

Money-Back Guarantee for Snowflake DSA-C03 Exam Questions



BTW, DOWNLOAD part of DumpsFree DSA-C03 dumps from Cloud Storage: <https://drive.google.com/open?id=11MzVS4pHAO-gWAWDyoBMi3pPdMZbHHx2>

Our DSA-C03 exam braindumps are set high standards for your experience. That is the reason why our DSA-C03 training questions gain well brand recognition and get attached with customers all these years around the world. Besides, our DSA-C03 learning questions are not only high effective but priced reasonably. Their prices are acceptable for everyone and help you qualify yourself as and benefit your whole life.

100% correct answers of SnowPro Advanced: Data Scientist Certification Exam flexible testing engine - unlimited exam practice! DSA-C03 exam learning materials has high pass rate. Test price is resonable and Snowflake certification exam dumps is updated. Exam actual practice test engine is for free. DSA-C03 Certification Book Torrent Download now! DSA-C03 Free pdf guide 365 days are updates.

>> **DSA-C03 Hot Questions <<**

DSA-C03 Online Lab Simulation, DSA-C03 Interactive EBook

There is almost no innovative and exam-oriented format that can be compared with the precision and relevance of the actual SnowPro Advanced: Data Scientist Certification Exam exam questions, you get with DumpsFree brain dumps PDF. As per the format of the DSA-C03 Exam, our experts have consciously created a questions and answers pattern. It saves your time by providing you direct and precise information that will help you cover the syllabus contents within no time.

Snowflake SnowPro Advanced: Data Scientist Certification Exam Sample Questions (Q147-Q152):

NEW QUESTION # 147

A data scientist is tasked with identifying customer segments for a new marketing campaign using transaction data stored in Snowflake. The transaction data includes features like transaction amount, frequency, recency, and product category. Which unsupervised learning algorithm would be MOST appropriate for this task, considering scalability and Snowflake's data processing

capabilities, and what preprocessing steps are crucial before applying the algorithm?

- A. K-Means clustering, after standardizing numerical features (transaction amount, frequency, recency) and using one-hot encoding for product category. This is highly scalable within Snowflake using UDFs and SQL.
- B. Hierarchical clustering, using the complete linkage method and Euclidean distance. No preprocessing is necessary, as hierarchical clustering can handle raw data.
- C. K-Means clustering, after applying min-max scaling to numerical features and converting categorical features to numerical representation. The optimal 'k' (number of clusters) should be determined using the elbow method or silhouette analysis.
- D. Principal Component Analysis (PCA) followed by K-Means. This reduces dimensionality and then clusters, improving the visualization of the cluster.
- E. DBSCAN, using raw data without any scaling or encoding. The algorithm's density-based nature will automatically handle the varying scales of the features.

Answer: C

Explanation:

K-Means clustering is a suitable algorithm for customer segmentation due to its scalability and efficiency. Min-max scaling is important to ensure that features with larger ranges don't dominate the distance calculations. Converting categorical features to numerical representation (e.g., one-hot encoding) is also essential for K-Means. The elbow method or silhouette analysis helps determine the optimal number of clusters. Options A, B, C, and D have flaws related to scaling requirements, algorithm suitability for large datasets, or lack of pre-processing.

NEW QUESTION # 148

A data scientist is analyzing website traffic data stored in Snowflake. The data includes daily page views for different pages. The data scientist suspects that the variance of page views for a particular page, 'home', has significantly increased recently. Which of the following steps and Snowflake SQL queries could be used to identify a potential change in the variance of 'home' page views over time (e.g., comparing variance before and after a specific date)? Select all that apply.

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

Answer: A,B,C,D

Explanation:

Options B, C, D and E are correct. Option B directly compares the variance before and after a date, allowing for a direct assessment of change. Option C uses a window function for a rolling variance calculation, revealing trends over time. Option D creates a histogram, which helps visualize the distribution and identify shifts in spread. Option E calculates standard deviation before and after a date. Option A, while calculating the overall variance, doesn't provide insight into changes over time.

NEW QUESTION # 149

A data scientist is developing a model within a Snowpark Python environment to predict customer churn. They have established a Snowflake session and loaded data into a Snowpark DataFrame named 'customer data'. The feature engineering pipeline requires a custom Python function, 'calculate engagement_score', to be applied to each row. This function takes several columns as input and returns a single score representing customer engagement. The data scientist wants to apply this function in parallel across the entire DataFrame using Snowpark's UDF capabilities. The following code snippet is used to define and register the UDF:

When the UDF is called the above error is observed. What change needs to be applied to make the UDF work as expected?

- A. Redefine the function to accept string arguments and cast them to the correct data types within the function.
- B. Wrap the Python function inside a stored procedure using '@F.sproc' and call that stored procedure instead of the plain python function.
- C. Change the function call to use the Snowpark DataFrame's 'select' function with column objects: 'customer_data.select(engagement_score_udf(F.col('num_transactions'), F.col('avg_transaction_value')))
- D. Remove argument from 'session.udf.register' call. Snowpark can infer the input types automatically.
- E. Add '@F.sproc' decorator before the function definition.

