

## Exam Questions DEA-C01

SnowPro Advanced: Data Engineer Certification Exam

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



### NEW QUESTION 1

Which use case would be BEST suited for the search optimization service?

- A. Analysts who need to perform aggregates over high cardinality columns
- B. Business users who need fast response times using highly selective filters
- C. Data Scientists who seek specific JOIN statements with large volumes of data
- D. Data Engineers who create clustered tables with frequent reads against clustering keys

**Answer: B**

#### Explanation:

The use case that would be best suited for the search optimization service is business users who need fast response times using highly selective filters. The search optimization service is a feature that enables faster queries on tables with high cardinality columns by creating inverted indexes on those columns. High cardinality columns are columns that have a large number of distinct values, such as customer IDs, product SKUs, or email addresses. Queries that use highly selective filters on high cardinality columns can benefit from the search optimization service because they can quickly locate the relevant rows without scanning the entire table. The other options are not best suited for the search optimization service. Option A is incorrect because analysts who need to perform aggregates over high cardinality columns will not benefit from the search optimization service, as they will still need to scan all the rows that match the filter criteria. Option C is incorrect because data scientists who seek specific JOIN statements with large volumes of data will not benefit from the search optimization service, as they will still need to perform join operations that may involve shuffling or sorting data across nodes. Option D is incorrect because data engineers who create clustered tables with frequent reads against clustering keys will not benefit from the search optimization service, as they already have an efficient way to organize and access data based on clustering keys.

### NEW QUESTION 2

Within a Snowflake account permissions have been defined with custom roles and role hierarchies. To set up column-level masking using a role in the hierarchy of the current user, what command would be used?

- A. CORRECT\_ROLE
- B. IKVOKER\_ROLE
- C. IS\_ROLE\_IN\_SESSION
- D. IS\_GRANTED\_TO\_INVOKER\_ROLE

**Answer: C**

#### Explanation:

The IS\_ROLE\_IN\_SESSION function is used to set up column-level masking using a role in the hierarchy of the current user. Column-level masking is a feature in Snowflake that allows users to apply dynamic data masking policies to specific columns based on the roles of the users who access them. The IS\_ROLE\_IN\_SESSION function takes a role name as an argument and returns true if the role is in the current user's session, or false otherwise. The function can be used in a masking policy expression to determine whether to mask or unmask a column value based on the role of the user. For example:

```
CREATE OR REPLACE MASKING POLICY email_mask AS (val string) RETURNS string -
> CASE WHEN IS_ROLE_IN_SESSION('HR') THEN val ELSE REGEXP_REPLACE(val, '(.)@.', '\1****\2') END;
```

In this example, the IS\_ROLE\_IN\_SESSION function is used to create a masking policy for an email column. The masking policy returns the original email value if the user has the HR role in their session, or returns a masked email value with asterisks if not.

### NEW QUESTION 3

Which methods will trigger an action that will evaluate a DataFrame? (Select TWO)

- A. DataFrame.random\_split ()
- B. DataFrame.collect ()
- C. DataFrame.select ()
- D. DataFrame.col ()
- E. DataFrame.show ()

**Answer: BE**

#### Explanation:

The methods that will trigger an action that will evaluate a DataFrame are DataFrame.collect() and DataFrame.show(). These methods will force the execution of any pending transformations on the DataFrame and return or display the results. The other options are not methods that will evaluate a DataFrame. Option A, DataFrame.random\_split(), is a method that will split a DataFrame into two or more DataFrames based on random weights. Option C, DataFrame.select(), is a method that will project a set of expressions on a DataFrame and return a new DataFrame. Option D, DataFrame.col(), is a method that will return a Column object based on a column name in a DataFrame.

### NEW QUESTION 4

Database XYZ has the data\_retention\_time\_in\_days parameter set to 7 days and table xyz.public.ABC has the data\_retention\_time\_in\_days set to 10 days. A Developer accidentally dropped the database containing this single table 8 days ago and just discovered the mistake. How can the table be recovered?

