

Amazon - Data-Engineer-Associate - AWS Certified Data Engineer - Associate (DEA-C01) High Hit-Rate Passing Score Feedback



AWS Certified Data Engineer - Associate

試験結果のお知らせ

受験者氏名: [REDACTED]	受験日: 2024/04, [REDACTED]
受験者 ID: [REDACTED]	登録番号: [REDACTED]
受験者スコア: 1000	結果: 合格

おめでとうございます。AWS Certified Data Engineer - Associate に無事合格し、AWS 認定されました。

AWS Certified Data Engineer - Associate (DEA-C01) のスケールリングスコアは 100~1,000 点です。合格に必要なスケールリングスコアは 720 点です。

P.S. Free & New Data-Engineer-Associate dumps are available on Google Drive shared by Pass4cram: https://drive.google.com/open?id=1fcUPRUS7JRAC5q_1NhP5BPah3vh3L4Hy

With "reliable credit" as the soul of our Data-Engineer-Associate study tool, "utmost service consciousness" as the management philosophy, we endeavor to provide customers with high quality service. Our customer service staff, who are willing to be your little helper and answer your any questions about our Data-Engineer-Associate qualification test, fully implement the service principle of customer-oriented service on our Data-Engineer-Associate Exam Questions. Any puzzle about our Data-Engineer-Associate test torrent will receive timely and effective response, just leave a message on our official website or send us an e-mail for our Data-Engineer-Associate study guide.

When you have adequately prepared for the AWS Certified Data Engineer - Associate (DEA-C01) (Data-Engineer-Associate) questions, only then you become capable of passing the Amazon exam. There is no purpose in attempting the Amazon Data-Engineer-Associate certification exam if you have not prepared with Pass4cram's Free Amazon Data-Engineer-Associate PDF Questions. It's time to get serious if you want to validate your abilities and earn the Amazon Data-Engineer-Associate Certification. If you hope to pass the AWS Certified Data Engineer - Associate (DEA-C01) exam on your first attempt, you must be studied with real Data-Engineer-Associate exam questions verified by Amazon Data-Engineer-Associate.

>> Passing Data-Engineer-Associate Score Feedback <<

Online Data-Engineer-Associate Lab Simulation, Data-Engineer-Associate Valid Test Bootcamp

In order to serve you better, we have a complete system for you if you choose us. We have free demo for Data-Engineer-Associate training materials for you to have a try. If you have decided to buy Data-Engineer-Associate exam dumps of us, just add them to your cart, and pay for it, our system will send the downloading link and password to you within ten minutes, and if you don't receive, just contact us, we will solve this problem for you as quickly as possible. For Data-Engineer-Associate Training Materials, we also have after-service, if you have questions about the exam dumps, you can contact us by email.

Amazon AWS Certified Data Engineer - Associate (DEA-C01) Sample Questions (Q71-Q76):

NEW QUESTION # 71

A data engineer has a one-time task to read data from objects that are in Apache Parquet format in an Amazon S3 bucket. The data engineer needs to query only one column of the data.

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

- A. Use S3 Select to write a SQL SELECT statement to retrieve the required column from the S3 objects.
- B. Prepare an AWS Glue DataBrew project to consume the S3 objects and to query the required column.
- C. Configure an AWS Lambda function to load data from the S3 bucket into a pandas dataframe- Write a SQL SELECT statement on the dataframe to query the required column.
- D. Run an AWS Glue crawler on the S3 objects. Use a SQL SELECT statement in Amazon Athena to query the required column.

Answer: A

Explanation:

Option B is the best solution to meet the requirements with the least operational overhead because S3 Select is a feature that allows you to retrieve only a subset of data from an S3 object by using simple SQL expressions.

S3 Select works on objects stored in CSV, JSON, or Parquet format. By using S3 Select, you can avoid the need to download and process the entire S3 object, which reduces the amount of data transferred and the computation time. S3 Select is also easy to use and does not require any additional services or resources.

Option A is not a good solution because it involves writing custom code and configuring an AWS Lambda function to load data from the S3 bucket into a pandas dataframe and query the required column. This option adds complexity and latency to the data retrieval process and requires additional resources and configuration.

Moreover, AWS Lambda has limitations on the execution time, memory, and concurrency, which may affect the performance and reliability of the data retrieval process.

Option C is not a good solution because it involves creating and running an AWS Glue DataBrew project to consume the S3 objects and query the required column. AWS Glue DataBrew is a visual data preparation tool that allows you to clean, normalize, and transform data without writing code. However, in this scenario, the data is already in Parquet format, which is a columnar storage format that is optimized for analytics.

Therefore, there is no need to use AWS Glue DataBrew to prepare the data. Moreover, AWS Glue DataBrew adds extra time and cost to the data retrieval process and requires additional resources and configuration.

