

# DSA-C03 Practice Exam Fee - DSA-C03 Authentic Exam Questions



2025 Latest It-Tests DSA-C03 PDF Dumps and DSA-C03 Exam Engine Free Share: [https://drive.google.com/open?id=1BbSECeC2wnUtRTWOUvpQMRKZz20z6l\\_n](https://drive.google.com/open?id=1BbSECeC2wnUtRTWOUvpQMRKZz20z6l_n)

Snowflake study material is designed to enhance your personal ability and professional skills to solve the actual problem. DSA-C03 exam certification will be the most important one. There are many study material online for you to choose. While, the DSA-C03 exam dumps provided by It-Tests site will be the best valid training material for you. DSA-C03 study pdf contains the questions which are all from the original question pool, together with verified answers. Besides, the explanations are very detail and helpful after the DSA-C03 questions where is needed. You can pass your test at first try with our DSA-C03 training pdf.

Our Snowflake is suitable for computer users with a Windows operating system. Snowflake DSA-C03 practice exam support team cooperates with users to tie up any issues with the correct equipment. If DSA-C03 Certification Exam material changes, It-Tests also issues updates free of charge for three months following the purchase of our DSA-C03 exam questions.

>> DSA-C03 Practice Exam Fee <<

## DSA-C03 Authentic Exam Questions | Latest DSA-C03 Test Camp

The DSA-C03 certificate stands out among the numerous certificates because its practicability and role to improve the clients' stocks of knowledge and practical ability. Owning a test DSA-C03 certificate equals owning a weighty calling card when the clients find jobs and the proof that the clients are the competent people. Our DSA-C03 Quiz prep is the great option for the clients to prepare for the test. Our DSA-C03 study materials boost high passing rate and hit rate. Our clients praise them highly after they use them and recognize them as the key tool to pass the DSA-C03 certification.

## Snowflake SnowPro Advanced: Data Scientist Certification Exam Sample Questions (Q43-Q48):

### NEW QUESTION # 43

You've trained a machine learning model using Scikit-learn and saved it as 'model.joblib'. You need to deploy this model to Snowflake. Which sequence of commands will correctly stage the model and create a Snowflake external function to use it for