- A. undrop database xyz;
- B. create table abc\_restore as select \* from xyz.public.abc at {offset => -60\*60\*24\*8};
- C. create table abc\_restore clone xyz.public.abc at (offset => -3\*60\*24\*3);
- D. Create a Snowflake Support case to restore the database and table from "a i-safe"

**Answer: A**

#### Explanation:

The table can be recovered by using the undrop database xyz; command. This command will restore the database that was dropped within the last 14 days, along with all its schemas and tables, including the customer table. The data\_retention\_time\_in\_days parameter does not affect this command, as it only applies to time travel queries that reference historical data versions of tables or databases. The other options are not valid ways to recover the table. Option B is incorrect

because creating a table as `select * from xyz.public.ABC at {offset => -6060248}` will not work, as this query will try to access a historical data version of the ABC table that does not exist anymore after dropping the database. Option C is incorrect because creating a table clone `xyz.public.ABC at {offset => -360024*3}` will not work, as this query will try to clone a historical data version of the ABC table that does not exist anymore after dropping the database. Option D is incorrect because creating a Snowflake Support case to restore the database and table from fail-safe will not work, as fail-safe is only available for disaster recovery scenarios and cannot be accessed by customers.

#### NEW QUESTION 5

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

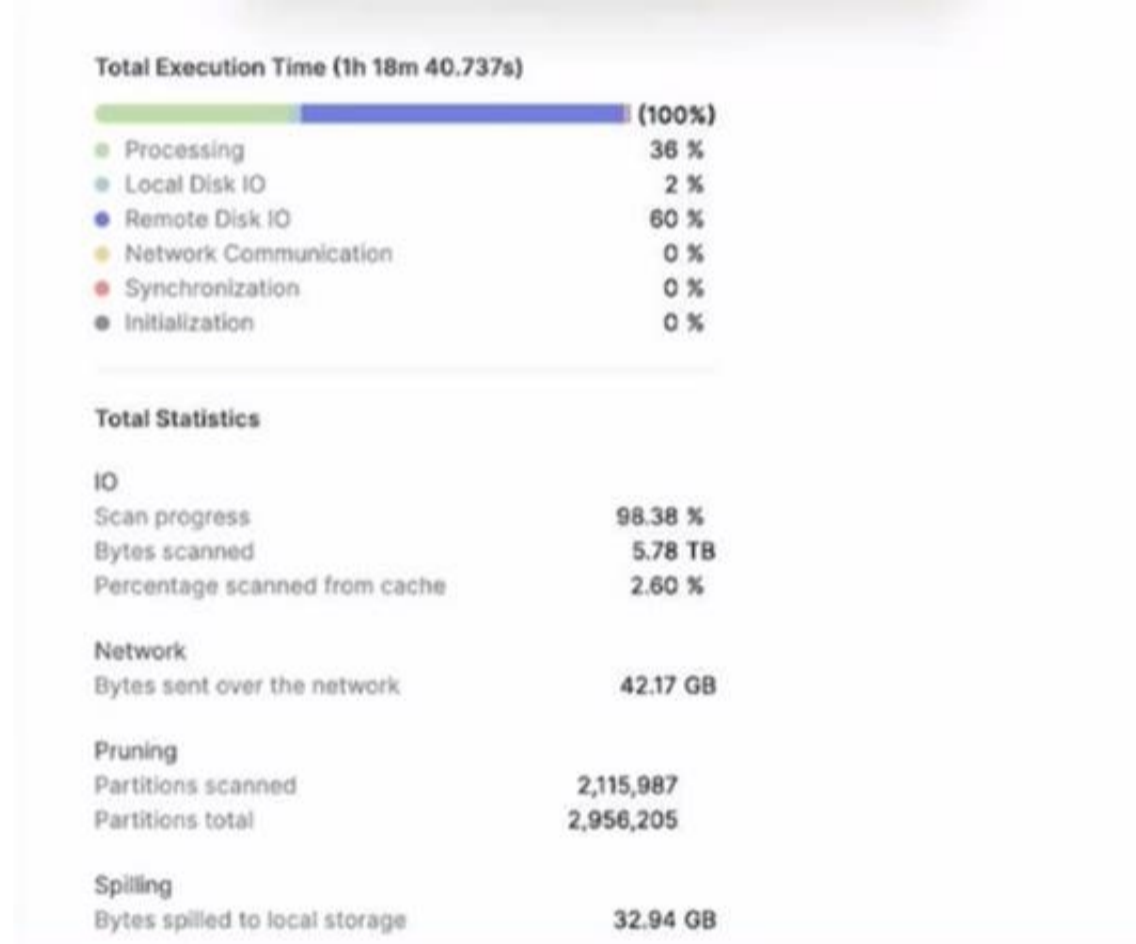
**Answer:** ABE

#### 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 6

A large table with 200 columns contains two years of historical data. When queried, the table is filtered on a single day. Below is the Query Profile:



Using a size 2XL virtual warehouse, this query took over an hour to complete. What will improve the query performance the MOST?

