


[Oct 05, 2023 DEA-C01 Exam Dumps - Snowflake Practice Test Questions [Q62-Q80]



[Oct 05, 2023] DEA-C01 Exam Dumps - Snowflake Practice Test Questions [Q62-Q80]

[Oct 05, 2023] DEA-C01 Exam Dumps - Snowflake Practice Test Questions
New Real DEA-C01 Exam Dumps Questions

Q62. By default, a newly-created Custom role is not assigned to any user, nor granted to any other role?

- * TRUE
- * FALSE

Q63. The smaller the average depth, the better clustered the table is with regards to the specified column?

- * TRUE
- * FALSE

Q64. If the data retention period for a table is less than 90 days, and a stream has not been consumed, Snowflake temporarily extends this period to prevent it from going stale?

- * TRUE
- * FALSE

Explanation

If the data retention period for a table is less than 14 days, and a stream has not been consumed, Snowflake temporarily extends this period to prevent it from going stale. The period is extended to the stream's offset, up to a maximum of 14 days by default, regardless of the Snowflake edition for your account. The maximum number of days for which Snowflake can extend the data retention period is determined by the `MAX_DATA_EXTENSION_TIME_IN_DAYS` parameter value. When the stream is consumed, the extended data retention period is reduced to the default period for the table.

Q65. In efforts to recover the dropped child tables within schema named `SCV_SCHEMA` by Data Engineer, She found that `DATA_RETENTION_TIME_IN_DAYS` parameter set with value 45 days at Schema level & the data retention period for child tables explicitly set at 85 days. What will happen when she will try to run `undrop table` command on Child tables to recover them on the 50th day assuming `SCV_SCHEMA` is already dropped on 45th day?

- * To honor the data retention period for child tables, She will be able to recover the child tables on 50th day as `DATA_RETENTION_TIME_IN_DAYS` is explicitly set with higher retention value.
- * When a schema is already dropped, the data retention period for child tables, if explicitly set to be different from the retention of the schema, is not honored. So `UNDROP` command will fail to run on

50th day for Child tables recovery.

- * Child tables can be recovered using Fail-Safe SQL commands.
- * Data Engineer needs to first recover the Schema & then Child tables will automatically be recovered irrespective of Retention Inheritance.

Explanation

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.

Q66. Which privilege are required on an object (i.e. user or role) with `USERADMIN` Role can modify the object properties?

- * `OPERATE`
- * `MANAGE GRANTS`
- * `OWNERSHIP`
- * `MODIFY`

Q67. Can the same column be specified in both a Dynamic data masking policy signature and a row access policy signature at the same time?

- * YES
- * NO

Q68. A SQL UDF evaluates an arbitrary SQL expression and returns the result(s) of the expression. Which value type it can return?

- * Single Value
- * A Set of Rows
- * Scaler or Tabular depend on input SQL expression
- * Regex

Q69. Which is the non-supportable JavaScript UDF data types?

- * Integers
- * String
- * Binary
- * Double

Q70. To troubleshoot data load failure in one of your Copy Statement, Data Engineer have Executed a COPY statement with the VALIDATION_MODE copy option set to RETURN_ALL_ERRORS with reference to the set of files he had attempted to load. Which below function can facilitate analysis of the problematic records on top of the Results produced? [Select 2]

- * RESULT_SCAN
- * LAST_QUERY_ID
- * Rejected_record
- * LOAD_ERROR

Explanation

LAST_QUERY_ID() Function

Returns the ID of a specified query in the current session. If no query is specified, the most recently executed query is returned.

RESULT_SCAN() Function

Returns the result set of a previous command (within 24 hours of when you executed the query) as if the result was a table.

The following example validates a set of files (SFfile.csv.gz) that contain errors. To facilitate analysis of the errors, a COPY INTO <location> statement then unloads the problematic records into a text file so they could be analyzed and fixed in the original data files. The statement queries the RESULT_SCAN table.