inference, assuming you already have a Snowflake stage named 'model\_stage'?

```

○ PUT file:///path/to/model.joblib @model_stage; CREATE OR REPLACE EXTERNAL FUNCTION predict(input VARCHAR) RETURNS VARCHAR LANGUAGE PYTHON
RUNTIME_VERSION = '3.8' PACKAGES = ('scikit-learn','joblib') HANDLER = 'main.predict' AS $$ import joblib import pandas as pd from
snowflake.snowpark.types import PandasDataFrame class PythonPredictor: def __init__(self): import_dir = os.path.dirname(os.path.abspath(__file__))
self._model = joblib.load(os.path.join(import_dir, 'model.joblib')) def predict(self, input): df = pd.DataFrame([input]) pred =
self._model.predict(df) return str(pred[0]) main = PythonPredictor() $$ ;

○ PUT file:///path/to/model.joblib @model_stage AUTO_COMPRESS=FALSE; CREATE OR REPLACE EXTERNAL FUNCTION predict(input VARCHAR) RETURNS VARCHAR LANGUAGE PYTHON
RUNTIME_VERSION = '3.8' PACKAGES = ('scikit-learn','joblib','pandas','snowflake-snowpark-python') HANDLER = 'main.predict' TARGET_PATH = '@model_stage/model.joblib' AS $$
import joblib import pandas as pd import os class PythonPredictor: def __init__(self): import_dir = os.path.dirname(os.path.abspath(__file__)) self._model =
joblib.load(os.path.join(import_dir, 'model.joblib')) def predict(self, input): df = pd.DataFrame([input]) pred = self._model.predict(df) return str(pred[0]) main =
PythonPredictor() $$ ;

○ PUT file:///path/to/model.joblib @model_stage AUTO_COMPRESS=FALSE; CREATE OR REPLACE EXTERNAL FUNCTION predict(input VARCHAR) RETURNS VARCHAR LANGUAGE PYTHON RUNTIME_VERSION = '3.8' PACKAGES
= ('scikit-learn','joblib','pandas','snowflake-snowpark-python') HANDLER = 'main.predict' IMPORTS = ('@model_stage/model.joblib') AS $$ import joblib import pandas as pd import os class
PythonPredictor: def __init__(self): import_dir = os.path.dirname(os.path.abspath(__file__)) self._model = joblib.load(os.path.join(import_dir, 'model.joblib')) def predict(self, input): df =
pd.DataFrame([input]) pred = self._model.predict(df) return str(pred[0]) main = PythonPredictor() $$ ;

○ PUT file:///path/to/model.joblib @model_stage AUTO_COMPRESS=FALSE; CREATE OR REPLACE EXTERNAL FUNCTION predict(input VARCHAR) RETURNS VARCHAR LANGUAGE PYTHON RUNTIME_VERSION = '3.8' PACKAGES = ('scikit-
learn','joblib','pandas','snowflake-snowpark-python') HANDLER = 'main.predict' IMPORTS = ('@model_stage/model.joblib') AS $$ import joblib import pandas as pd import os import snowflake.snowpark class PythonPredictor: def
__init__(self): import_dir = os.path.dirname(os.path.abspath(__file__)) self._model = joblib.load(os.path.join(import_dir, 'model.joblib')) def predict(self, input): df = pd.DataFrame([input]) pred =
self._model.predict(df) return str(pred[0]) main = PythonPredictor() $$ ;

○ PUT file:///path/to/model.joblib @model_stage OVERWRITE=TRUE; CREATE OR REPLACE EXTERNAL FUNCTION predict(input VARCHAR) RETURNS VARCHAR LANGUAGE PYTHON RUNTIME_VERSION = '3.8' PACKAGES = ('scikit-learn','joblib','pandas','snowflake-snowpark-
python') HANDLER = 'main.predict' IMPORTS = ('@model_stage/model.joblib') AS $$ import joblib import pandas as pd import os import snowflake.snowpark class PythonPredictor: def
__init__(self): import_dir = os.path.dirname(os.path.abspath(__file__)) self._model = joblib.load(os.path.join(import_dir, 'model.joblib')) def predict(self, input): df = pd.DataFrame([input]) pred = self._model.predict(df) return str(pred[0]) main = PythonPredictor() $$ ;

```

- A. Option B
- **B. Option E**
- C. Option A
- D. Option D
- E. Option C

**Answer: B**

**Explanation:**

Option E is correct. `AUTO\_COMPRESS=FALSE` isn't strictly needed for joblib files, and adding `OVERWRITE=TRUE` ensures the model is updated in case of re-deployment. Importantly, the IMPORTS parameter correctly specifies the location of the model within the stage, and includes necessary packages (pandas and snowflake-snowpark-python) and the code correctly loads the model from the stage within the Python handler. Also snowflake.snowpark need to be imported to avoid errors with snowpark library. Option A lacks pandas, snowflake-snowpark-python, imports section. Option B has wrong syntax TARGET\_PATH which is incorrect syntax. Option C doesn't import snowpark, so might give error. Option D doesn't OVERWRITE flag, and can give error if trying to overwrite files.

**NEW QUESTION # 44**

You are tasked with training a complex machine learning model using scikit-learn and need to leverage Snowflake's data for training outside of Snowflake using an external function. The training data resides in a Snowflake table named 'CUSTOMER DATA'. Due to data governance policies, you must ensure minimal data movement and secure communication. You choose to implement the external function using AWS Lambda'. Which of the following steps are crucial to achieve secure and efficient model training outside of Snowflake?

