

Realistic DSA-C03 New Braindumps Sheet–Pass DSA-C03 First Attempt

How to Pass AWS Certified Solutions Architect - Associate (SAA-C03) in First Attempt 2025?

Tips To Pass:

- Understand the Exam Blueprint
- Master AWS Core Services
- Study from the Right Resources
- Smart Exam Strategies
- Last 7 Days Revision Plan

100% Success In First Attempt

VISIT NOW

WWW.CLEARCATNET.COM

DOWNLOAD the newest DumpsTests DSA-C03 PDF dumps from Cloud Storage for free: https://drive.google.com/open?id=1f3jvLx0NHFHO_o9XH8pMZh_05HELcO

Working in IT industry, IT people most want to attend Snowflake certification exam. As a widely recognized certification examination, Snowflake certification exams are becoming more and more popular. Among them, Snowflake DSA-C03 certification test is the most important exam. Having DSA-C03 certificate proves you have high skills. Owing to its importance, it is very difficult to pass Snowflake DSA-C03 exam successfully. Although to pass the exam is hard, you also don't need to worry about it. DumpsTests exam dumps will help you sail through DSA-C03 test.

After clients pay successfully for our SnowPro Advanced: Data Scientist Certification Exam guide torrent, they will receive our mails sent by our system in 5-10 minutes. Then they can click the mail and log in to use our software to learn immediately. For that time is extremely important for the learners, everybody hope that they can get the efficient learning. So clients can use our DSA-C03 test torrent immediately is the great merit of our product. We have set strict computer procedure to protect the client's privacy about purchasing DSA-C03 Study Tool and there is no one which can see the privacy information through online or other illegal channels except us. We have set the rigorous interception procedure to protect others from stealing the client's personal privacy information.

>> DSA-C03 New Braindumps Sheet <<

Latest DSA-C03 Braindumps Questions & Latest Study DSA-C03 Questions

One way to makes yourself competitive is to pass the DSA-C03 certification exams. Hence, if you need help to get certified, you are in the right place. DumpsTests offers the most comprehensive and updated braindumps for DSA-C03's certifications. To ensure that our products are of the highest quality, we have tapped the services of DSA-C03 experts to review and evaluate our DSA-C03 certification test materials. In fact, we continuously provide updates to every customer to ensure that our DSA-C03 products can cope with the fast changing trends in DSA-C03 certification programs.

Snowflake SnowPro Advanced: Data Scientist Certification Exam Sample Questions (Q129-Q134):

NEW QUESTION # 129

You are working on a fraud detection model and need to prepare transaction data. You have two tables: 'transactions' (transaction_id, customer_id, transaction_date, amount, merchant_id) and (merchant_id, city, state). You need to perform the following data cleaning and feature engineering steps using Snowpark: 1. Remove duplicate transactions based on 'transaction_id'. 2. Join the 'transactions' table with the 'merchant_locations' table to add city and state information to each transaction. 3. Create a new feature called 'amount_category' based on the transaction amount, categorized as 'Low', 'Medium', or 'High'. 4. The categorization thresholds are defined as follows: 'LoW': amount < 50 'Medium': 50 amount < 200 'High': amount >= 200 Which of the following statements about performing these operations using Snowpark are accurate?

- A. A LEFT JOIN should be used to join the 'transactions' and 'merchant_location' tables to ensure that all transactions are included, even if some merchant IDs are not present in the 'merchant_location' table.
- B. The construct in Snowpark can be used to create the 'amount_category' feature directly within the DataFrame transformation without needing a UDF
- C. Removing duplicate transactions can be efficiently done using the method on the Snowpark DataFrame, specifying 'transaction_id' as the subset. Creating the amount categories can be completed using the 'when' clause with multiple 'otherwise' clauses.
- D. Removing duplicate transactions can be efficiently done using the method on the Snowpark DataFrame, specifying 'transaction_id' as the subset. Creating the amount categories requires use of a User-Defined Function (UDF) as the logic can't be efficiently embedded in a single 'when' clause.
- E. You can register SQL UDF to calculate the 'amount_category' using 'CASE WHEN' statement

