

ITPassLeader



Pass Your Next Certification Exam Fast!

Select a vendor... Select an test... Your email address [Free Download Demo](#)



Instant Download



365 Days Free Updates



Money Back Guarantee



Security & Privacy

Choose the version that fits your needs

PDF Version

Desktop Test Engine

Online Test Engine

Latest and Up-to-Date exam dumps with real exam questions answers.



Get 12-Months free updates without any extra charges.



Experience same exam environment before appearing in the certification exam.



100% exam passing guarantee in the first attempt.



20% discount on more than one license and 30% discount on 5+ license purchases.



100% secure purchase on SSL.



Completely private purchase without sharing your personal info with anyone.



<http://www.itpassleader.com>

High-praise Exam Dumps Questions grant you success by high pass rate - ITPassLeader

Exam : **SnowPro-Core**

Title : **SnowPro Core Certification Exam**

Vendor : **Snowflake**

Version : **DEMO**

QUESTION NO: 1

While working with unstructured data, which file function generates a Snowflake-hosted file URL to a staged file using the stage name and relative file path as inputs?

- A. GET_PRESIGNED_URL
- B. GET_ABSOLUTE_PATH
- C. BUILD_STAGE_FILE_URL
- D. BUILD SCOPED FILE URL

Answer: C

Explanation:

The BUILD_STAGE_FILE_URL function generates a Snowflake-hosted file URL to a staged file using the stage name and relative file path as inputs.

QUESTION NO: 2

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: B E

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

QUESTION NO: 3

What privilege should a user be granted to change permissions for new objects in a managed access schema?

- A. Grant the OWNERSHIP privilege on the schema.
- B. Grant the OWNERSHIP privilege on the database.
- C. Grant the MANAGE GRANTS global privilege.
- D. Grant ALL privileges on the schema.

Answer: C

Explanation:

To change permissions for new objects in a managed access schema, a user should be granted the MANAGE GRANTS global privilege. This privilege allows the user to manage access control through grants on all securable objects within Snowflake. References: [COF-C02] SnowPro Core Certification Exam Study Guide

QUESTION NO: 4

The following settings are configured:

The `MIN_DATA_RETENTION_TIME_IN_DAYS` is set to 5 at the account level.
The `DATA_RETENTION_TIME_IN_DAYS` is set to 2 at the object level.

For how many days will the data be retained at the object level?

- A. 2
- B. 3
- C. 5
- D. 7

Answer: A

Explanation:

The settings shown in the image indicate that the data retention time in days is configured at two different levels: the account level and the object level. At the account level, the `MIN_DATA_RETENTION_TIME_IN_DAYS` is set to 5 days, and at the object level, the `DATA_RETENTION_TIME_IN_DAYS` is set to 2 days. Since the object level setting has a lower value, it takes precedence over the account level setting for the specific object. Therefore, the data will be retained for 2 days at the object level. References: Snowflake Documentation on Data Retention Policies

QUESTION NO: 5

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: B D

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

QUESTION NO: 6

In which hierarchy is tag inheritance possible?

- A. Organization » Account » Role
- B. Account » User » Schema
- C. Database » View » Column
- D. Schema » Table » Column

Answer: D

Explanation:

In Snowflake, tag inheritance is a feature that allows tags, which are key-value pairs assigned to objects for the purpose of data governance and metadata management, to be

inherited within a hierarchy. The hierarchy in which tag inheritance is possible is from Schema to Table to Column. This means that a tag applied to a schema can be inherited by the tables within that schema, and a tag applied to a table can be inherited by the columns within that table. References: Snowflake Documentation on Tagging and Object Hierarchy

QUESTION NO: 7

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

QUESTION NO: 8

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

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

Answer: B C

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

QUESTION NO: 9

What happens when a cloned table is replicated to a secondary database? (Select TWO)

- A. A read-only copy of the cloned tables is stored.
- B. The replication will not be successful.
- C. The physical data is replicated
- D. Additional costs for storage are charged to a secondary account
- E. Metadata pointers to cloned tables are replicated

Answer: C E

Explanation:

When a cloned table is replicated to a secondary database in Snowflake, the following occurs:

* C. The physical data is replicated: The actual data of the cloned table is physically

replicated to the secondary database. This ensures that the secondary database has its own copy of the data, which can be used for read-only purposes or failover scenarios¹.

