

Exam Snowflake DEA-C01 Labs - Training DEA-C01 Tools

Snowflake DEA-C01 Practice Questions

SnowPro Advanced Data Engineer Certification

Order our DEA-C01 Practice Questions Today and Get Ready to Pass with Flying Colors!



DEA-C01 Practice Exam Features | QuestionsTube

- Latest & Updated Exam Questions
- Subscribe to FREE Updates
- Both PDF & Exam Engine
- Download Directly Without Waiting

<https://www.questiontube.com/exam/dea-c01/>

At QuestionsTube, you can read DEA-C01 free demo questions in pdf file, so you can check the questions and answers before deciding to download the Snowflake DEA-C01 practice questions. These free demo questions are parts of the DEA-C01 exam questions. Download and read them carefully, you will find that the DEA-C01 test questions of QuestionsTube will be your great learning materials online. Share some DEA-C01 exam online questions below.

1. sum(ifnull(quantity * (s.price - m.item_price), 0)) as profit

BTW, DOWNLOAD part of ITPassLeader DEA-C01 dumps from Cloud Storage: <https://drive.google.com/open?id=11VpIecPnFhyugBIKSMVrxyTbvOtuXiXx>

Nowadays the requirements for jobs are higher than any time in the past. The job-hunters face huge pressure because most jobs require both working abilities and profound major knowledge. Passing DEA-C01 exam can help you find the ideal job. If you buy our DEA-C01 test prep you will pass the DEA-C01 Exam easily and successfully, and you will realize you dream to find an ideal job and earn a high income. Our DEA-C01 training braindump is of high quality and the passing rate and the hit rate are both high as more than 98%.

Snowflake DEA-C01 Exam Syllabus Topics:

Topic	Details
Topic 1	<ul style="list-style-type: none">• Storage and Data Protection: The topic tests the implementation of data recovery features and the understanding of Snowflake's Time Travel and micro-partitions. Engineers are evaluated on their ability to create new environments through cloning and ensure data protection, highlighting essential skills for maintaining Snowflake data integrity and accessibility.

Topic 2	<ul style="list-style-type: none"> • Data Transformation: The SnowPro Advanced: Data Engineer exam evaluates skills in using User-Defined Functions (UDFs), external functions, and stored procedures. It assesses the ability to handle semi-structured data and utilize Snowpark for transformations. This section ensures Snowflake engineers can effectively transform data within Snowflake environments, critical for data manipulation tasks.
Topic 3	<ul style="list-style-type: none"> • Data Movement: Snowflake Data Engineers and Software Engineers are assessed on their proficiency to load, ingest, and troubleshoot data in Snowflake. It evaluates skills in building continuous data pipelines, configuring connectors, and designing data sharing solutions.
Topic 4	<ul style="list-style-type: none"> • Performance Optimization: This topic assesses the ability to optimize and troubleshoot underperforming queries in Snowflake. Candidates must demonstrate knowledge in configuring optimal solutions, utilizing caching, and monitoring data pipelines. It focuses on ensuring engineers can enhance performance based on specific scenarios, crucial for Snowflake Data Engineers and Software Engineers.
Topic 5	<ul style="list-style-type: none"> • Security: The Security topic of the DEA-C01 test covers the principles of Snowflake security, including the management of system roles and data governance. It measures the ability to secure data and ensure compliance with policies, crucial for maintaining secure data environments for Snowflake Data Engineers and Software Engineers.

>> Exam Snowflake DEA-C01 Labs <<

Training DEA-C01 Tools, Exam DEA-C01 Fees

Our company deeply knows that product quality is very important, so we have been focusing on ensuring the development of a high quality of our DEA-C01 test torrent. All customers who have purchased our products have left deep impression on our DEA-C01 guide torrent. Of course, the customer not only has left deep impression on the high quality of our products but also the efficiency of our products. Our DEA-C01 Exam Questions can help you save much time, if you use our DEA-C01 study prep, you just need to spend 20-30 hours on learning, and you will pass your DEA-C01 exam successfully.

Snowflake SnowPro Advanced: Data Engineer Certification Exam Sample Questions (Q329-Q334):

NEW QUESTION # 329

You as Data engineer might want to consider disabling auto-suspend for a warehouse if?

