

Latest DEA-C01 Exam Camp, DEA-C01 Test Dumps Demo



DEA-C01
SnowPro
Advanced
Data Engineer

Certification Questions & Exams Dumps

www.edurely.com

The banner features a blue border and a central illustration of a graduation cap on top of a stack of books, with a laptop and a smartphone nearby.

P.S. Free & New DEA-C01 dumps are available on Google Drive shared by Exams-boost: https://drive.google.com/open?id=1Bb2dy9705uVFu9N5digUmKQoq_HbkFxE

You should not register for the Snowflake SnowPro Advanced: Data Engineer Certification Exam certification exam without proper preparation. Passing the SnowPro Advanced: Data Engineer Certification Exam exam is quite a challenging task. This difficult task becomes easier if you use valid Snowflake DEA-C01 Exam Dumps of Exams-boost. Don't forget that the SnowPro Advanced: Data Engineer Certification Exam (DEA-C01) test registration fee is hefty and your money will go to waste if you don't crack this exam.

Snowflake DEA-C01 Exam Syllabus Topics:

Topic	Details
Topic 1	<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 2	<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.
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"> • 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 5	<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.

>> Latest DEA-C01 Exam Camp <<

Snowflake DEA-C01 Test Dumps Demo | Free DEA-C01 Braindumps

There are rare products which can rival with our products and enjoy the high recognition and trust by the clients like our products. Our products provide the DEA-C01 test guide to clients and help they pass the test DEA-C01 certification which is highly authorized and valuable. Our company is a famous company which bears the world-wide influences and our DEA-C01 Test Prep is recognized as the most representative and advanced study materials among the same kinds of products. Whether the qualities and functions or the service of our product, are leading and we boost the most professional expert team domestically.

Snowflake SnowPro Advanced: Data Engineer Certification Exam Sample Questions (Q57-Q62):

NEW QUESTION # 57

To view/monitor the clustering metadata for a table, Snowflake provides which of the following system functions?

- **A. SYSTEM\$CLUSTERING_DEPTH**
- B. SYSTEM\$CLUSTERING_DEPTH_KEY
- C. SYSTEM\$CLUSTERING_KEY_INFORMATION (including clustering depth)
- **D. SYSTEM\$CLUSTERING_INFORMATION (including clustering depth)**

Answer: A,D

Explanation:

Explanation

SYSTEM\$CLUSTERING_DEPTH:

Computes the average depth of the table according to the specified columns (or the clustering key defined for the table). The average depth of a populated table (i.e. a table containing data) is always 1 or more. The smaller the average depth, the better clustered the table is with regards to the specified columns.

Calculate the clustering depth for a table using two columns in the table:

```
SELECT SYSTEM$CLUSTERING_DEPTH('TPCH_PRODUCT', '(C2, C9');
```

SYSTEM\$CLUSTERING_INFORMATION:

Returns clustering information, including average clustering depth, for a table based on one or more columns in the table.

```
SELECT SYSTEM$CLUSTERING_INFORMATION('SAMPLE_TABLE', '(col1, col3');
```

NEW QUESTION # 58

A company maintains a data warehouse in an on-premises Oracle database. The company wants to build a data lake on AWS. The company wants to load data warehouse tables into Amazon S3 and synchronize the tables with incremental data that arrives from the data warehouse every day.

Each table has a column that contains monotonically increasing values. The size of each table is less than 50 GB. The data warehouse tables are refreshed every night between 1 AM and 2 AM.

A business intelligence team queries the tables between 10 AM and 8 PM every day.

Which solution will meet these requirements in the MOST operationally efficient way?

- A. Use an AWS Glue Java Database Connectivity (JDBC) connection. Configure a job bookmark for a column that contains monotonically increasing values. Write custom logic to append the daily incremental data to a full-load copy that is in Amazon

S3.

