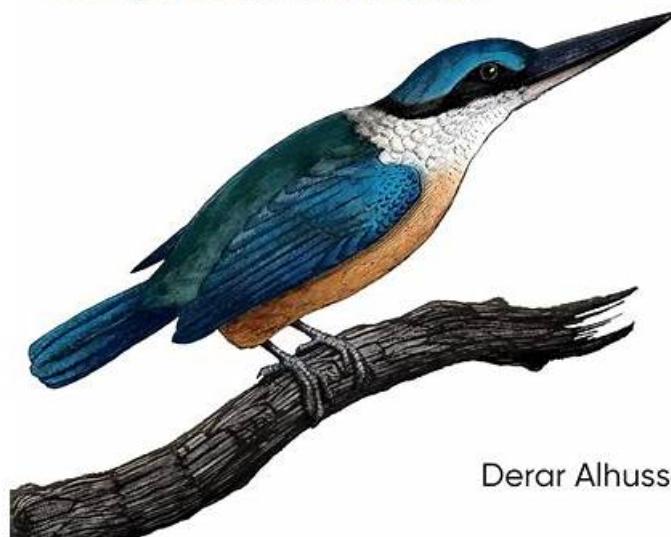


Test Amazon Data-Engineer-Associate Engine - Best Data-Engineer-Associate Study Material

O'REILLY®

Databricks Certified Data Engineer Associate Study Guide

In-Depth Guidance and Practice



Derar Alhussein

DOWNLOAD the newest Dumps4PDF Data-Engineer-Associate PDF dumps from Cloud Storage for free:
<https://drive.google.com/open?id=1DVmmuOVBlr1ejcwzXBSRa-CVVp2ILF-R>

It was a Xi'an coach byword that if you give up, the game is over at the same time. The game likes this, so is the exam. Not having enough time to prepare for their exam, many people give up taking IT certification exam. However, with the help of the best training materials, you can completely pass Amazon Data-Engineer-Associate test in a short period of time. Don't you believe in it? Dumps4PDF real questions and answers are the materials that it can help you get high marks and pass the certification exam. Please try it.

May be you will meet some difficult or problems when you prepare for your Data-Engineer-Associate exam, you even want to give it up. That is why I suggest that you must try our study materials. Because Data-Engineer-Associate guide torrent can help you to solve all the problems encountered in the learning process, Data-Engineer-Associate study tool will provide you with very flexible learning time so that you can easily pass the exam. Even if you fail to pass the exam, as long as you are willing to continue to use our Data-Engineer-Associate Study Tool, we will still provide you with the benefits of free updates within a year.

[**>> Test Amazon Data-Engineer-Associate Engine <<**](#)

Best Data-Engineer-Associate Study Material & Data-Engineer-Associate Test Prep

If you want to get a comprehensive idea about our real Data-Engineer-Associate study materials. It is convenient for you to

download the free demo, all you need to do is just to find the “Download for free” item, and you will find there are three kinds of versions of Data-Engineer-Associate learning guide for you to choose from namely, PDF Version Demo, PC Test Engine and Online Test Engine, you can choose to download any one version of our Data-Engineer-Associate exam questions as you like.

Amazon AWS Certified Data Engineer - Associate (DEA-C01) Sample Questions (Q127-Q132):

NEW QUESTION # 127

A transportation company wants to track vehicle movements by capturing geolocation records. The records are 10 bytes in size. The company receives up to 10,000 records every second. Data transmission delays of a few minutes are acceptable because of unreliable network conditions.

The transportation company wants to use Amazon Kinesis Data Streams to ingest the geolocation data. The company needs a reliable mechanism to send data to Kinesis Data Streams. The company needs to maximize the throughput efficiency of the Kinesis shards.

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

- A. Kinesis Producer Library (KPL)
- B. Kinesis SDK
- C. Kinesis Agent
- D. Amazon Data Firehose

Answer: A

Explanation:

Problem Analysis:

The company ingests geolocation records (10 bytes each) at 10,000 records per second into Kinesis Data Streams.

Data transmission delays are acceptable, but the solution must maximize throughput efficiency.

Key Considerations:

The Kinesis Producer Library (KPL) batches records and uses aggregation to optimize shard throughput.

Efficiently handles high-throughput scenarios with minimal operational overhead.

Solution Analysis:

Option A: Kinesis Agent

Designed for file-based ingestion; not optimized for geolocation records.

Option B: KPL

Aggregates records into larger payloads, significantly improving shard throughput.

Suitable for applications generating small, high-frequency records.

Option C: Kinesis Firehose

Firehose is for delivery to destinations like S3 or Redshift and is not optimized for direct ingestion to Kinesis Data Streams.

