

Exam DSA-C03 Score & Latest DSA-C03 Dumps Questions

QUESTION	SCALED SCORE			
	SET A EASY	SET B MID	SET C HARD	SET ...
0	0	0	0	More SAA-C03 Sets...
1	225	234	240	
...	...	...	...	
29	709	730	751	
30	719	751	560	
31	730	756	780	
32	751	760	747	
...	...	...	...	
50	1000	1000	1000	

Small 0-50 Question Architecture Questions

An Oversimplified Representation of the Scaled Scoring Model used in the AWS Certification tests

TD

TUTORIALS DOJO

What's more, part of that FreeDumps DSA-C03 dumps now are free: https://drive.google.com/open?id=1Ilpov_zQ1CVTgMPMskWkLjkn3xILFkIM

FreeDumps have the latest Snowflake certification DSA-C03 exam training materials. The industrious FreeDumps's IT experts through their own expertise and experience continuously produce the latest Snowflake DSA-C03 training materials to facilitate IT professionals to pass the Snowflake Certification DSA-C03 Exam. The certification of Snowflake DSA-C03 more and more valuable in the IT area and a lot people use the products of FreeDumps to pass Snowflake certification DSA-C03 exam. Through so many feedbacks of these products, our FreeDumps products prove to be trusted.

Before clients buy our DSA-C03 questions torrent they can download them and try out them freely. The pages of our product provide the demo and the aim is to let the client know part of our titles before their purchase and what form our DSA-C03 guide torrent is. The pages introduce the quantity of our questions and answers of our DSA-C03 Guide Torrent. After you try out the free demo you could decide whether our DSA-C03 exam torrent is worthy to buy or not. So you needn't worry that you will waste your money or our DSA-C03 exam torrent is useless and boosts no values.

>> Exam DSA-C03 Score <<

Quiz 2025 Snowflake High-quality DSA-C03: Exam SnowPro Advanced: Data Scientist Certification Exam Score

Our online resources and events enable you to focus on learning just what you want on your timeframe. You get access to every DSA-C03 exams files and there continuously update our DSA-C03 Study Materials; these exam updates are supplied free of charge to our valued customers. Get the best DSA-C03 exam Training; as you study from our exam-files.

Snowflake SnowPro Advanced: Data Scientist Certification Exam Sample Questions (Q134-Q139):

NEW QUESTION # 134

You are developing a Python stored procedure in Snowflake to train a machine learning model using scikit-learn. The training data resides in a Snowflake table named 'SALES DATA'. You need to pass the feature columns (e.g., 'PRICE', 'QUANTITY') and the target column ('REVENUE') dynamically to the stored procedure. Which of the following approaches is the MOST secure and efficient way to achieve this, preventing SQL injection vulnerabilities and ensuring data integrity within the stored procedure?

- Pass the column names directly as strings in the SQL call to the stored procedure and use string formatting within the Python code to construct the SELECT statement. E.g., `CALL TRAIN_MODEL('PRICE, QUANTITY', 'REVENUE');` and then build the query `SELECT {feature_cols}, {target_col} FROM SALES_DATA`.
- Pass the column names as a VARIANT array in the SQL call to the stored procedure, and then access the elements of the array within the Python code to dynamically construct and execute the SELECT statement using Snowflake's cursor execute method with parameterized queries. E.g., `CALL TRAIN_MODEL(ARRAY_CONSTRUCT('PRICE', 'QUANTITY'), 'REVENUE');` and then use `cursor.execute('SELECT {}, {} FROM SALES_DATA', [feature_cols[0], feature_cols[1]])` after parsing the array in python.
- Pass the complete SELECT statement as a string in the SQL call to the stored procedure. E.g., `CALL TRAIN_MODEL('SELECT PRICE, QUANTITY, REVENUE FROM SALES_DATA');` and then execute this SQL statement directly using Snowflake's cursor. This relies on the caller to ensure the statement is valid.
- Use Snowflake's dynamic data masking policies to mask sensitive data columns before passing the data to the stored procedure, even though the column names are passed as strings. Then pass the column names directly as strings in the SQL call to the stored procedure and construct the query. E.g., `CALL TRAIN_MODEL('PRICE, QUANTITY', 'REVENUE');`
- Define a Snowflake view that selects only the necessary feature and target columns and then pass the view name to the stored procedure. The stored procedure selects all columns from the view using `SELECT FROM`. This avoids passing column names directly.