1.#copy into Snowtable

2.from @SFstage/SFfile.csv.gz

3.validation_mode=return_all_errors;

4.#set qid=last_query_id();

5.#copy into @SFstage/errors/load_errors.txt from (select rejected_record from table(result_scan(\$qid))); Note: Other options are not valid functions.

Q71. As Data Engineer, you have been asked to access data held in AWS Glacier Deep Archive storage class for Historical Data Analysis, which one is the correct statement to recommend?

- * You cannot access data held in archival cloud storage classes that requires restoration before it can be retrieved.
- * Loading data from AWS cloud storage services is supported regardless of the cloud platform that hosts your Snowflake account.
- * Data can be accessed from External stage using AWS Private link in this case.
- * We can simply access AWS Glacier Deep Archive storage External Stage data using PUT command.
- * Upload (i.e. stage) files to your cloud storage account using the tools provided by the cloud storage service.

Explanation

External stage

References data files stored in a location outside of Snowflake. Currently, the following cloud storage services are supported:

Amazon S3 buckets

Google Cloud Storage buckets

Microsoft Azure containers

The storage location can be either private/protected or public.

You cannot access data held in archival cloud storage classes that requires restoration before it can be retrieved. These archival storage classes include, for example, the Amazon S3 Glacier Flexible Retrieval or Glacier Deep Archive storage class, or Microsoft Azure Archive Storage.

Q72. Jeff, a Data Engineer, accessing elements in JSON object in its 3 data loading scripts, he unknow-ingly use the upper case while accessing the elements. e.g.

Script 1 >> fruits:apple.sweet

Script 2 >> FRUITS:apple.sweet

Script 3 >> FRUITS:Apple.Sweet

Which are the correct statements?

- * Script 1 & Script 2 traverse path twill be treated same, but Script 2 will not.
- * Script 1,2,3 traverse path will treat in same way.
- * Script 1 & 3 traverse path will be treated in same way.
- * Script 2&3 traverse path will be same.

Explanation

There are two ways to access elements in a JSON object:

Dot Notation

Bracket Notation

Regardless of which notation you use, the column name is case-insensitive but element names are case-sensitive.

Q73. While creating External function, Which Database object required with at least ACCOUNTAD-MIN privileges?

- * STORAGE Integration
- * SECURITY Integration
- * API Integration
- * None of the above required.

Q74. To view/monitor the clustering metadata for a table, Snowflake provides which of the following system functions?

- * SYSTEM\$CLUSTERING_DEPTH_KEY
- * SYSTEM\$CLUSTERING_KEY_INFORMATION (including clustering depth)
- * SYSTEM\$CLUSTERING_DEPTH
- * SYSTEM\$CLUSTERING_INFORMATION (including clustering depth)

Explanation

SYSTEM\$CLUSTERING_DEPTH:

Computes the average depth of the table according to the specified columns (or the clustering key defined for the table). The average depth of a populated table (i.e. a table containing data) is always 1 or more. The smaller the average depth, the better clustered the table is with regards to the specified columns.

Calculate the clustering depth for a table using two columns in the table:

```
SELECT SYSTEM$CLUSTERING_DEPTH(TPCH_PRODUCT, (C2, C9));
```

SYSTEM\$CLUSTERING_INFORMATION:

Returns clustering information, including average clustering depth, for a table based on one or more columns in the table.

```
SELECT SYSTEM$CLUSTERING_INFORMATION(SAMPLE_TABLE, (col1, col3));
```

Q75. Search optimization works best to improve the performance of a query when the following conditions are true:[Select All that apply]

- * The table is not clustered.
- * The table is frequently queried on columns other than the primary cluster key.
- * Search Query uses Equality predicates (for example, <column_name> = <constant>) OR Predicates that use IN.
- * Search Query uses Sort Operations.

Explanation

Materialized Views works best for search query performance in case of Sort Operations. For Rest of the points Search optimization works best to improve query performance.

Q76. Snowflake does not provide which of following set of SQL functions to support retrieving information about tasks?

