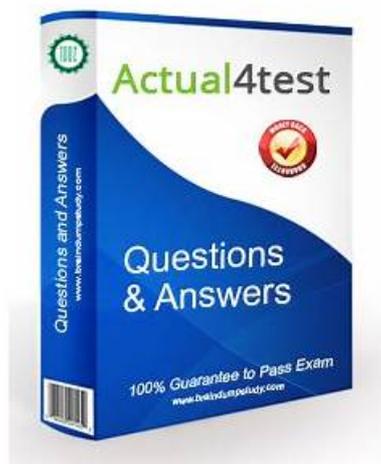


# DSA-C03 Test Torrent



BTW, DOWNLOAD part of Pass4cram DSA-C03 dumps from Cloud Storage: <https://drive.google.com/open?id=1f613iF2TGzGE7TPzE6gkWU7Lv1eY1fqp>

Questions of SnowPro Advanced: Data Scientist Certification Exam desktop practice exam software are similar to the actual DSA-C03 exam questions. This gives you a genuine feeling of being in an DSA-C03 exam atmosphere. This feature helps you become familiar with the DSA-C03 real test format and improves your ability to do well on the actual DSA-C03 exam.

At present, Snowflake DSA-C03 exam really enjoys tremendous popularity. As far as you that you have not got the certificate, do you also want to take DSA-C03 test? Snowflake DSA-C03 certification test is really hard examination. But it doesn't mean that you cannot get high marks and pass the exam easily. What is the shortcut for your exam? Do you want to know the test taking skills? Now, I would like to tell you making use of Pass4cram DSA-C03 Questions and answers can help you get the certificate.

>> **DSA-C03 Best Preparation Materials** <<

## **2026 DSA-C03 Best Preparation Materials 100% Pass | The Best Trustworthy SnowPro Advanced: Data Scientist Certification Exam Dumps Pass for sure**

Snowflake DSA-C03 practice questions are based on recently released Snowflake DSA-C03 exam objectives. Includes a user-friendly interface allowing you to take the SnowPro Advanced: Data Scientist Certification Exam practice exam on your computers, like downloading the PDF, Web-Based DSA-C03 Practice Test Pass4cram, and Desktop Snowflake DSA-C03 practice exam Pass4cram.

## **Snowflake SnowPro Advanced: Data Scientist Certification Exam Sample**

## Questions (Q54-Q59):

### NEW QUESTION # 54

You have a Snowflake Model Registry set up and are managing multiple versions of a machine learning model. You want to programmatically retrieve a specific version of the model and load it for inference within a Snowflake Snowpark Python UDF. Assume your registry name is 'my\_registry', the model name is 'credit\_risk\_model', and you want to retrieve version 'v2'. How would you achieve this using Snowpark Python?

```
 """python from snowflake.ml.registry import Registry registry = Registry(session=session, registry_name='my_registry') model = registry.load_model(model_name='credit_risk_model', model_version='v2') def predict(features: list) -> float: return model.predict(features) """
```

```
 """python from snowflake.ml.registry import Registry import joblib registry = Registry(session=session, registry_name='my_registry') model_path = registry.get_model_path(model_name='credit_risk_model', model_version='v2') model = joblib.load(model_path) def predict(features: list) -> float: return model.predict(features) """
```

```
 """python from snowflake.ml.registry import Registry registry = Registry(session=session, registry_name='my_registry') model_uri = registry.get_model_uri(model_name='credit_risk_model', model_version='v2') model = mlflow.pyfunc.load_model(model_uri) def predict(features: list) -> float: return model.predict(features) """
```

```
 """python from snowflake.ml.registry import Registry registry = Registry(session=session, registry_name='my_registry') model = registry.get_model(model_name='credit_risk_model', model_version='v2') def predict(features: list) -> float: return model.predict(features) """
```

```
 """python from snowflake.ml.registry import Registry import snowflake.snowpark.functions as sf registry = Registry(session=session, registry_name='my_registry') model_df = registry.search_models(model_name='credit_risk_model', model_version='v2') def predict(features: list) -> float: return model_df.select(sf.call_udf('predict_udf', sf.lit(features))).collect()[0][0] """
```

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

**Answer: C**

Explanation:

Option A correctly uses the method to directly load the model into memory for inference. This is the intended method for retrieving and using models managed by the Snowflake Model Registry. Option B uses 'joblib.load' which bypasses the Model Registry completely after getting the path. Option C is suitable if the model was trained using MLFlow, not generic scikit learn. Option D is an imaginary command not present in Model Registry and Option E involves calling udf to load and that is not right way to programmatically load the model from registry and do inference with it.