- A. You require the warehouse to be available with delay.
- **B. You have a heavy, steady workload for the warehouse.**
- C. You have a low, fluctuating workload for the warehouse.
- **D. You require the warehouse to be available with no delay or lag time.**

Answer: B,D

Explanation:

Explanation

Automating Warehouse Suspension

Data Engineer might want to consider disabling auto-suspend for a warehouse if:

He/She have a heavy, steady workload for the warehouse.

He/She require the warehouse to be available with no delay or lag time. Warehouse provisioning is generally very fast (e.g. 1 or 2 seconds); however, depending on the size of the warehouse and the availability of compute resources to provision, it can take longer.

If he/she chose to disable auto-suspend, He/she must carefully consider the costs associated with running a warehouse continually, even when the warehouse is not processing queries. The costs can be significant, especially for larger warehouses (X-Large, 2X-Large, etc.).

To disable auto-suspend, Engineer must explicitly select Never in the web interface, or specify 0 or NULL in SQL.

NEW QUESTION # 330

A global ecommerce company processes customer transactions, inventory updates, and user activity logs across multiple AWS services. The company needs a scalable, fully managed, and event-driven orchestration solution to coordinate complex extract, transform, and load (ETL) workflows. The solution must use AWS Glue and Amazon EMR to process data. The data will be stored in Amazon Redshift and Amazon S3. The solution must support dependency management, automated retries, and data pipeline monitoring. Which solution will meet these requirements?

- A. Use AWS Step Functions to define an express workflow that invokes the data transformation and loading tasks across Amazon EMR and AWS Glue.
- B. Create an AWS Lambda function that runs each step of the workflow. Create an Amazon EventBridge scheduled rule to invoke the function every day.
- **C. Use Apache Airflow on Amazon Managed Workflows for Apache Airflow (Amazon MWAA) to create Directed Acyclic Graphs (DAGs) to manage ETL workflows.**
- D. Create AWS Lambda functions for each step of the workflow. Configure Amazon EventBridge to invoke AWS Glue jobs. Configure the Lambda functions to process and move data through the pipeline.

Answer: C

Explanation:

Amazon Managed Workflows for Apache Airflow is a fully managed orchestration service built specifically for coordinating complex ETL workflows with dependencies, retries, scheduling, and monitoring. Airflow DAGs can natively orchestrate AWS Glue and Amazon EMR jobs and manage data movement into Amazon Redshift and Amazon S3, making it the best fit for a scalable, event-driven ETL orchestration requirement.

NEW QUESTION # 331

What are Common Query Problems a Data Engineer can identify using Query Profiler?

- **A. Queries Too Large to Fit in Memory**
- **B. "Exploding" Joins i.e Joins resulting due to a "Cartesian product"**
- C. Ineffective Data Sharing
- **D. Inefficient Pruning**

Answer: A,B,D

Explanation:

Explanation

"Exploding" Joins

One of the common mistakes SQL users make is joining tables without providing a join condition (resulting in a "Cartesian product"), or providing a condition where records from one table match multiple records from another table. For such queries, the Join operator produces significantly (often by orders of magnitude) more tuples than it consumes.

This can be observed by looking at the number of records produced by a Join operator in the profile interface, and typically is also reflected in Join operator consuming a lot of time.

Queries Too Large to Fit in Memory

For some operations (e.g. duplicate elimination for a huge data set), the amount of memory available for the compute resources used to execute the operation might not be sufficient to hold intermediate results. As a result, the query processing engine will start spilling the data to local disk. If the local disk space is not sufficient, the spilled data is then saved to remote disks.

This spilling can have a profound effect on query performance (especially if remote disk is used for spilling).

Spilling statistics can be checked in Query Profile Interface.

Inefficient Pruning

Snowflake collects rich statistics on data allowing it not to read unnecessary parts of a table based on the query filters. However, for this to have an effect, the data storage order needs to be correlated with the query filter attributes.

The efficiency of pruning can be observed by comparing Partitions scanned and Partitions total statistics in the TableScan operators. If the former is a small fraction of the latter, pruning is efficient. If not, the pruning did not have an effect.