Answer: B,C,E

Explanation:

Options C, D and E are correct. Option C is correct because Snowpark's construct allows creating new features based on conditional logic directly within DataFrame transformations, avoiding the need for a UDF for simple categorizations like this. Option D is correct because SQL UDF can be used to create a function that returns Option E is also correct because the method efficiently removes duplicates, and the 'when' clauses enables easy categorization of in snowflake. Option A is incorrect, the categorization doesn't necessarily require UDF. Option B is incorrect since a RIGHT or INNER join is valid as well.

NEW QUESTION # 130

A financial institution wants to use Snowflake Cortex to analyze customer reviews and feedback extracted from various online sources to gauge customer sentiment towards their new mobile banking application. The goal is to identify positive, negative, and neutral sentiments, and also extract key phrases that drive these sentiments. Which of the following steps represent a viable workflow for achieving this using Snowflake Cortex and related functionalities?

- A. 1. Ingest the customer reviews into a Snowflake table. 2. Use Snowflake's built-in 'NLP_SENTIMENT' function (if available) or a similar UDF based on a pre-trained sentiment analysis model to get the sentiment score. 3. Use regular expressions in SQL to extract key phrases based on frequency and context.
- B. 1. Create a Streamlit application hosted externally that connects to the Snowflake database. 2. The Streamlit app uses a Python library like 'transformers' to perform sentiment analysis and key phrase extraction on the customer reviews read from Snowflake. 3. The results are then written back to a separate Snowflake table.
- C. 1. Ingest the customer reviews into a Snowflake table. 2. Create a custom JavaScript UDF that calls the Snowflake Cortex 'COMPLETE' endpoint with a prompt that asks for both sentiment and key phrases. 3. Store the results in a new Snowflake table.
- D. 1. Ingest the customer reviews into a Snowflake table. 2. Use the 'SNOWFLAKML.PREDICT' function with the appropriate task-specific model to determine the sentiment score for each review. 3. Further fine-tune the sentiment model with customer review data to improve the score and accuracy.
- E. 1. Ingest the customer reviews into a Snowflake table. 2. Use the 'SNOWFLAKE.ML.PREDICT' function with a sentiment analysis model to determine the overall sentiment score for each review. 3. Apply a separate key phrase extraction model via 'SNOWFLAKE.ML.PREDICT' to identify important keywords in the reviews.

Answer: E

Explanation:

Option A is the most viable workflow. It leverages Snowflake Cortex directly to perform both sentiment analysis and key phrase extraction. By using the 'SNOWFLAKE.ML.PREDICT' function with appropriate models, it keeps the processing within the Snowflake environment and avoids the need for external dependencies or custom coding (as in options B and C). The rest of the options are less effective because they involve use third party components when Snowflake Cortex readily has modules that can do what is required.

NEW QUESTION # 131

You are working on a customer churn prediction project. One of the features you want to normalize is 'customer_age'. However, a Snowflake table constraint ensures that all 'customer_age' values are between 0 and 120 (inclusive). Furthermore, you want to avoid using any stored procedures and prefer a pure SQL approach for data transformation. Considering these constraints, which normalization technique and associated SQL query is the most appropriate in Snowflake for this scenario, guaranteeing that the scaled values remain within a predictable range?

