Desktop-Based Amazon Data-Engineer-Associate Practice Exam Software

P.S. Free & New Data-Engineer-Associate dumps are available on Google Drive shared by ValidVCE: https://drive.google.com/open?id=1_o_bv3b7hQd6kiaY09IkIBwfZ76YV0z2

Our industry experts are constantly adding new content to Data-Engineer-Associate test dumps based on constantly changing syllabus and industry development breakthroughs. We also hired dedicated IT staff to continuously update our question bank daily, so no matter when you buy Data-Engineer-Associate Study Materials, what you learn is the most advanced. Even if you fail to pass the exam, as long as you are willing to continue to use our Data-Engineer-Associate test answers, we will still provide you with the benefits of free updates within a year.

Our Data-Engineer-Associate training quiz is provided by PDF, Software/PC, and App/Online, which allows you to choose a suitable way to study anytime and anywhere. The PDF versions of Data-Engineer-Associate study materials can be printed into a paper file, more convenient to read and take notes. You can also try the simulated exam environment with Data-Engineer-Associate software on PC. Anyway, you can practice the key knowledge repeatedly with our Data-Engineer-Associate test prep, and at the same time, you can consolidate your weaknesses more specifically.

>> Data-Engineer-Associate Latest Exam Simulator <<

New Data-Engineer-Associate Study Guide | Data-Engineer-Associate Lab Questions

If you buy the Amazon Data-Engineer-Associate practice materials within one year you can enjoy free updates. Being the most competitive and advantageous company in the market, our AWS Certified Data Engineer - Associate (DEA-C01) Data-Engineer-Associate exam questions have help tens of millions of exam candidates, realized their dreams all these years. What you can harvest is not only certificate but of successful future from now on just like our former clients.

Amazon AWS Certified Data Engineer - Associate (DEA-C01) Sample Questions (Q90-Q95):

NEW QUESTION # 90
A company uses Amazon Redshift as a data warehouse solution. One of the datasets that the company stores in Amazon Redshift contains data for a vendor.
Recently, the vendor asked the company to transfer the vendor ' s data into the vendor ' s Amazon S3 bucket once each week.
Which solution will meet this requirement?

- A. Configure Amazon Redshift Spectrum to use the vendor ' s S3 bucket as destination. Enable data querying in both directions.
- B. Use the Amazon Redshift data sharing feature. Set the vendor ' s S3 bucket as the destination.
  Configure the source to be as a custom SQL query that selects the required data.
- C. Create an AWS Glue job to connect to the Redshift data warehouse. Configure the AWS Glue job to use the Redshift UNLOAD command to load the required data to the vendor ' s S3 bucket on a schedule.

- D. Create an AWS Lambda function to connect to the Redshift data warehouse. Configure the Lambda function to use the Redshift COPY command to copy the required data to the vendor ' s S3 bucket on a schedule.

Answer: C

Explanation:
The Redshift UNLOAD command is specifically designed to export query results to Amazon S3, and AWS Glue can orchestrate this as part of a scheduled job. This is the cleanest and most appropriate approach for recurring weekly data transfers:
"Use the Redshift UNLOAD command with AWS Glue to export data to Amazon S3. This pattern enables routine exports of selected data to external locations."
- Ace the AWS Certified Data Engineer - Associate Certification - version 2 - apple.pdf This avoids complexities of Redshift Spectrum or unsupported use of COPY commands in Lambda.

NEW QUESTION # 91
A company receives a daily file that contains customer data in .xls format. The company stores the file in Amazon S3. The daily file is approximately 2 GB in size.
A data engineer concatenates the column in the file that contains customer first names and the column that contains customer last names. The data engineer needs to determine the number of distinct customers in the file.
Which solution will meet this requirement with the LEAST operational effort?

- A. Create and run an Apache Spark job in Amazon EMR Serverless to calculate the number of distinct customers.
- B. Create and run an Apache Spark job in an AWS Glue notebook. Configure the job to read the S3 file and calculate the number of distinct customers.
- C. Use AWS Glue DataBrew to create a recipe that uses the COUNT_DISTINCT aggregate function to calculate the number of distinct customers.
- D. Create an AWS Glue crawler to create an AWS Glue Data Catalog of the S3 file. Run SQL queries from Amazon Athena to calculate the number of distinct customers.

