

2026 Answers DEA-C01 Free: SnowPro Advanced: Data Engineer Certification Exam - Latest Snowflake Reliable DEA-C01 Exam Review



What's more, part of that VCE4Dumps DEA-C01 dumps now are free: <https://drive.google.com/open?id=18FtV0FEz0XNEDCJsp8c9ZFmkpGETGLrm>

In this age of anxiety, everyone seems to have great pressure. If you are better, you will have a more relaxed life. DEA-C01 guide materials allow you to increase the efficiency of your work. You can spend more time doing other things. Our study materials allow you to pass the DEA-C01 exam in the shortest possible time. You will stand at a higher starting point than others. Why are DEA-C01 Practice Questions worth your choice? I hope you can spend a little time free downloading our demo of our DEA-C01 exam questions, then you will know the advantages of our DEA-C01 study materials!

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 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 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"> 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.

Reliable DEA-C01 Exam Review & DEA-C01 Reliable Test Preparation

The DEA-C01 practice exam we offered is designed with the real questions that will help you in enhancing your knowledge about the DEA-C01 certification exam. Our online test engine will improve your ability to solve the difficulty of DEA-C01 Real Questions and get used to the atmosphere of the formal test. Our experts created the valid DEA-C01 study guide for most of candidates to help them get good result with less time and money.

Snowflake SnowPro Advanced: Data Engineer Certification Exam Sample Questions (Q160-Q165):

NEW QUESTION # 160

The following is returned from `SYSTEMCLUSTERING_INFORMATION ()` for a table named `orders` with a date column named `O_ORDERDATE`:

What does the `total_constant_partition_count` value indicate about this table?

- A. The data in `O_ORDERDATE` does not change very often as there are 493 micro-partitions containing rows where that column has not been modified since the row was created
- B. The data in `O_ORDERDATE` has a very low cardinality as there are 493 micro-partitions where there is only a single distinct value in that column for all rows in the micro-partition
- C. The table is clustered very well on `O_ORDERDATE`, as there are 493 micro-partitions that could not be significantly improved by reclustering
- **D. The table is not clustered well on `O_ORDERDATE`, as there are 493 micro-partitions where the range of values in that column overlap with every other micro-partition in the table.**

Answer: D

Explanation:

Explanation

The `total_constant_partition_count` value indicates the number of micro-partitions where the clustering key column has a constant value across all rows in the micro-partition. However, this does not necessarily mean that the table is clustered well on that column, as there could be other micro-partitions where the range of values in that column overlap with each other. This is the case for the `orders` table, as the clustering depth is 1, which means that every micro-partition overlaps with every other micro-partition on `O_ORDERDATE`. This indicates that the table is not clustered well on `O_ORDERDATE` and could benefit from reclustering.

NEW QUESTION # 161

A data engineer uses Amazon Kinesis Data Streams to ingest and process records that contain user behavior data from an application every day.

The data engineer notices that the data stream is experiencing throttling because hot shards receive much more data than other shards in the data stream.

How should the data engineer resolve the throttling issue?

- A. Use a random partition key to distribute the ingested records.
- B. Limit the number of records that are sent each second by the producer to match the capacity of the stream.
- **C. Increase the number of shards in the data stream. Distribute the records across the shards.**
- D. Decrease the size of the records that the producer sends to match the capacity of the stream.

Answer: C

Explanation:

Amazon Kinesis Data Streams distribute data across multiple shards, with each shard having its own capacity for read and write operations. Throttling occurs when one or more shards, referred to as "hot shards," receive significantly more data than they can handle. To resolve this, increasing the number of shards in the data stream and redistributing the records across the shards is the appropriate solution. This approach ensures that the workload is spread more evenly, thereby preventing throttling on individual shards.

NEW QUESTION # 162

A gaming company uses a NoSQL database to store customer information. The company is planning to migrate to AWS. The company needs a fully managed AWS solution that will handle high online transaction processing (OLTP) workload, provide single-digit millisecond performance, and provide high availability around the world. Which solution will meet these requirements with the LEAST operational overhead?

- A. Amazon DocumentDB (with MongoDB compatibility)
- **B. Amazon DynamoDB**
- C. Amazon Timestream
- D. Amazon Keyspaces (for Apache Cassandra)

Answer: B

NEW QUESTION # 163

A company uses AWS Glue jobs to implement several data pipelines. The pipelines are critical to the company. The company needs to implement a monitoring mechanism that will alert stakeholders if the pipelines fail. Which solution will meet these requirements with the LEAST operational overhead?

- A. Create an Amazon EventBridge rule to match AWS Glue job failure events. Configure the rule to target an AWS Lambda function to process events. Configure the function to send notifications to an Amazon Simple Notification Service (Amazon SNS) topic.
- B. Configure an Amazon CloudWatch Logs log group for the AWS Glue jobs. Create an Amazon EventBridge rule to match new log creation events in the log group. Configure the rule to send notifications to an Amazon Simple Notification Service (Amazon SNS) topic.
- C. Configure an Amazon CloudWatch Logs log group for the AWS Glue jobs. Create an Amazon EventBridge rule to match new log creation events in the log group. Configure the rule to target an AWS Lambda function that reads the logs and sends notifications to an Amazon Simple Notification Service (Amazon SNS) topic if AWS Glue job failure logs are present.
- **D. Create an Amazon EventBridge rule to match AWS Glue job failure events. Define an Amazon CloudWatch metric based on the EventBridge rule. Set up a CloudWatch alarm based on the metric to send notifications to an Amazon Simple Notification Service (Amazon SNS) topic.**

Answer: D

Explanation:

By creating an EventBridge rule that matches Glue job state-change events where the state is FAILED, you can emit a custom CloudWatch metric without writing any code. Then you define a CloudWatch alarm on that metric to notify an SNS topic upon breach. This approach requires no Lambda functions or log parsing (just the rule, the metric, and the alarm) minimizing operational overhead.

NEW QUESTION # 164

To support Time Travel, Which of the following SQL extensions/parameters/commands have been implemented?

- A. ONSET (time difference in seconds from the present time)
- **B. UNDROP command for tables, schemas, and databases.**
- **C. OFFSET (time difference in seconds from the present time)**
- D. STATEMENT_ID (identifier for statement, e.g. query ID)
- **E. STATEMENT (identifier for statement, e.g. query ID)**
- **F. AT | BEFORE clause which can be specified in the CREATE ... CLONE commands.**

Answer: B,C,E,F

NEW QUESTION # 165

.....

VCE4Dumps SnowPro Advanced: Data Engineer Certification Exam (DEA-C01) exam dumps save your study and preparation time. Our experts have added hundreds of SnowPro Advanced: Data Engineer Certification Exam (DEA-C01) questions similar to the real exam. You can prepare for the SnowPro Advanced: Data Engineer Certification Exam (DEA-C01) exam dumps during your job. You don't need to visit the market or any store because VCE4Dumps Snowflake DEA-C01 exam questions are easily