- A. Box-Cox transformation:

```
SELECT customer_age
```

- B. Z-score standardization after clipping values outside 1 and 99 percentile:

```
WITH Percentiles AS (SELECT APPROX_PERCENTILE(customer_age, 0.01) AS p01, APPROX_PERCENTILE(customer_age, 0.99) AS p99 FROM your_table),  
ClippedAges AS (SELECT LEAST(GREATEST(customer_age, (SELECT p01 FROM Percentiles)), (SELECT p99 FROM Percentiles)) AS clipped_age FROM your_table)  
SELECT (clipped_age - AVG(clipped_age) OVER ()) / STDDEV(clipped_age) OVER () FROM ClippedAges;
```

- C. Min-Max scaling to the range [0, 1]:

```
(customer_age - MIN(customer_age) OVER ()) / (MAX(customer_age) OVER () - MIN(customer_age) OVER ())
```

- D. Z-score standardization:

```
SELECT (customer_age - AVG(customer_age) OVER ()) / STDDEV(customer_age) OVER ()
```

- E. Min-Max scaling directly to the range [0, 1] using the known bounds (0 and 120):

```
(customer_age - 0) / (120 - 0)
```

Answer: E

Explanation:

Option D is the most appropriate. Given the existing constraint on 'customer_age' (0-120), and the requirement to avoid stored procedures, directly scaling to the range [0, 1] using the known minimum and maximum values is efficient and guarantees the output remains within a predictable range. This approach avoids data-dependent calculations (like MIN and MAX over the entire dataset) which are unnecessary given the constraint. Option A won't guarantee values within [0, 1]. Option B is correct but option D is the efficient solution to get the expected outcome and avoid cost and complexity. Option C would not scale to between 0 and 1 and adds complexity. Option E is not a normalization technique.

NEW QUESTION # 132

You are building an automated model retraining pipeline for a sales forecasting model in Snowflake using Snowflake Tasks and Stored Procedures. After retraining, you want to validate the new model against a champion model already deployed. You need to define a validation strategy using the following models: champion model deployed as UDF 'FORECAST_UDF', and contender model deployed as UDF 'FORECAST_UDF_NEW'. Given the following objectives: (1) Minimal impact on production latency, (2) Ability to compare predictions on a large volume of real-time data, (3) A statistically sound comparison metric. Which of the following SQL statements best represents how to efficiently compare the forecasts of the two models on a sample dataset and calculate the Root Mean Squared Error (RMSE) to validate the new model?

- A.

```
CREATE OR REPLACE TEMPORARY TABLE model_comparison AS  
SELECT  
  AVG(POWER(FORECAST_UDF(feature1, feature2) - FORECAST_UDF_NEW(feature1, feature2), 2)) AS mse,  
  SQRT(mse) AS rmse  
FROM sales_data;  
  
SELECT rmse FROM model_comparison;
```



```

CREATE OR REPLACE TEMPORARY TABLE model_comparison AS
SELECT
    sales_data.actual_sales,
    FORECAST_UDF(feature1, feature2) AS champion_forecast,
    FORECAST_UDF_NEW(feature1, feature2) AS challenger_forecast
FROM sales_data
SAMPLE BERNOULLI(1);

SELECT
    SQRT(AVG(POWER(actual_sales - challenger_forecast, 2))) AS challenger_rmse,
    SQRT(AVG(POWER(actual_sales - champion_forecast, 2))) AS champion_rmse,
    challenger_rmse - champion_rmse AS rmse_diff
FROM model_comparison;

```

- B.
- C.

```

CREATE OR REPLACE TEMPORARY TABLE model_comparison AS
SELECT
    sales_data.actual_sales,
    FORECAST_UDF(feature1, feature2) AS champion_forecast,
    FORECAST_UDF_NEW(feature1, feature2) AS challenger_forecast
FROM sales_data
SAMPLE BERNOULLI(10);
SELECT
    SQRT(AVG(POWER(actual_sales - challenger_forecast, 2))) AS challenger_rmse,
    SQRT(AVG(POWER(actual_sales - champion_forecast, 2))) AS champion_rmse
FROM model_comparison;

```

```

CREATE OR REPLACE TEMPORARY TABLE model_comparison AS
SELECT
    sales_data.actual_sales,
    FORECAST_UDF(feature1, feature2) AS champion_forecast,
    FORECAST_UDF_NEW(feature1, feature2) AS challenger_forecast
FROM sales_data
SAMPLE BERNOULLI(10);

SELECT
    SQRT(AVG(POWER(actual_sales - challenger_forecast, 2))) AS challenger_rmse,
    SQRT(AVG(POWER(actual_sales - champion_forecast, 2))) AS champion_rmse,
FROM model_comparison
WHERE challenger_rmse < champion_rmse;