- A. Utilize Snowflake's data masking policies on the table to anonymize sensitive information before sending it to the external function for training. This ensures data privacy and compliance with regulations.
- B. In the Lambda function, establish a direct connection to the Snowflake database using the Snowflake JDBC driver and Snowflake user credentials stored in the Lambda environment variables. This allows the Lambda function to directly query the 'CUSTOMER DATA' table.
- **C. Create an API integration object in Snowflake that points to your AWS API Gateway endpoint, configured to invoke the Lambda function. This API integration must use a service principal and access roles for secure authentication.**
- **D. Create an external function in Snowflake that accepts a JSON payload containing the necessary parameters for model training, such as features to use and model hyperparameters. This function will call the API integration to invoke the Lambda function.**
- E. Grant usage privilege on the API integration object to the role that will be calling the external function, ensuring only authorized users can trigger the model training.

**Answer: C,D,E**

**Explanation:**

Options A, B, and D are correct. An API integration is essential for securely communicating with the external function via API Gateway, offering authentication and authorization. Granting usage privileges ensures only authorized roles can execute the external function. The external function serves as the bridge between Snowflake and the external service. Option C is incorrect because storing Snowflake credentials directly in Lambda environment variables is a security risk. Option E, while good practice, is not strictly related to the external function configuration for training model.

### NEW QUESTION # 45

You are building a machine learning model using Snowpark Python to predict house prices. The dataset contains a feature column named 'location' which contains free-form text descriptions of house locations. You want to leverage a pre-trained Large Language Model (LLM) hosted externally to extract structured location features like city, state, and zip code from the free-form text within Snowpark. You want to minimize the data transferred out of Snowflake. Which approach is most efficient and secure?

- A. Use to load the 'location' column data into a Pandas DataFrame, call the external LLM API in your Python script to enrich the location data and then use to store the enriched data back into a Snowflake table.
- B. Use Snowpark's 'createOrReplaceStage' to create an external stage pointing to the LLM API endpoint. Load the 'location' data into this stage and call the LLM API directly from the Snowflake stage using SQL.
- C. Create a Snowpark User-Defined Function (UDF) that calls the external LLM API. Pass the 'location' column data to the UDF and retrieve the structured location features. Then apply the UDF directly on the Snowpark DataFrame.
- D. Use the Snowflake Connector for Python to directly query the 'location' column and call the external LLM API from the connector. Then write the updated data into a new table.
- E. Create a Snowflake External Function that calls the external LLM API. Pass the 'location' column data to the External Function and retrieve the structured location features. Then apply the External Function directly on the Snowpark DataFrame.

**Answer: E**

Explanation:

Using a Snowflake External Function is the most efficient and secure way to interact with an external LLM API for this task. Here's why: Efficiency: External Functions allow Snowflake to directly call the external service in parallel, leveraging Snowflake's compute resources. This minimizes data transfer between Snowflake and the client environment. Security: External Functions support secure communication with external services using API integration objects, which handle authentication and authorization. Data Governance: Keeps all processing within Snowflake's secure environment, reducing the risk of data leakage. Options A, C, and E involve transferring the data outside of Snowflake, which is less secure and less performant. Option D is not a valid approach for integrating with an external LLM API.

### NEW QUESTION # 46

You are designing a feature engineering pipeline using Snowpark Feature Store for a fraud detection model. You have a transaction table in Snowflake. One crucial feature is the 'average\_transaction\_amount\_last\_7\_days' for each customer. You want to implement this feature using Snowpark Python and materialize it in the Feature Store. You have the following Snowpark DataFrame 'transactions\_df' containing 'customer\_id' and 'transaction\_amount'. Which of the following code snippets correctly defines and registers this feature in the Snowpark Feature Store, ensuring efficient computation and storage?

- A.

```
from snowflake.snowpark import functions as F
from snowflake.ml.feature_store.feature_group import FeatureGroup

def define_avg_transaction(transactions_df):
    avg_transactions = transactions_df.groupBy('customer_id').agg(F.avg('transaction_amount')).alias('average_transaction_amount_last_7_days').select('customer_id', 'average_transaction_amount_last_7_days')
    return avg_transactions

fg = FeatureGroup(
    name='customer_transaction_features',
    version=1,
    features=['average_transaction_amount_last_7_days'],
    source_data=define_avg_transaction(transactions_df),
    primary_keys=['customer_id']
)

fg.create_or_update(blocking=True)
```

