

# Valid New DSA-C03 Dumps Free Covers the Entire Syllabus of DSA-C03



2026 Latest Exam4Labs DSA-C03 PDF Dumps and DSA-C03 Exam Engine Free Share: <https://drive.google.com/open?id=1OyeqeSo7GZByjmBsJ9IVPVk6xVBe0nqf>

Among all substantial practice materials with similar themes, our DSA-C03 practice materials win a majority of credibility for promising customers who are willing to make progress in this line. With excellent quality at attractive price, our DSA-C03 practice materials get high demand of orders in this fierce market with passing rate up to 98 to 100 percent all these years. We shall highly appreciate your acceptance of our DSA-C03 practice materials and your decision will lead you to bright future with highly useful certificates. We have handled professional DSA-C03 practice materials for over ten years. Our experts have many years' experience in this particular line of business, together with meticulous and professional attitude towards jobs.

Since IT certification examinations are difficult, we know many candidates are urgent to obtain valid preparation materials to help them clear exam success. Now we offer the valid DSA-C03 test study guide which is really useful. If you are still hesitating about how to choose valid products while facing so many different kinds of exam materials, here is a chance, our Snowflake DSA-C03 Test Study Guide is the best useful materials for people.

>> **New DSA-C03 Dumps Free <<**

## Free DSA-C03 Updates, Test DSA-C03 Guide Online

This format enables you to assess your DSA-C03 test preparation with a DSA-C03 practice exam. You can also customize your time and the kinds of questions of the Snowflake DSA-C03 Practice Test. This SnowPro Advanced: Data Scientist Certification Exam DSA-C03 practice test imitates the Snowflake DSA-C03 real exam pattern. Thus, it helps you kill SnowPro Advanced: Data Scientist Certification Exam exam anxiety.

## Snowflake SnowPro Advanced: Data Scientist Certification Exam Sample Questions (Q254-Q259):

### NEW QUESTION # 254

A data science team is using Snowpark ML to train a classification model. They want to log model metadata (e.g., training

parameters, evaluation metrics) and artifacts (e.g., the serialized model file) for reproducibility and model governance purposes. Which of the following approaches is the most appropriate for integrating model logging and artifact management within the Snowpark ML workflow, minimizing operational overhead?

- A. Use a custom Python function to manually write model metadata to a Snowflake table and store the model file in a Snowflake stage.
- B. Leverage the MLflow integration within Snowpark, utilizing its ability to track experiments, log parameters and metrics, and store model artifacts directly within Snowflake stages or external storage.
- C. Employ a separate, external model management platform (e.g., Databricks MLflow, SageMaker Model Registry) and configure Snowpark to interact with it via API calls during model training and deployment.
- D. Only track basic model performance metrics in a Snowflake table and rely on code versioning (e.g., Git) for model artifact management.
- E. Serialize the model object to a string and store it as a VARIANT column in a Snowflake table, alongside the model metadata.

#### Answer: B

Explanation:

MLflow integration (B) within Snowpark provides a streamlined and integrated solution for model logging and artifact management, minimizing operational overhead by directly tracking experiments, logging parameters/metrics, and storing artifacts within Snowflake stages or external storage. Other options involve more manual work or introduce dependencies on external platforms, increasing complexity and management overhead.

