Data-Engineer-Associate試験問題、Data-Engineer-Associate日本語試験対策

ちなみに、MogiExam Data-Engineer-Associateの一部をクラウドストレージからダウンロードできます：https://drive.google.com/open?id=1yETMTKvgJE48pyFgskC0duF15zmwngze

時間は何もありません。 タイミングが全てだ。 heしないでください。 Data-Engineer-Associate VCEダンプは、試験をクリアする時間を節約するのに役立ちます。 有効な試験ファイルを選択した場合、試験は一発で合格します。 Amazon VCEダンプで最短時間で認定資格を取得できます。 今すぐ上級職に就くと、他の人よりも絶対に有利になります。 これで、時間を無駄にせずに、Data-Engineer-Associate VCEダンプから始めてください。 優れた有効なVCEダンプは、あなたの夢を実現し、他の仲間よりも先に人生のピークを迎えます。

我々MogiExamはお客様に満足させるための最高のサービスを提供します。あなたに安心させるため、我々はData-Engineer-Associate問題集のサンプルを無料で提供して、あなたはダウンロードしてやってみることができます。あなたはData-Engineer-Associate模擬問題集をご購入になってから、我々は一年間の無料更新サービスを提供します。

>> Data-Engineer-Associate試験問題 <<

試験の準備方法-素晴らしいData-Engineer-Associate試験問題試験-最新のData-Engineer-Associate日本語試験対策

Data-Engineer-Associate学習教材があれば、あなたは自分の夢を叶えます。Data-Engineer-Associate学習教材はすごく人気があります。全世界のお客様からいい評価をもらいました。なんといっても、自分はData-Engineer-Associate学習教材を利用したら、その資料のよさを感じることができます。大切なのは、Data-Engineer-Associate学習教材の合格率が高いので、多くのお客様はData-Engineer-Associate認定試験資格証明書を取得したということです。

Amazon AWS Certified Data Engineer - Associate (DEA-C01) 認定 Data-Engineer-Associate 試験問題 (Q152-Q157):

質問 # 152
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. 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.
- B. 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.
- C. Use AWS Application Auto Scaling to schedule higher provisioned capacity for peak usage times. Schedule lower

capacity during off-peak times.

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

正解：C

解説：
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)

質問 # 153
A company uses an Amazon Redshift provisioned cluster as its database. The Redshift cluster has five reserved ra3.4xlarge nodes and uses key distribution.
A data engineer notices that one of the nodes frequently has a CPU load over 90%. SQL Queries that run on the node are queued. The other four nodes usually have a CPU load under 15% during daily operations.
The data engineer wants to maintain the current number of compute nodes. The data engineer also wants to balance the load more evenly across all five compute nodes.
Which solution will meet these requirements?

- A. Upgrade the reserved node from ra3.4xlarqe to ra3.16xlarqe.
- B. Change the sort key to be the data column that is most often used in a WHERE clause of the SQL SELECT statement.
- C. Change the distribution key to the table column that has the largest dimension.
- D. Change the primary key to be the data column that is most often used in a WHERE clause of the SQL SELECT statement.

正解：C

解説：
Changing the distribution key to the table column that has the largest dimension will help to balance the load more evenly across all five compute nodes. The distribution key determines how the rows of a table are distributed among the slices of the cluster. If the distribution key is not chosen wisely, it can cause data skew, meaning some slices will have more data than others, resulting in uneven CPU load and query performance. By choosing the table column that has the largest dimension, meaning the column that has the most distinct values, as the distribution key, the data engineer can ensure that the rows are distributed more uniformly across the slices, reducing data skew and improving query performance.
The other options are not solutions that will meet the requirements. Option A, changing the sort key to be the data column that is most often used in a WHERE clause of the SQL SELECT statement, will not affect the data distribution or the CPU load. The sort key determines the order in which the rows of a table are stored on disk, which can improve the performance of range-restricted queries, but not the load balancing. Option C, upgrading the reserved node from ra3.4xlarge to ra3.16xlarge, will not maintain the current number of compute nodes, as it will increase the cost and the capacity of the cluster. Option D, changing the primary key to be the data column that is most often used in a WHERE clause of the SQL SELECT statement, will not affect the data distribution or the CPU load either. The primary key is a constraint that enforces the uniqueness of the rows in a table, but it does not influence the data layout or the query optimization. Reference:
Choosing a data distribution style
Choosing a data sort key
Working with primary keys

質問 # 154
A data engineer needs to build an enterprise data catalog based on the company's Amazon S3 buckets and Amazon RDS databases. The data catalog must include storage format metadata for the data in the catalog.
Which solution will meet these requirements with the LEAST effort?

- A. Use scripts to scan data elements and to assign data classifications based on the format of the data.
- B. Use an AWS Glue crawler to build a data catalog. Use AWS Glue crawler classifiers to recognize the format of data and store the format in the catalog.
- C. Use an AWS Glue crawler to scan the S3 buckets and RDS databases and build a data catalog. Use data stewards to inspect the data and update the data catalog with the data format.
- D. Use Amazon Macie to build a data catalog and to identify sensitive data elements. Collect the data format information from Macie.

