



Up-to-date Questions and Answers from authentic resources to improve knowledge and pass the exam at very first attempt. ----- Guaranteed.



DEA-C01 MCQs
DEA-C01 TestPrep
DEA-C01 Study Guide
DEA-C01 Practice Test
DEA-C01 Exam Questions



killexams.com

SnowFlake

DEA-C01

SnowPro Advanced Data Engineer

ORDER FULL VERSION

<https://killexams.com/pass4sure/exam-detail/DEA-C01>



Question: 62

Each micro partition contains between 50 mb and 500 MB of uncompressed data
A. TRUE
B. FALSE

Answer: A

Explanation

What are Micro-partitions?

All data in Snowflake tables is automatically divided into micro-partitions, which are contiguous units of storage. Each micro-partition contains between 50 MB and 500 MB of uncompressed data (note that the actual size in Snowflake is smaller because data is always stored compressed). Groups of rows in tables are mapped into individual micro-partitions, organized in a columnar fashion. This size and structure allows for extremely granular pruning of very large tables, which can be comprised of millions, or even hundreds of millions, of micro-partitions.

Snowflake stores metadata about all rows stored in a micro-partition, including:

Question: 63

An existing clustering key is copied in which of the below scenarios
A. CREATE TABLE...CLONE
B. CREATE TABLE...LIKE
C. CREATE TABLE...AS SELECT

Answer: A

Explanation

https://docs.snowflake.com/en/sql-reference/functions/system_estimate_search_optimization_costs.html#out

put

BuildCosts

This object describes the predicted costs of building the search access path for the table. If search optimization has already been added to the table, this object contains no cost information. StorageCosts

This object describes the predicted amount of storage space (in TB) needed for the search access path for the table. If search optimization has already been added to the table, this object shows the current amount of space used by the search access path.

Benefit

This object does not contain any cost information at this time.

MaintenanceCosts

This object describes the predicted costs of maintaining the search access path for the table. If this table has been created recently, this object does not contain any cost information.

Question: 64

CREATE OR REPLACE TABLE TIME_TRAVEL_SCHEMA.TIME_TRAVEL_TABLE (ID NUMBER) DATA_RETENTION_TIME_IN_DAYS =20;

Later you dropped the schema. In this scenario what data retention value will be honored for the table, if we need to retrieve the table data

A. 10

B. 20

C. 30

Answer: A

Explanation

<https://docs.snowflake.com/en/user-guide/data-time-travel.html#dropped-containers-and-object-retention-inheritance>

Dropped Containers and Object Retention Inheritance

Currently, when a database is dropped, the data retention period for child schemas or tables, if explicitly set to be different from the retention of the database, is not honored. The child schemas or tables are retained for the same period of time as the database.

Similarly, when a schema is dropped, the data retention period for child tables, if explicitly set to be different from the retention of the schema, is not honored. The child tables are retained for the same period of time as the schema.

To honor the data retention period for these child objects (schemas or tables), drop them explicitly before you drop the database or schema.

Question: 65

Which system table will you use to get the total credit consumption over a specific time period?

- A. WAREHOUSE_METERING_HISTORY
- B. WAREHOUSE_CREDIT_USAGE_HISTORY
- C. WAREHOUSE_USAGE_HISTORY

Answer: A

Explanation

The WAREHOUSE_METERING_HISTORY table in the ACCOUNT_USAGE Schema can be used to get the desired information. Run the below query to try this out.

```
SELECT WAREHOUSE_NAME, SUM(CREDITS_USED_COMPUTE) AS CREDITS_USED_COMPUTE_SUM
FROM ACCOUNT_USAGE.WAREHOUSE_METERING_HISTORY
GROUP BY 1
ORDER BY 2 DESC;
```

Question: 66

Snowpipe supports loading from both internal and external stage

- A. FALSE
- B. TRUE

Answer: B

Explanation

Step 1: Create a Stage (If Needed)

Snowpipe supports loading from the following stage types:

Question: 67

Each micro partition contains between 50 mb and 500 MB of uncompressed data

- A. TRUE
- B. FALSE

Answer: A

Explanation

What are Micro-partitions?