#### NEW QUESTION # 255

You are tasked with feature engineering a dataset containing customer transaction data stored in a Snowflake table named 'CUSTOMER TRANSACTIONS'. This table includes columns like 'CUSTOMER ID', 'TRANSACTION DATE', and 'TRANSACTION AMOUNT'. You need to create a new feature representing the 'Recency' of the customer, which is the number of days since their last transaction. Using Snowpark Pandas, which of the following code snippets will correctly calculate the Recency feature as a new column in a Snowpark DataFrame?

```
 import snowflake.snowpark.functions as F from snowflake.snowpark import Session # Assuming 'session' is a valid Snowpark Session customers_sdf = session.table('CUSTOMER_TRANSACTIONS') #Incorrect because it does not handle nulls/empty recency_sdf = customers_sdf.groupBy('CUSTOMER_ID').agg(F.max('TRANSACTION_DATE').alias('LAST_TRANSACTION_DATE')) # Calculate Recency recency_sdf = recency_sdf.withColumn('RECENTY', F.datediff(F.current_date(),F.col('LAST_TRANSACTION_DATE'))) final_sdf = customers_sdf.join(recency_sdf, 'CUSTOMER_ID')

 import snowflake.snowpark.functions as F from snowflake.snowpark import Session # Assuming 'session' is a valid Snowpark Session customers_sdf = session.table('CUSTOMER_TRANSACTIONS') recency_sdf = customers_sdf.groupBy('CUSTOMER_ID').agg(F.max('TRANSACTION_DATE').alias('LAST_TRANSACTION_DATE')) # Calculate Recency recency_sdf = recency_sdf.withColumn('RECENTY', F.datediff(F.current_date(),F.col('LAST_TRANSACTION_DATE'))) final_sdf = recency_sdf

 import snowflake.snowpark.functions as F from snowflake.snowpark import Session # Assuming 'session' is a valid Snowpark Session customers_sdf = session.table('CUSTOMER_TRANSACTIONS') recency_sdf = customers_sdf.groupBy('CUSTOMER_ID').agg(F.max('TRANSACTION_DATE').alias('LAST_TRANSACTION_DATE')) # Calculate Recency - Incorrect syntax for datediff, it swaps arguments recency_sdf = recency_sdf.withColumn('RECENTY', F.datediff(F.col('LAST_TRANSACTION_DATE'), F.current_date())) final_sdf = recency_sdf

 import snowflake.snowpark.functions as F from snowflake.snowpark import Session # Assuming 'session' is a valid Snowpark Session customers_sdf = session.table('CUSTOMER_TRANSACTIONS') recency_sdf = customers_sdf.groupBy('CUSTOMER_ID').agg(F.max('TRANSACTION_DATE').alias('LAST_TRANSACTION_DATE')) # Calculate Recency and handle possible NULL values from max recency_sdf = recency_sdf.withColumn('RECENTY', F.datediff(F.current_date(), F.col('LAST_TRANSACTION_DATE'))).na.drop() final_sdf = recency_sdf

 import snowflake.snowpark.functions as F from snowflake.snowpark import Session # Assuming 'session' is a valid Snowpark Session customers_sdf = session.table('CUSTOMER_TRANSACTIONS') recency_sdf = customers_sdf.groupBy('CUSTOMER_ID').agg(F.max('TRANSACTION_DATE').alias('LAST_TRANSACTION_DATE')) # Calculate Recency recency_sdf = recency_sdf.withColumn('RECENTY', F.datediff(F.current_date(),F.to_date(F.col('LAST_TRANSACTION_DATE')))) final_sdf = recency_sdf
```

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

#### Answer: C

Explanation:

Option E is the only fully correct approach. It correctly groups by 'CUSTOMER\_ID' and finds the maximum transaction date. It calculates the Recency by using 'datediff', and casting 'LAST\_TRANSACTION\_DATE' with 'Without the cast to', it is possible to

run into error in 'datediff' function. 'datediff' function will cause issues when used on a timestamp. The 'recency\_sdf' data frame will only have customer\_id and recency.

### NEW QUESTION # 256

You are developing a model to predict equipment failure in a factory using sensor data stored in Snowflake. The data is partitioned by 'EQUIPMENT\_ID' and 'TIMESTAMP'. After initial model training and cross-validation using the following code snippet:

```
-- Assume TRAINING_DATA contains preprocessed sensor data
CREATE OR REPLACE MODEL equipment_failure_model
  INPUT_DATA => TABLE TRAINING_DATA
  TARGET_COL => 'FAILURE_FLAG'
  MODEL_TYPE => 'REGRESSION'
  PARTITION_COLS => ['EQUIPMENT_ID']
  -- Initial training parameters (simplified)
;
```

You observe significant performance variations across different equipment groups when evaluating on out-of-sample data'. Which of the following strategies could you employ to address this issue within the Snowflake environment to improve the model's generalization ability across all equipment?