正解：B

解説：
To build an enterprise data catalog with metadata for storage formats, the easiest and most efficient solution is using an AWS Glue crawler. The Glue crawler can scan Amazon S3 buckets and Amazon RDS databases to automatically create a data catalog that includes metadata such as the schema and storage format (e.g., CSV, Parquet, etc.). By using AWS Glue crawler classifiers, you can configure the crawler to recognize the format of the data and store this information directly in the catalog.
* Option B: Use an AWS Glue crawler to build a data catalog. Use AWS Glue crawler classifiers to recognize the format of data and store the format in the catalog.This option meets the requirements with the least effort because Glue crawlers automate the discovery and cataloging of data from multiple sources, including S3 and RDS, while recognizing various file formats via classifiers. Other options (A, C, D) involve additional manual steps, like having data stewards inspect the data, or using services like Amazon

Macie that focus more on sensitive data detection rather than format cataloging.
References:
* AWS Glue Crawler Documentation
* AWS Glue Classifiers

質問 # 155
A data engineer needs to use AWS Step Functions to design an orchestration workflow. The workflow must parallel process a large collection of data files and apply a specific transformation to each file.
Which Step Functions state should the data engineer use to meet these requirements?

- A. Wait state
- B. Parallel state
- C. Choice state
- D. Map state

正解：D

解説：
Option C is the correct answer because the Map state is designed to process a collection of data in parallel by applying the same transformation to each element. The Map state can invoke a nested workflow for each element, which can be another state machine or a Lambda function. The Map state will wait until all the parallel executions are completed before moving to the next state.
Option A is incorrect because the Parallel state is used to execute multiple branches of logic concurrently, not to process a collection of data. The Parallel state can have different branches with different logic and states, whereas the Map state has only one branch that is applied to each element of the collection.
Option B is incorrect because the Choice state is used to make decisions based on a comparison of a value to a set of rules. The Choice state does not process any data or invoke any nested workflows.
Option D is incorrect because the Wait state is used to delay the state machine from continuing for a specified time. The Wait state does not process any data or invoke any nested workflows.
Reference:
AWS Certified Data Engineer - Associate DEA-C01 Complete Study Guide, Chapter 5: Data Orchestration, Section 5.3: AWS Step Functions, Pages 131-132 Building Batch Data Analytics Solutions on AWS, Module 5: Data Orchestration, Lesson 5.2: AWS Step Functions, Pages 9-10 AWS Documentation Overview, AWS Step Functions Developer Guide, Step Functions Concepts, State Types, Map State, Pages 1-3

質問 # 156
A company stores time-series data that is collected from streaming services in an Amazon S3 bucket. The company must ensure that only workloads that are deployed within the company's VPC can access the data.
Which solution will meet this requirement?

- A. Create an S3 bucket policy that uses a condition to allow access only to traffic that originates from the company's VPC.
- B. Apply a security group to the S3 bucket that allows connections only from the company's VPC CIDR block.
- C. Use a network ACL on the VPC subnets to allow only specific resources to access the S3 bucket.
- D. Define an IAM policy that denies access to all users unless the request originates from within the company's VPC.

正解：A

解説：
The best practice to restrict Amazon S3 access to specific VPCs is to use a bucket policy with a StringEquals or StringLike condition on aws:SourceVpc. This ensures only requests from a specified VPC are allowed.
* IAM policies (option C) control who can access the resource but are not suitable alone to restrict by VPC.
* Security Groups and NACLs (options B and D) do not apply to Amazon S3 because it is a global service and not VPC-bound.
"You can restrict access to your S3 bucket so that only requests coming from a specific VPC endpoint are allowed." Source: AWS Documentation - Amazon S3 Bucket Policies for VPC Endpoints

質問 # 157
.....

MogiExamの専門家チームは彼らの経験と知識を利用して長年の研究をわたって多くの人は待ちに待ったAmazonのData-Engineer-Associate「AWS Certified Data Engineer - Associate (DEA-C01)」認証試験について教育資料が完成してから、大変にお客様に歓迎されます。MogiExamの模擬試験は真実の試験問題はとても似ている専門家チームの勤労の結果としてとても値打ちがあります。

Data-Engineer-Associate日本語試験対策: https://www.mogiexam.com/Data-Engineer-Associate-exam.html

Amazon Data-Engineer-Associate試験問題 それでは、この問題集はあなたに適用するかどうかを確認した後、購入することを決定します、Amazon Data-Engineer-Associate試験問題 これもあなたの意志が強いことを表示する方法です、Amazon Data-Engineer-Associate試験問題 ふさわしいアシスタントを選ぶのは一番重要なのです、周知のように、すべての人は支払い後、できるだけ早くData-Engineer-Associate練習問題を受けたいと思います、弊社のData-Engineer-Associate模擬問題集は多くのIT専門家の協力で開発されています、Data-Engineer-Associateの無料のトレーニング資料は、他の人よりも労働市場で目立つようにし、より多くの機会があなたを見つけるためのイニシアチブを取ります、Amazon Data-Engineer-Associate 試験問題 一部の厳選された顧客でさえ、彼らの高品質と正確さの実践をやめることはできません。

