

Valid Amazon Data-Engineer-Associate Exam Review - Reliable Data-Engineer-Associate Test Tutorial



DOWNLOAD the newest DumpsTorrent Data-Engineer-Associate PDF dumps from Cloud Storage for free:
<https://drive.google.com/open?id=1FiOecMtC0qta-5AyW74Z3-mueiUZWxvI>

Although it is not an easy thing for some candidates to pass the exam, but our Data-Engineer-Associate question torrent can help aggressive people to achieve their goals. This is the reason why we need to recognize the importance of getting the test Data-Engineer-Associate certification. If you have any doubt about our products that will bring a lot of benefits for you. The trial demo of our Data-Engineer-Associate question torrent must be a good choice for you. By the trial demo provided by our company, you will have the opportunity to closely contact with our Data-Engineer-Associate exam torrent, and it will be possible for you to have a view of our products.

If you want to be an excellent elites in this line, you need to get the Data-Engineer-Associate certification, thus it can be seen through the importance of qualification examination. Only through qualification examination, has obtained the corresponding qualification certificate, we will be able to engage in related work, so the Data-Engineer-Associate Test Torrent is to help people in a relatively short period of time a great important tool to pass the qualification test. Choose our Data-Engineer-Associate study tool, can help users quickly analysis in the difficult point, and pass the Data-Engineer-Associate exam successfully.

>> Valid Amazon Data-Engineer-Associate Exam Review <<

Reliable Data-Engineer-Associate Test Tutorial - Reliable Data-Engineer-Associate Exam Cost

There is an irreplaceable trend that an increasingly amount of clients are picking up Data-Engineer-Associate study materials from tremendous practice materials in the market. There are unconquerable obstacles ahead of us if you get help from our Data-Engineer-Associate Exam Questions. So many exam candidates feel privileged to have our Data-Engineer-Associate practice braindumps. And our website is truly very famous for the hot hit in the market and easy to be found on the internet.

Amazon AWS Certified Data Engineer - Associate (DEA-C01) Sample Questions (Q251-Q256):

NEW QUESTION # 251

A data engineer is using Amazon QuickSight to build a dashboard to report a company's revenue in multiple AWS Regions. The data engineer wants the dashboard to display the total revenue for a Region, regardless of the drill-down levels shown in the visual. Which solution will meet these requirements?

- A. Create a level-aware calculation - aggregate (LAC-A) function.
- B. Create a table calculation.
- C. Create a simple calculated field.
- D. Create a level-aware calculation - window (LAC-W) function.

Answer: A

Explanation:

Option C (LAC-A) is the correct choice because the requirement is to always show the Region-level total even when the visual is drilled down into lower levels (for example, country # city # store). A simple calculated field (Option B) is computed at the row level and then aggregated by the visual, so it will change as the drill-down changes the grain. A table calculation (Option A) is evaluated based on the current visual layout and can vary with the fields placed in the visual, which makes it unreliable for enforcing a fixed "Region total regardless of drill-down."

A level-aware calculation (aggregate) is specifically intended to "lock" an aggregation to a chosen dimensional level (here: Region). That means you can compute revenue aggregated at the Region level and reuse that value across lower drill levels without it recalculating at city/store granularity. A window LAC (LAC-W) is primarily for windowed analytics (running totals, period-over-period, rank, moving averages) over a partition/order, not for enforcing a fixed dimensional aggregation level. Therefore, LAC-A best matches the requirement.

NEW QUESTION # 252

A data engineer needs to build an extract, transform, and load (ETL) job. The ETL job will process daily incoming .csv files that users upload to an Amazon S3 bucket. The size of each S3 object is less than 100 MB. Which solution will meet these requirements MOST cost-effectively?

- A. Write a PySpark ETL script. Host the script on an Amazon EMR cluster.
- B. Write a custom Python application. Host the application on an Amazon Elastic Kubernetes Service (Amazon EKS) cluster.
- C. Write an AWS Glue PySpark job. Use Apache Spark to transform the data.
- **D. Write an AWS Glue Python shell job. Use pandas to transform the data.**

Answer: D

Explanation:

AWS Glue is a fully managed serverless ETL service that can handle various data sources and formats, including .csv files in Amazon S3. AWS Glue provides two types of jobs: PySpark and Python shell. PySpark jobs use Apache Spark to process large-scale data in parallel, while Python shell jobs use Python scripts to process small-scale data in a single execution environment. For this requirement, a Python shell job is more suitable and cost-effective, as the size of each S3 object is less than 100 MB, which does not require distributed processing. A Python shell job can use pandas, a popular Python library for data analysis, to transform the .csv data as needed. The other solutions are not optimal or relevant for this requirement. Writing a custom Python application and hosting it on an Amazon EKS cluster would require more effort and resources to set up and manage the Kubernetes environment, as well as to handle the data ingestion and transformation logic. Writing a PySpark ETL script and hosting it on an Amazon EMR cluster would also incur more costs and complexity to provision and configure the EMR cluster, as well as to use Apache Spark for processing small data files. Writing an AWS Glue PySpark job would also be less efficient and economical than a Python shell job, as it would involve unnecessary overhead and charges for using Apache Spark for small data files. References:

AWS Glue

Working with Python Shell Jobs

pandas

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

NEW QUESTION # 253

A data engineer is using Amazon Athena to analyze sales data that is in Amazon S3. The data engineer writes a query to retrieve sales amounts for 2023 for several products from a table named sales_data. However, the query does not return results for all of the products that are in the sales_data table. The data engineer needs to troubleshoot the query to resolve the issue.

The data engineer's original query is as follows:

```
SELECT product_name, sum(sales_amount)
```

```
FROM sales_data
```

```
WHERE year = 2023
```

```
GROUP BY product_name
```

How should the data engineer modify the Athena query to meet these requirements?

- **A. Change WHERE year = 2023 to WHERE extract(year FROM sales_data) = 2023.**
- B. Add HAVING sum(sales amount) > 0 after the GROUP BY clause.
- C. Replace sum(sales amount) with count(*) for the aggregation.
- D. Remove the GROUP BY clause

Answer: A

Explanation:

The original query does not return results for all of the products because the year column in the sales_data table is not an integer, but a timestamp. Therefore, the WHERE clause does not filter the data correctly, and only returns the products that have a null value for the year column. To fix this, the data engineer should use the extract function to extract the year from the timestamp and compare it with 2023. This way, the query will return the correct results for all of the products in the sales_data table. The other options are either incorrect or irrelevant, as they do not address the root cause of the issue. Replacing sum with count does not change the filtering condition, adding HAVING clause does not affect the grouping logic, and removing the GROUP BY clause does not solve the problem of missing products. References:

Troubleshooting JSON queries - Amazon Athena (Section: JSON related errors) When I query a table in Amazon Athena, the TIMESTAMP result is empty (Section: Resolution) AWS Certified Data Engineer - Associate DEA-C01 Complete Study Guide (Chapter 7, page 197)

NEW QUESTION # 254

A data engineer needs to create an AWS Lambda function that converts the format of data from .csv to Apache Parquet. The Lambda function must run only if a user uploads a .csv file to an Amazon S3 bucket. Which solution will meet these requirements with the LEAST operational overhead?

- A. Create an S3 event notification that has an event type of s3:ObjectCreated:*. Use a filter rule to generate notifications only when the suffix includes .csv. Set an Amazon Simple Notification Service (Amazon SNS) topic as the destination for the event notification. Subscribe the Lambda function to the SNS topic.
- **B. Create an S3 event notification that has an event type of s3:ObjectCreated:*. Use a filter rule to generate notifications only when the suffix includes .csv. Set the Amazon Resource Name (ARN) of the Lambda function as the destination for the event notification.**
- C. Create an S3 event notification that has an event type of s3:*. Use a filter rule to generate notifications only when the suffix includes .csv. Set the Amazon Resource Name (ARN) of the Lambda function as the destination for the event notification.
- D. Create an S3 event notification that has an event type of s3:ObjectTagging:* for objects that have a tag set to .csv. Set the Amazon Resource Name (ARN) of the Lambda function as the destination for the event notification.

Answer: B

Explanation:

Option A is the correct answer because it meets the requirements with the least operational overhead. Creating an S3 event notification that has an event type of s3:ObjectCreated:* will trigger the Lambda function whenever a new object is created in the S3 bucket. Using a filter rule to generate notifications only when the suffix includes .csv will ensure that the Lambda function only runs for .csv files. Setting the ARN of the Lambda function as the destination for the event notification will directly invoke the Lambda function without any additional steps.

Option B is incorrect because it requires the user to tag the objects with .csv, which adds an extra step and increases the operational overhead.

Option C is incorrect because it uses an event type of s3:*, which will trigger the Lambda function for any S3 event, not just object creation. This could result in unnecessary invocations and increased costs.

Option D is incorrect because it involves creating and subscribing to an SNS topic, which adds an extra layer of complexity and operational overhead.