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

Answer: E

Explanation:

Passing the column names as a VARIANT array and using parameterized queries is the safest and most efficient approach. This avoids SQL injection vulnerabilities, as the column names are treated as data rather than code. It also allows Snowflake to optimize the query execution plan. Options A and C are vulnerable to SQL injection. Option D doesn't address the core problem of dynamically specifying columns and security. Option E introduces an extra layer of abstraction (the view) but doesn't inherently solve the dynamic column specification or SQL injection risks if the view definition is itself dynamically constructed.

NEW QUESTION # 135

You've deployed a regression model in Snowflake to predict product sales. After a month, you observe that the RMSE on your validation dataset has increased significantly compared to the initial deployment. Analyzing the prediction errors, you notice a pattern: the model consistently underestimates sales for products with a recent surge in social media mentions. Which of the following actions would be MOST effective in addressing this issue and improving the model's RMSE?

- A. Increase the regularization strength of the model to prevent overfitting to the original training data.
- B. Retrain the model using only the most recent data (e.g., last week) to adapt to the changing sales patterns.
- C. Decrease the learning rate of the optimization algorithm during retraining to avoid overshooting the optimal weights.
- D. Implement a moving average smoothing technique on the target variable (sales) before retraining the model.
- E. Incorporate a feature representing the number of social media mentions for each product into the model and retrain.

Answer: E

Explanation:

Incorporating the social media mentions feature directly addresses the observed pattern in the errors. While other options might have some impact, adding the missing information is the most targeted and effective approach. Option A might help prevent overfitting, but doesn't address the missing information. Option B could lead to instability if the recent data isn't representative. Option D affects training but isn't specific to the issue. Option E smooths the target but doesn't explicitly account for social media influence.

NEW QUESTION # 136

You have a regression model deployed in Snowflake predicting customer churn probability, and you're using RMSE to monitor its performance. The current production RMSE is consistently higher than the RMSE you observed during initial model validation. You suspect data drift is occurring. Which of the following are effective strategies for monitoring, detecting, and mitigating this data drift to improve RMSE? (Select TWO)

- A. Implement a process to continuously calculate and track the RMSE on a holdout dataset representing the most recent data, alerting you when the RMSE exceeds a predefined threshold.
- B. Disable model monitoring, because the increased RMSE shows that the model is adapting to new patterns.
- C. Randomly sample a large subset of the production data and manually compare it to the original training data to identify any

differences.

- D. Use Snowflake's data lineage features to identify any changes in the upstream data sources feeding the model and assess their potential impact.
- E. Regularly re-train the model on the entire historical dataset to ensure it captures all possible data patterns.

Answer: A,D

Explanation:

Option A provides a proactive approach to monitoring the model's performance on new data and triggering alerts when the RMSE deteriorates. Option C helps identify changes in the input data that could be causing the drift. Option B is not ideal, as retraining on all historical data might not effectively adapt to recent drifts. Option D is inefficient and impractical for large datasets. Option E is incorrect because a high RMSE indicates poor performance and warrants investigation, not ignoring.

NEW QUESTION # 137

You are building a predictive model on customer churn using Snowflake data'. You observe that the distribution of 'TIME SINCE LAST PURCHASE' is heavily left-skewed. Which of the following strategies would be MOST appropriate to handle this skewness before feeding the data into a linear regression model to improve its performance? (Select TWO)

- A. Remove all records with 'TIME SINCE LAST PURCHASE' values below the mean.
- B. Apply a square root transformation to the 'TIME SINCE LAST PURCHASE' column.
- C. Apply a logarithmic transformation to the 'TIME SINCE LAST PURCHASE' column.
- D. Standardize the 'TIME SINCE LAST PURCHASE' column using Z-score normalization.
- E. Use a winsorization technique to cap extreme values in the 'TIME SINCE LAST PURCHASE' column at a predefined percentile (e.g., 99th percentile).