Answer: C

NEW QUESTION # 92
A company is migrating its database servers from Amazon EC2 instances that run Microsoft SQL Server to Amazon RDS for Microsoft SQL Server DB instances. The company's analytics team must export large data elements every day until the migration is complete. The data elements are the result of SQL joins across multiple tables. The data must be in Apache Parquet format. The analytics team must store the data in Amazon S3.
Which solution will meet these requirements in the MOST operationally efficient way?

- A. Use a SQL query to create a view in the EC2 instance-based SQL Server databases that contains the required data elements. Create and run an AWS Glue crawler to read the view. Create an AWS Glue job that retrieves the data and transfers the data in Parquet format to an S3 bucket. Schedule the AWS Glue job to run every day.
- B. Schedule SQL Server Agent to run a daily SQL query that selects the desired data elements from the EC2 instance-based SQL Server databases. Configure the query to direct the output .csv objects to an S3 bucket. Create an S3 event that invokes an AWS Lambda function to transform the output format from .csv to Parquet.
- C. Create an AWS Lambda function that queries the EC2 instance-based databases by using Java Database Connectivity (JDBC). Configure the Lambda function to retrieve the required data, transform the data into Parquet format, and transfer the data into an S3 bucket. Use Amazon EventBridge to schedule the Lambda function to run every day.
- D. Create a view in the EC2 instance-based SQL Server databases that contains the required data elements.
Create an AWS Glue job that selects the data directly from the view and transfers the data in Parquet format to an S3 bucket.
Schedule the AWS Glue job to run every day.

Answer: D

Explanation:
Option A is the most operationally efficient way to meet the requirements because it minimizes the number of steps and services involved in the data export process. AWS Glue is a fully managed service that can extract, transform, and load (ETL) data from various sources to various destinations, including Amazon S3. AWS Glue can also convert data to different formats, such as Parquet, which is a columnar storage format that is optimized for analytics. By creating a view in the SQL Server databases that contains the required data elements, the AWS Glue job can select the data directly from the view without having to perform any joins or transformations on the source data. The AWS Glue job can then transfer the data in Parquet format to an S3 bucket and run

on a daily schedule.
Option B is not operationally efficient because it involves multiple steps and services to export the data. SQL Server Agent is a tool that can run scheduled tasks on SQL Server databases, such as executing SQL queries.
However, SQL Server Agent cannot directlyexport data to S3, so the query output must be saved as .csv objects on the EC2 instance. Then, an S3 event must be configured to trigger an AWS Lambda function that can transform the .csv objects to Parquet format and upload them to S3. This option adds complexity and latency to the data export process and requires additional resources and configuration.
Option C is not operationally efficient because it introduces an unnecessary step of running an AWS Glue crawler to read the view. 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. However, in this scenario, the schema and format of the data elements are already known and fixed, so there is no need to run a crawler to discover them. The AWS Glue job can directly select the data from the view without using the Data Catalog. Running a crawler adds extra time and cost to the data export process.
Option D is not operationally efficient because it requires custom code and configuration to query the databases and transform the data. An AWS Lambda function is a service that can run code in response to events or triggers, such as Amazon EventBridge. Amazon EventBridge is a service that can connect applications and services with event sources, such as schedules, and route them to targets, such as Lambda functions. However, in this scenario, using a Lambda function to query the databases and transform the data is not the best option because it requires writing and maintaining code that uses JDBC to connect to the SQL Server databases, retrieve the required data, convert the data to Parquet format, and transfer the data to S3.
This option also has limitations on the execution time, memory, and concurrency of the Lambda function, which may affect the performance and reliability of the data export process.
:
AWS Certified Data Engineer - Associate DEA-C01 Complete Study Guide
AWS Glue Documentation
Working with Views in AWS Glue
Converting to Columnar Formats

NEW QUESTION # 93
A company stores data from an application in an Amazon DynamoDB table that operates in provisioned capacity mode. The workloads of the application have predictable throughput load on a regular schedule.
Every Monday, there is an immediate increase in activity early in the morning. The application has very low usage during weekends.
The company must ensure that the application performs consistently during peak usage times.
Which solution will meet these requirements in the MOST cost-effective way?