All data in Snowflake tables is automatically divided into micro-partitions, which are contiguous units of storage. Each micro-partition contains between 50 MB and 500 MB of uncompressed data (note that the actual size in Snowflake is smaller because data is always stored compressed). Groups of rows in tables are mapped into individual micro-partitions, organized in a columnar fashion. This size and structure allows for extremely granular pruning of very large tables, which can be comprised of millions, or even hundreds of millions, of micro-partitions.

Snowflake stores metadata about all rows stored in a micro-partition, including:

Question: 68

Which of the below statements are true?

- A. ACCOUNT USAGE includes dropped objects but INFORMATION SCHEMA does not
- B. INFORMATION SCHEMA includes dropped objects but ACCOUNT USAGE does not
- C. BOTH includes dropped object
- D. BOTH does not include dropped object

Answer: A

Explanation

<https://docs.snowflake.com/en/sql-reference/account-usage.html#differences-between-account-usage-and-information-schema>

Dropped Object Records

Account usage views include records for all objects that have been dropped. An additional DELETED column displays the timestamp when the object was dropped.

In addition, because objects can be dropped and recreated with the same name, to differentiate between objects records that have the same name, the account usage views include ID columns, where appropriate, that display the internal IDs generated and assigned to each record by the system.

Question: 69

Which of the below statements are true?

- A. ACCOUNT USAGE includes dropped objects but INFORMATION SCHEMA does not
- B. INFORMATION SCHEMA includes dropped objects but ACCOUNT USAGE does not
- C. BOTH includes dropped object
- D. BOTH does not include dropped object

Answer: A

Explanation

<https://docs.snowflake.com/en/sql-reference/account-usage.html#differences-between-account-usage-and-information-schema>

Dropped Object Records

Account usage views include records for all objects that have been dropped. An additional DELETED column displays the timestamp when the object was dropped.

In addition, because objects can be dropped and recreated with the same name, to differentiate between objects records that have the same name, the account usage views include ID columns, where appropriate, that display the internal IDs generated and assigned to each record by the system.

Question: 70

CREATE OR REPLACE TABLE TIME_TRAVEL_SCHEMA.TIME_TRAVEL_TABLE (ID NUMBER) DATA_RETENTION_TIME_IN_DAYS =20;

Later you dropped the schema. In this scenario what data retention value will be honored for the table, if we need to retrieve the table data

- A. 10
- B. 20
- C. 30

Answer: A

Explanation

<https://docs.snowflake.com/en/user-guide/data-time-travel.html#dropped-containers-and-object-retention-inheritance>

Dropped Containers and Object Retention Inheritance

Currently, when a database is dropped, the data retention period for child schemas or tables, if explicitly set to be different from the retention of the database, is not honored. The child schemas or tables are retained for the same period of time as the database.

Similarly, when a schema is dropped, the data retention period for child tables, if explicitly set to be different from the retention of the schema, is not honored. The child tables are retained for the same period of time as the schema.

To honor the data retention period for these child objects (schemas or tables), drop them explicitly before you drop the database or schema.

Question: 71

Time travel cannot be disabled for an account, but it can be disabled for individual databases, schemas and tables by specifying `DATA_RETENTION_TIME_IN_DAYS` with a value of 0 for the object

- A. TRUE
- B. FALSE

Answer: A

Explanation

<https://docs.snowflake.com/en/user-guide/data-time-travel.html#enabling-and-disabling-time-travel>

Enabling and Disabling Time Travel

No tasks are required to enable Time Travel. It is automatically enabled with the standard, 1-day retention period.

However, you may wish to upgrade to Snowflake Enterprise Edition to enable configuring longer data retention periods of up to 90 days for databases, schemas, and tables. Note that extended data retention requires additional storage which will be reflected in your monthly storage charges. For more information about storage charges, see [Storage Costs for Time Travel and Fail-safe](#).

