

Reliable DEA-C01 Dumps Book, DEA-C01 Valid Test Question

AWS Data Engineer Associate (DEA-C01) Exam Dumps 2024
 AWS Data Engineer Associate (DEA-C01) Practice Tests 2024. Contains 390+ exam questions to pass the exam in first attempt.
 SkillCertPro offers real-exam questions for practice for all major IT certifications.

- For a full set of 390+ questions, Go to <https://skillcertpro.com/product/aws-data-engineer-associate-sca-c01-exam-questions>
- SkillCertPro offers detailed explanations to each question which helps to understand the concepts better.
- It is recommended to score above 85% in SkillCertPro exams before attempting a real exam.
- SkillCertPro updates exam questions every 2 weeks.
- You will get life time access and life time free updates
- SkillCertPro assures 100% pass guarantee in first attempt.

Below are the free 10 sample questions.

Question 1:

A data engineer working for an analytics company is working on a consumer to a Kinesis Data Streams application. They have written the consumer using Kinesis Client Library (KCL), however, currently they are receiving an ExpiredIteratorException when reading records from Kinesis Data Streams. What would you recommend to the engineer to solve their issue?

- A. Change the capacity mode of the Kinesis Data Stream to on-demand.
- B. Increase WCU in DynamoDB checkpointing table.
- C. Increase the amount of shards in Kinesis Data Streams.
- D. Increase RCU in DynamoDB checkpointing table.

Answer: B

P.S. Free 2026 Snowflake DEA-C01 dumps are available on Google Drive shared by DumpsMaterials:
https://drive.google.com/open?id=1WtAA2O8zkYwTJb_0Po7aPT8HlBk9TxbR

DumpsMaterials is an invisible assent that can give your advantage and get better life higher than your current situation and help you stand out among the average with the best and most accurate DEA-C01 study braindumps. For the great merit of our DEA-C01 Exam Guide is too many to count. Our experts have been dedicated in this area for more than ten years on compiling the content of our DEA-C01 training guide and keeping updating it to the latest.

Snowflake DEA-C01 Exam Syllabus Topics:

Topic	Details
Topic 1	<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 2	<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 3	<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 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"> • 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.

>> **Reliable DEA-C01 Dumps Book** <<

Free download Snowflake certification DEA-C01 exam practice questions and answers

It is well acknowledged that people who have a chance to participate in the simulation for the real DEA-C01 exam, they must have a fantastic advantage over other people to get good grade in the DEA-C01 exam. Now, it is so lucky for you to meet this opportunity once in a blue. We offer you the simulation test with the Software version of our DEA-C01 Preparation dumps in order to let you be familiar with the environment of test as soon as possible.

Snowflake SnowPro Advanced: Data Engineer Certification Exam Sample Questions (Q44-Q49):

NEW QUESTION # 44

A data engineer is configuring an AWS Glue Apache Spark extract, transform, and load (ETL) job. The job contains a sort-merge join of two large and equally sized DataFrames.

The job is failing with the following error: No space left on device.

Which solution will resolve the error?

- **A. Use the AWS Glue Spark shuffle manager.**
- B. Convert the sort-merge join in the job to be a broadcast join.
- C. Deploy an Amazon Elastic Block Store (Amazon EBS) volume for the job to use.
- D. Convert the DataFrames to DynamicFrames, and perform a DynamicFrame join in the job.

Answer: A

Explanation:

The error "No space left on device" during a sort-merge join typically results from insufficient disk space during the shuffle phase of a Spark job. Using the AWS Glue Spark shuffle manager (optimized shuffle manager) helps reduce disk usage and improves shuffle performance by optimizing how intermediate data is written and read. This directly addresses the issue without requiring job logic changes or manual resource management.

NEW QUESTION # 45

A company extracts approximately 1 TB of data every day from data sources such as SAP HANA, Microsoft SQL Server, MongoDB, Apache Kafka, and Amazon DynamoDB. Some of the data sources have undefined data schemas or data schemas that change.

A data engineer must implement a solution that can detect the schema for these data sources.

The solution must extract, transform, and load the data to an Amazon S3 bucket. The company has a service level agreement (SLA) to load the data into the S3 bucket within 15 minutes of data creation.

Which solution will meet these requirements with the LEAST operational overhead?

- A. Create a stored procedure in Amazon Redshift to detect the schema and to extract, transform, and load the data into a Redshift Spectrum table. Access the table from Amazon S3.
- B. Use Amazon EMR to detect the schema and to extract, transform, and load the data into the S3 bucket. Create a pipeline in Apache Spark.
- C. Use AWS Glue to detect the schema and to extract, transform, and load the data into the S3 bucket. Create a pipeline in Apache Spark.
- D. Create a PySpark program in AWS Lambda to extract, transform, and load the data into the S3 bucket.

Answer: C

NEW QUESTION # 46

A data engineer maintains a materialized view that is based on an Amazon Redshift database.

The view has a column named load_date that stores the date when each row was loaded.

The data engineer needs to reclaim database storage space by deleting all the rows from the materialized view.

Which command will reclaim the MOST database storage space?

- A. VACUUM table_name where load_date<=current_date materializedview
- B. DELETE FROM materialized_view_name where 1=1
- C. TRUNCATE materialized_view_name
- D. DELETE FROM materialized_view_name where load_date<=current_date

Answer: C

NEW QUESTION # 47

A Data Engineer wants to centralize grant management to maximize security. A user needs ownership on a table m a new schema

However, this user should not have the ability to make grant decisions What is the correct way to do this?

- A. Revoke grant decisions from the user on the schema.
- B. Grant ownership to the user on the table
- C. Revoke grant decisions from the user on the table
- D. Add the with managed access parameter on the schema

Answer: D

Explanation:

Explanation

The with managed access parameter on the schema enables the schema owner to control the grant and revoke privileges on the objects within the schema. This way, the user who owns the table cannot make grant decisions, but only the schema owner can. This is the best way to centralize grant management and maximize security.

NEW QUESTION # 48

A Data Engineer is working on a continuous data pipeline which receives data from Amazon Kinesis Firehose and loads the data into a staging table which will later be used in the data transformation process The average file size is 300-500 MB.

The Engineer needs to ensure that Snowpipe is performant while minimizing costs.

How can this be achieved?

- A. Split the files before loading them and set the SIZE_LIMIT option to 250 MB.
- B. Increase the size of the virtual warehouse used by Snowpipe.
- C. Decrease the buffer size to trigger delivery of files sized between 100 to 250 MB in Kinesis Firehose
- D. Change the file compression size and increase the frequency of the Snowpipe loads

Answer: A

Explanation:

Explanation

This option is the best way to ensure that Snowpipe is performant while minimizing costs. By splitting the files before loading them, the Data Engineer can reduce the size of each file and increase the parallelism of loading.

By setting the SIZE_LIMIT option to 250 MB, the Data Engineer can specify the maximum file size that can be loaded