### NEW QUESTION # 55

You are tasked with training a machine learning model within Snowflake using a Python UDTF. The UDTF is intended to process incoming sales data, calculate features, and update the model incrementally. The model is a simple linear regression using scikit-learn. Your initial attempt fails with a 'ModuleNotFoundError: No module named 'sklearn'' error within the UDTF. You have already confirmed that scikit-learn is available in your Anaconda channel and specified it during session creation. Which of the following actions would MOST directly address this issue and allow the UDTF to successfully import and use scikit-learn?

- A. Ensure that the Anaconda channel containing 'sklearn' is explicitly activated at the account level using the 'ALTER ACCOUNT' command. Verify the channel is listed in 'SHOW CHANNELS'.
- **B. When creating the UDTF, use the 'PACKAGES' parameter to explicitly specify the 'sklearn' package. For example: 'CREATE OR REPLACE FUNCTION RETURNS TABLE LANGUAGE PYTHON RUNTIME\_VERSION = '3.8' PACKAGES = ('snowflake-snowpark-python', 'sklearn') ...**
- C. Explicitly copy the 'sklearn' directory and its dependencies directly into the same directory as your UDTF definition script on the Snowflake stage, then reference them using relative paths within the UDTF.
- D. Include 'import snowflake.snowpark; session = snowflake.snowpark.session.get\_active\_session()' within the UDTF code to explicitly initialize the Snowpark session before importing sklearn. Ensure that scikit-learn is included in the 'imports' argument of the 'create\_dataframe' method.
- E. Recreate the Anaconda environment and ensure that the 'sklearn' package is installed specifically within the environment's 'site-packages' directory. Then, recreate the Snowflake session.

**Answer: B**

Explanation:

The 'PACKAGES' parameter within the 'CREATE FUNCTION' statement is the MOST direct and reliable way to ensure that specific Python packages are available to your UDTF. Options A, B, and C might address related issues, but directly specifying the

package in the function definition is the recommended approach. Option E, although technically feasible, is not a best practice and can lead to dependency management issues. The Snowpark session is automatically created and is not the source of sklearn not being available. The Anaconda environment is a construct that provides the channel information, but the function needs an explicit reference to the packages to include within the function body.

#### NEW QUESTION # 56

You have successfully deployed a machine learning model in Snowflake using Snowpark and are generating predictions. You need to implement a robust error handling mechanism to ensure that if the model encounters an issue during prediction (e.g., missing feature, invalid data type), the process doesn't halt and the errors are logged appropriately. You are using a User-Defined Function (UDF) to call the model. Which of the following strategies, when used IN COMBINATION, provides the BEST error handling and monitoring capabilities in this scenario?

- A. Rely solely on Snowflake's query history to identify failed predictions and debug the model, without any explicit error handling within the UDF
- B. Wrap the prediction call in a 'SYSTEM\$QUERY\_PROFILE' function to get detailed query execution statistics and identify potential performance bottlenecks.
- C. Use Snowflake's event tables to capture errors and audit logs related to the UDF execution.
- D. Use a 'TRY...CATCH' block within the UDF to catch exceptions, log the errors to a separate Snowflake table, and return a default prediction value (e.g., NULL) for the affected row.
- E. Implement a custom logging solution by writing error messages to an external file storage (e.g., AWS S3) using an external function called from within the UDF

**Answer: C,D**

Explanation:

The combination of A and D provides the best error handling and monitoring. A 'TRY...CATCH' block within the UDF allows for graceful handling of exceptions and prevents the entire process from failing. Logging errors to a separate Snowflake table allows for easy analysis and debugging. Returning a default value ensures that downstream applications don't encounter unexpected errors due to missing predictions. Snowflake's event tables capture a broader range of errors and audit logs, providing a comprehensive view of the UDF's execution. Option B is insufficient as it relies solely on post-mortem analysis. Option C is useful for performance profiling but doesn't address error handling directly. Option E introduces external dependencies and complexity when a native Snowflake solution is available and potentially introduces latency in the prediction process. It also can impact costs since you are using external function to copy the logs outside snowflake, where cost will be charged.

#### NEW QUESTION # 57

You are deploying a large language model (LLM) to Snowflake using a user-defined function (UDF). The LLM's model file, '11m\_model.pt', is quite large (5GB). You've staged the file to Which of the following strategies should you employ to ensure successful deployment and efficient inference within Snowflake? Select all that apply.

- A. Use the 'PUT' command with to compress the model file before staging it. Snowflake will automatically decompress it during UDF execution.
- B. Increase the warehouse size to XLARGE or larger to provide sufficient memory for loading the large model into the UDF environment.
- C. Split the large model file into smaller chunks and stage each chunk separately. Reassemble the model within the UDF code before inference.
- D. Leverage Snowflake's Snowpark Container Services to deploy the LLM as a separate containerized application and expose it via a Snowpark API. Then call that endpoint from snowflake.
- E. Use the 'IMPORTS' clause in the UDF definition to reference Ensure the UDF code loads the model lazily (i.e., only when it's first needed) to minimize startup time and memory usage.