- A. increase the size of the virtual warehouse.
- B. Increase the number of clusters in the virtual warehouse
- C. Implement the search optimization service on the table
- D. Add a date column as a cluster key on the table

**Answer:** D

#### Explanation:

Adding a date column as a cluster key on the table will improve the query performance by reducing the number of micro-partitions that need to be scanned. Since the table is filtered on a single day, clustering by date will make the query more selective and efficient.

#### NEW QUESTION 7

A Data Engineer ran a stored procedure containing various transactions. During the execution, the session abruptly disconnected, preventing one transaction from committing or rolling back. The transaction was left in a detached state and created a lock on resources. ...must the Engineer take to immediately run a new transaction?

- A. Call the system function `SYSTEM$ABORT_TRANSACTION`.
- B. Call the system function `SYSTEM$CANCEL_TRANSACTION`.
- C. Set the `LOCK_TIMEOUT` to `FALSE` in the stored procedure.
- D. Set the transaction abort on error to true in the stored procedure.

**Answer:** A

**Explanation:**

The system function SYSTEM\$ABORT\_TRANSACTION can be used to abort a detached transaction that was left in an open state due to a session disconnect or termination. The function takes one argument: the transaction ID of the detached transaction. The function will abort the transaction and release any locks held by it. The other options are incorrect because they do not address the issue of a detached transaction. The system function SYSTEM\$CANCEL\_TRANSACTION can be used to cancel a running transaction, but not a detached one. The LOCK\_TIMEOUT parameter can be used to set a timeout period for acquiring locks on resources, but it does not affect existing locks. The TRANSACTION\_ABORT\_ON\_ERROR parameter can be used to control whether a transaction should abort or continue when an error occurs, but it does not affect detached transactions.

**NEW QUESTION 8**

A stream called TRANSACTIONS\_STM is created on top of a transactions table in a continuous pipeline running in Snowflake. After a couple of months, the TRANSACTIONS table is renamed transactiok3\_raw to comply with new naming standards. What will happen to the TRANSACTIONS\_STM object?

- A. TRANSACTIONS\_STM will keep working as expected
- B. TRANSACTIONS\_STM will be stale and will need to be re-created
- C. TRANSACTIONS\_STM will be automatically renamed TRANSACTIONS\_RAW\_STM.
- D. Reading from the traksactioks\_3T>: stream will succeed for some time after the expected STALE\_TIME.

**Answer:** B

**Explanation:**

A stream is a Snowflake object that records the history of changes made to a table. A stream is associated with a specific table at the time of creation, and it cannot be altered to point to a different table later. Therefore, if the source table is renamed, the stream will become stale and will need to be re-created with the new table name. The other options are not correct because:

? TRANSACTIONS\_STM will not keep working as expected, as it will lose track of the changes made to the renamed table.

? TRANSACTIONS\_STM will not be automatically renamed TRANSACTIONS\_RAW\_STM, as streams do not inherit the name changes of their source tables.

? Reading from the transactions\_stm stream will not succeed for some time after the expected STALE\_TIME, as streams do not have a STALE\_TIME property.

**NEW QUESTION 9**

Which stages support external tables?

- A. Internal stages only; within a single Snowflake account
- B. internal stages only from any Snowflake account in the organization
- C. External stages only from any region, and any cloud provider
- D. External stages only, only on the same region and cloud provider as the Snowflake account

**Answer:** C

**Explanation:**

External stages only from any region, and any cloud provider support external tables. External tables are virtual tables that can query data from files stored in external stages without loading them into Snowflake tables. External stages are references to locations outside of Snowflake, such as Amazon S3 buckets, Azure Blob Storage containers, or Google Cloud Storage buckets. External stages can be created from any region and any cloud provider, as long as they have a valid URL and credentials. The other options are incorrect because internal stages do not support external tables. Internal stages are locations within Snowflake that can store files for loading or unloading data. Internal stages can be user stages, table stages, or named stages.

**NEW QUESTION 10**

What is a characteristic of the use of external tokenization?

- A. Secure data sharing can be used with external tokenization
- B. External tokenization cannot be used with database replication
- C. Pre-loading of unmasked data is supported with external tokenization
- D. External tokenization allows (he preservation of analytical values after de-identification)

**Answer:** D

**Explanation:**

External tokenization is a feature in Snowflake that allows users to replace sensitive data values with tokens that are generated and managed by an external service. External tokenization allows the preservation of analytical values after de-identification, such as preserving the format, length, or range of the original values. This way, users can perform analytics on the tokenized data without compromising the security or privacy of the sensitive data.

**NEW QUESTION 10**

What is a characteristic of the operations of streams in Snowflake?