- B. Use AWS Glue to load a full copy of the data warehouse tables into Amazon S3 every day. Overwrite the previous day's full-load copy every day.
- C. Use an AWS Database Migration Service (AWS DMS) full load plus CDC job to load tables that contain monotonically increasing data columns from the on-premises data warehouse to Amazon S3. Use custom logic in AWS Glue to append the daily incremental data to a full-load copy that is in Amazon S3.
- D. Use an AWS Database Migration Service (AWS DMS) full load migration to load the data warehouse tables into Amazon S3 every day. Overwrite the previous day's full-load copy every day.

Answer: C

NEW QUESTION # 59

A data engineer must build an extract, transform, and load (ETL) pipeline to process and load data from 10 source systems into 10 tables that are in an Amazon Redshift database. All the source systems generate .csv, JSON, or Apache Parquet files every 15 minutes. The source systems all deliver files into one Amazon S3 bucket. The file sizes range from 10 MB to 20 GB.

The ETL pipeline must function correctly despite changes to the data schema.

Which data pipeline solutions will meet these requirements? (Choose two.)

- A. Configure an AWS Lambda function to invoke an AWS Glue job when a file is loaded into the S3 bucket. Configure the AWS Glue job to read the files from the S3 bucket into an Apache Spark DataFrame. Configure the AWS Glue job to also put smaller partitions of the DataFrame into an Amazon Kinesis Data Firehose delivery stream. Configure the delivery stream to load data into the Amazon Redshift tables.
- B. Use an Amazon EventBridge rule to run an AWS Glue job every 15 minutes. Configure the AWS Glue job to process and load the data into the Amazon Redshift tables.
- C. Use an Amazon EventBridge rule to invoke an AWS Glue workflow job every 15 minutes. Configure the AWS Glue workflow to have an on-demand trigger that runs an AWS Glue crawler and then runs an AWS Glue job when the crawler finishes running successfully. Configure the AWS Glue job to process and load the data into the Amazon Redshift tables.
- D. Configure an AWS Lambda function to invoke an AWS Glue workflow when a file is loaded into the S3 bucket. Configure the AWS Glue workflow to have an on-demand trigger that runs an AWS Glue crawler and then runs an AWS Glue job when the crawler finishes running successfully. Configure the AWS Glue job to process and load the data into the Amazon Redshift tables.
- E. Configure an AWS Lambda function to invoke an AWS Glue crawler when a file is loaded into the S3 bucket. Configure an AWS Glue job to process and load the data into the Amazon Redshift tables. Create a second Lambda function to run the AWS Glue job. Create an Amazon EventBridge rule to invoke the second Lambda function when the AWS Glue crawler finishes running successfully.

Answer: C,D

NEW QUESTION # 60

A gaming company uses Amazon Kinesis Data Streams to collect clickstream data. The company uses Amazon Data Firehose delivery streams to store the data in JSON format in Amazon S3.

Data scientists at the company use Amazon Athena to query the most recent data to obtain business insights.

The company wants to reduce Athena costs but does not want to recreate the data pipeline.

Which solution will meet these requirements with the LEAST management effort?

- A. Change the Firehose output format to Apache Parquet. Provide a custom S3 object YYYYMMDD prefix expression and specify a large buffer size. For the existing data, create an AWS Glue extract, transform, and load (ETL) job. Configure the ETL job to combine small JSON files, convert the JSON files to large Parquet files, and add the YYYYMMDD prefix. Use the ALTER TABLE ADD PARTITION statement to reflect the partition on the existing Athena table.
- B. Create an Apache Spark job that combines JSON files and converts the JSON files to Apache Parquet files. Launch an Amazon EMR ephemeral cluster every day to run the Spark job to create new Parquet files in a different S3 location. Use the ALTER TABLE SET LOCATION statement to reflect the new S3 location on the existing Athena table.
- C. Create a Kinesis data stream as a delivery destination for Firehose. Use Amazon Managed Service for Apache Flink (previously known as Amazon Kinesis Data Analytics) to run Apache Flink on the Kinesis data stream. Use Flink to aggregate the data and save the data to Amazon S3 in Apache Parquet format with a custom S3 object YYYYMMDD prefix. Use the ALTER TABLE ADD PARTITION statement to reflect the partition on the existing Athena table.

