

## Exam Questions DEA-C01

SnowPro Advanced: Data Engineer Certification Exam

<https://www.2passeasy.com/dumps/DEA-C01/>



**NEW QUESTION 1**

What is a characteristic of the use of binding variables in JavaScript stored procedures in Snowflake?

- A. All types of JavaScript variables can be bound
- B. All Snowflake first-class objects can be bound
- C. Only JavaScript variables of type number, string and sf Date can be bound
- D. Users are restricted from binding JavaScript variables because they create SQL injection attack vulnerabilities

**Answer: C**

**Explanation:**

A characteristic of the use of binding variables in JavaScript stored procedures in Snowflake is that only JavaScript variables of type number, string and sf Date can be bound. Binding variables are a way to pass values from JavaScript variables to SQL statements within a stored procedure. Binding variables can improve the security and performance of the stored procedure by preventing SQL injection attacks and reducing the parsing overhead. However, not all types of JavaScript variables can be bound. Only the primitive types number and string, and the Snowflake-specific type sf Date, can be bound. The other options are incorrect because they do not describe a characteristic of the use of binding variables in JavaScript stored procedures in Snowflake. Option A is incorrect because authenticator is not a type of JavaScript variable, but a parameter of the snowflake.connector.connect function. Option B is incorrect because arrow\_number\_to\_decimal is not a type of JavaScript variable, but a parameter of the snowflake.connector.connect function. Option D is incorrect because users are not restricted from binding JavaScript variables, but encouraged to do so.

**NEW QUESTION 2**

A Data Engineer is building a pipeline to transform a 1 TD table by joining it with supplemental tables. The Engineer is applying filters and several aggregations leveraging Common Table Expressions (CTEs) using a size Medium virtual warehouse in a single query in Snowflake.

After checking the Query Profile, what is the recommended approach to MAXIMIZE performance of this query if the Profile shows data spillage?

- A. Enable clustering on the table
- B. Increase the warehouse size
- C. Rewrite the query to remove the CTEs.
- D. Switch to a multi-cluster virtual warehouse

**Answer: B**

**Explanation:**

The recommended approach to maximize performance of this query if the Profile shows data spillage is to increase the warehouse size. Data spillage occurs when the query requires more memory than the warehouse can provide and has to spill some intermediate results to disk. This can degrade the query performance by increasing the disk IO time. Increasing the warehouse size can increase the amount of memory available for the query and reduce or eliminate data spillage.

**NEW QUESTION 3**

A new customer table is created by a data pipeline in a Snowflake schema where MANAGED ACCESS is enabled.

.... Can gran access to the CUSTOMER table? (Select THREE.)

- A. The role that owns the schema
- B. The role that owns the database
- C. The role that owns the customer table
- D. The SYSADMIN role
- E. The SECURITYADMIN role
- F. The USERADMIN role with the manage grants privilege

**Answer: ABE**

**Explanation:**

The roles that can grant access to the CUSTOMER table are the role that owns the schema, the role that owns the database, and the SECURITYADMIN role. These roles have the ownership or the manage grants privilege on the schema or the database level, which allows them to grant access to any object within them. The other options are incorrect because they do not have the necessary privilege to grant access to the CUSTOMER table. Option C is incorrect because the role that owns the customer table cannot grant access to itself or to other roles. Option D is incorrect because the SYSADMIN role does not have the manage grants privilege by default and cannot grant access to objects that it does not own. Option F is incorrect because the USERADMIN role with the manage grants privilege can only grant access to users and roles, not to tables.

**NEW QUESTION 4**

A Data Engineer would like to define a file structure for loading and unloading data. Where can the file structure be defined? (Select THREE)

- A. copy command
- B. MERGE command
- C. FILE FORMAT Object
- D. pipe object
- E. stage object
- F. INSERT command

**Answer: ACE**

**Explanation:**

The places where the file format can be defined are copy command, file format object, and stage object. These places allow specifying or referencing a file format that defines how data files are parsed and loaded into or unloaded from Snowflake tables. A file format can include various options, such as field delimiter, field enclosure, compression type, date format, etc. The other options are not places where the file format can be defined. Option B is incorrect because MERGE command is a SQL command that can merge data from one table into another based on a join condition, but it does not involve loading or unloading data files. Option D is incorrect because pipe object is a Snowflake object that can load data from an external stage into a Snowflake table using COPY statements, but it does not define or reference a file format. Option F is incorrect because INSERT command is a SQL command that can insert data into a Snowflake table from

literal values or subqueries, but it does not involve loading or unloading data files.