- A. Use AWS Application Auto Scaling to schedule higher provisioned capacity for peak usage times. Schedule lower capacity during off-peak times.
- B. Change the capacity mode from provisioned to on-demand. Configure the table to scale up and scale down based on the load on the table.
- C. Increase the provisioned capacity to the maximum capacity that is currently present during peak load times.
- D. Divide the table into two tables. Provision each table with half of the provisioned capacity of the original table. Spread queries evenly across both tables.

Answer: A

Explanation:
Amazon DynamoDB is a fully managed NoSQL database service that provides fast and predictable performance with seamless scalability. DynamoDB offers two capacity modes for throughput capacity:
provisioned and on-demand. In provisioned capacity mode, you specify the number of read and write capacity units per second that you expect your application to require. DynamoDB reserves the resources to meet your throughput needs with consistent performance. In on-demand capacity mode, you pay per request and DynamoDB scales the resources up and down automatically based on the actual workload. On-demand capacity mode is suitable for unpredictable workloads that can vary significantly over time1.
The solution that meets the requirements in the most cost-effective way is to use AWS Application Auto Scaling to schedule higher provisioned capacity for peak usage times and lower capacity during off-peak times. This solution has the following advantages:
* It allows you to optimize the cost and performance of your DynamoDB table by adjusting the provisioned capacity according to your predictable workload patterns. You can use scheduled scaling to specify the date and time for the scaling actions, and the new minimum and maximum capacity limits. For example, you can schedule higher capacity for every Monday morning and lower capacity for weekends2.
* It enables you to take advantage of the lower cost per unit of provisioned capacity mode compared to on-demand capacity mode. Provisioned capacity mode charges a flat hourly rate for the capacity you reserve, regardless of how much you use. On-demand

capacity mode charges for each read and write request you consume, with no minimum capacity required. For predictable workloads, provisioned capacity mode can be more cost-effective than on-demand capacity mode1.
* It ensures that your application performs consistently during peak usage times by having enough capacity to handle the increased load. You can also use auto scaling to automatically adjust the provisioned capacity based on the actual utilization of your table, and set a target utilization percentage for your table or global secondary index. This way, you can avoid under-provisioning or over-provisioning your table2.
Option A is incorrect because it suggests increasing the provisioned capacity to the maximum capacity that is currently present during peak load times. This solution has the following disadvantages:
* It wastes money by paying for unused capacity during off-peak times. If you provision the same high capacity for all times, regardless of the actual workload, you are over-provisioning your table and paying for resources that you don't need1.
* It does not account for possible changes in the workload patterns over time. If your peak load times increase or decrease in the future, you may need to manually adjust the provisioned capacity to match the new demand. This adds operational overhead and complexity to your application2.
Option B is incorrect because it suggests dividing the table into two tables and provisioning each table with half of the provisioned capacity of the original table. This solution has the following disadvantages:
* It complicates the data model and the application logic by splitting the data into two separate tables.
You need to ensure that the queries are evenly distributed across both tables, and that the data is consistent and synchronized between them. This adds extra development and maintenance effort to your application3.
* It does not solve the problem of adjusting the provisioned capacity according to the workload patterns.
You still need to manually or automatically scale the capacity of each table based on the actual utilization and demand. This may result in under-provisioning or over-provisioning your tables2.
Option D is incorrect because it suggests changing the capacity mode from provisioned to on-demand. This solution has the following disadvantages:
* It may incur higher costs than provisioned capacity mode for predictable workloads. On-demand capacity mode charges for each read and write request you consume, with no minimum capacity required. For predictable workloads, provisioned capacity mode can be more cost-effective than on- demand capacity mode, as you can reserve the capacity you need at a lower rate1.
* It may not provide consistent performance during peak usage times, as on-demand capacity mode may take some time to scale up the resources to meet the sudden increase in demand. On-demand capacity mode uses adaptive capacity to handle bursts of traffic, but it may not be able to handle very large spikes or sustained high throughput. In such cases, you may experience throttling or increased latency.
:
1: Choosing the right DynamoDB capacity mode - Amazon DynamoDB
2: Managing throughput capacity automatically with DynamoDB auto scaling - Amazon DynamoDB
3: Best practices for designing and using partition keys effectively - Amazon DynamoDB
[4]: On-demand mode guidelines - Amazon DynamoDB
[5]: How to optimize Amazon DynamoDB costs - AWS Database Blog
[6]: DynamoDB adaptive capacity: How it works and how it helps - AWS Database Blog
[7]: Amazon DynamoDB pricing - Amazon Web Services (AWS)