- A. Whenever a stream is queried, the offset is automatically advanced.
- B. When a stream is used to update a target table the offset is advanced to the current time.
- C. Querying a stream returns all change records and table rows from the current offset to the current time.
- D. Each committed and uncommitted transaction on the source table automatically puts a change record in the stream.

**Answer:** C

**Explanation:**

A stream is a Snowflake object that records the history of changes made to a table. A stream has an offset, which is a point in time that marks the beginning of the change records to be returned by the stream. Querying a stream returns all change records and table rows from the current offset to the current time. The offset is not automatically advanced by querying the stream, but it can be manually

advanced by using the ALTER STREAM command. When a stream is used to update a target table, the offset is advanced to the current time only if the ON UPDATE clause is specified in the stream definition. Each committed transaction on the source table automatically puts a change record in the stream, but uncommitted transactions do not.

**NEW QUESTION 12**

A Data Engineer has developed a dashboard that will issue the same SQL select clause to Snowflake every 12 hours. ---will Snowflake use the persisted query results from the result cache provided that the underlying data has not changed^

- A. 12 hours
- B. 24 hours
- C. 14 days
- D. 31 days

**Answer: C**

**Explanation:**

Snowflake uses the result cache to store the results of queries that have been executed recently. The result cache is maintained at the account level and is shared across all sessions and users. The result cache is invalidated when any changes are made to the tables or views referenced by the query. Snowflake also has a retention policy for the result cache, which determines how long the results are kept in the cache before they are purged. The default retention period for the result cache is 24 hours, but it can be changed at the account, user, or session level. However, there is a maximum retention period of 14 days for the result cache, which cannot be exceeded. Therefore, if the underlying data has not changed, Snowflake will use the persisted query results from the result cache for up to 14 days.

**NEW QUESTION 15**

The following code is executed in a Snowflake environment with the default settings:

```
create table customer;  
  
insert into customer values ('1', 'John');  
  
select $1 from customer;
```

What will be the result of the select statement?

- A. SQL compilation error object CUSTOMER' does not exist or is not authorized.
- B. John
- C. 1
- D. 1John

**Answer: C**

**NEW QUESTION 16**

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 21**

What are characteristics of Snowpark Python packages? (Select THREE).

Third-party packages can be registered as a dependency to the Snowpark session using the session, import () method.

- A. Python packages can access any external endpoints
- B. Python packages can only be loaded in a local environment
- C. Third-party supported Python packages are locked down to prevent hitting
- D. The SQL command DESCRIBE FUNCTION will list the imported Python packages of the Python User-Defined Function (UDF).

E. Querying information schema .packages will provide a list of supported Python packages and versions

**Answer:** ADE

**Explanation:**

The characteristics of Snowpark Python packages are:

- ? Third-party packages can be registered as a dependency to the Snowpark session using the session.import() method.
- ? The SQL command DESCRIBE FUNCTION will list the imported Python packages of the Python User-Defined Function (UDF).
- ? Querying information\_schema.packages will provide a list of supported Python packages and versions.

These characteristics indicate how Snowpark Python packages can be imported, inspected, and verified in Snowflake. The other options are not characteristics of Snowpark Python packages. Option B is incorrect because Python packages can be loaded in both local and remote environments using Snowpark. Option C is incorrect because third-party supported Python packages are not locked down to prevent hitting external endpoints, but rather restricted by network policies and security settings.

**NEW QUESTION 24**

A company is building a dashboard for thousands of Analysts. The dashboard presents the results of a few summary queries on tables that are regularly updated. The query conditions vary by tope according to what data each Analyst needs Responsiveness of the dashboard queries is a top priority, and the data cache should be preserved.

How should the Data Engineer configure the compute resources to support this dashboard?

- A. Assign queries to a multi-cluster virtual warehouse with economy auto-scaling Allow the system to automatically start and stop clusters according to demand.
- B. Assign all queries to a multi-cluster virtual warehouse set to maximized mode Monitor to determine the smallest suitable number of clusters.
- C. Create a virtual warehouse for every 250 Analysts Monitor to determine how many of these virtual warehouses are being utilized at capacity.
- D. Create a size XL virtual warehouse to support all the dashboard queries Monitor query runtimes to determine whether the virtual warehouse should be resized.