- A. Implement cross-validation at the partition level by splitting 'TRAINING\_DATA' into train and test sets before creating the model, and then using the 'FIT' command to train on the train set and 'PREDICT' to evaluate on the test set, repeating for each partition.
- B. Implement a hyperparameter search using 'SYSTEM\$OPTIMIZE\_MODEL' with a wider range of parameters for each 'EQUIPMENT\_ID' individually, creating a separate model for each 'EQUIPMENT\_ID'.
- C. Retrain the model with additional feature engineering to create interaction terms between 'EQUIPMENT\_ID' and other relevant sensor features to capture equipment-specific patterns. For instance, you can one hot encode and add to model and include in 'INPUT DATA'.
- D. Increase the overall size of the 'TRAINING\_DATA' to include more historical data for all equipment, assuming this will balance the representation of each EQUIPMENT\_ID'
- E. Create separate models per equipment ID. For each equipment ID, split data into training and testing data. For each equipment ID, use 'SYSTEM\$OPTIMIZE\_MODEL' to perform hyper parameter search individually. Train and Deploy the model at equipment ID Level.

**Answer: C,E**

**Explanation:**

Options C and E are the most effective strategies. Option C (Feature Engineering): By creating interaction terms between EQUIPMENT\_ID and other sensor features, the model can learn equipment-specific patterns. This enables the model to account for the unique characteristics of each equipment group, improving its ability to generalize across all equipment. For example, the optimal temperature threshold for triggering a failure might differ significantly between EQUIPMENT\_ID groups, and this can be captured using interaction terms. Option E (Separate models per Equipment ID): Hyperparameter tuning and training separate models per equipment ID enables you to optimize and customize the model specific to each equipment ID. The downside is that we need to create and manage more models. Options A and D are less effective or may have limitations: Option A (Increase Training Data Size): While increasing the training data size can sometimes improve model performance, it doesn't guarantee that the model will learn to differentiate between the equipment groups effectively, especially if some groups have significantly different data characteristics. This can also consume a lot of resources unnecessarily. Option D (Custom cross Validation) : While it's valid, it is difficult to implement and the built in Snowflake cross validation features is much more performant and easier to use.

### NEW QUESTION # 257

You are tasked with preparing customer data for a churn prediction model in Snowflake. You have two tables: 'customers' (customer\_id, name, signup\_date, plan\_id) and 'usage' (customer\_id, usage\_date, data\_used\_gb). You need to create a Snowpark DataFrame that calculates the total data usage for each customer in the last 30 days and joins it with customer information. However, the 'usage' table contains potentially erroneous entries with negative values, which should be treated as zero. Also, some customers might not have any usage data in the last 30 days, and these customers should be included in the final result with a total data usage of 0. Which of the following Snowpark Python code snippets will correctly achieve this?

- A.

```
from snowflake.snowpark.functions import sum,lit,coalesce
from snowflake.snowpark import functions as F

customers_df = session.table('customers')
usage_df = session.table('usage')

recent_usage_df = usage_df.filter(F.datediff('day', F.col('usage_date'), F.current_date()) <= 30)

aggregated_usage = recent_usage_df.with_column('data_used_gb', F.when(F.col('data_used_gb') < 0, lit(0)).otherwise(F.col('data_used_gb'))).groupBy('customer_id').agg(sum('data_used_gb').alias('total_data_usage'))

final_df = customers_df.join(aggregated_usage, customers_df['customer_id'] == aggregated_usage['customer_id'], how='right').select(customers_df['customer_id'], 'name', 'signup_date', 'plan_id', coalesce(F.col('total_data_usage'),lit(0)).alias('total_data_usage'))
```