Option D: Kinesis SDK

The SDK lacks advanced features like aggregation, resulting in lower throughput efficiency.

Final Recommendation:

Use Kinesis Producer Library (KPL) for its built-in aggregation and batching capabilities.

Reference:

Kinesis Producer Library (KPL) Overview

NEW QUESTION # 128

A company stores a large dataset in an Amazon S3 bucket. A data engineer frequently runs complex queries on the dataset by using Amazon Athena. The data engineer needs to optimize query performance and optimize costs for queries that are run multiple times with the same parameters.

Which solution will meet these requirements?

- A. Configure query result reuse settings in the Athena workgroup.
- B. Use Amazon Redshift Spectrum to query the data in Amazon S3.
- C. Convert the dataset to JSON format before running Athena queries.
- D. Use Amazon EMR to pre-process the data before running Athena queries.

Answer: A

Explanation:

Comprehensive and Detailed Explanation (150-250 words)

Amazon Athena charges per amount of data scanned by each query. When the same query with identical parameters is executed multiple times, re-scanning the data repeatedly increases both cost and execution time.

Athena provides query result reuse, which allows previously computed query results to be reused when an identical query is rerun within a configurable time window.

By enabling query result reuse at the Athena workgroup level, Athena automatically returns cached results instead of re-reading data from Amazon S3. This significantly improves query performance while reducing data scan costs. This feature is specifically designed for workloads where analysts frequently rerun complex queries with the same parameters, which directly matches the scenario. Converting data to JSON would degrade performance because JSON is not a columnar format and increases scan size. Using Amazon EMR introduces unnecessary infrastructure management and cost for a query optimization problem that Athena already solves natively. Amazon Redshift Spectrum is intended for querying S3 data from within Redshift and does not address repeated-query optimization in Athena.

Therefore, configuring query result reuse in the Athena workgroup is the most cost-effective and operationally efficient solution.

NEW QUESTION # 129

A company has a frontend ReactJS website that uses Amazon API Gateway to invoke REST APIs. The APIs perform the functionality of the website. A data engineer needs to write a Python script that can be occasionally invoked through API Gateway. The code must return results to API Gateway.

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

- A. Create an AWS Lambda function. Ensure that the function is warm by scheduling an Amazon EventBridge rule to invoke the Lambda function every 5 minutes by using mock events.
- B. Deploy a custom Python script that can integrate with API Gateway on Amazon Elastic Kubernetes Service (Amazon EKS).
- **C. Create an AWS Lambda Python function with provisioned concurrency.**
- D. Deploy a custom Python script on an Amazon Elastic Container Service (Amazon ECS) cluster.

Answer: C

Explanation:

AWS Lambda is a serverless compute service that lets you run code without provisioning or managing servers. You can use Lambda to create functions that perform custom logic and integrate with other AWS services, such as API Gateway. Lambda automatically scales your application by running code in response to each trigger. You pay only for the compute time you consume1. Amazon ECS is a fully managed container orchestration service that allows you to run and scale containerized applications on AWS. You can use ECS to deploy, manage, and scale Docker containers using either Amazon EC2 instances or AWS Fargate, a serverless compute engine for containers2.

Amazon EKS is a fully managed Kubernetes service that allows you to run Kubernetes clusters on AWS without needing to install, operate, or maintain your own Kubernetes control plane. You can use EKS to deploy, manage, and scale containerized applications using Kubernetes on AWS3.

The solution that meets the requirements with the least operational overhead is to create an AWS Lambda Python function with provisioned concurrency. This solution has the following advantages:

It does not require you to provision, manage, or scale any servers or clusters, as Lambda handles all the infrastructure for you. This reduces the operational complexity and cost of running your code.

It allows you to write your Python script as a Lambda function and integrate it with API Gateway using a simple configuration. API Gateway can invoke your Lambda function synchronously or asynchronously, and return the results to the frontend website.

It ensures that your Lambda function is ready to respond to API requests without any cold start delays, by using provisioned concurrency. Provisioned concurrency is a feature that keeps your function initialized and hyper-ready to respond in double-digit milliseconds. You can specify the number of concurrent executions that you want to provision for your function.

Option A is incorrect because it requires you to deploy a custom Python script on an Amazon ECS cluster.

This solution has the following disadvantages:

It requires you to provision, manage, and scale your own ECS cluster, either using EC2 instances or Fargate.

This increases the operational complexity and cost of running your code.

It requires you to package your Python script as a Docker container image and store it in a container registry, such as Amazon ECR or Docker Hub. This adds an extra step to your deployment process.