Option D is not a good solution because it involves running an AWS Glue crawler on the S3 objects and using a SQL SELECT statement in Amazon Athena to query the required column. An AWS Glue crawler is a service that can scan data sources and create metadata tables in the AWS Glue Data Catalog. The Data Catalog is a central repository that stores information about the data sources, such as schema, format, and location. Amazon Athena is a serverless interactive query service that allows you to analyze data in S3 using standard SQL. However, in this scenario, the schema and format of the data are already known and fixed, so there is no need to run a crawler to discover them. Moreover, running a crawler and using Amazon Athena adds extra time and cost to the data retrieval process and requires additional services and configuration.

AWS Certified Data Engineer - Associate DEA-C01 Complete Study Guide

S3 Select and Glacier Select - Amazon Simple Storage Service

AWS Lambda - FAQs

What Is AWS Glue DataBrew? - AWS Glue DataBrew

Populating the AWS Glue Data Catalog - AWS Glue

What is Amazon Athena? - Amazon Athena

NEW QUESTION # 72

A retail company uses Amazon Aurora PostgreSQL to process and store live transactional data. The company uses an Amazon Redshift cluster for a data warehouse.

An extract, transform, and load (ETL) job runs every morning to update the Redshift cluster with new data from the PostgreSQL database. The company has grown rapidly and needs to cost optimize the Redshift cluster.

A data engineer needs to create a solution to archive historical data. The data engineer must be able to run analytics queries that effectively combine data from live transactional data in PostgreSQL, current data in Redshift, and archived historical data. The solution must keep only the most recent 15 months of data in Amazon Redshift to reduce costs.

Which combination of steps will meet these requirements? (Select TWO.)

- A. Schedule a monthly job to copy data that is older than 15 months to Amazon S3 by using the UNLOAD command. Delete the old data from the Redshift cluster. Configure Amazon Redshift Spectrum to access historical data in Amazon S3.
- B. Schedule a monthly job to copy data that is older than 15 months to Amazon S3 Glacier Flexible Retrieval by using the UNLOAD command. Delete the old data from the Redshift cluster. Configure Redshift Spectrum to access historical data from S3 Glacier Flexible Retrieval.
- C. Configure the Amazon Redshift Federated Query feature to query live transactional data that is in the PostgreSQL database.
- D. Configure Amazon Redshift Spectrum to query live transactional data that is in the PostgreSQL database.
- E. Create a materialized view in Amazon Redshift that combines live, current, and historical data from different sources.

Answer: A,C

Explanation:

The goal is to archive historical data from an Amazon Redshift data warehouse while combining live transactional data from Amazon Aurora PostgreSQL with current and historical data in a cost-efficient manner. The company wants to keep only the last 15 months of data in Redshift to reduce costs.

* Option A: "Configure the Amazon Redshift Federated Query feature to query live transactional data that is in the PostgreSQL database." Redshift Federated Query allows querying live transactional data directly from Aurora PostgreSQL without having to move it into Redshift, thereby enabling seamless integration of the current data in Redshift and live data in PostgreSQL. This is a cost-effective approach, as it avoids unnecessary data duplication.

* Option C: "Schedule a monthly job to copy data that is older than 15 months to Amazon S3 by using the UNLOAD command. Delete the old data from the Redshift cluster. Configure Amazon Redshift Spectrum to access historical data in Amazon S3." This option uses Amazon Redshift Spectrum, which enables Redshift to query data directly in S3 without moving it into Redshift. By unloading older data (older than 15 months) to S3, and then using Spectrum to access it, this approach reduces storage costs significantly while still allowing the data to be queried when necessary.

* Option B (Redshift Spectrum for live PostgreSQL data) is not applicable, as Redshift Spectrum is intended for querying data in Amazon S3, not live transactional data in Aurora.

* Option D (S3 Glacier Flexible Retrieval) is not suitable because Glacier is designed for long-term archival storage with infrequent access, and querying data in Glacier for analytics purposes would incur higher retrieval times and costs.

* Option E (materialized views) would not meet the need to archive data or combine it from multiple sources; it is best suited for combining frequently accessed data already in Redshift.

References:

* Amazon Redshift Federated Query

* Amazon Redshift Spectrum Documentation

* Amazon Redshift UNLOAD Command

NEW QUESTION # 73

A data engineer needs to join data from multiple sources to perform a one-time analysis job. The data is stored in Amazon DynamoDB, Amazon RDS, Amazon Redshift, and Amazon S3.

Which solution will meet this requirement MOST cost-effectively?