- B. None of the above

- C.

```
from snowflake.snowpark.functions import sum,lit,coalesce
from snowflake.snowpark import functions as F

customers_df = session.table('customers')
usage_df = session.table('usage')

recent_usage_df = usage_df.filter(F.datediff('day', F.col('usage_date'), F.current_date()) <= 30)

aggregated_usage = recent_usage_df.with_column('data_used_gb', F.when(F.col('data_used_gb') < 0, lit(0)).otherwise(F.col('data_used_gb'))).groupBy('customer_id').agg(sum('data_used_gb').alias('total_data_usage'))

final_df = customers_df.join(aggregated_usage, customers_df['customer_id'] == aggregated_usage['customer_id'], how='inner').select(customers_df['customer_id'], 'name', 'signup_date', 'plan_id', coalesce(F.col('total_data_usage'),lit(0)).alias('total_data_usage'))
```

- D.

```
from snowflake.snowpark.functions import sum,lit,coalesce
from snowflake.snowpark import functions as F

customers_df = session.table('customers')
usage_df = session.table('usage')

recent_usage_df = usage_df.filter(F.datediff('day', F.col('usage_date'), F.current_date()) <= 30)

aggregated_usage = recent_usage_df.with_column('data_used_gb', F.when(F.col('data_used_gb') < 0, lit(0)).otherwise(F.col('data_used_gb'))).groupBy('customer_id').agg(sum('data_used_gb').alias('total_data_usage'))

final_df = customers_df.join(aggregated_usage, customers_df['customer_id'] == aggregated_usage['customer_id'], how='left').select(customers_df['customer_id'], 'name', 'signup_date', 'plan_id', coalesce(F.col('total_data_usage'),lit(0)).alias('total_data_usage'))
```

- E.

```
from snowflake.snowpark.functions import sum
customers_df = session.table('customers')
usage_df = session.table('usage')

recent_usage_df = usage_df.filter(F.datediff('day', F.col('usage_date'), F.current_date()) <= 30)

aggregated_usage = recent_usage_df.groupBy('customer_id').agg(sum(F.iff(F.col('data_used_gb') < 0, 0, F.col('data_used_gb'))).alias('total_data_usage'))

final_df = customers_df.join(aggregated_usage, customers_df['customer_id'] == aggregated_usage['customer_id'], how='left').select(customers_df['customer_id'], 'name', 'signup_date', 'plan_id', aggregated_usage['total_data_usage'])
```

**Answer: D**

**Explanation:**

Option A correctly addresses all requirements: Filters usage data for the last 30 days. Corrects negative values by setting them to 0 using and ' Calculates the sum of for each customer. Uses a 'LEFT JOIN' to include all customers, even those without recent usage data. Uses 'coalesce()' to set the to 0 for customers with no usage data after the join. Option B uses an ' INNER JOIN' , which would exclude customers without any recent usage data, violating the requirement to include all customers. Option C does not treat negative usage values correctly. Option D uses a 'RIGHT JOIN' which would return incorrect results. Option E isn't right as option A correctly addresses all the scenarios.

## NEW QUESTION # 258

You have developed a customer churn prediction model using Python and deployed it as a Snowflake UDE. You are monitoring its performance and notice a significant drop in accuracy over time. To address this, you need to implement automated model retraining with regular validation. Which of the following steps and validation techniques are MOST critical for ensuring the retrained model is effective and avoids overfitting to recent data? (Select THREE)

- A. Implement a data drift detection mechanism. Monitor the distribution of input features over time and trigger retraining if significant drift is detected using tools such as Snowflake's Anomaly Detection features or custom drift metrics calculated in SQL.
- B. Use cross-validation techniques (e.g., k-fold cross-validation) during the retraining process to estimate the model's performance on unseen data and prevent overfitting. Evaluate on a held-out validation set.
- C. Update the UDF in place using 'CREATE OR REPLACE FUNCTION' immediately after retraining completes, regardless of the validation results.
- D. Monitor the model's performance on a live dataset and trigger retraining only when the performance drops below a predefined threshold, using metrics like accuracy, precision, or recall. Save Model Performance to 'MODEL\_PERFORMANCE'.
- E. Retrain the model using the entire available dataset, as this will maximize the amount of data the model learns from.

