

Amazon Data-Engineer-Associate日本語受験攻略、 Data-Engineer-Associate資格取得



BONUS！！！ Tech4Exam Data-Engineer-Associateダンプの一部を無料でダウンロード：<https://drive.google.com/open?id=1XRCVcRBxPBKgnwxjL8BHz8hZH5cO3OPR>

国際的に認められているData-Engineer-Associate認定は、特定の分野の知識を十分に活用できることを意味するのは当然です。ワークロードに圧倒され、息を吸うことができない場合、Data-Engineer-Associate準備トレントを選択してみませんか？私たちは、最も信頼性が高く正確な試験資料をお客様に提供することに特化しており、お客様が満足のいくスコアを達成することで試験に合格できるよう支援しています。Data-Engineer-Associate練習教材を使用すると、試験は簡単になります。

競争が常に激化している世界では、特定の分野で優れた能力と深い知識を所有することで、高い社会的地位を獲得し、社会での地位を確立することができます。Tech4ExamテストData-Engineer-Associate認定に合格すると、目標を実現し、理想的な仕事を見つけるのに役立ちます。Data-Engineer-Associateの最新の質問を購入すると、Data-Engineer-Associate試験に合格できます。Data-Engineer-Associate試験問題の無料デモを試してみてください。Data-Engineer-Associate学習資料を気に入っていることでしょう。

>> [Amazon Data-Engineer-Associate日本語受験攻略](#) <<

Data-Engineer-Associate試験の準備方法 | 最新のData-Engineer-Associate日本語受験攻略試験 | 効果的なAWS Certified Data Engineer - Associate (DEA-C01)資格取得

Tech4Examの商品は100%の合格率を保証いたします。Tech4ExamはData-Engineer-Associateに対応性研究続けて、高品質で低価格な問題集が開発いたしました。Tech4Examの商品の最大の特徴は20時間だけ育成課程を通して樂々に合格できます。

Amazon AWS Certified Data Engineer - Associate (DEA-C01) 認定 Data-Engineer-Associate 試験問題 (Q213-Q218):

質問 # 213

A company uses Amazon RDS for MySQL as the database for a critical application. The database workload is mostly writes, with a small number of reads.

A data engineer notices that the CPU utilization of the DB instance is very high. The high CPU utilization is slowing down the application. The data engineer must reduce the CPU utilization of the DB Instance.

Which actions should the data engineer take to meet this requirement? (Choose two.)

- A. Modify the database schema to include additional tables and indexes.
- B. **Implement caching to reduce the database query load.**
- C. Reboot the RDS DB instance once each week.
- D. Upgrade to a larger instance size.
- E. **Use the Performance Insights feature of Amazon RDS to identify queries that have high CPU utilization. Optimize the problematic queries.**

正解： B、 E

解説：

Amazon RDS is a fully managed service that provides relational databases in the cloud. Amazon RDS for MySQL is one of the supported database engines that you can use to run your applications. Amazon RDS provides various features and tools to monitor and optimize the performance of your DB instances, such as Performance Insights, Enhanced Monitoring, CloudWatch metrics and alarms, etc.

Using the Performance Insights feature of Amazon RDS to identify queries that have high CPU utilization and optimizing the problematic queries will help reduce the CPU utilization of the DB instance. Performance Insights is a feature that allows you to analyze the load on your DB instance and determine what is causing performance issues. Performance Insights collects, analyzes, and displays database performance data using an interactive dashboard. You can use Performance Insights to identify the top SQL statements, hosts, users, or processes that are consuming the most CPU resources. You can also drill down into the details of each query and see the execution plan, wait events, locks, etc. By using Performance Insights, you can pinpoint the root cause of the high CPU utilization and optimize the queries accordingly. For example, you can rewrite the queries to make them more efficient, add or remove indexes, use prepared statements, etc.

Implementing caching to reduce the database query load will also help reduce the CPU utilization of the DB instance. Caching is a technique that allows you to store frequently accessed data in a fast and scalable storage layer, such as Amazon ElastiCache. By using caching, you can reduce the number of requests that hit your database, which in turn reduces the CPU load on your DB instance. Caching also improves the performance and availability of your application, as it reduces the latency and increases the throughput of your data access.

You can use caching for various scenarios, such as storing session data, user preferences, application configuration, etc. You can also use caching for read-heavy workloads, such as displaying product details, recommendations, reviews, etc.