- * SYSTEM\$CURRENT_USER_TASK_NAME
- * TASK_HISTORY
- * TASK_DEPENDENTS
- * TASK_QUERY_HISTORY
- * SYSTEM\$TASK_DEPENDENTS_ENABLE

Explanation

SYSTEM\$CURRENT_USER_TASK_NAME

Returns the name of the task currently executing when invoked from the statement or stored procedure defined by the task.

SYSTEM\$TASK_DEPENDENTS_ENABLE

Recursively resumes all dependent tasks tied to a specified root task.

TASK_DEPENDENTS

This table function returns the list of child tasks for a given root task in a DAG of tasks.

TASK_HISTORY

This table function can be used to query the history of task usage within a specified date range.

Q77. Which of the below concepts/functions helps while implementing advanced Column-level Security?

- * CURRENT_ROLE
- * INVOKER_ROLE
- * Role Hierarchy
- * CURRENT_CLIENT

Explanation

Column-level Security supports using Context Functions in the conditions of the masking policy body to enforce whether a user has authorization to see data. To determine whether a user can see data in a given SQL statement, it is helpful to consider:

Masking policy conditions using CURRENT_ROLE target the role in use for the current session.

Masking policy conditions using INVOKER_ROLE target the executing role in a SQL statement.

Role hierarchy

Determine if a specified role in a masking policy condition (e.g. ANALYST custom role) is a lower privilege role in the CURRENT_ROLE or INVOKER_ROLE role hierarchy. If so, then the role returned by the CURRENT_ROLE or INVOKER_ROLE functions inherits the privileges of the specified role.

Q78. Data Engineer wants to analyze query performance & looking out for profiling information, He went to Query/Operator Details also called Profile Overview of Query Profile Interface & searching for statistics attributes around I/O. Which of the following information he can get from there?

- * Percentage scanned from cache; the percentage of data scanned from the local disk cache.
- * Bytes written; bytes written (e.g. when loading into a table).
- * External bytes scanned; bytes read from an external object, e.g. a stage.
- * Bytes sent over the wireframe; amount of data sent over the wireframe
- * Bytes read from result; bytes read from the result object.

Explanation

To help you analyze query performance, Query/Operator Details panel also called Profile overview panel provides two classes of profiling information:

Execution time, broken down into categories

Detailed statistics

Apart from Option a Bytes sent over the wireframe; amount of data sent over the wireframe, Rest of the Statistics Information provided by Query/Operator details in the Query Profile Inter-face.

To Know More about the Query/Operator Details options, please refer the link:

<https://docs.snowflake.com/en/user-guide/ui-query-profile#query-operator-details>

Q79. Pascal, a Data Engineer, have requirement to retrieve the 10 most recent executions of a specified task (completed, still running, or scheduled in the future) scheduled within the last hour, which of the following is the correct SQL Code ?

- * 1.select *
- 2.from table(information_schema.task_history(

3.scheduled_time_range_start=>dateadd(#{hour};,-1,current_timestamp()),

4.result_limit => 10,

5.task_name=>MYTASK;) WHERE query_id IS NOT NULL);

* 1.select *

2.from table(information_schema.task_history(

3.scheduled_time_range_start=>dateadd(#{hour};,-1,current_timestamp()),

4.result_limit => 11,

5.task_name=>MYTASK;) WHERE query_id IS NOT NULL);

* 1.select *

2.from table(information_schema.task_history(

3.scheduled_time_range_start=>dateadd(#{hour};,-1,current_timestamp()),

4.result_limit => 10,query_id IS NOT NULL

5.task_name=>MYTASK;));

* 1.select *

2.from table(information_schema.task_history(

3.scheduled_time_range_start=>dateadd(#{hour};,-1,current_timestamp()),

4.result_limit => 10,

5.task_name=>MYTASK;));

Explanation

To retrieve only tasks that are completed or still running, filter the query using WHERE query_id IS NOT NULL.

Q80. Which Role inherits the privileges of the USERADMIN role via the system role hierarchy?

- * SYSADMIN
- * SECURITYADMIN
- * PUBLIC
- * CUSTOM ROLE

Pass Your DEA-C01 Exam Easily with Accurate PDF Questions: https://www.braindumpsit.com/DEA-C01_real-exam.html