- D. Integrate an AWS Lambda function with Firehose to convert source records to Apache Parquet and write them to Amazon S3. In parallel, run an AWS Glue extract, transform, and load (ETL) job to combine the JSON files and convert the JSON files to large Parquet files. Create a custom S3 object YYYYMMDD prefix. Use the ALTER TABLE ADD PARTITION statement to reflect the partition on the existing Athena table.

Answer: A

Explanation:

Changing the Firehose output format to Apache Parquet reduces Amazon Athena query costs because Parquet is a columnar storage format, which is much more efficient for queries compared to row-based formats like JSON. Parquet helps reduce the amount of data scanned by Athena, thus lowering costs.

Configuring Firehose to use a large buffer size ensures fewer, larger files, which also improves query performance in Athena.

For the existing JSON data, using an AWS Glue ETL job to convert the files into Parquet and apply the necessary partitioning (e.g., YYYYMMDD) will help optimize future queries without needing to recreate the data pipeline. Using the ALTER TABLE ADD PARTITION statement in Athena will allow the existing table to reflect the new partitions.

This solution minimizes management effort by reconfiguring Amazon Kinesis Data Firehose and using AWS Glue for handling existing data, without requiring significant changes to the existing pipeline.

Running an Apache Spark job on Amazon EMR requires more operational management, as you would need to launch and manage ephemeral EMR clusters daily, which increases complexity.

Amazon Managed Service for Apache Flink introduces additional components (Kinesis Data Streams and Flink), which adds complexity and management overhead compared to using the existing Firehose and simply changing its output format.

Using Lambda for real-time transformation and combining that with a Glue ETL job introduces unnecessary complexity. Firehose already supports direct output to Parquet format, making the Lambda function redundant.

NEW QUESTION # 61

A company runs a multi-tenant Amazon EMR cluster on Amazon EC2 instances. Multiple teams perform interactive query analyses and data transformations on the data in the EMR cluster. The teams can access the cluster only through EMR Studio workspaces and EMR steps.

The teams need to use EMR steps to run Apache Spark jobs to fetch data from an Amazon DynamoDB table. The DynamoDB table contains confidential data that must be accessible to only one specific team. The company needs to ensure that only the appropriate team can access the confidential data in the EMR cluster. Which solution will meet these requirements?

- A. Set up a DynamoDB resource-based policy.
- B. Set up AWS Lake Formation permissions.
- **C. Set up runtime roles for EMR steps.**
- D. Set up IAM roles for EMR File System (EMRFS) requests.

Answer: C

Explanation:

Runtime roles for EMR steps allow each submitted step (such as a Spark job) to assume a distinct IAM role at execution time. By granting DynamoDB read permissions only to the role used by the authorized team's steps, the confidential table data becomes accessible only to that team's Spark jobs, while other teams' steps run with roles that lack access.

NEW QUESTION # 62

.....

The modern job market is becoming more competitive with every passing moment. You have to be ready for it and learn in-demand skills with the SnowPro Advanced: Data Engineer Certification Exam Exam DEA-C01 certification exam. If you are not doing this you are going to end up in a normal company with low pay. Be smart in your decision and get registered for the SnowPro Advanced: Data Engineer Certification Exam DEA-C01 certification exam and put all your efforts, commitment and dedication to crack the SnowPro Advanced: Data Engineer Certification Exam DEA-C01 exam. Once you pass the SnowPro Advanced: Data Engineer Certification Exam DEA-C01 certification exam you will get personal and professional benefits throughout your career. Do you have the plan to accept this challenge and enroll in the DEA-C01 Certification Exam? Looking for a simple, quick, and smart way to pass the SnowPro Advanced: Data Engineer Certification Exam DEA-C01 exam? If your answer is yes then you do not need to get worried about it. Just visit Exams-boost and explore the top features of Snowflake DEA-C01 PDF Questions and practice tests. The Exams-boost is quite confident that you will crack the DEA-C01 exam shortly.

DEA-C01 Test Dumps Demo: <https://www.exams-boost.com/DEA-C01-valid-materials.html>