- A. Use an Amazon EMR provisioned cluster to read from all sources. Use Apache Spark to join the data and perform the analysis.
- B. Use Redshift Spectrum to query data from DynamoDB, Amazon RDS, and Amazon S3 directly from Redshift.
- **C. Use Amazon Athena Federated Query to join the data from all data sources.**
- D. Copy the data from DynamoDB, Amazon RDS, and Amazon Redshift into Amazon S3. Run Amazon Athena queries directly on the S3 files.

Answer: C

Explanation:

Amazon Athena Federated Query is a feature that allows you to query data from multiple sources using standard SQL. You can use Athena Federated Query to join data from Amazon DynamoDB, Amazon RDS, Amazon Redshift, and Amazon S3, as well as other data sources such as MongoDB, Apache HBase, and Apache Kafka¹. Athena Federated Query is a serverless and interactive service, meaning you do not need to provision or manage any infrastructure, and you only pay for the amount of data scanned by your queries.

Athena Federated Query is the most cost-effective solution for performing a one-time analysis job on data from multiple sources, as it eliminates the need to copy or move data, and allows you to query data directly from the source.

The other options are not as cost-effective as Athena Federated Query, as they involve additional steps or costs. Option A requires you to provision and pay for an Amazon EMR cluster, which can be expensive and time-consuming for a one-time job. Option B requires you to copy or move data from DynamoDB, RDS, and Redshift to S3, which can incur additional costs for data transfer and storage, and also introduce latency and complexity. Option D requires you to have an existing Redshift cluster, which can be costly and may not be necessary for a one-time job. Option D also does not support querying data from RDS directly, so you would need to use Redshift Federated Query to access RDS data, which adds another layer of complexity².

References:

Amazon Athena Federated Query

Redshift Spectrum vs Federated Query

NEW QUESTION # 74

A financial company recently added more features to its mobile app. The new features required the company to create a new topic in an existing Amazon Managed Streaming for Apache Kafka (Amazon MSK) cluster.

A few days after the company added the new topic, Amazon CloudWatch raised an alarm on the RootDiskUsed metric for the MSK cluster.

How should the company address the CloudWatch alarm?

- A. Expand the storage of the MSK broker. Configure the MSK cluster storage to expand automatically.
- B. Update the MSK broker instance to a larger instance type. Restart the MSK cluster.
- C. Expand the storage of the Apache ZooKeeper nodes.
- D. Specify the Target-Volume-in-GiB parameter for the existing topic.

Answer: A

Explanation:

The RootDiskUsed metric for the MSK cluster indicates that the storage on the broker is reaching its capacity. The best solution is to expand the storage of the MSK broker and enable automatic storage expansion to prevent future alarms.

Expand MSK Broker Storage:

AWS Managed Streaming for Apache Kafka (MSK) allows you to expand the broker storage to accommodate growing data volumes. Additionally, auto-expansion of storage can be configured to ensure that storage grows automatically as the data increases.

Reference:

Alternatives Considered:

B (Expand Zookeeper storage): Zookeeper is responsible for managing Kafka metadata and not for storing data, so increasing Zookeeper storage won't resolve the root disk issue.

C (Update instance type): Changing the instance type would increase computational resources but not directly address the storage problem.

D (Target-Volume-in-GiB): This parameter is irrelevant for the existing topic and will not solve the storage issue.

Amazon MSK Storage Auto Scaling

NEW QUESTION # 75

A company maintains multiple extract, transform, and load (ETL) workflows that ingest data from the company's operational databases into an Amazon S3 based data lake. The ETL workflows use AWS Glue and Amazon EMR to process data.

The company wants to improve the existing architecture to provide automated orchestration and to require minimal manual effort. Which solution will meet these requirements with the LEAST operational overhead?

- A. Amazon Managed Workflows for Apache Airflow (Amazon MWAA) workflows
- B. AWS Glue workflows
- C. AWS Lambda functions
- D. AWS Step Functions tasks

Answer: B

Explanation:

AWS Glue workflows are a feature of AWS Glue that enable you to create and visualize complex ETL pipelines using AWS Glue components, such as crawlers, jobs, triggers, and development endpoints. AWS Glue workflows provide automated orchestration and require minimal manual effort, as they handle dependency resolution, error handling, state management, and resource allocation for your ETL workflows. You can use AWS Glue workflows to ingest data from your operational databases into your Amazon S3 based data lake, and then use AWS Glue and Amazon EMR to process the data in the data lake. This solution will meet the requirements with the least operational overhead, as it leverages the serverless and fully managed nature of AWS Glue, and the scalability and flexibility of Amazon EMR².

The other options are not optimal for the following reasons:

B . AWS Step Functions tasks. AWS Step Functions is a service that lets you coordinate multiple AWS services into serverless workflows. You can use AWS Step Functions tasks to invoke AWS Glue and Amazon EMR jobs as part of your ETL workflows, and use AWS Step Functions state machines to define the logic and flow of your workflows. However, this option would require more manual effort than AWS Glue workflows, as you would need to write JSON code to define your state machines, handle errors and retries, and monitor the execution history and status of your workflows³.