ハズが が不合格になるのも死活問題として困る、必ずや〈精霊の た存在でData-Engineer-Associateある、それでは、この問題集はあなたに適用するかどうかを確認した後、購入することを決定します、これもあなたの意志が強いことを表示する方法です。

Data-Engineer-Associate試験の準備方法｜効率的なData-Engineer-Associate試験問題試験｜実際的なAWS Certified Data Engineer - Associate (DEA-C01)日本語試験対策

ふさわしいアシスタントを選ぶのは一番重要なのです、周知のように、すべての人は支払い後、できるだけ早くData-Engineer-Associate練習問題を受けたいと思います、弊社のData-Engineer-Associate模擬問題集は多くのIT専門家の協力で開発されています。

- Data-Engineer-Associate無料問題 □ Data-Engineer-Associate復習攻略問題 □ Data-Engineer-Associate復習攻略問題 □ ✔ www.passtest.jp □✔□サイトで《 Data-Engineer-Associate 》の最新問題が使えるData-Engineer-Associate日本語版テキスト内容
- Data-Engineer-Associateクラムメディア □ Data-Engineer-Associate模擬モード □ Data-Engineer-Associate全真問題集 □ ➽ www.goshiken.com □に移動し、⇒ Data-Engineer-Associate ⇐を検索して、無料でダウンロード可能な試験資料を探しますData-Engineer-Associate認定試験
- 最新のData-Engineer-Associate試験問題試験ツールの保証購入の安全性Data-Engineer-Associate日本語試験対策 □（ www.jptestking.com ）サイトにて最新【 Data-Engineer-Associate 】問題集をダウンロードData-Engineer-Associateクラムメディア
- Data-Engineer-Associate日本語 □ Data-Engineer-Associate無料問題 □ Data-Engineer-Associate認定デベロッパー □ [www.goshiken.com]には無料の（ Data-Engineer-Associate ）問題集がありますData-Engineer-Associate日本語認定対策
- Data-Engineer-Associate模擬試験サンプル □ Data-Engineer-Associate勉強の資料 □ Data-Engineer-Associate日本語 □（ jp.fast2test.com ）で□ Data-Engineer-Associate □を検索し、無料でダウンロードしてくださいData-Engineer-Associate日本語認定対策
- Data-Engineer-Associate試験の準備方法｜正確的なData-Engineer-Associate試験問題試験｜有効的なAWS Certified Data Engineer - Associate (DEA-C01)日本語試験対策 □ □ www.goshiken.com □から{ Data-Engineer-Associate }を検索して、試験資料を無料でダウンロードしてくださいData-Engineer-Associate日本語版トレーリング
- Data-Engineer-Associate勉強の資料 □ Data-Engineer-Associate参考資料 □ Data-Engineer-Associate日本語 □ サイト➥ www.it-passports.com □で{ Data-Engineer-Associate }問題集をダウンロードData-Engineer-Associate日本語版トレーリング
- Data-Engineer-Associate日本語版テキスト内容 □ Data-Engineer-Associate無料問題 □ Data-Engineer-Associate認定試験 □ ✔ www.goshiken.com □✔□から簡単に《 Data-Engineer-Associate 》を無料でダウンロードできますData-Engineer-Associate合格資料
- Data-Engineer-Associate参考資料 □ Data-Engineer-Associate日本語版テキスト内容 □ Data-Engineer-Associate日本語 □ ➥ www.xhs1991.com □にて限定無料の“ Data-Engineer-Associate ”問題集をダウンロードせよData-Engineer-Associate模擬モード
- Data-Engineer-Associate模擬試験サンプル □ Data-Engineer-Associate模擬試験サンプル □ Data-Engineer-Associate模擬試験サンプル □ 時間限定無料で使える{ Data-Engineer-Associate }の試験問題は□ www.goshiken.com □サイトで検索Data-Engineer-Associate全真問題集
- 最高のAmazonのData-Engineer-Associate試験テストソフトウェア □ サイト“ www.japancert.com ”で□ Data-Engineer-Associate □問題集をダウンロードData-Engineer-Associate試験合格攻略

- www.stes.tyc.edu.tw, www.stes.tyc.edu.tw, www.klemminghundar.se, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, www.aliusa.net, www.stes.tyc.edu.tw, www.stes.tyc.edu.tw, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, www.stes.tyc.edu.tw, Disposable vapes

P.S.MogiExamがGoogle Driveで共有している無料の2026 Amazon Data-Engineer-Associateダンプ：https://drive.google.com/open?id=1yETMTKvgJE48pyFgskC0duF15zmwngze