Answer: B,E

Explanation:

For left-skewed data, a square root transformation (Option B) can help reduce the impact of smaller values and bring the distribution closer to normal. Winsorization (Option C) can mitigate the influence of extreme values on the left tail of the distribution, making the data more suitable for linear regression. Logarithmic transformation is more suitable for right-skewed data (Option A). Z-score normalization (Option D) centers the data around zero but doesn't change the skewness. Removing records below the mean (Option E) is generally not a good practice as it can introduce bias and lose valuable information.

NEW QUESTION # 138

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. 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.
- B. 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.
- C. 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'.

- D. 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'.
- E. 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.

Answer: A,D

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 # 139

.....

Overall we can say that SnowPro Advanced: Data Scientist Certification Exam (DSA-C03) certification can provide you with several benefits that can assist you to advance your career and achieve your professional goals. Are you ready to gain all these personal and professional benefits? Looking for a sample, is smart and quick for SnowPro Advanced: Data Scientist Certification Exam (DSA-C03) exam dumps preparation? If your answer is yes then you do not need to go anywhere, just download FreeDumps DSA-C03 Questions and start SnowPro Advanced: Data Scientist Certification Exam (DSA-C03) exam preparation with complete peace of mind and satisfaction.

Latest DSA-C03 Dumps Questions: <https://www.freedumps.top/DSA-C03-real-exam.html>

To buy real DSA-C03 Exam Dumps, Snowflake Exam DSA-C03 Score Our training program includes simulation test before the formal examination, specific training course and the current exam which has 95% similarity with the real exam, Snowflake Exam DSA-C03 Score Failing to address these issues can result in wasted time and money, The test questions cover the practical questions in the test DSA-C03 certification and these possible questions help you explore varied types of questions which may appear in the DSA-C03 test and the approaches you should adapt to answer the questions.

This chapter also shows the tools' changes from earlier versions, This DSA-C03 is probably evident to anyone who has had to configure something on multiple machines or perform the same configuration numerous times.

Free PDF 2025 Professional Snowflake Exam DSA-C03 Score

To buy real DSA-C03 Exam Dumps, Our training program includes simulation test before the formal examination, specific training course and the current exam which has 95% similarity with the real exam.

Failing to address these issues can result in wasted time and money, The test questions cover the practical questions in the test DSA-C03 certification and these possible questions help you explore varied types of questions which may appear in the DSA-C03 test and the approaches you should adapt to answer the questions.

Do you want to get more recognition and employment opportunities?

- Free PDF 2025 DSA-C03: Latest Exam SnowPro Advanced: Data Scientist Certification Exam Score ☐ Search for 《 DSA-C03 》 and obtain a free download on (www.pdf.dumps.com) ☐ DSA-C03 Hot Spot Questions
- High-quality Exam DSA-C03 Score | DSA-C03 100% Free Latest Dumps Questions ☐ Open ☒ www.pdfvce.com ☐ ☒ and search for ☐ DSA-C03 ☐ to download exam materials for free ☐ DSA-C03 Hot Spot Questions
- Free PDF Exam DSA-C03 Score | Amazing Pass Rate For DSA-C03 Exam | First-Grade DSA-C03: SnowPro Advanced: Data Scientist Certification Exam ☐ [【 www.passcollection.com 】](http://www.passcollection.com) is best website to obtain [【 DSA-C03 】](#) for free download ☐ DSA-C03 Updated Test Cram

- [illegible]

P.S. Free & New DSA-C03 dumps are available on Google Drive shared by FreeDumps: https://drive.google.com/open?id=1Ilpov_zQ1CVTgMPMskWkLjkn3xILFkIM