NEW QUESTION # 94
A company ' s application needs to search and analyze data in near real time. The application must handle up to 1,000 requests each second with low query latency. The company wants a solution that individual data teams can own and configure to meet each team ' s cost and performance optimization requirements.
Which solution will meet these requirements?

- A. Use Amazon Aurora clusters that run on Aurora I/O-Optimized instances. Assign each data team a separate Aurora cluster to configure for storage and queries.
- B. Use Amazon OpenSearch Service clusters with indexing to query the data. Assign each data team a separate cluster to configure for storage and queries.
- C. Use streams in Amazon Kinesis Data Streams and Amazon Managed Service for Apache Flink to query and analyze the data. Assign each data team a separate stream to manage and consume.
- D. Use Amazon S3 buckets to store the data. Use Amazon Athena to query and analyze the data. Assign each data team a separate S3 bucket prefix to optimize queries.

Answer: B

Explanation:
Option C is correct because Amazon OpenSearch Service is designed for search and analytics in near real time with low-latency queries. AWS documentation states that Amazon OpenSearch Service makes it easy to deploy, secure, operate, and scale OpenSearch to search, analyze, and visualize data in real time, and that it provides real-time analytics for use cases such as log

analytics, full-text search, application monitoring, and clickstream analytics. That aligns directly with the requirement for handling many requests per second with low query latency.
This option also best satisfies the requirement that individual data teams can own and configure their own environments for cost and performance optimization. By assigning each team a separate OpenSearch cluster, each team can independently tune indexing, storage, shard layout, retention, and scaling policies. Option A is less suitable because Amazon Athena is serverless SQL over data in S3 and is not the best fit for low-latency, high-request-rate interactive search workloads. Option B focuses on streaming data processing, not primary low-latency indexed search. Option D uses a relational database, which is not the native AWS choice for large-scale search and analytics. Therefore, OpenSearch Service is the most appropriate and scalable solution.

NEW QUESTION # 95
.....

The AWS Certified Data Engineer - Associate (DEA-C01) (Data-Engineer-Associate) product can be easily accessed just after purchasing it from ValidVCE. You can receive free Amazon Dumps updates for up to 1 year after buying material. The 24/7 support system is also available for you, which helps you every time you get stuck somewhere. Many students have studied from the ValidVCE Amazon Data-Engineer-Associate practice material and rated it positively because they have passed the AWS Certified Data Engineer - Associate (DEA-C01) (Data-Engineer-Associate) certification exam on the first try.

New Data-Engineer-Associate Study Guide: https://www.validvce.com/Data-Engineer-Associate-exam-collection.html

Our AWS Certified Data Engineer - Associate (DEA-C01) (Data-Engineer-Associate) exam dumps comes in three formats: Amazon Data-Engineer-Associate PDF dumps file, desktop-based practice test software, and a web-based practice exam, Although you will take each Data-Engineer-Associate online test one at a time - each one builds upon the previous, In addition, except Data-Engineer-Associate, many other certification exams are also useful, You can download Data-Engineer-Associate certkingdom pdf demo for a try.

Archiving Classes with Jar, Little did I know, my Data-Engineer-Associate own life was encased in a bubble of deception that was about to burst, Our AWS Certified Data Engineer - Associate (DEA-C01) (Data-Engineer-Associate) exam dumps comes in three formats: Amazon Data-Engineer-Associate PDF dumps file, desktop-based practice test software, and a web-based practice exam.

Quiz Amazon - Data-Engineer-Associate - AWS Certified Data Engineer - Associate (DEA-C01) Pass-Sure Latest Exam Simulator

Although you will take each Data-Engineer-Associate online test one at a time - each one builds upon the previous, In addition, except Data-Engineer-Associate, many other certification exams are also useful.

You can download Data-Engineer-Associate certkingdom pdf demo for a try, How to find such good learning material software?