**Answer: A,B,D**

Explanation:

B, C, and D are the most critical steps. Option B is essential because data drift can significantly impact model performance. Detecting and addressing data drift is crucial for maintaining accuracy over time. Option C is vital for preventing overfitting and ensuring the model generalizes well to unseen data. Cross-validation provides a more robust estimate of model performance than a single train-test split. Option D is necessary to ensure that the retraining process is only triggered when the model's performance degrades. Monitoring live data and using performance metrics as triggers is a key component of automated retraining. Option A is incorrect because retraining on the entire dataset without validation can lead to overfitting. Option E is dangerous, as it deploys the retrained model without confirming its effectiveness.

#### NEW QUESTION # 259

.....

If you buy Exam4Labs's Snowflake certification DSA-C03 exam practice questions and answers, you can not only pass Snowflake certification DSA-C03 exam, but also enjoy a year of free update service. If you fail your exam, Exam4Labs will full refund to you. You can free download part of practice questions and answers about Snowflake Certification DSA-C03 Exam as a try to test the reliability of Exam4Labs's products.

**Free DSA-C03 Updates:** <https://www.exam4labs.com/DSA-C03-practice-torrent.html>

The DSA-C03 authorized training exams provided by Exam4Labs helps you to clear about your strengths and weaknesses before you take the exam, Snowflake New DSA-C03 Dumps Free As old saying goes, genuine gold fears no fire, What's more, when you have shown your talent with Free DSA-C03 Updates - SnowPro Advanced: Data Scientist Certification Exam certification in relating field, naturally, you will have the chance to enlarge your friends circle with a lot of distinguished persons who may influence you career life profoundly, Our DSA-C03 training materials are popular because of high quality.

What's more, our DSA-C03 learning materials are committed to grasp the most knowledgeable points with the fewest problems, As anyone who's recently visited a newly designed corporate office knows, they look a lot like coworking spaces.

### Free PDF New DSA-C03 Dumps Free & Accurate Free DSA-C03 Updates Ensure You a High Passing Rate

The DSA-C03 authorized training exams provided by Exam4Labs helps you to clear about your strengths and weaknesses before you take the exam, As old saying goes, genuine gold fears no fire.

What's more, when you have shown your talent DSA-C03 with SnowPro Advanced: Data Scientist Certification Exam certification in relating field, naturally, you will have the chance to enlarge your friends circle with New DSA-C03 Dumps Free a lot of distinguished persons who may influence you career life profoundly.

Our DSA-C03 training materials are popular because of high quality, Many companies offer job opportunities to qualified candidates, but they have specific DSA-C03 certification criteria to select qualified candidates.

- 2026 New DSA-C03 Dumps Free | Pass-Sure Snowflake DSA-C03: SnowPro Advanced: Data Scientist Certification Exam 100% Pass ☐ ⇒ [www.pdfdumps.com](http://www.pdfdumps.com) ⇌ is best website to obtain ➔ DSA-C03 ☐☐☐ for free download ☐Valid Braindumps DSA-C03 Ppt
- Pass Guaranteed Quiz 2026 Snowflake DSA-C03: SnowPro Advanced: Data Scientist Certification Exam First-grade New Dumps Free ☐ Search for ✓ DSA-C03 ☐✓☐ and download exam materials for free through ⇒ [www.pdfvce.com](http://www.pdfvce.com) ⇌ ☐

□Reliable DSA-C03 Exam Papers

BTW, DOWNLOAD part of Exam4Labs DSA-C03 dumps from Cloud Storage: <https://drive.google.com/open?id=1OyeqeSo7GZByjmBsJ9IVPVk6xVBe0nqf>