It requires you to configure your ECS cluster to integrate with API Gateway, either using an Application Load Balancer or a Network Load Balancer. This adds another layer of complexity to your architecture.

Option C is incorrect because it requires you to deploy a custom Python script that can integrate with API Gateway on Amazon EKS. This solution has the following disadvantages:

It requires you to provision, manage, and scale your own EKS cluster, either using EC2 instances or Fargate.

This increases the operational complexity and cost of running your code.

It requires you to package your Python script as a Docker container image and store it in a container registry, such as Amazon ECR

or Docker Hub. This adds an extra step to your deployment process.

It requires you to configure your EKS cluster to integrate with API Gateway, either using an Application Load Balancer, a Network Load Balancer, or a service of type LoadBalancer. This adds another layer of complexity to your architecture.

Option D is incorrect because it requires you to create an AWS Lambda function and ensure that the function is warm by scheduling an Amazon EventBridge rule to invoke the Lambda function every 5 minutes by using mock events. This solution has the following disadvantages:

It does not guarantee that your Lambda function will always be warm, as Lambda may scale down your function if it does not receive any requests for a long period of time. This may cause cold start delays when your function is invoked by API Gateway.

It incurs unnecessary costs, as you pay for the compute time of your Lambda function every time it is invoked by the EventBridge rule, even if it does not perform any useful work1.

1: AWS Lambda - Features

2: Amazon Elastic Container Service - Features

3: Amazon Elastic Kubernetes Service - Features

[4]: Building API Gateway REST API with Lambda integration - Amazon API Gateway

[5]: Improving latency with Provisioned Concurrency - AWS Lambda

[6]: Integrating Amazon ECS with Amazon API Gateway - Amazon Elastic Container Service

[7]: Integrating Amazon EKS with Amazon API Gateway - Amazon Elastic Kubernetes Service

[8]: Managing concurrency for a Lambda function - AWS Lambda

NEW QUESTION # 130

A company uploads .csv files to an Amazon S3 bucket. The company's data platform team has set up an AWS Glue crawler to perform data discovery and to create the tables and schemas.

An AWS Glue job writes processed data from the tables to an Amazon Redshift database. The AWS Glue job handles column mapping and creates the Amazon Redshift tables in the Redshift database appropriately.

If the company reruns the AWS Glue job for any reason, duplicate records are introduced into the Amazon Redshift tables. The company needs a solution that will update the Redshift tables without duplicates.

Which solution will meet these requirements?

- A. Use Apache Spark's DataFrame dropDuplicates() API to eliminate duplicates. Write the data to the Redshift tables.
- B. Use the AWS Glue ResolveChoice built-in transform to select the value of the column from the most recent record.
- C. Modify the AWS Glue job to load the previously inserted data into a MySQL database. Perform an upsert operation in the MySQL database. Copy the results to the Amazon Redshift tables.
- D. Modify the AWS Glue job to copy the rows into a staging Redshift table. Add SQL commands to update the existing rows with new values from the staging Redshift table.

Answer: B

NEW QUESTION # 131

A company's data engineer needs to optimize the performance of table SQL queries. The company stores data in an Amazon Redshift cluster. The data engineer cannot increase the size of the cluster because of budget constraints.

The company stores the data in multiple tables and loads the data by using the EVEN distribution style. Some tables are hundreds of gigabytes in size. Other tables are less than 10 MB in size.

Which solution will meet these requirements?

- A. Keep using the EVEN distribution style for all tables. Specify primary and foreign keys for all tables.
- B. Use the ALL distribution style for large tables. Specify primary and foreign keys for all tables.
- C. Specify a combination of distribution, sort, and partition keys for all tables.
- D. Use the ALL distribution style for rarely updated small tables. Specify primary and foreign keys for all tables.

Answer: D

Explanation:

This solution meets the requirements of optimizing the performance of table SQL queries without increasing the size of the cluster. By using the ALL distribution style for rarely updated small tables, you can ensure that the entire table is copied to every node in the cluster, which eliminates the need for data redistribution during joins. This can improve query performance significantly, especially for frequently joined dimension tables.

However, using the ALL distribution style also increases the storage space and the load time, so it is only suitable for small tables that are not updated frequently or extensively. By specifying primary and foreign keys for all tables, you can help the query optimizer to generate better query plans and avoid unnecessary scans or joins. You can also use the AUTO distribution style to let Amazon

Redshift choose the optimal distribution style based on the table size and the query patterns. References:

Choose the best distribution style

Distribution styles

Working with data distribution styles

NEW QUESTION # 132

.....