C . AWS Lambda functions. AWS Lambda is a service that lets you run code without provisioning or managing servers. You can use AWS Lambda functions to trigger AWS Glue and Amazon EMR jobs as part of your ETL workflows, and use AWS Lambda event sources and destinations to orchestrate the flow of your workflows. However, this option would also require more manual effort than AWS Glue workflows, as you would need to write code to implement your business logic, handle errors and retries, and

monitor the invocation and execution of your Lambda functions. Moreover, AWS Lambda functions have limitations on the execution time, memory, and concurrency, which may affect the performance and scalability of your ETL workflows.

D . Amazon Managed Workflows for Apache Airflow (Amazon MWAA) workflows. Amazon MWAA is a managed service that makes it easy to run open source Apache Airflow on AWS. Apache Airflow is a popular tool for creating and managing complex ETL pipelines using directed acyclic graphs (DAGs). You can use Amazon MWAA workflows to orchestrate AWS Glue and Amazon EMR jobs as part of your ETL workflows, and use the Airflow web interface to visualize and monitor your workflows. However, this option would have more operational overhead than AWS Glue workflows, as you would need to set up and configure your Amazon MWAA environment, write Python code to define your DAGs, and manage the dependencies and versions of your Airflow plugins and operators.

Reference:

1: AWS Glue Workflows

2: AWS Glue and Amazon EMR

3: AWS Step Functions

: AWS Lambda

: Amazon Managed Workflows for Apache Airflow

NEW QUESTION # 76

.....

We know that every user has their favorite. Therefore, we have provided three versions of Data-Engineer-Associate practice guide: the PDF, the Software and the APP online. You can choose according to your actual situation. If you like to use computer to learn, you can use the Software and the APP online versions of the Data-Engineer-Associate Exam Questions. If you like to write your own experience while studying, you can choose the PDF version of the Data-Engineer-Associate study materials. Our PDF version can be printed and you can take notes as you like.

Online Data-Engineer-Associate Lab Simulation: https://www.pass4cram.com/Data-Engineer-Associate_free-download.html

Amazon Passing Data-Engineer-Associate Score Feedback The money will be refund to your account and no extra questions will be asked, What's more, we have achieved breakthroughs in Data-Engineer-Associate certification training application as well as interactive sharing and after-sales service, Data-Engineer-Associate AWS Certified Data Engineer - Associate (DEA-C01) exam questions & answers are codified by Amazon qualified experts, Amazon Passing Data-Engineer-Associate Score Feedback Obtaining certification will make you stand out from other people and make a big difference in your work.

The firm is targeting large-scale corporate users as its main Data-Engineer-Associate Valid Test Bootcamp audience, selling an Appmail Server for enterprises to host these applications themselves, We have included this chapter to provide additional background information and tips Data-Engineer-Associate that we hope will help you if you want to develop a more advanced Cocoon application, such as an Internet portal.

Pass Guaranteed Quiz Amazon - Valid Data-Engineer-Associate - Passing AWS Certified Data Engineer - Associate (DEA-C01) Score Feedback

The money will be refund to your account and no extra questions will be asked, What's more, we have achieved breakthroughs in Data-Engineer-Associate certification training application as well as interactive sharing and after-sales service.

Data-Engineer-Associate AWS Certified Data Engineer - Associate (DEA-C01) exam questions & answers are codified by Amazon qualified experts, Obtaining certification will make you stand out from other people and make a big difference in your work.

Do you want to choose a lifetime Data-Engineer-Associate Study Dumps of mediocrity or become better and pursue your dreams?

- 2026 Data-Engineer-Associate: Authoritative Passing AWS Certified Data Engineer - Associate (DEA-C01) Score Feedback Search for [Data-Engineer-Associate] and easily obtain a free download on www.vce4dumps.com Test Data-Engineer-Associate Valid
- Excellent Passing Data-Engineer-Associate Score Feedback - Trustable Source of Data-Engineer-Associate Exam Easily obtain free download of 《 Data-Engineer-Associate 》 by searching on (www.pdfvce.com) Latest Test Data-Engineer-Associate Simulations
- Passing Data-Engineer-Associate Score Feedback Key Data-Engineer-Associate Concepts Latest Test Data-Engineer-Associate Simulations The page for free download of [Data-Engineer-Associate](#) on www.vce4dumps.com will open immediately Data-Engineer-Associate Authentic Exam Hub
- Valid Data-Engineer-Associate Real Test Data-Engineer-Associate Latest Learning Materials Data-Engineer-Associate New Cram Materials Open website www.pdfvce.com and search for “ Data-Engineer-Associate ”