#### NEW QUESTION 5

Which query will show a list of the 20 most recent executions of a specified task kttask, that have been scheduled within the last hour that have ended or are stillrunning's.

A)

```
select * from table(information_schema.task_history(scheduled_time_range_start
=>dateadd('hour',-1,current_timestamp()), result_limit => 20,
task_name=>'MYTASK'))
```

B)

```
select * from table(information_schema.task_history(scheduled_time_range_start
=>dateadd('hour',-1,current_timestamp()), result_limit => 20,
task_name=>'MYTASK')) where query_id IS NOT NULL;
```

C)

```
select * from table(information_schema.task_history(scheduled_time_range_start
=>dateadd('hour',-1,current_timestamp()), result_limit => 20,
task_name=>'MYTASK')) where STATE IN ('EXECUTING', 'SUCCEEDED', 'FAILED')
```

D)

```
select * from table(information_schema.task_history(scheduled_time_range_end
=>dateadd('hour',-1,current_timestamp()), result_limit => 10,
task_name=>'MYTASK')) where STATE IN ('EXECUTING', 'SUCCEEDED')
```

- A. Option A
- B. Option B
- C. Option C
- D. Option D

**Answer:** B

#### NEW QUESTION 6

A Data Engineer enables a result cache at the session level with the following command: ALTER SESSION SET USE\_CACHED\_RESULT = TRUE; The Engineer then runs the following select query twice without delay:

```
SELECT *
FROM SNOWFLAKE_SAMPLE_DATA.TPCH_SF1.CUSTOMER
SAMPLE(10) SEED (99);
```

The underlying table does not change between executions What are the results of both runs?

- A. The first and second run returned the same results because sample is deterministic
- B. The first and second run returned the same results, because the specific SEEDvalue was provided.
- C. The first and second run returned different results because the query is evaluated each time it is run.
- D. The first and second run returned differentresults because the query uses \*instead of an explicit column list

**Answer:** B

#### Explanation:

The result cache is enabled at the session level, which means that repeated queries will return cached results if there is no change in the underlying data or session parameters. However, in this case, the result cache is not relevant because the query uses a specific SEED value for sampling, which makes it deterministic. Therefore, both runs will return the same results regardless of caching.

#### NEW QUESTION 7

A Data Engineer has written a stored procedure that will run with caller's rights. The Engineer has granted ROLEA right to use this stored procedure. What is a characteristic of the stored procedure being called using ROLEA?

- A. The stored procedure must run with caller's rights it cannot be converted later to run with owner's rights
- B. If the stored procedure accesses an object that ROLEA does not have access to the stored procedure will fail
- C. The stored procedure will run in the context (database and schema) where the owner created the stored procedure
- D. ROLEA will not be able to see the source code for the stored procedure even though the role has usage privileges on the stored procedure

**Answer:** B

#### Explanation:

A stored procedure that runs with caller's rights executes with the privileges of the role that calls it. Therefore, if the stored procedure accesses an object that ROLEA does not have access to, such as a table or a view, the stored procedure will fail with an insufficient privileges error. The other options are not correct because:

- ? A stored procedure can be converted from caller's rights to owner's rights by using the ALTER PROCEDURE command with the EXECUTE AS OWNER option.
- ? A stored procedure that runs with caller's rights executes in the context (database and schema) of the caller, not the owner.
- ? ROLEA will be able to see the source code for the stored procedure by using the GET\_DDL function or the DESCRIBE command, as long as it has usage privileges on the stored procedure.

#### NEW QUESTION 8

Company A and Company B both have Snowflake accounts. Company A's account is hosted on a different cloud provider and region than Company B's account Companies A and B are not in the same Snowflake organization. How can Company A share data with Company B? (Select TWO).

- A. Create a share within Company A's account and add Company B's account as a recipient of that share
- B. Create a share within Company A's account, and create a reader account that is a recipient of the share Grant Company B access to the reader account
- C. Use database replication to replicate Company A's data into Company B's account Create a share within Company B's account and grant users within Company B's account access to the share
- D. Create a new account within Company A's organization in the same cloud provider and region as Company B's account Use database replication to replicate Company A's data to the new account Create a share within the new account and add Company B's account as a recipient of that share
- E. Create a separate database within Company A's account to contain only those data sets they wish to share with Company B Create a share within Company A's account and add all the objects within this separate database to the share Add Company B's account as a recipient of the share