- B.

```
from snowflake.snowpark import functions as F
from snowflake.ml.feature_store.feature_group import FeatureGroup

def define_avg_transaction(transactions_df):
    avg_transactions = transactions_df.groupBy('customer_id').agg(F.mean('transaction_amount')).alias('average_transaction_amount_last_7_days')
    return avg_transactions

fg = FeatureGroup(
    name='customer_transaction_features',
    version=1,
    features=['average_transaction_amount_last_7_days'],
    source_data=define_avg_transaction,
    primary_keys=['customer_id']
)

fg.create_or_update(blocking=True)
```

- C.

```

from snowflake.snowpark import functions as F
from snowflake.ml.feature_store.feature_group import FeatureGroup

def define_avg_transaction(transactions_df):
    avg_transactions = transactions_df.groupBy('customer_id').agg(F.mean('transaction_amount').alias('average_transaction_amount_last_7_days'))
    return avg_transactions

fg = FeatureGroup(
    name='customer_transaction_features',
    version=1,
    features=['average_transaction_amount_last_7_days'],
    source_data=define_avg_transaction(transactions_df),
    primary_keys=['customer_id']
)

fg.create_or_update()

```

• D.

```

from snowflake.snowpark import functions as F
from snowflake.ml.feature_store.feature_group import FeatureGroup

def define_avg_transaction(transactions_df):
    avg_transactions = transactions_df.groupBy('customer_id').agg(F.avg('transaction_amount').alias('average_transaction_amount_last_7_days'))
    return avg_transactions.select('customer_id', 'average_transaction_amount_last_7_days')

fg = FeatureGroup(
    name='customer_transaction_features',
    version=1,
    features=['average_transaction_amount_last_7_days'],
    source_data=define_avg_transaction(transactions_df),
    primary_keys=['customer_id']
)

fg.create_or_update()

```

• E.

```

from snowflake.snowpark import functions as F
from snowflake.ml.feature_store.feature_group import FeatureGroup

def define_avg_transaction(transactions_df):
    avg_transactions = transactions_df.groupBy('customer_id').agg(F.avg('transaction_amount').alias('average_transaction_amount_last_7_days'))
    return avg_transactions

g = FeatureGroup(
    name='customer_transaction_features',
    version=1,
    features=['average_transaction_amount_last_7_days'],
    source_data=define_avg_transaction(transactions_df),
    primary_keys=['customer_id']
)

g.create_or_update(blocking=True)

```

**Answer: A**

**Explanation:**

Option E is correct. It uses 'F.avg' for calculating the average, selects only the required columns ('customer\_id', 'average\_transaction\_amount\_last\_7\_days') in and then sets on to ensure the feature group is fully materialized before proceeding. 'blocking=True' is important for production pipelines to avoid race conditions.

### NEW QUESTION # 47

You have successfully trained a binary classification model using Snowpark ML and deployed it as a UDF in Snowflake. The UDF takes several input features and returns the predicted probability of the positive class. You need to continuously monitor the model's performance in production to detect potential data drift or concept drift. Which of the following methods and metrics, when used together, would provide the MOST comprehensive and reliable assessment of model performance and drift in a production environment? (Select TWO)

- A. Monitor the average predicted probability score over time. A significant shift in the average score indicates data drift.
- B. Continuously calculate and track performance metrics like AUC, precision, recall, and F1-score on a representative sample of labeled production data over regular intervals. Compare these metrics to the model's performance on the holdout set during training.
- C. Calculate the Kolmogorov-Smirnov (KS) statistic between the distribution of predicted probabilities in the training data and the production data over regular intervals. Track any substantial changes in the KS statistic.
- D. Check for null values in the input features passed to the UDF. A sudden increase in null values indicates a problem with

data quality.

- E. Monitor the volume of data processed by the UDF per day. A sudden drop in volume indicates a problem with the data pipeline.