The other options are not as effective as using Performance Insights and caching. Modifying the database schema to include additional tables and indexes may or may not improve the CPU utilization, depending on the nature of the workload and the queries. Adding more tables and indexes may increase the complexity and overhead of the database, which may negatively affect the performance. Rebooting the RDS DB instance once each week will not reduce the CPU utilization, as it will not address the underlying cause of the high CPU load. Rebooting may also cause downtime and disruption to your application. Upgrading to a larger instance size may reduce the CPU utilization, but it will also increase the cost and complexity of your solution.

Upgrading may also not be necessary if you can optimize the queries and reduce the database load by using caching. References:

* Amazon RDS

* Performance Insights

* Amazon ElastiCache

* [AWS Certified Data Engineer - Associate DEA-C01 Complete Study Guide], Chapter 3: Data Storage and Management, Section 3.1: Amazon RDS

質問 # 214

A company manages an Amazon Redshift data warehouse. The data warehouse is in a public subnet inside a custom VPC A security group allows only traffic from within itself. An ACL is open to all traffic.

The company wants to generate several visualizations in Amazon QuickSight for an upcoming sales event.

The company will run QuickSight Enterprise edition in a second AW5 account inside a public subnet within a second custom VPC. The new public subnet has a security group that allows outbound traffic to the existing Redshift cluster.

A data engineer needs to establish connections between Amazon Redshift and QuickSight. QuickSight must refresh dashboards by querying the Redshift cluster.

Which solution will meet these requirements?

- A. **Configure the Redshift security group to allow inbound traffic on the Redshift port from the QuickSight security group.**

- B. Assign Elastic IP addresses to the QuickSight visualizations. Configure the QuickSight security group to allow inbound traffic on the Redshift port from the Elastic IP addresses.
- C. Create a QuickSight gateway endpoint in the Redshift VPC. Attach an endpoint policy to the gateway endpoint to ensure only specific QuickSight accounts can use the endpoint.
- D. Confirm that the CIDR ranges of the Redshift VPC and the QuickSight VPC are the same. If CIDR ranges are different, reconfigure one CIDR range to match the other. Establish network peering between the VPCs.

正解: A

質問 # 215

A company uses Amazon RDS to store transactional data. The company runs an RDS DB instance in a private subnet. A developer wrote an AWS Lambda function with default settings to insert, update, or delete data in the DB instance.

The developer needs to give the Lambda function the ability to connect to the DB instance privately without using the public internet. Which combination of steps will meet this requirement with the LEAST operational overhead? (Choose two.)

- A. Update the security group of the DB instance to allow only Lambda function invocations on the database port.
- B. **Configure the Lambda function to run in the same subnet that the DB instance uses.**
- C. **Attach the same security group to the Lambda function and the DB instance. Include a self-referencing rule that allows access through the database port.**
- D. Turn on the public access setting for the DB instance.
- E. Update the network ACL of the private subnet to include a self-referencing rule that allows access through the database port.

正解: B、C

解説:

To enable the Lambda function to connect to the RDS DB instance privately without using the public internet, the best combination of steps is to configure the Lambda function to run in the same subnet that the DB instance uses, and attach the same security group to the Lambda function and the DB instance. This way, the Lambda function and the DB instance can communicate within the same private network, and the security group can allow traffic between them on the database port. This solution has the least operational overhead, as it does not require any changes to the public access setting, the network ACL, or the security group of the DB instance.

The other options are not optimal for the following reasons:

* A. Turn on the public access setting for the DB instance. This option is not recommended, as it would expose the DB instance to the public internet, which can compromise the security and privacy of the data. Moreover, this option would not enable the Lambda function to connect to the DB instance privately, as it would still require the Lambda function to use the public internet to access the DB instance.

* B. Update the security group of the DB instance to allow only Lambda function invocations on the database port. This option is not sufficient, as it would only modify the inbound rules of the security group of the DB instance, but not the outbound rules of the security group of the Lambda function.

Moreover, this option would not enable the Lambda function to connect to the DB instance privately, as it would still require the Lambda function to use the public internet to access the DB instance.

* E. Update the network ACL of the private subnet to include a self-referencing rule that allows access through the database port. This option is not necessary, as the network ACL of the private subnet already allows all traffic within the subnet by default.

Moreover, this option would not enable the Lambda function to connect to the DB instance privately, as it would still require the Lambda function to use the public internet to access the DB instance.

References:

- * 1: Connecting to an Amazon RDS DB instance
- * 2: Configuring a Lambda function to access resources in a VPC
- * 3: Working with security groups
- * : Network ACLs

質問 # 216

A data engineer needs to schedule a workflow that runs a set of AWS Glue jobs every day. The data engineer does not require the Glue jobs to run or finish at a specific time.

Which solution will run the Glue jobs in the MOST cost-effective way?