* E. Metadata pointers to cloned tables are replicated: Along with the physical data, the metadata pointers that refer to the cloned tables are also replicated. This metadata includes information about the structure of the table and any associated properties².

It's important to note that while the physical data and metadata are replicated, the secondary database is typically read-only and cannot be used for write operations. Additionally, while there may be additional storage costs associated with the secondary account, this is not a direct result of the replication process but rather a consequence of storing additional data.

References:

* SnowPro Core Exam Prep - Answers to Snowflake's LEVEL UP: Backup and Recovery

* Snowflake SnowPro Core Certification Exam Questions Set 10

QUESTION NO: 10

What is a directory table in Snowflake?

- A. A separate database object that is used to store file-level metadata
- B. An object layered on a stage that is used to store file-level metadata
- C. A database object with grantable privileges for unstructured data tasks
- D. A Snowflake table specifically designed for storing unstructured files

Answer: B

Explanation:

A directory table in Snowflake is an object layered on a stage that is used to store file-level metadata. It is not a separate database object but is conceptually similar to an external table because it stores metadata about the data files in the stage⁵.

QUESTION NO: 11

A company needs to allow some users to see Personally Identifiable Information (PII) while limiting other users from seeing the full value of the PII.

Which Snowflake feature will support this?

- A. Row access policies
- B. Data masking policies
- C. Data encryption
- D. Role based access control

Answer: B

Explanation:

Data masking policies in Snowflake allow for the obfuscation of specific data within a field, enabling some users to see the full data while limiting others. This feature is particularly useful for handling PII, ensuring that sensitive information is only visible to authorized users¹.

QUESTION NO: 12

What are common issues found by using the Query Profile? (Choose two.)

- A. Identifying queries that will likely run very slowly before executing them
- B. Locating queries that consume a high amount of credits
- C. Identifying logical issues with the queries

- D. Identifying inefficient micro-partition pruning
- E. Data spilling to a local or remote disk

Answer: D E

Explanation:

The Query Profile in Snowflake is used to identify performance issues with queries. Common issues that can be found using the Query Profile include identifying inefficient micro-partition pruning (D) and data spilling to a local or remote disk (E). Micro-partition pruning is related to the efficiency of query execution, and data spilling occurs when the memory is insufficient, causing the query to write data to disk, which can slow down the query performance¹.

QUESTION NO: 13

How can a data provider ensure that a data consumer is going to have access to the required objects?

- A. Enable the data sharing feature in the account and validate the view.
- B. Use the CURRENT_ROLE and CURRENT_USER functions to validate secure views.
- C. Use the CURRENT_ function to authorize users from a specific account to access rows in a base table.
- D. Set the SIMULATED DATA SHARING CONSUMER session parameter to the name of the consumer account for which access is being simulated.

Answer: A

Explanation:

To ensure a data consumer has access to the required objects, a data provider can enable the data sharing feature and validate that the consumer can access the views or tables shared with them. References: Based on general data sharing practices in cloud services as of 2021.

QUESTION NO: 14

Which command sets the Virtual Warehouse for a session?

- A. COPY WAREHOUSE FROM <<config file>>;
- B. SET WAREHOUSE = <<warehouse name>>;
- C. USE WAREHOUSE <<warehouse name>>;
- D. USE VIRTUAL_WAREHOUSE <<warehouse name>>;

Answer: C

Explanation:

The command USE WAREHOUSE <<warehouse name>>; is used to set the virtual warehouse for the current session in Snowflake. This command specifies which virtual warehouse to use for executing queries in that session¹.

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

QUESTION NO: 15

Which command is used to start configuring Snowflake for Single Sign-On (SSO)?

- A. CREATE SESSION POLICY
- B. CREATE NETWORK RULE
- C. CREATE SECURITY INTEGRATION

D. CREATE PASSWORD POLICY

Answer: C

Explanation:

To start configuring Snowflake for Single Sign-On (SSO), the CREATE SECURITY INTEGRATION command is used. This command sets up a security integration object in Snowflake, which is necessary for enabling SSO with external identity providers using SAML 2.01.

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

QUESTION NO: 16

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

QUESTION NO: 17

Which type of role can be granted to a share?

