Data-Engineer-Associate Dumps Download, Data-Engineer-Associate Actual Exam Dumps

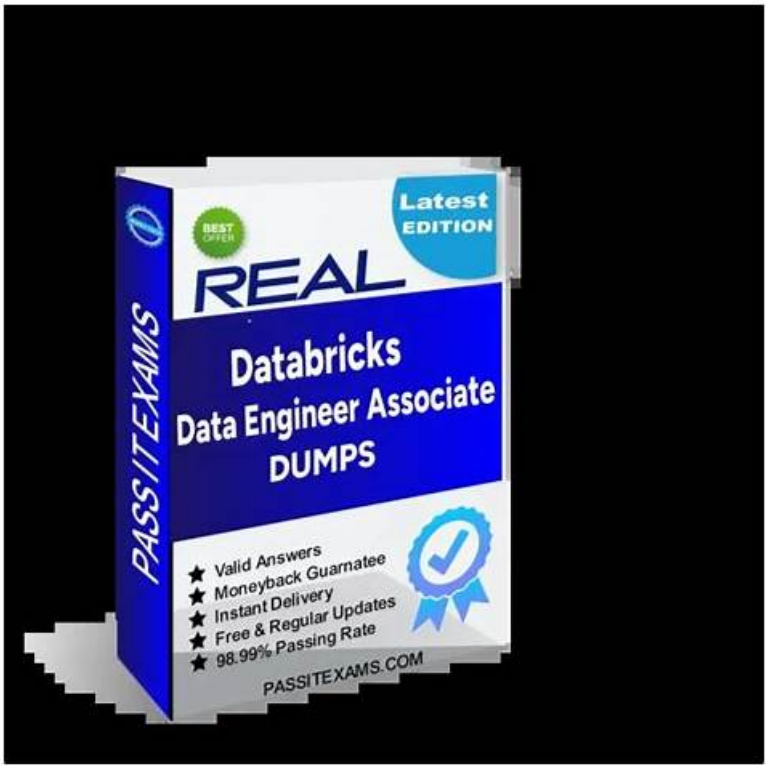

P.S. Free 2026 Amazon Data-Engineer-Associate dumps are available on Google Drive shared by TestkingPass: https://drive.google.com/open?id=1Z_BTXi8W1IH2yrJwgY3thi9_c-Y9DENi

You can also be part of successful Data-Engineer-Associate exam candidates. To do this you just need to enroll in Data-Engineer-Associate exam and strive hard to get success in the Amazon Data-Engineer-Associate certification exam. In this journey, the Data-Engineer-Associate Dumps can help you perfectly. The AWS Certified Data Engineer - Associate (DEA-C01) Data-Engineer-Associate Exam Questions are the real, updated AWS Certified Data Engineer - Associate (DEA-C01) Data-Engineer-Associate exam practice Test that will assist you in Amazon Data-Engineer-Associate exam preparation and enable you to pass the final Amazon Data-Engineer-Associate exam easily.

The study system of our company will provide all customers with the best study materials. If you buy the Data-Engineer-Associate latest questions of our company, you will have the right to enjoy all the Data-Engineer-Associate certification training materials from our company. More importantly, there are a lot of experts in our company; the first duty of these experts is to update the study system of our company day and night for all customers. By updating the study system of the Data-Engineer-Associate Training Materials, we can guarantee that our company can provide the newest information about the Data-Engineer-Associate exam for all people.

>> Data-Engineer-Associate Dumps Download <<

Data-Engineer-Associate Actual Exam Dumps & Data-Engineer-Associate Valid Test Braindumps

The top features of TestkingPass Data-Engineer-Associate exam questions are the availability of Amazon certification exam in three different formats, real, valid, and updated Data-Engineer-Associate exam questions, subject matter experts verified Data-Engineer-Associate Exam Questions, free demo download facility, 1 year updated Data-Engineer-Associate exam questions download facility, affordable price and 100 percent Amazon Data-Engineer-Associate exam passing money back guarantee.

Amazon AWS Certified Data Engineer - Associate (DEA-C01) Sample

Questions (Q100-Q105):

NEW QUESTION # 100
An airline company is collecting metrics about flight activities for analytics. The company is conducting a proof of concept (POC) test to show how analytics can provide insights that the company can use to increase on-time departures.
The POC test uses objects in Amazon S3 that contain the metrics in .csv format. The POC test uses Amazon Athena to query the data. The data is partitioned in the S3 bucket by date.
As the amount of data increases, the company wants to optimize the storage solution to improve query performance.
Which combination of solutions will meet these requirements? (Choose two.)