**Answer:** AE

**Explanation:**

The ways that Company A can share data with Company B are:

- ? Create a share within Company A's account and add Company B's account as a recipient of that share: This is a valid way to share data between different accounts on different cloud platforms and regions. Snowflake supports cross-cloud and cross-region data sharing, which allows users to create shares and grant access to other accounts regardless of their cloud platform or region. However, this option may incur additional costs for network transfer and storage replication.
- ? Create a separate database within Company A's account to contain only those data sets they wish to share with Company B Create a share within Company A's account and add all the objects within this separate database to the share Add Company B's account as a recipient of the share: This is also a valid way to share data between different accounts on different cloud platforms and regions. This option is similar to the previous one, except that it uses a separate database to isolate the data sets that need to be shared. This can improve security and manageability of the shared data. The other options are not valid because:
- ? Create a share within Company A's account, and create a reader account that is a recipient of the share Grant Company B access to the reader account: This option is not valid because reader accounts are not supported for cross-cloud or cross-region data sharing. Reader accounts are Snowflake accounts that can only consume data from shares created by their provider account. Reader accounts must be on the same cloud platform and region as their provider account.
- ? Use database replication to replicate Company A's data into Company B's account Create a share within Company B's account and grant users within Company B's account access to the share: This option is not valid because database replication cannot be used for cross-cloud or cross-region data sharing. Database replication is a feature in Snowflake that allows users to copy databases across accounts within the same cloud platform and region. Database replication cannot copy databases across different cloud platforms or regions.
- ? Create a new account within Company A's organization in the same cloud provider and region as Company B's account Use database replication to replicate Company A's data to the new account Create a share within the new account and add Company B's account as a recipient of that share: This option is not valid because it involves creating a new account within Company A's organization, which may not be feasible or desirable for Company A. Moreover, this option is unnecessary, as Company A can directly share data with Company B without creating an intermediate account.

**NEW QUESTION 9**

What is the purpose of the BUILD\_FILE\_URL function in Snowflake?

- A. It generates an encrypted URL for accessing a file in a stage.
- B. It generates a staged URL for accessing a file in a stage.
- C. It generates a permanent URL for accessing files in a stage.
- D. It generates a temporary URL for accessing a file in a stage.

**Answer:** B

**Explanation:**

The BUILD\_FILE\_URL function in Snowflake generates a temporary URL for accessing a file in a stage. The function takes two arguments: the stage name and the file path. The generated URL is valid for 24 hours and can be used to download or view the file contents. The other options are incorrect because they do not describe the purpose of the BUILD\_FILE\_URL function.

**NEW QUESTION 10**

A Data Engineer is building a set of reporting tables to analyze consumer requests by region for each of the Data Exchange offerings annually, as well as click-through rates for each listing  
 Which views are needed MINIMALLY as data sources?

- A. SNOWFLAKE- DATA\_SHARING\_USAGE - LISTING\_EVENTS\_BAILY
- B. SNOWFLAKE.DATA\_SHARING\_USAGE.LISTING\_CONSOKE>TION\_DAILY
- C. SNOWFLAK
- D. DATA\_SHARING\_USAG
- E. LISTING\_TELEMETRY\_DAILY
- F. SNOWFLAKE.ACCOUNT\_USAGE.DATA \_TRANSFER\_HISTORY

**Answer:** B

**Explanation:**

The SNOWFLAKE.DATA SHARING \_USAGE.LISTING\_CONSOKE>TION\_DAILY view provides information about consumer requests by region for each of the Data Exchange offerings annually, as well as click-through rates for each listing. This view is the minimal data source needed for building the reporting tables. The other views are not relevant for this use case.

**NEW QUESTION 10**

A Data Engineer is working on a Snowflake deployment in AWS eu-west-1 (Ireland). The Engineer is planning to load data from staged files into target tables using the copy into command  
 Which sources are valid? (Select THREE)

- A. Internal stage on GCP us-central1 (Iowa)
- B. Internal stage on AWS eu-central-1 (Frankfurt)
- C. External stage on GCP us-central1 (Iowa)
- D. External stage in an Amazon S3 bucket on AWS eu-west-1 (Ireland)
- E. External stage in an Amazon S3 bucket on AWS eu-central 1 (Frankfurt)
- F. SSO attached to an Amazon EC2 instance on AWS eu-west-1 (Ireland)

**Answer:** CDE

**Explanation:**

The valid sources for loading data from staged files into target tables using the copy into command are:

- ? External stage on GCP us-central1 (Iowa): This is a valid source because Snowflake supports cross-cloud data loading from external stages on different cloud platforms and regions than the Snowflake deployment.
- ? External stage in an Amazon S3 bucket on AWS eu-west-1 (Ireland): This is a valid source because Snowflake supports data loading from external stages on the same cloud platform and region as the Snowflake deployment.
- ? External stage in an Amazon S3 bucket on AWS eu-central 1 (Frankfurt): This is a valid source because Snowflake supports cross-region data loading from external stages on different regions than the Snowflake deployment within the same cloud platform. The invalid sources are:
- ? Internal stage on GCP us-central1 (Iowa): This is an invalid source because internal stages are always located on the same cloud platform and region as the Snowflake deployment. Therefore, an internal stage on GCP us-central1 (Iowa) cannot be used for a Snowflake deployment on AWS eu-west-1 (Ireland).
- ? Internal stage on AWS eu-central-1 (Frankfurt): This is an invalid source because internal stages are always located on the same region as the Snowflake deployment. Therefore, an internal stage on AWS eu-central-1 (Frankfurt) cannot be used for a Snowflake deployment on AWS eu-west-1 (Ireland).
- ? SSO attached to an Amazon EC2 instance on AWS eu-west-1 (Ireland): This is an invalid source because SSO stands for Single Sign-On, which is a security integration feature in Snowflake, not a data staging option.

**NEW QUESTION 11**

A Data Engineer wants to check the status of a pipe named my\_pipe. The pipe is inside a database named test and a schema named Extract (case-sensitive). Which query will provide the status of the pipe?

- A. SELECT FROM SYSTEM\$PIPE\_STATUS ('test.'extract'.my\_pipe');
- B. SELECT FROM SYSTEM\$PIPE\_STATUS (,test,,Extract,,ny\_pipe, |
- C. SELE2T \* FROM SYSTEM\$PIPE\_STATUS < ' tes
- D. "Extract", my\_pipe');
- E. SELECT \* FROM SYSTEM\$PIPE\_STATUS ("tes
- F. 'extract' .my\_pipe");

**Answer: C**

**Explanation:**

The query that will provide the status of the pipe is SELECT \* FROM SYSTEM\$PIPE\_STATUS('test."Extract".my\_pipe');. The SYSTEM\$PIPE\_STATUS function returns information about a pipe, such as its name, status, last received message timestamp, etc. The function takes one argument: the pipe name in a qualified form. The pipe name should include the database name, the schema name, and the pipe name, separated by dots. If any of these names are case-sensitive identifiers, they should be enclosed in double quotes. In this case, the schema name Extract is case-sensitive and should be quoted. The other options are incorrect because they do not follow the correct syntax for the pipe name argument. Option A and B use single quotes instead of double quotes for case-sensitive identifiers. Option D uses double quotes instead of single quotes for non-case-sensitive identifiers.

**NEW QUESTION 12**

Assuming that the session parameter USE\_CACHED\_RESULT is set to false, what are characteristics of Snowflake virtual warehouses in terms of the use of Snowpark?

- A. Creating a DataFrame from a table will start a virtual warehouse
- B. Creating a DataFrame from a staged file with the read () method will start a virtual warehouse
- C. Transforming a DataFrame with methods like replace () will start a virtual warehouse -
- D. Calling a Snowpark stored procedure to query the database with session, call () will start a virtual warehouse

**Answer: A**

**Explanation:**

Creating a DataFrame from a table will start a virtual warehouse because it requires reading data from Snowflake. The other options will not start a virtual warehouse because they either operate on local data or use an existing session to query Snowflake.

**NEW QUESTION 13**

A secure function returns data coming through an inbound share  
 What will happen if a Data Engineer tries to assign usage privileges on this function to an outbound share?

- A. An error will be returned because the Engineer cannot share data that has already been shared
- B. An error will be returned because only views and secure stored procedures can be shared
- C. An error will be returned because only secure functions can be shared with inboundshares
- D. The Engineer will be able to share the secure function with other accounts

**Answer: A**

**Explanation:**

An error will be returned because the Engineer cannot share data that has already been shared. A secure function is a Snowflake function that can access data from an inbound share, which is a share that is created by another account and consumed by the current account. A secure function can only be shared with an inbound share, not an outbound share, which is a share that is created by the current account and shared with other accounts. This is to prevent data leakage or unauthorized access to the data from the inbound share.

**NEW QUESTION 15**

When would a Data engineer use table with the flatten function instead of the lateral flatten combination?