- A. Use the Spot Instance type in Glue job properties.
- B. Choose the STANDARD execution class in the Glue job properties.

- C. Choose the latest version in the GlueVersion field in the Glue job properties.
- D. Choose the FLEX execution class in the Glue job properties.

正解: D

解説:

The FLEX execution class allows you to run AWS Glue jobs on spare compute capacity instead of dedicated hardware. This can reduce the cost of running non-urgent or non-time sensitive data integration workloads, such as testing and one-time data loads. The FLEX execution class is available for AWS Glue 3.0 Spark jobs. The other options are not as cost-effective as FLEX, because they either use dedicated resources (STANDARD) or do not affect the cost at all (Spot Instance type and GlueVersion). Reference: Introducing AWS Glue Flex jobs: Cost savings on ETL workloads Serverless Data Integration - AWS Glue Pricing AWS Certified Data Engineer - Associate DEA-C01 Complete Study Guide (Chapter 5, page 125)

質問 # 217

A data engineer needs to use Amazon Neptune to develop graph applications.

Which programming languages should the engineer use to develop the graph applications? (Select TWO.)

- A. SQL
- B. Gremlin
- C. Spark SQL
- D. ANSI SQL
- E. SPARQL

正解: B、E

解説:

Amazon Neptune supports graph applications using Gremlin and SPARQL as query languages. Neptune is a fully managed graph database service that supports both property graph and RDF graph models.

Option A: Gremlin

Gremlin is a query language for property graph databases, which is supported by Amazon Neptune. It allows the traversal and manipulation of graph data in the property graph model.

Option D: SPARQL

SPARQL is a query language for querying RDF graph data in Neptune. It is used to query, manipulate, and retrieve information stored in RDF format.

Other options:

SQL (Option B) and ANSI SQL (Option C) are traditional relational database query languages and are not used for graph databases.

Spark SQL (Option E) is related to Apache Spark for big data processing, not for querying graph databases.

Reference:

Amazon Neptune Documentation

Gremlin Documentation

SPARQL Documentation

質問 # 218

.....

Data-Engineer-Associateテストガイドは、時間の無駄を避けるために、できるだけ早くこれらの資料を学習できることを保証できます。 AWS Certified Data Engineer - Associate (DEA-C01) Study Questionは、不明瞭な概念を簡素化することにより、学習方法を最適化するのに役立ちます。 Data-Engineer-Associate試験問題は、アフターサービスを完璧にするための努力をspareしません。

Data-Engineer-Associate資格取得: <https://www.tech4exam.com/Data-Engineer-Associate-pass-shiken.html>

そして、Tech4ExamのAmazonのData-Engineer-Associate問題集を購入したら、私たちは一年間で無料更新サービスを提供することができます、Data-Engineer-Associate試験を恐れることはありません、Tech4Exam Data-Engineer-Associate資格取得を選ぶかどうか状況があれば、弊社の無料なサンプルをダウンロードしてから、決めて大丈夫です、AmazonのData-Engineer-Associate問題集もCredit Cardで支払われることができます、Tech4ExamのAmazonのData-Engineer-Associate試験トレーニング資料はTech4Examの実力と豊富な経験を持っているIT専門家が研究し

たもので、本物のAmazonのData-Engineer-Associate試験問題とほぼ同じです、Amazon Data-Engineer-Associate日本語受験攻略やってみて第一步を進める勇気があります。

近所に遊び仲間のA子がいた、でもその前に一言ルーファスは指摘されて、目についたイチゴをパクリと口に変態イチゴ男、そして、Tech4ExamのAmazonのData-Engineer-Associate問題集を購入したら、私たちは一年間で無料更新サービスを提供することができます。

真実のData-Engineer-Associate | 効率的なData-Engineer-Associate日本語受験攻略試験 | 試験の準備方法 AWS Certified Data Engineer - Associate (DEA-C01) 資格取得

Data-Engineer-Associate試験を恐れることはありません、Tech4Examを選ぶかどうか状況があれば、弊社の無料なサンプルをダウンロードしてから、決めて大丈夫です、AmazonのData-Engineer-Associate問題集もCredit Cardで支払うことができます。

Tech4Examの AmazonのData-Engineer-Associate試験 トレーニング 資料は Tech4Examの実力と豊富な経験を持ついるIT専門家が研究したもので、本物のAmazonのData-Engineer-Associate試験問題とほぼ同じです。

P.S. Tech4ExamがGoogle Driveで共有している無料かつ新しいData-Engineer-Associateダンプ：<https://drive.google.com/open?id=1XRCVcRBxPBKgmwxjL8BH2hZH5cO3OPR>