- A. Account role
- B. Custom role
- C. Database role
- D. Secondary role

Answer: B

Explanation:

In Snowflake, shares are used to share data between Snowflake accounts. When creating a share, it is possible to grant access to the share to roles within the Snowflake account that is creating the share. The type of role that can be granted to a share is a Custom role. Custom roles are user-defined roles that account administrators can create to manage access control in a more granular way. Unlike predefined roles such as ACCOUNTADMIN or SYSADMIN, custom roles can be tailored with specific privileges to meet the security and access requirements of different groups within an organization.

Granting a custom role access to a share enables users associated with that role to access the shared data if the share is received by another Snowflake account. It is important to carefully manage the privileges granted to custom roles to ensure that data sharing aligns with organizational policies and data governance standards.

References:

- * Snowflake Documentation on Shares: Shares
- * Snowflake Documentation on Roles: Access Control

QUESTION NO: 18

Which Snowflake database object can be used to track data changes made to table data?

- A. Tag
- B. Task
- C. Stream
- D. Stored procedure

Answer: C

Explanation:

A Stream object in Snowflake is used for change data capture (CDC), which records data manipulation language (DML) changes made to tables, including inserts, updates, and deletes³.

QUESTION NO: 19

What does the TableScan operator represent in the Query Profile?

- A. The access to a single table
- B. The access to data stored in stage objects
- C. The list of values provided with the VALUES clause
- D. The records generated using the TABLE (GENERATOR (...)) construct

Answer: A

Explanation:

In the Query Profile of Snowflake, the TableScan operator represents the access to a single table. This operator indicates that the query execution involved reading data from a table stored in Snowflake. TableScan is a fundamental operation in query execution plans, showing how the database engine retrieves data directly from tables as part of processing a query.

References:

- * Snowflake Documentation: Understanding the Query Profile

QUESTION NO: 20

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 rekeying¹. References: [COF-C02] SnowPro Core Certification Exam Study Guide

QUESTION NO: 21

In the Snowflake access control model, which entity owns an object by default?

- A. The user who created the object
- B. The SYSADMIN role
- C. Ownership depends on the type of object
- D. The role used to create the object

Answer: D

Explanation:

In Snowflake's access control model, the default owner of an object is the role that was used to create the object. This role has the OWNERSHIP privilege on the object and can grant access to other roles¹

QUESTION NO: 22

What is the MAXIMUM Time Travel retention period for a transient table?

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

Answer: B

Explanation:

The maximum Time Travel retention period for a transient table in Snowflake is 1 day. This is the default and maximum duration for which Snowflake maintains the historical data for transient tables, allowing users to query data as it appeared at any point within the past 24 hours².

QUESTION NO: 23

What is a responsibility of Snowflake's virtual warehouses?

- A. Infrastructure management
- B. Metadata management
- C. Query execution
- D. Query parsing and optimization
- E. Management of the storage layer

Answer: C

Explanation:

The primary responsibility of Snowflake's virtual warehouses is to execute queries. Virtual warehouses are one of the key components of Snowflake's architecture, providing the compute power required to perform data processing tasks such as running SQL queries, performing joins, aggregations, and other data manipulations.

References:

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

* Snowflake Documentation on Virtual Warehouses¹

QUESTION NO: 24

What Snowflake role must be granted for a user to create and manage accounts?

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

Answer: A

Explanation:

The ACCOUNTADMIN role is required for a user to create and manage accounts in Snowflake. This role has the highest level of privileges and is responsible for managing all aspects of the Snowflake account, including the ability to create and manage other user accounts¹.

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

QUESTION NO: 25

By default, which Snowflake role is required to create a share?

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

Answer: D

Explanation:

By default, the Snowflake role required to create a share is ACCOUNTADMIN (D). This role has the necessary privileges to perform administrative tasks, including creating shares for data sharing purposes

QUESTION NO: 26

What action can a user take to address query concurrency issues?

- A. Enable the query acceleration service.
- B. Enable the search optimization service.
- C. Add additional clusters to the virtual warehouse
- D. Resize the virtual warehouse to a larger instance size.

Answer: C

Explanation:

To address query concurrency issues, a user can add additional clusters to the virtual warehouse. This allows for the distribution of queries across multiple clusters, reducing the load on any single cluster and improving overall query performance².

QUESTION NO: 27

Which Snowflake table supports unstructured data?