- Data-Engineer-Associate Exam Paper Pdf □ New Data-Engineer-Associate Test Syllabus ✏ Reliable Data-Engineer-Associate Dumps Pdf □ Easily obtain （ Data-Engineer-Associate ） for free download through 「 www.troytecdumps.com 」 □Valid Braindumps Data-Engineer-Associate Free
- Current Data-Engineer-Associate Exam Content □ Valid Data-Engineer-Associate Exam Online □ New Data-Engineer-Associate Exam Practice 🈸 Easily obtain free download of ➥ Data-Engineer-Associate □ by searching on " www.pdfvce.com " □Data-Engineer-Associate Authentic Exam Hub
- Data-Engineer-Associate Valid Exam Question □ Passing Data-Engineer-Associate Score Feedback □ Data-Engineer-Associate Valid Exam Question □ Easily obtain free download of （ Data-Engineer-Associate ） by searching on □ www.vce4dumps.com □ □Data-Engineer-Associate Valid Test Cost
- Reliable Data-Engineer-Associate Dumps Pdf □ Valid Braindumps Data-Engineer-Associate Free □ Data-Engineer-Associate Exam Paper Pdf □ Search for ➢ Data-Engineer-Associate □ on ➠ www.pdfvce.com □ immediately to obtain a free download □New Data-Engineer-Associate Test Syllabus
- Data-Engineer-Associate Valid Exam Question □ Data-Engineer-Associate Exam Paper Pdf □ Data-Engineer-Associate Practice Exams □ Download ➽ Data-Engineer-Associate □ for free by simply searching on ➠ www.torrentvce.com □ □ □Data-Engineer-Associate Authentic Exam Hub
- New Data-Engineer-Associate Exam Practice □ Valid Dumps Data-Engineer-Associate Free □ Data-Engineer-Associate New Braindumps □ Open □ www.pdfvce.com □ and search for { Data-Engineer-Associate } to download exam materials for free □Data-Engineer-Associate Authentic Exam Hub
- Data-Engineer-Associate Valid Test Cost □ Valid Data-Engineer-Associate Test Materials □ Data-Engineer-Associate New Braindumps □ Easily obtain 《 Data-Engineer-Associate 》 for free download through 「 www.prepawayete.com 」 □Data-Engineer-Associate Authentic Exam Hub
- Free PDF Quiz Amazon - Unparalleled Data-Engineer-Associate Latest Exam Simulator □ Search for { Data-Engineer-

Associate } on ▸ www.pdfvce.com ◂ immediately to obtain a free download □Data-Engineer-Associate New Braindumps

- 100% Pass-Rate Data-Engineer-Associate Latest Exam Simulator - Leading Offer in Qualification Exams - First-Grade Amazon AWS Certified Data Engineer - Associate (DEA-C01) □ Copy URL ➥ www.examcollectionpass.com □ open and search for ▹ Data-Engineer-Associate ◃ to download for free □Valid Data-Engineer-Associate Exam Online
- Data-Engineer-Associate Latest Exam Simulator - How to Download for New Data-Engineer-Associate Study Guide free □ □ Open 【 www.pdfvce.com 】 and search for （ Data-Engineer-Associate ） to download exam materials for free □ □Real Data-Engineer-Associate Exams
- 100% Pass-Rate Data-Engineer-Associate Latest Exam Simulator - Leading Offer in Qualification Exams - First-Grade Amazon AWS Certified Data Engineer - Associate (DEA-C01) □ Search for [Data-Engineer-Associate] and easily obtain a free download on □ www.dumpsmaterials.com □ □Data-Engineer-Associate Valid Exam Question
- hamzahocft274048.wiki-cms.com, www.impactio.com, emmaklewis.sites.gettysburg.edu, rsaoxsy076792.wikitron.com, hotbookmarkings.com, scrapbookmarket.com, estelleokmq422821.blogacep.com, caoimhelhcy036213.59bloggers.com, get-social-now.com, arrangvmn104771.vblogetin.com, Disposable vapes

What's more, part of that ValidVCE Data-Engineer-Associate dumps now are free: https://drive.google.com/open?id=1_o_bv3b7hQd6kiaY09IkIBwfZ76YV0z2