- A. Use an S3 bucket that is in the same account that uses Athena to query the data.
- B. Add a randomized string to the beginning of the keys in Amazon S3 to get more throughput across partitions.
- C. Preprocess the .csv data to Apache Parquet format by fetching only the data blocks that are needed for predicates.
- D. Use an S3 bucket that is in the same AWS Region where the company runs Athena queries.
- E. Preprocess the .csv data to JSON format by fetching only the document keys that the query requires.

Answer: C,D

Explanation:
Using an S3 bucket that is in the same AWS Region where the company runs Athena queries can improve query performance by reducing data transfer latency and costs. Preprocessing the .csv data to Apache Parquet format can also improve query performance by enabling columnar storage, compression, and partitioning, which can reduce the amount of data scanned and fetched by the query. These solutions can optimize the storage solution for the POC test without requiring much effort or changes to the existing data pipeline. The other solutions are not optimal or relevant for this requirement. Adding a randomized string to the beginning of the keys in Amazon S3 can improve the throughput across partitions, but it can also make the data harder to query and manage. Using an S3 bucket that is in the same account that uses Athena to query the data does not have any significant impact on query performance, as long as the proper permissions are granted.
Preprocessing the .csv data to JSON format does not offer any benefits over the .csv format, as both are row- based and verbose formats that require more data scanning and fetching than columnar formats like Parquet.
References:
* Best Practices When Using Athena with AWS Glue
* Optimizing Amazon S3 Performance
* AWS Certified Data Engineer - Associate DEA-C01 Complete Study Guide

NEW QUESTION # 101
A company is developing an application that runs on Amazon EC2 instances. Currently, the data that the application generates is temporary. However, the company needs to persist the data, even if the EC2 instances are terminated.
A data engineer must launch new EC2 instances from an Amazon Machine Image (AMI) and configure the instances to preserve the data.
Which solution will meet this requirement?

- A. Launch new EC2 instances by using an AMI that is backed by an EC2 instance store volume that contains the application data. Apply the default settings to the EC2 instances.
- B. Launch new EC2 instances by using an AMI that is backed by a root Amazon Elastic Block Store (Amazon EBS) volume that contains the application data. Apply the default settings to the EC2 instances.
- C. Launch new EC2 instances by using an AMI that is backed by an EC2 instance store volume. Attach an Amazon Elastic Block Store (Amazon EBS) volume to contain the application data. Apply the default settings to the EC2 instances.
- D. Launch new EC2 instances by using an AMI that is backed by an Amazon Elastic Block Store (Amazon EBS) volume. Attach an additional EC2 instance store volume to contain the application data. Apply the default settings to the EC2 instances.

Answer: C

Explanation:
Amazon EC2 instances can use two types of storage volumes: instance store volumes and Amazon EBS volumes. Instance store volumes are ephemeral, meaning they are only attached to the instance for the duration of its life cycle. If the instance is stopped, terminated, or fails, the data on the instance store volume is lost. Amazon EBS volumes are persistent, meaning they can be detached from the instance and attached to another instance, and the data on the volume is preserved. To meet the requirement of persisting the data even if the EC2 instances are terminated, the data engineer must use Amazon EBS volumes to store the application data. The solution is to launch new EC2 instances by using an AMI that is backed by an EC2 instance store volume, which is the default

option for most AMIs. Then, the data engineer must attach an Amazon EBS volume to each instance and configure the application to write the data to the EBS volume. This way, the data will be saved on the EBS volume and can be accessed by another instance if needed. The data engineer can apply the default settings to the EC2 instances, as there is no need to modify the instance type, security group, or IAM role for this solution. The other options are either not feasible or not optimal. Launching new EC2 instances by using an AMI that is backed by an EC2 instance store volume that contains the application data (option A) or by using an AMI that is backed by a root Amazon EBS volume that contains the application data (option B) would not work, as the data on the AMI would be outdated and overwritten by the new instances. Attaching an additional EC2 instance store volume to contain the application data (option D) would not work, as the data on the instance store volume would be lost if the instance is terminated.
References:
* Amazon EC2 Instance Store
* Amazon EBS Volumes
* AWS Certified Data Engineer - Associate DEA-C01 Complete Study Guide, Chapter 2: Data Store Management, Section 2.1: Amazon EC2

NEW QUESTION # 102
A company has a data warehouse that contains a table that is named Sales. The company stores the table in Amazon Redshift The table includes a column that is named city_name. The company wants to query the table to find all rows that have a city_name that starts with "San" or "El." Which SQL query will meet this requirement?

- A. Select * from Sales where city_name - '$(San|El)";
- B. Select * from Sales where city_name - '$(San&El)";
- C. Select * from Sales where city_name -,