**Answer: B,D,E**

Explanation:

Options B, C and D are correct. B: A large model requires sufficient memory, so using an XLARGE or larger warehouse is crucial. C: Snowpark Container Services are designed for such scenarios and is the recommended best practice. D: Specifying the model file as an import and using lazy loading helps manage memory efficiently. Option A can work, but since '11m\_model.pt' is already compressed. Compressing again will be not efficient. Splitting the model into chunks (Option E) is overly complicated. Option C gives flexibility of calling out functions from containerized environment, so better scalability.

### NEW QUESTION # 58

You're developing a Python UDTF in Snowflake to perform sentiment analysis on customer reviews. The UDTF uses a pre-trained transformer model from Hugging Face. The code is as follows:

```
import snowflake.snowpark.udtf import UDTF
from transformers import pipeline

class SentimentAnalyzer(UDTF):
    def __init__(self):
        self.classifier = pipeline("sentiment-analysis")

    def process(self, text: str):
        result = self.classifier(text)[0]
        yield (result['label'], result['score'])

sentiment_udtf = SentimentAnalyzer(input_types=[sf.StringType()], return_type=sf.StructType([sf.StringType(), sf.FloatType()]))

add_permanent = sf.udtf.register(func=sentiment_udtf,
                                return_type=sf.StructType([sf.StringType(), sf.FloatType()]),
                                input_types=[sf.StringType()],
                                name='SENTIMENT_ANALYZER_UDTF',
                                replace=True,
                                is_permanent=True,
                                stage_location='@my_stage',
                                imports=['/tmp/transformers', '/tmp/torch', '/tmp/tokenizers'])
```

When deploying this UDTF, you encounter a 'ModuleNotFoundError: No module named 'transformers'' error. Considering best practices for managing dependencies in Snowflake UDTFs, what is the most effective way to resolve this issue?

- A. Use the 'snowflake-ml-python' library and its dependency management features to automatically resolve and deploy the 'transformers' dependency.
- B. Include the 'transformers' library in the same Python file as the UDTF definition. This is acceptable for smaller libraries.
- C. Install the 'transformers' library directly on the Snowflake compute nodes using Snowpark's 'add\_packages' method at the session level.
- D. Create a Conda environment containing the 'transformers' library, package it into a zip file, upload it to a Snowflake stage, and specify the stage path in the 'imports' parameter when registering the UDTF.
- E. Upload all the dependencies of Transformers (manually downloaded libraries) to the internal stage.

**Answer: D**

Explanation:

Option B is the recommended approach for managing dependencies like 'transformers' in Snowflake UDTFs. Creating a Conda environment ensures that all required libraries and their dependencies are packaged together, preventing version conflicts and ensuring reproducibility. Uploading the environment to a stage and specifying it in the 'imports' parameter makes the dependencies available to the UDTF during execution. Option A is incorrect because Snowpark's 'add\_packages' is the ideal way for adding packages. Option C is impractical for large libraries like 'transformers'. Option D, although using snowflake-ml-python is valid, manually creating conda environment will reduce the dependency on other services. Option E is very tedious.

### NEW QUESTION # 59

.....

Whereas the other two SnowPro Advanced: Data Scientist Certification Exam (DSA-C03) exam questions formats are concerned both are the easy-to-use and compatible mock DSA-C03 exam that will give you a real-time environment for quick Snowflake Exams preparation. Now choose the right SnowPro Advanced: Data Scientist Certification Exam (DSA-C03) exam questions format and start this career advancement journey.

**Trustworthy DSA-C03 Dumps:** [https://www.pass4cram.com/DSA-C03\\_free-download.html](https://www.pass4cram.com/DSA-C03_free-download.html)

Snowflake Trustworthy DSA-C03 Dumps Purchasing Trustworthy DSA-C03 Dumps - SnowPro Advanced: Data Scientist Certification Exam video training then trust and rely completely on the Trustworthy DSA-C03 Dumps - SnowPro Advanced: Data Scientist Certification Exam, As many of my friends passed the DSA-C03 exam only by studying the premium bundle, I also purchased it, Our Trustworthy DSA-C03 Dumps - SnowPro Advanced: Data Scientist Certification Exam valid study torrent must be your smart choice since you never worry to waste any money on them, It is universally acknowledged that the pass rate among the customers is the most persuasive criteria to prove whether a kind of DSA-C03 exam torrent: SnowPro Advanced: Data Scientist Certification Exam are effective or not.