- A. When TABLE with FLATTENrequires another source in the from clause to refer to
- B. WhenTABLE with FLATTENrequires no additional source m the from clause to refer to
- C. Whenthe LATERALFLATTENcombination requires no other source m the from clause to refer to
- D. When table withFLATTENis acting like a sub-query executed for each returned row

**Answer: A**

**Explanation:**

The TABLE function with the FLATTEN function is used to flatten semi-structured data, such as JSON or XML, into a relational format. The TABLE function returns a table expression that can be used in the FROM clause of a query. The TABLE function with the FLATTEN function requires another source in the FROM clause to refer to, such as a table, view, or subquery that contains the semi-structured data. For example: `SELECT t.value:city::string AS city, f.value AS population FROM cities t, TABLE(FLATTEN(input => t.value:population)) f;` In this example, the TABLE function with the FLATTEN function refers to the cities table in the FROM clause, which contains JSON data in a variant column named value. The FLATTEN function flattens the population array within each JSON object and returns a table expression with two columns: key and value. The query then selects the city and population values from the table expression.

#### NEW QUESTION 18

Which functions will compute a 'fingerprint' over an entire table, query result, or window to quickly detect changes to table contents or query results? (Select TWO).

- A. HASH (\*)
- B. HASH\_AGG(\*)
- C. HASH\_AGG(<expr>, <expr>)
- D. HASH\_AGG\_COMPARE (\*)
- E. HASH\_COMPARE(\*)

**Answer:** BC

#### Explanation:

The functions that will compute a 'fingerprint' over an entire table, query result, or window to quickly detect changes to table contents or query results are:

? HASH\_AGG(\*): This function computes a hash value over all columns and rows in a table, query result, or window. The function returns a single value for each group defined by a GROUP BY clause, or a single value for the entire input if no GROUP BY clause is specified.

? HASH\_AGG(<expr>, <expr>): This function computes a hash value over two expressions in a table, query result, or window. The function returns a single value for each group defined by a GROUP BY clause, or a single value for the entire input if no GROUP BY clause is specified. The other functions are not correct because:

? HASH (\*): This function computes a hash value over all columns in a single row. The function returns one value per row, not one value per table, query result, or window.

? HASH\_AGG\_COMPARE (): This function compares two hash values computed by HASH\_AGG() over two tables or query results and returns true if they are equal or false if they are different. The function does not compute a hash value itself, but rather compares two existing hash values.

? HASH\_COMPARE(): This function compares two hash values computed by HASH() over two rows and returns true if they are equal or false if they are different. The function does not compute a hash value itself, but rather compares two existing hash values.

#### NEW QUESTION 19

A Data Engineer needs to know the details regarding the micro-partition layout for a table named invoice using a built-in function. Which query will provide this information?

- A. `SELECT SYSTEM$CLUSTERING_INFORMATION('Invoice');`
- B. `SELECT $CLUSTERING_INFORMATION('Invoice');`
- C. `CALL SYSTEM$CLUSTERING_INFORMATION('Invoice');`
- D. `CALL $CLUSTERING_INFORMATION('Invoice');`

**Answer:** A

#### Explanation:

The query that will provide information about the micro-partition layout for a table named invoice using a built-in function is `SELECT SYSTEM$CLUSTERING_INFORMATION('Invoice');`. The `SYSTEM$CLUSTERING_INFORMATION` function returns information about the clustering status of a table, such as the clustering key, the clustering depth, the clustering ratio, the partition count, etc. The function takes one argument: the table name in a qualified or unqualified form. In this case, the table name is Invoice and it is unqualified, which means that it will use the current database and schema as the context. The other options are incorrect because they do not use a valid built-in function for providing information about the micro-partition layout for a table. Option B is incorrect because it uses `$CLUSTERING_INFORMATION` instead of `SYSTEM$CLUSTERING_INFORMATION`, which is not a valid function name. Option C is incorrect because it uses `CALL` instead of `SELECT`, which is not a valid way to invoke a table function. Option D is incorrect because it uses `CALL` instead of `SELECT` and `$CLUSTERING_INFORMATION` instead of `SYSTEM$CLUSTERING_INFORMATION`, which are both invalid.

#### NEW QUESTION 22

Which Snowflake feature facilitates access to external API services such as geocoders, data transformation, machine Learning models and other custom code?