```

- D.

```

CREATE OR REPLACE TEMPORARY TABLE model_comparison AS
SELECT
    sales_data.actual_sales,
    FORECAST_UDF(feature1, feature2) AS champion_forecast,
    FORECAST_UDF_NEW(feature1, feature2) AS challenger_forecast
FROM sales_data
LIMIT 1000;

SELECT
    SQRT(AVG(POWER(champion_forecast - challenger_forecast, 2)))
FROM model_comparison;

```

Answer: C

Explanation:

Option E is the best approach. It samples the data using 'SAMPLE BERNOULLI(IO)' for minimal impact on production. Then, it calculates both the challenger RMSE (new model) and the champion RMSE on this sample data. This provides a direct comparison of the model performance against actual sales and also allows to minimize runtime to compute this metric compared to option C which computes a difference without evaluating if the new model has a better score. Sampling helps with minimal impact while comparison metric in this case needs the actual_sales column. This provides a statistically relevant comparison within Snowflake, minimizing external processing. Option A does not compare the model to the ground truth (actual sales). Option B only compares the challenger and champion models' predictions against each other on a small, limited dataset (1000 records), which may not be representative. Option C calculates the RMSE difference directly and has a SAMPLE size of 1, which is unlikely to reflect the reality and Option D filters based on RMSE, which makes the approach bias and makes it harder to evaluate if the RMSE is statistically significant.

NEW QUESTION # 133

You are performing exploratory data analysis on a dataset containing customer transaction data in Snowflake. The dataset has a column named 'transaction_amount' and a column named 'customer_segment'. You want to analyze the distribution of transaction amounts for each customer segment using Snowflake's statistical functions. Which of the following approaches would BEST achieve this, providing insights into the central tendency and spread of the data?

- Using 'APPROX_COUNT_DISTINCT(transaction_amount)' and 'AVG(transaction_amount)' functions grouped by 'customer_segment'.
- Using 'STDDEV(transaction_amount)', 'VARIANCE(transaction_amount)', and 'MEDIAN(transaction_amount)' functions grouped by 'customer_segment'.
- Using 'MIN(transaction_amount)', 'MAX(transaction_amount)', and 'COUNT(*)' functions grouped by 'customer_segment'.
- Using 'CORR(transaction_amount, customer_segment)' and 'COVAR_POP(transaction_amount, customer_segment)' functions.
- Using 'AVG(transaction_amount)', 'MEDIAN(transaction_amount)', 'STDDEV(transaction_amount)', and 'QUANTILE(transaction_amount, 0.25, 0.5, 0.75)' functions grouped by 'customer_segment'.

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

Answer: C

Explanation:

Option E is the best approach. It uses to calculate the mean, to calculate the median (robust to outliers), to calculate the standard deviation (measure of spread), and 'QUANTILE(transaction_amount, 0.25, 0.5, 0.75)' to calculate the quartiles (25th, 50th, and 75th percentiles), all grouped by 'customer_segment'. This provides a comprehensive view of the distribution. Option A only provides an approximate count of distinct transaction amounts and the average. Option B provides standard deviation, variance, and median but lacks the mean and quartiles. Option C provides the range and count, which are useful but not as comprehensive. Option D calculates correlation and covariance, which are useful for understanding the relationship between transaction amount and

myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, www.stes.tyc.edu.tw, eishkul.com,
www.wcs.edu.eu, daotao.wisebusiness.edu.vn, a.gdds.top, www.stes.tyc.edu.tw, Disposable vapes

2025 Latest DumpsTests DSA-C03 PDF Dumps and DSA-C03 Exam Engine Free Share: https://drive.google.com/open?id=1f3jvLx0NHFHO_o9XHz8pMZh_05HELcO