**Answer: B,C**

Explanation:

Options B and D provide the most comprehensive assessment of model performance and drift. Option D, by continuously calculating key performance metrics (AUC, precision, recall, F1 -score) on labeled production data, directly assesses how well the model is performing on real- world data. Comparing these metrics to the holdout set provides insights into potential overfitting or degradation over time (concept drift). Option B, calculating the KS statistic between the predicted probability distributions of training and production data, helps to identify data drift, indicating that the input data distribution has changed. Option A can be an indicator but is less reliable than the KS statistic. Option C monitors data pipeline health, not model performance. Option E focuses on data quality, which is important but doesn't directly assess model performance drift.

## NEW QUESTION # 48

.....

Our company is considerably cautious in the selection of talent and always hires employees with store of specialized knowledge and skills on our DSA-C03 exam questions. All the members of our experts and working staff maintain a high sense of responsibility, which is why there are so many people choose our DSA-C03 Exam Materials and to be our long-term partner. Believe in our DSA-C03 study guide, and you will have a brighter future!

**DSA-C03 Authentic Exam Questions:** <https://www.it-tests.com/DSA-C03.html>

As a candidate for the DSA-C03 Authentic Exam Questions - SnowPro Advanced: Data Scientist Certification Exam certification, you should insist on and never give up for a higher pursue no matter how difficult it is to conquer, Snowflake DSA-C03 Practice Exam Fee The software boosts varied self-learning and self-assessment functions to check the results of the learning. If you are worried about how to prepare for the DSA-C03 certification exam, just download It-Tests real DSA-C03 Dumps PDF and study well to crack it.

And keep attempting the same question several time so you can overcome DSA-C03 Authentic Exam Questions the exam anxiety, The first clues you have to the success of your backup are the messages displayed at the end of the backup.

## 2026 DSA-C03 – 100% Free Practice Exam Fee | Reliable SnowPro Advanced: Data Scientist Certification Exam Authentic Exam Questions

As a candidate for the SnowPro Advanced: Data Scientist Certification Exam certification, DSA-C03 Testing Center you should insist on and never give up for a higher pursue no matter how difficult it is to conquer, The software boosts varied DSA-C03 self-learning and self-assessment functions to check the results of the learning.

If you are worried about how to prepare for the DSA-C03 certification exam, just download It-Tests real DSA-C03 Dumps PDF and study well to crack it, We offer money back guarantee for our customers.

We want to make you fully prepared and confident for the SnowPro Advanced: Data Scientist Certification Exam DSA-C03 Practice Exam Fee Exam, and that is why we created this preparation material that fulfills all the requirements of your preparation.

- DSA-C03 Practice Exam Fee and Snowflake DSA-C03 Authentic Exam Questions: SnowPro Advanced: Data Scientist Certification Exam Pass for Sure  Enter [www.testkingpass.com](http://www.testkingpass.com)  and search for { DSA-C03 } to download for free  DSA-C03 Reliable Mock Test
- Snowflake DSA-C03 Practice Exam Fee: SnowPro Advanced: Data Scientist Certification Exam - Pdfvce Reliable Platform  Simply search for “ DSA-C03 ” for free download on “ [www.pdfvce.com](http://www.pdfvce.com) ”  DSA-C03 Reliable Test Camp
- 2026 The Best 100% Free DSA-C03 – 100% Free Practice Exam Fee | DSA-C03 Authentic Exam Questions  Search for [ DSA-C03 ] and download it for free on [www.testkingpass.com](http://www.testkingpass.com)  website  DSA-C03 Valuable Feedback
- Snowflake DSA-C03 Practice Exam Fee: SnowPro Advanced: Data Scientist Certification Exam - Pdfvce Reliable Platform  The page for free download of  DSA-C03  on [www.pdfvce.com](http://www.pdfvce.com)  will open immediately  Test DSA-C03 Price
- DSA-C03 Latest Test Bootcamp  Valid Test DSA-C03 Test  Reliable DSA-C03 Exam Guide  Search for 《 DSA-C03 》 and download it for free on 《 [www.vce4dumps.com](http://www.vce4dumps.com) 》 website  Valid DSA-C03 Cram Materials
- Reliable DSA-C03 Exam Voucher  Latest DSA-C03 Exam Papers  Exam DSA-C03 Learning  Enter

