflake account

**NO.297** Which of the following terms best describes Snowflake's database architecture?

- \* Columnar shared nothing

- \* Shared disk
- \* Multi-cluster, shared data
- \* Cloud-native shared memory

<https://www.snowflake.com/product/architecture/>

Built from the ground up for the cloud, Snowflake's unique multi-cluster shared data architecture delivers the performance, scale, elasticity, and concurrency today's organizations require.

**Get SnowPro-Core Actual Free Exam Q&As to Prepare Certification:**

<https://www.vceprep.com/SnowPro-Core-latest-vce-prep.html>