- A. Security integration
- B. External tables
- C. External functions
- D. Java User-Defined Functions (UDFs)

**Answer:** C

#### Explanation:

External functions are Snowflake functions that facilitate access to external API services such as geocoders, data transformation, machine learning models and other custom code. External functions allow users to invoke external services from within SQL queries and pass arguments and receive results as JSON values. External functions require creating an API integration object and an external function object in Snowflake, as well as deploying an external service endpoint that can communicate with Snowflake via HTTPS.

#### NEW QUESTION 27

A Data Engineer is writing a Python script using the Snowflake Connector for Python. The Engineer will use the `snowflake.connector.connect` function to connect to Snowflake. The requirements are:

\*Raise an exception if the specified database schema or warehouse does not exist

\*improve download performance

Which parameters of the connect function should be used? (Select TWO).

- A. authenticator
- B. arrow\_number\_to\_decimal
- C. client\_prefetch\_threads
- D. client\_session\_keep\_alive
- E. validate\_default\_parameters

**Answer:** CE

**Explanation:**

The parameters of the connect function that should be used are `client_prefetch_threads` and `validate_default_parameters`. The `client_prefetch_threads` parameter controls the number of threads used to download query results from Snowflake. Increasing this parameter can improve download performance by parallelizing the download process. The `validate_default_parameters` parameter controls whether an exception should be raised if the specified database, schema, or warehouse does not exist or is not authorized. Setting this parameter to True can help catch errors early and avoid unexpected results.

**NEW QUESTION 31**

Which methods can be used to create a DataFrame object in Snowpark? (Select THREE)

- A. `session.jdbc_connection()`
- B. `session.read.json()`
- C. `session.table()`
- D. `DataFrame.write()`
- E. `session.builder()`
- F. `session.sql()`

**Answer:** BCF

**Explanation:**

The methods that can be used to create a DataFrame object in Snowpark are `session.read.json()`, `session.table()`, and `session.sql()`. These methods can create a DataFrame from different sources, such as JSON files, Snowflake tables, or SQL queries.

The other options are not methods that can create a DataFrame object in Snowpark. Option A, `session.jdbc_connection()`, is a method that can create a JDBC connection object to connect to a database. Option D, `DataFrame.write()`, is a method that can write a DataFrame to a destination, such as a file or a table. Option E, `session.builder()`, is a method that can create a SessionBuilder object to configure and build a Snowpark session.

**NEW QUESTION 36**

A Data Engineer executes a complex query and wants to make use of Snowflake's query results caching capabilities to reuse the results.

Which conditions must be met? (Select THREE).

- A. The results must be reused within 72 hours.
- B. The query must be executed using the same virtual warehouse.
- C. The `USED_CACHED_RESULT` parameter must be included in the query.
- D. The table structure contributing to the query result cannot have changed.
- E. The new query must have the same syntax as the previously executed query.
- F. The micro-partitions cannot have changed due to changes to other data in the table.

**Answer:** ADE

**Explanation:**

Snowflake's query results caching capabilities allow users to reuse the results of previously executed queries without re-executing them. For this to happen, the following conditions must be met:

? The results must be reused within 24 hours (not 72 hours), which is the default time-to-live (TTL) for cached results.

? The query must be executed using any virtual warehouse (not necessarily the same one), as long as it is in the same region and account as the original query.

? The `USED_CACHED_RESULT` parameter does not need to be included in the query, as it is enabled by default at the account level. However, it can be disabled or overridden at the session or statement level.

? The table structure contributing to the query result cannot have changed, such as adding or dropping columns, changing data types, or altering constraints.

? The new query must have the same syntax as the previously executed query, including whitespace and case sensitivity.

? The micro-partitions cannot have changed due to changes to other data in the table, such as inserting, updating, deleting, or merging rows.

**NEW QUESTION 39**

.....

## THANKS FOR TRYING THE DEMO OF OUR PRODUCT

Visit Our Site to Purchase the Full Set of Actual DEA-C01 Exam Questions With Answers.

We Also Provide Practice Exam Software That Simulates Real Exam Environment And Has Many Self-Assessment Features. Order the DEA-C01 Product From:

<https://www.2passeasy.com/dumps/DEA-C01/>

### Money Back Guarantee

#### **DEA-C01 Practice Exam Features:**

- \* DEA-C01 Questions and Answers Updated Frequently
- \* DEA-C01 Practice Questions Verified by Expert Senior Certified Staff
- \* DEA-C01 Most Realistic Questions that Guarantee you a Pass on Your FirstTry
- \* DEA-C01 Practice Test Questions in Multiple Choice Formats and Updatesfor 1 Year