Time Travel cannot be disabled for an account; however, it can be disabled for individual databases, schemas, and tables by specifying `DATA_RETENTION_TIME_IN_DAYS` with a value of 0 for the object. Also, users with the `ACCOUNTADMIN` role can set `DATA_RETENTION_TIME_IN_DAYS` to 0 at the account level, which means that all databases (and subsequently all schemas and tables) created in the account have no retention period by default; however, this default can be overridden at any time for any database, schema, or table.

Question: 72

```
CREATE OR REPLACE TABLE TIME_TRAVEL_SCHEMA.TIME_TRAVEL_TABLE (ID NUMBER) DATA_RETENTION_TIME_IN_DAYS =20;
```

Later you dropped the schema. In this scenario what data retention value will be honored for the table, if we need to retrieve the table data

- A. 10
- B. 20
- C. 30

Answer: A

Explanation

<https://docs.snowflake.com/en/user-guide/data-time-travel.html#dropped-containers-and-object-retention-inheritance>

Dropped Containers and Object Retention Inheritance

Currently, when a database is dropped, the data retention period for child schemas or tables, if explicitly set to be different from the retention of the database, is not honored. The child schemas or tables are retained for the same period of time as the database.

Similarly, when a schema is dropped, the data retention period for child tables, if explicitly set to be different from the retention of the schema, is not honored. The child tables are retained for the same period of time as the schema.

To honor the data retention period for these child objects (schemas or tables), drop them explicitly before you drop the database or schema.

Question: 73

Snowpipe supports loading from both internal and external stage

- A. FALSE
- B. TRUE

Answer: B

Explanation

Step 1: Create a Stage (If Needed)

Snowpipe supports loading from the following stage types:

Question: 74

Snowflake charges a per-byte fee when users transfer data from your snowflake account into cloud storages in another region on the same cloud platform or into cloud storage in another cloud platform

- A. TRUE
- B. FALSE

Answer: A

Explanation

<https://docs.snowflake.com/en/user-guide/billing-data-transfer.html#understanding-snowflake-data-transfer-billing>

Cloud providers apply data egress charges in either of the following use cases:

Data is transferred from one region to another within the same cloud platform.

Data is transferred out of the cloud platform.

To recover these expenses, Snowflake charges a per-byte fee when users transfer data from your Snowflake account (hosted on AWS, Google Cloud Platform, or Microsoft Azure) into cloud storage in another region on the same cloud platform, or into cloud storage in another cloud platform.

The amount charged per byte depends on the region where your Snowflake account is hosted. For data transfer pricing, see the pricing guide (on the Snowflake website):

Question: 75

```
;  
  
call sp1();  
  
SELECT * FROM EMPLOYEE ORDER BY ID;  
A. 1 MOHAN  
2 RON  
B. 1 MOHAN  
2 RON  
3 RANJAN  
C. 1 MOHAN  
3 RANJAN  
D. 1 MOHAN
```

Answer: A

Explanation

<https://docs.snowflake.com/en/sql-reference/transactions.html#scoped-transactions>

Scoped Transactions

A stored procedure that contains a transaction can be called from within another transaction. For example, a transaction inside a stored procedure can include a call to another stored procedure that contains a transaction.

Snowflake does not treat the inner transaction as nested; instead, the inner transaction is a separate transaction. Snowflake calls these “autonomous scoped transactions” (or simply “scoped transactions”). The starting point and ending point of each scoped transaction determine which statements are included in the transaction. The start and end can be explicit or implicit. Each SQL statement is part of only one transaction. An enclosing ROLLBACK or COMMIT does not undo an enclosed COMMIT or ROLLBACK.

Question: 76

Which of the below statements are true for API integration object?