Answer: C

Explanation:

The error message 'UDFArgumentException: Invalid argument types for function 'calculate_engagement_score_udf. Expected arguments: [ONT, FLOAT, INT], actual arguments: [COLUMN_NAME, COLUMN_NAME, COLUMN_NAME]' indicates that the UDF is receiving column objects instead of the actual data values. This is because when calling the UDF on a Snowpark DataFrame, you need to explicitly reference the columns using The correct way to apply the UDF to the DataFrame is to use the 'select' function with 'F.col()' to pass the column objects as arguments to the UDF.

NEW QUESTION # 150

You are a data scientist working for an e-commerce company. You have a table named 'sales_data' with columns 'product_id', 'customer_id', 'transaction_date', and 'sale_amount'. You need to identify the top 5 products by total sale amount for each month. Which of the following Snowflake SQL queries is the MOST efficient and correct way to achieve this, while also handling potential ties in sale amounts?

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

Answer: A,D

Explanation:

Options C and E are correct. Both use a subquery to calculate the rank of each product within each month's sales, then filter for the top 5 products. The main difference is that option C uses DENSE_RANK(), which assigns consecutive ranks even if there are ties in sales amount (resulting in more than 5 products being selected if there are ties for the 5th position), while option E uses RANK(), which assigns the same rank to tied values but can skip ranks. Option A is incorrect because it attempts to filter using HAVING on a ranking calculated within the same query level, which is not allowed in many SQL implementations (and can be logically incorrect). Options B and D are incorrect as they employ ROW_NUMBER() and NTILE(5) respectively. ROW_NUMBER will not handle ties correctly, while NTILE just divides the data into 5 groups without explicitly identifying the 'top' 5. Option A uses a rank function inside the HAVING clause which is often syntactically invalid.

NEW QUESTION # 151

A marketing team is using Snowflake to store customer data including demographics, purchase history, and website activity. They want to perform customer segmentation using hierarchical clustering. Considering performance and scalability with very large datasets, which of the following strategies is the MOST suitable approach?

- A. Utilize a SQL-based affinity propagation method directly within Snowflake. This removes the need for feature scaling and specialized hardware.
- B. Randomly sample a small subset of the customer data and perform hierarchical clustering on this subset using an external tool like R or Python with scikit-learn. Assume that results generalize well to the entire dataset. Avoid using Snowflake for this purpose.
- C. Perform mini-batch K-means clustering using Snowflake's compute resources through a Snowpark DataFrame. Take a large sample of each mini-batch and perform hierarchical clustering on each mini-batch and then create clusters of clusters.
- D. Directly apply an agglomerative hierarchical clustering algorithm with complete linkage to the entire dataset within Snowflake, using SQL. This is computationally feasible due to SQL's efficiency.
- E. **Employ BIRCH clustering with Snowflake Python UDF. Configure Snowflake resources accordingly. Optimize the clustering process. And tune parameters.**

Answer: E

Explanation:

Hierarchical clustering has a high time complexity, making it impractical for large datasets. While mini-batch K-means provides the most efficient option for large datasets. BIRCH is more suited for huge datasets and can be applied as a Snowflake Python UDF with Snowpark DataFrames to provide scalability and high performance as its better than other clustering such as affinity propagation. Options A and E are impractical due to the computational cost of hierarchical clustering in SQL or affinity propagation in SQL. Sampling (Option C) can lead to inaccurate results.

NEW QUESTION # 152

Immediately after you have made a purchase for our DSA-C03 practice test, you can download our exam study materials to make preparations for the exams. It is universally acknowledged that time is a key factor in terms of the success of exams. The more time you spend in the preparation for DSA-C03 training materials, the higher possibility you will pass the exam. And with our DSA-C03 study torrent, you can make full use of those time originally spent in waiting for the delivery of exam files. There is why our DSA-C03 test prep exam is well received by the general public.

DSA-C03 Online Lab Simulation: <https://www.dumpsfree.com/DSA-C03-valid-exam.html>

DSA-C03 soft test engine can simulate the real test, thus you can take a simulation test in advance, If you are determined to pass exams as soon as possible, the wise choice is to select our DSA-C03 exam preparation, Quality of DSA-C03 learning quiz you purchased is of prior importance for consumers, Yes, we provide you with the comprehensive and most valid DSA-C03 study torrent.

Creating a New HyperTerminal Connectoid, In addition to DSA-C03 the memory used by the object itself, each thread has two execution call stacks allocated for it by the JavaVM.

DSA-C03 soft test engine can simulate the real test, thus you can take a simulation test in advance. If you are determined to pass exams as soon as possible, the wise choice is to select our DSA-C03 Exam Preparation.

Three Easy-to-Use and Compatible Formats of DumpsFree Snowflake DSA-C03 Practice Test

Quality of DSA-C03 learning quiz you purchased is of prior importance for consumers, Yes, we provide you with the comprehensive and most valid DSA-C03 study torrent.

Self-assessment is a very important element of any test preparation process, so we provide you with one of our best products – Snowflake DSA-C03 practice test.

www.stes.tyc.edu.tw, Disposable vapes

2026 Latest DumpsFree DSA-C03 PDF Dumps and DSA-C03 Exam Engine Free Share: <https://drive.google.com/open?id=11MzVS4pHAO-gWAWDyoBMi3pPdMZbHHx2>