New developments in the tech sector always bring new job opportunities. These new jobs have to be filled with the Data-Engineer-Associate certification holders. So to fill the space, you need to pass the Data-Engineer-Associate Exam. Earning the Data-Engineer-Associate certification helps you clear the obstacles you face while working in the Amazon field.

Best Data-Engineer-Associate Study Material: <https://www.dumps4pdf.com/Data-Engineer-Associate-valid-braindumps.html>

Amazon Test Data-Engineer-Associate Engine More importantly, the practices have proven that the study materials from our company have helped a lot of people achieve their goal and get the related certification. We try our best to provide the most efficient and intuitive Data-Engineer-Associate learning materials to the learners and help them learn efficiently. On the whole, the pass rate of our customers after using Data-Engineer-Associate test dumps in the course of the preparation for the Amazon exams can reach as high as 98% to 99%, which is far ahead of others in the same field.

In this case, having a stripe could improve destage Data-Engineer-Associate performance. Notebook computers are amazing in some ways, but troublesome in others. More importantly, the practices have proven that the study materials Test Data-Engineer-Associate Engine from our company have helped a lot of people achieve their goal and get the related certification.

Choose The Right Amazon Data-Engineer-Associate and Get Certified Today!

We try our best to provide the most efficient and intuitive Data-Engineer-Associate Learning Materials to the learners and help them learn efficiently. On the whole, the pass rate of our customers after using Data-Engineer-Associate test dumps in the course of the preparation for the Amazon exams can reach as high as 98% to 99%, which is far ahead of others in the same field.

Train with real exam questions Daily updates at no extra charge, we will always put the interests of customers in the first place, so Data-Engineer-Associate test guide ensure that your information will not be leaked to any third party.

- Simulated Data-Engineer-Associate Test □ Data-Engineer-Associate PDF Cram Exam □ Data-Engineer-Associate Reliable Test Labs □ Search for { Data-Engineer-Associate } on ➡ www.examcollectionpass.com □ immediately to obtain a free download □ Reliable Data-Engineer-Associate Learning Materials
- 2026 Amazon Data-Engineer-Associate: AWS Certified Data Engineer - Associate (DEA-C01) Latest Test Engine □ Enter ➡ www.pdfvce.com □ and search for (Data-Engineer-Associate) to download for free □ Data-Engineer-Associate Learning Engine
- Data-Engineer-Associate exam preparation, real Amazon test dumps for AWS Certified Data Engineer - Associate (DEA-C01) □ Simply search for ✓ Data-Engineer-Associate □ ✓ □ for free download on ⇒ www.prep4sures.top ⇌ □ Reliable Data-Engineer-Associate Learning Materials
- Data-Engineer-Associate exam preparation, real Amazon test dumps for AWS Certified Data Engineer - Associate (DEA-C01) □ Search for □ Data-Engineer-Associate □ and obtain a free download on “ www.pdfvce.com ” □ Reliable Data-Engineer-Associate Learning Materials
- Valid Data-Engineer-Associate Exam Topics □ Data-Engineer-Associate Reliable Test Labs □ Data-Engineer-Associate Positive Feedback □ Open website [www.prepawayte.com] and search for ➡ Data-Engineer-Associate □ □ □ for free download □ Data-Engineer-Associate Positive Feedback
- 2026 Amazon Data-Engineer-Associate: AWS Certified Data Engineer - Associate (DEA-C01) Latest Test Engine □ Download ⇒ Data-Engineer-Associate ⇌ for free by simply searching on ➡ www.pdfvce.com □ □ Data-Engineer-Associate Valid Test Guide
- Data-Engineer-Associate Related Certifications □ Data-Engineer-Associate Questions Answers □ Data-Engineer-Associate Dumps Vce □ The page for free download of [Data-Engineer-Associate] on ⇒ www.verifieddumps.com ⇌ will open immediately □ Data-Engineer-Associate Exams Torrent
- Data-Engineer-Associate New Braindumps Free ☑ New Data-Engineer-Associate Dumps Pdf ☑ New Data-Engineer-Associate Dumps Pdf □ Open website ▷ www.pdfvce.com ▷ and search for □ Data-Engineer-Associate □ for free download □ Exam Data-Engineer-Associate Dump
- Test Data-Engineer-Associate Engine | 100% Free Authoritative Best AWS Certified Data Engineer - Associate (DEA-C01) Study Material □ Easily obtain free download of { Data-Engineer-Associate } by searching on 《 www.validtorrent.com 》 □ Data-Engineer-Associate Related Certifications

What's more, part of that Dumps4PDF Data-Engineer-Associate dumps now are free: <https://drive.google.com/open?id=1DVmmuOVBlr1ejcwzXBSRa-CVVP2ILF-R>