References:

AWS Certified Data Engineer - Associate DEA-C01 Complete Study Guide, Chapter 3: Data Ingestion and Transformation, Section 3.2: S3 Event Notifications and Lambda Functions, Pages 67-69 Building Batch Data Analytics Solutions on AWS, Module 4: Data Transformation, Lesson 4.2: AWS Lambda, Pages 4-8 AWS Documentation Overview, AWS Lambda Developer Guide, Working with AWS Lambda Functions, Configuring Function Triggers, Using AWS Lambda with Amazon S3, Pages 1-5

NEW QUESTION # 255

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. 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.
- D. Increase the provisioned capacity to the maximum capacity that is currently present during peak load times.

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.

Reference:

[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 # 256

.....

Our web-based practice test is accessible from anywhere with an internet connection, which means you can take it at your convenience. This Amazon Data-Engineer-Associate Practice Test is designed to simulate the actual exam and help you become familiar with the test format. You can access the web-based practice exam from anywhere with an internet connection to study on the go or from the comfort of your own home. You can receive your mock exam result instantly.

Reliable Data-Engineer-Associate Test Tutorial: <https://www.dumpstorrent.com/Data-Engineer-Associate-exam-dumps-torrent.html>

Amazon Valid Data-Engineer-Associate Exam Review It was never so easy to make your way to the world's most worthwhile professional qualification as it has become now, It is well known that Reliable Data-Engineer-Associate Test Tutorial- AWS Certified Data Engineer - Associate (DEA-C01) exam is an international recognition certification test, which is very important for people who are engaged in this field, We offer you free demo for Data-Engineer-Associate exam braindumps, and we recommend you have a try before buying.

This newly forged interaction between people and AI is another fascinating Reliable Data-Engineer-Associate Test Tutorial part of this space, Use a password manager By now, you may be sick of hearing about how you should be using a password manager.

Start Preparation With Amazon Data-Engineer-Associate Latest Dumps Today

It was never so easy to make your way to the world's Valid Data-Engineer-Associate Exam Review most worthwhile professional qualification as it has become now, It is well known that AWS Certified Data Engineer - Associate (DEA-C01) exam is an international recognition Data-Engineer-Associate Certification test, which is very important for people who are engaged in this field.

We offer you free demo for Data-Engineer-Associate exam braindumps, and we recommend you have a try before buying, Considering the current plea of our exam candidates we make up our mind to fight for your satisfaction and wish to pass the Data-Engineer-Associate exam.

If so, you have come to the right place.

- Flexible Data-Engineer-Associate Learning Mode Data-Engineer-Associate Exam Torrent Data-Engineer-Associate Valid Exam Preparation Search for Data-Engineer-Associate and obtain a free download on www.practicevce.com Valid Data-Engineer-Associate Exam Notes
- Quiz Amazon - Data-Engineer-Associate - The Best Valid AWS Certified Data Engineer - Associate (DEA-C01) Exam Review Search for [Data-Engineer-Associate] and download it for free immediately on www.pdfvce.com Data-Engineer-Associate Exam Cram Questions
- Valid Data-Engineer-Associate Exam Notes Data-Engineer-Associate Latest Dump Data-Engineer-Associate Latest Dumps Pdf Search for Data-Engineer-Associate and obtain a free download on www.validtorrent.com Practical Data-Engineer-Associate Information
- Amazon Data-Engineer-Associate Exam | Valid Data-Engineer-Associate Exam Review - High Pass Rate Reliable Data-Engineer-Associate Test Tutorial Search for 【 Data-Engineer-Associate 】 and download it for free immediately on www.pdfvce.com Data-Engineer-Associate Free Dump Download
- Amazon certification Data-Engineer-Associate the latest examination questions and answers come out Open www.practicevce.com and search for Data-Engineer-Associate to download exam materials for free Cost Effective Data-Engineer-Associate Dumps
- Quiz Amazon - Data-Engineer-Associate - The Best Valid AWS Certified Data Engineer - Associate (DEA-C01) Exam Review Search for (Data-Engineer-Associate) and download exam materials for free through [www.pdfvce.com] Practical Data-Engineer-Associate Information
- Data-Engineer-Associate Brain Dump Free Data-Engineer-Associate Exam Torrent Data-Engineer-Associate Test Cram Download Data-Engineer-Associate for free by simply entering www.examcollectionpass.com website Data-Engineer-Associate Trustworthy Practice
- Data-Engineer-Associate Actual Exams Data-Engineer-Associate Brain Dump Free Data-Engineer-Associate Exam Cram Questions www.pdfvce.com is best website to obtain Data-Engineer-Associate for free download Data-Engineer-Associate Latest Dump
- Desktop Amazon Data-Engineer-Associate Practice Test Software By www.dumpsmaterials.com Easily obtain [Data-Engineer-Associate] for free download through www.dumpsmaterials.com Data-Engineer-Associate Trustworthy Practice