- A. Directory
- B. Transient
- C. Temporary
- D. Permanent

Answer: D

Explanation:

While Snowflake primarily deals with structured and semi-structured data, it also has the capability to handle unstructured data. Unstructured data can be stored in Snowflake using variants of SQL data types in tables, which can be permanent tables. These permanent tables, while traditionally used for structured or semi-structured data (like JSON, Avro, or Parquet), can also accommodate unstructured data in the form of binary formats or strings, offering flexibility in data storage and analysis. However, the management and querying of unstructured data in Snowflake may require additional considerations compared to structured data. References: Snowflake Documentation on Data Types

QUESTION NO: 28

Which privilege is required to use the search optimization service in Snowflake?

- A. GRANT SEARCH OPTIMIZATION ON SCHEMA <schema_name> TO ROLE <role>
- B. GRANT SEARCH OPTIMIZATION ON DATABASE <database_name> TO ROLE <role>
- C. GRANT ADD SEARCH OPTIMIZATION ON SCHEMA <schema_name> TO ROLE <role>
- D. GRANT ADD SEARCH OPTIMIZATION ON DATABASE <database name> TO ROLE <role>

Answer: C

Explanation:

To utilize the search optimization service in Snowflake, the correct syntax for granting privileges to a role involves specific commands that include adding search optimization capabilities:

* Option C: GRANT ADD SEARCH OPTIMIZATION ON SCHEMA <schema_name> TO ROLE

<role>. This command grants the specified role the ability to implement search optimization at the schema level, which is essential for enhancing search capabilities within that schema. Options A and B do not include the correct verb "ADD," which is necessary for this specific type of grant command in Snowflake. Option D incorrectly mentions the database level, as search optimization privileges are typically configured at the schema level, not the database level. References: Snowflake documentation on the use of GRANT statements for configuring search optimization.

QUESTION NO: 29

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. References: Understanding and viewing Fail-safe | Snowflake Documentation

QUESTION NO: 30

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.

QUESTION NO: 31

A marketing co-worker has requested the ability to change a warehouse size on their medium virtual warehouse called mktg__WH.

Which of the following statements will accommodate this request?

- A. ALLOW RESIZE ON WAREHOUSE MKTG__WH TO USER MKTG__LEAD;
- B. GRANT MODIFY ON WAREHOUSE MKTG WH TO ROLE MARKETING;
- C. GRANT MODIFY ON WAREHOUSE MKTG__WH TO USER MKTG__LEAD;
- D. GRANT OPERATE ON WAREHOUSE MKTG WH TO ROLE MARKETING;

Answer: B

Explanation:

The correct statement to accommodate the request for a marketing co-worker to change the size of their medium virtual warehouse called mktg__WH is to grant the MODIFY privilege on the warehouse to the ROLE MARKETING. This privilege allows the role to change the warehouse size among other properties.

References:

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

* Snowflake Documentation on Access Control Privileges1

QUESTION NO: 32

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

QUESTION NO: 33

At what level is the MIN_DATA_RETENTION_TIME_IN_DAYS parameter set?

- A. Account
- B. Database
- C. Schema
- D. Table

Answer: A

Explanation:

The MIN_DATA_RETENTION_TIME_IN_DAYS parameter is set at the account level. This parameter determines the minimum number of days Snowflake retains historical data for Time Travel operations

QUESTION NO: 34

Which of the following describes the Snowflake Cloud Services layer?

- A. Coordinates activities in the Snowflake account
- B. Executes queries submitted by the Snowflake account users
- C. Manages quotas on the Snowflake account storage
- D. Manages the virtual warehouse cache to speed up queries

Answer: A

Explanation:

The Snowflake Cloud Services layer is a collection of services that coordinate activities across Snowflake, tying together all the different components to process user requests, from login to query dispatch1.

References = [COF-C02] SnowPro Core Certification Exam Study Guide, Snowflake Documentation1

QUESTION NO: 35

When is the result set cache no longer available? (Select TWO)

- A. When another warehouse is used to execute the query
- B. When another user executes the query
- C. When the underlying data has changed
- D. When the warehouse used to execute the query is suspended
- E. When it has been 24 hours since the last query

Answer: C E

Explanation:

The result set cache in Snowflake is invalidated and no longer available when the underlying data of the query results has changed, ensuring that queries return the most current data. Additionally, the cache expires after 24 hours to maintain the efficiency and accuracy of data

retrieval1.