Of course, pruning can only help for queries that actually filter out a significant amount of data. If the pruning statistics do not show data reduction, but there is a Filter operator above TableScan which filters out a number of records, this might signal that a different data organization might be beneficial for this query.

NEW QUESTION # 332

A company needs to use an AWS Glue PySpark job to read specific data from an Amazon DynamoDB table. The company knows the partition key values for the required records. The existing processing logic of the AWS Glue PySpark job requires the data to be

in DynamicFrame format. The company needs a solution to ensure that the job reads only the specified data. Which solution will meet this requirement with the MINIMUM number of read capacity units (RCUs)?

- A. Perform a scan on the DynamoDB table in the Aws Glue job. Put the data into a DynamicFrame. Filter the DynamicFrame on the partition key.
- **B. Perform a query on the DynamoDB table in the AWS Glue job. Use the partition key in the key condition expression. Put the data into a DynamicFrame.**
- C. Use the AWS Glue DynamoDB ETL connector to read the DynamoDB table. Use the filter option to read the required partition key.
- D. Perform a query on the DynamoDB table in the AWS Glue job by using only the sort key in the key condition expression. Load the data into a DynamicFrame.

Answer: B

Explanation:

A DynamoDB Query that uses the partition key in the key condition expression reads only the items in the specified partition(s), which minimizes consumed RCUs compared with any Scan- based approach. The job can then convert the query results into a DynamicFrame to satisfy the existing processing logic.

NEW QUESTION # 333

A company receives a daily file that contains customer data in .xls format. The company stores the file in Amazon S3. The daily file is approximately 2 GB in size.

A data engineer concatenates the column in the file that contains customer first names and the column that contains customer last names. The data engineer needs to determine the number of distinct customers in the file.

Which solution will meet this requirement with the LEAST operational effort?

- **A. Use AWS Glue DataBrew to create a recipe that uses the COUNT_DISTINCT aggregate function to calculate the number of distinct customers.**
- B. Create and run an Apache Spark job in an AWS Glue notebook. Configure the job to read the S3 file and calculate the number of distinct customers.
- C. Create and run an Apache Spark job in Amazon EMR Serverless to calculate the number of distinct customers.
- D. Create an AWS Glue crawler to create an AWS Glue Data Catalog of the S3 file. Run SQL queries from Amazon Athena to calculate the number of distinct customers.

Answer: A

Explanation:

AWS Glue DataBrew: AWS Glue DataBrew is a visual data preparation tool that allows data engineers and data analysts to clean and normalize data without writing code. Using DataBrew, a data engineer could create a recipe that includes the concatenation of the customer first and last names and then use the COUNT_DISTINCT function. This would not require complex code and could be performed through the DataBrew user interface, representing a lower operational effort.

NEW QUESTION # 334

.....

All of our DEA-C01 pdf torrent are up-to-date and reviewed by our IT experts and professionals. We have written our DEA-C01 study guide in such a way that you don't need to prepare anything else after practice our DEA-C01 Exam Questions. You can pass the real exam easily with our latest DEA-C01 vce dumps and this is the only smartest way to get success. Just contact us if you have any questions.

Training DEA-C01 Tools: <https://www.itpassleader.com/Snowflake/DEA-C01-dumps-pass-exam.html>

- Test DEA-C01 Topics Pdf New DEA-C01 Exam Price Reliable DEA-C01 Test Cram Search for ▷ DEA-C01 on [www.vce4dumps.com] immediately to obtain a free download Valid Test DEA-C01 Tips
- Test DEA-C01 Topics Pdf Valid Test DEA-C01 Tips Sure DEA-C01 Pass The page for free download of > DEA-C01 on ➡ www.pdfvce.com will open immediately DEA-C01 Exam Guide Materials
- 2026 Exam DEA-C01 Labs | Perfect 100% Free Training DEA-C01 Tools Search for DEA-C01 and download exam materials for free through (www.dumpsmaterials.com) DEA-C01 Exam Tests
- 100% Pass Quiz 2026 DEA-C01: Perfect Exam SnowPro Advanced: Data Engineer Certification Exam Labs Simply