- A. Only Snowflake users who have the ACCOUNTADMIN role or who have a role with the global CREATE INTEGRATION privilege can execute CREATE API INTEGRATION.
- B. Only Snowflake roles with OWNERSHIP or USAGE privileges on the API integration can use the API integration directly (e.g. by creating an external function that specifies that API integration).
- C. An API integration object is tied to a specific cloud platform account and role within that account, but not to a specific HTTPS proxy URL. You can create more than one instance of an HTTPS proxy service in a cloud provider account, and you can use the same API integration to authenticate to multiple proxy services in that account.
- D. Your Snowflake account can have multiple API integration objects, for example, for different cloud platform accounts.
- E. Multiple external functions can use the same API integration object, and thus the same HTTPS proxy service.
- F. ALL OF THE ABOVE

Answer: F

Explanation

API integration is an important topic to focus on for the certification. Please read this topic thoroughly. <https://docs.snowflake.com/en/sql-reference/sql/create-api-integration.html#create-api-integration> Usage Notes

Only Snowflake users who have the ACCOUNTADMIN role or who have a role with the global CREATE INTEGRATION privilege can execute CREATE API INTEGRATION.

Only Snowflake roles with OWNERSHIP or USAGE privileges on the API integration can use the API integration directly (e.g. by creating an external function that specifies that API integration).

An API integration object is tied to a specific cloud platform account and role within that account, but not to a specific HTTPS proxy URL. You can create more than one instance of an HTTPS proxy service in a cloud provider account, and you can use the same API integration to authenticate to multiple proxy services in that account.

Your Snowflake account can have multiple API integration objects, for example, for different cloud platform accounts.

Multiple external functions can use the same API integration object, and thus the same HTTPS proxy service.

Question: 77

Time travel cannot be disabled for an account, but it can be disabled for individual databases, schemas and tables by specifying DATA_RETENTION_TIME_IN_DAYS with a value of 0 for the object

- A. TRUE
- B. FALSE

Answer: A

Explanation

<https://docs.snowflake.com/en/user-guide/data-time-travel.html#enabling-and-disabling-time-travel>

Enabling and Disabling Time Travel

No tasks are required to enable Time Travel. It is automatically enabled with the standard, 1-day retention period.

However, you may wish to upgrade to Snowflake Enterprise Edition to enable configuring longer data retention periods of up to 90 days for databases, schemas, and tables. Note that extended data retention requires additional storage which will be reflected in your monthly storage charges. For more information about storage charges, see Storage Costs for Time Travel and Fail-safe.

Time Travel cannot be disabled for an account; however, it can be disabled for individual databases, schemas, and tables by specifying DATA_RETENTION_TIME_IN_DAYS with a value of 0 for the object. Also, users with the ACCOUNTADMIN role can set DATA_RETENTION_TIME_IN_DAYS to 0 at the account level, which means that all databases (and subsequently all schemas and tables) created in the account have no retention period by default; however, this default can be overridden at any time for any database, schema, or table.

Killexams.com is a leading online platform specializing in high-quality certification exam preparation. Offering a robust suite of tools, including MCQs, practice tests, and advanced test engines, Killexams.com empowers candidates to excel in their certification exams. Discover the key features that make Killexams.com the go-to choice for exam success.



Exam Questions:

Killexams.com provides exam questions that are experienced in test centers. These questions are updated regularly to ensure they are up-to-date and relevant to the latest exam syllabus. By studying these questions, candidates can familiarize themselves with the content and format of the real exam.

Exam MCQs:

Killexams.com offers exam MCQs in PDF format. These questions contain a comprehensive collection of questions and answers that cover the exam topics. By using these MCQs, candidate can enhance their knowledge and improve their chances of success in the certification exam.

Practice Test:

Killexams.com provides practice test through their desktop test engine and online test engine. These practice tests simulate the real exam environment and help candidates assess their readiness for the actual exam. The practice test cover a wide range of questions and enable candidates to identify their strengths and weaknesses.

Guaranteed Success:

Killexams.com offers a success guarantee with the exam MCQs. Killexams claim that by using this materials, candidates will pass their exams on the first attempt or they will get refund for the purchase price. This guarantee provides assurance and confidence to individuals preparing for certification exam.

Updated Contents:

Killexams.com regularly updates its question bank of MCQs to ensure that they are current and reflect the latest changes in the exam syllabus. This helps candidates stay up-to-date with the exam content and increases their chances of success.