**Answer:** B

**Explanation:**

This option is the best way to configure the compute resources to support this dashboard. By assigning all queries to a multi-cluster virtual warehouse set to maximized mode, the Data Engineer can ensure that there is enough compute capacity to handle thousands of concurrent queries from different analysts. A multi-cluster virtual warehouse can scale up or down by adding or removing clusters based on the load. A maximized scaling policy ensures that there is always at least one cluster running and that new clusters are added as soon as possible whenneeded. By monitoring the utilization and performance of the virtual warehouse, the Data Engineer can determine the smallest suitable number of clusters that can meet the responsiveness requirement and minimize costs.

**NEW QUESTION 26**

A Data Engineer is trying to load the following rows from a CSV file into a table in Snowflake with the following structure:

```

MERID, ADDRESS, REGISTERDT
30 Ford Walk, Dante, Rhode Island, 366",2014-02-08
14 Monroe Street, Kersey, Nevada, 6384",2021-04-19
33 Gate Ave, Edgewater, New York, 1757",2020-07-03
    
```

	type
MERID	NUMBER(38,0)
SS	VARCHAR(255)
ERDT	DATE

...engineer is using the following COPY INTO statement:

```

copy into stgCustomer
from @csv_stage/address.csv.gz
file_format = (type = CSV skip_header = 1);
    
```

However, the following error is received.

```

Number of columns in file (6) does not match that of the corresponding table (3). use file format option
error_on_column_count_mismatch=false to ignore this error File 'address.csv.gz', line 3, character 1 Row 1 starts at line 2, column
"STGCUSTOMER"(6) If you would like to continue loading when an error is encountered, use other values such as 'SKIP_FILE' or
'CONTINUE' for the ON_ERROR option.
    
```

Which file format option should be used to resolve the error and successfully load all the data into the table?

- A. ESC&PE\_UNENGL09ED\_FIELD = '\\'
- B. ERROR\_ON\_COLUMN\_COUKT\_MISMATCH = FALSE
- C. FIELD\_DELIMITER = ","
- D. FIELD\_OPTIONALLY\_ENCLOSED\_BY = " "

**Answer:** D

**Explanation:**

The file format option that should be used to resolve the error and successfully load all the data into the table is FIELD\_OPTIONALLY\_ENCLOSED\_BY = "". This option specifies that fields in the file may be enclosed by double quotes, which allows for fields that contain commas or newlines within them. For example, in row 3 of the file, there is a field that contains a comma within double quotes: "Smith Jr., John". Without specifying this option, Snowflake will treat this field as two separate fields and cause an error due to column count mismatch. By specifying this option, Snowflake will treat this field as one field and load it correctly into the table.

**NEW QUESTION 30**

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 querywill provide the status of the pipe?

- A. SELECT FROM SYSTEM\$PIPE\_STATUS ('test.extract.my\_pipe');
- B. SELECT FROM SYSTEM\$PIPE\_STATUS (,test,,Extract,,my\_pipe, i l
- 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 34**

A Data Engineer needs to ingest invoice data in PDF format into Snowflake so that the data can be queried and used in a forecasting solution. .... recommended way to ingest this data?

- A. Use Snowpipe to ingest the files that land in an external stage into a Snowflake table
- B. Use a COPY INTO command to ingest the PDF files in an external stage into a Snowflake table with a VARIANT column.
- C. Create an external table on the PDF files that are stored in a stage and parse the data nto structured data
- D. Create a Java User-Defined Function (UDF) that leverages Java-based PDF parser libraries to parse PDF data into structured data

**Answer:** D

**Explanation:**

The recommended way to ingest invoice data in PDF format into Snowflake is to create a Java User-Defined Function (UDF) that leverages Java-based PDF parser libraries to parse PDF data into structured data. This option allows for more flexibility and control over how the PDF data is extracted and transformed. The other options are not suitable for ingesting PDF data into Snowflake. Option A and B are incorrect because Snowpipe and COPY INTO commands can only ingest files that are in supported file formats, such as CSV, JSON, XML, etc. PDF files are not supported by Snowflake and will cause errors or unexpected results. Option C is incorrect because external tables can only query files that are in supported file formats as well. PDF files cannot be parsed by external tables and will cause errors or unexpected results.

**NEW QUESTION 36**

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 40**

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 42**

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 45

At what isolation level are Snowflake streams?

- A. Snapshot
- B. Repeatable read
- C. Read committed
- D. Read uncommitted

**Answer:** B

#### Explanation:

The isolation level of Snowflake streams is repeatable read, which means that each transaction sees a consistent snapshot of data that does not change during its execution. Streams use time travel internally to provide this isolation level and ensure that queries on streams return consistent results regardless of concurrent transactions on their source tables.

#### NEW QUESTION 48

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. `DataFraas.writeO`
- 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 53

.....

## 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