

SAA-C03日本語解説集 & SAA-C03受験記対策



BONUS！！！ Tech4Exam SAA-C03ダンプの一部を無料でダウンロード：<https://drive.google.com/open?id=1jx39xbGLGZAmmgq6SuD7DZLdx6mdsoJJ>

私たちはSAA-C03試験参考書について多くのユーザーと話し合ったので、あなたの希望の資料だと思います。ユーザーのニーズに合わせるために、私たちのSAA-C03試験参考書が絶えず改善されています。私たちのSAA-C03試験参考書は世界からのユーザーを引きつけてきます。そして私たちのSAA-C03試験参考書は理解しやすいです。SAA-C03試験参考書を作る専門家は問題集の内容の研究に取り組んでいます。

SAA-C03試験は、EC2、S3、RDSなどのAWSのコアサービスや、AWS Lambda、Amazon Elastic Container Service (ECS)、Amazon API Gatewayなどの高度なサービスを含む幅広いトピックをカバーしています。この試験は、高可用性でスケーラブルなシステムを設計・展開する能力や、AWSプラットフォーム上のセキュリティとコンプライアンスに関する知識も評価します。この試験に合格することで、個人がAWSプラットフォーム上で高可用性、耐障害性、スケーラブルなシステムを設計・展開するために必要なスキルと知識を持っていることが証明され、クラウドへの移行や既存のAWSインフラストラクチャの最適化を目指す組織にとって貴重な資産となります。

Amazon SAA-C03、またはAmazon AWS Certified Solutions Architect - Associate (SAA-C03) 試験は、AWS上でスケーラブルで耐障害性のあるシステムを設計・展開する技術的な専門知識を検証する資格試験です。この認定資格は、ITアーキテクチャ、クラウドコンピューティング、およびAWSサービスの経験があるプロフェッショナルに適しています。

Amazon SAA-C03試験は、複数の選択肢と複数の回答方式の問題で構成され、AWS認定ポータルを通じてオンラインで実施されます。試験時間は130分で、候補者は少なくとも1000点中720点を得る必要があります。試験は、AWSストレージソリューション、コンピューティングとネットワーキングサービス、セキュリティ、アプリケーション統合など、幅広いトピックをカバーしています。候補者は、これらのトピックについて十分な理解を持っている必要があります。

SAA-C03受験記対策、SAA-C03復習解答例

Tech4ExamのSAA-C03には何か品質問題があることを見つければ、あるいは試験に合格しなかったのなら、弊社が無条件で全額返金することを約束します。Tech4Examは専門的にAmazonのSAA-C03試験の最新問題と解答を提供するサイトで、SAA-C03についての知識をほとんどカバーしています。

Amazon AWS Certified Solutions Architect - Associate 認定 SAA-C03 試験問題 (Q658-Q663):

質問 # 658

A company needs to keep user transaction data in an Amazon DynamoDB table.

The company must retain the data for 7 years.

What is the MOST operationally efficient solution that meets these requirements?

- A. Use AWS Backup to create backup schedules and retention policies for the table.
- B. Create an on-demand backup of the table by using the DynamoDB console. Store the backup in an Amazon S3 bucket. Set an S3 Lifecycle configuration for the S3 bucket.
- C. Create an Amazon EventBridge (Amazon CloudWatch Events) rule to invoke an AWS Lambda function. Configure the Lambda function to back up the table and to store the backup in an Amazon S3 bucket. Set an S3 Lifecycle configuration for the S3 bucket.
- D. Use DynamoDB point-in-time recovery to back up the table continuously.

正解: B

質問 # 659

A company uses an on-premises network-attached storage (NAS) system to provide file shares to its high performance computing (HPC) workloads. The company wants to migrate its latency-sensitive HPC workloads and its storage to the AWS Cloud. The company must be able to provide NFS and SMB multi-protocol access from the file system.

Which solution will meet these requirements with the LEAST latency? (Select TWO.)

- A. Deploy compute optimized EC2 instances into a cluster placement group.
- B. Attach the EC2 instances to an Amazon FSx for NetApp ONTAP file system
- C. Attach the EC2 instances to an Amazon FSx for Lustre file system
- D. Attach the EC2 instances to an Amazon FSx for OpenZFS file system
- E. Deploy compute optimized EC2 instances into a partition placement group.

正解: A、B

解説:

A cluster placement group is a logical grouping of EC2 instances within a single Availability Zone that are placed close together to minimize network latency. This is suitable for latency-sensitive HPC workloads that require high network performance. A compute optimized EC2 instance is an instance type that has a high ratio of vCPUs to memory, which is ideal for compute-intensive applications. Amazon FSx for NetApp ONTAP is a fully managed service that provides NFS and SMB multi-protocol access from the file system, as well as features such as data deduplication, compression, thin provisioning, and snapshots. This solution will meet the requirements with the least latency, as it leverages the low-latency network and storage performance of AWS.

Reference:

1 explains how cluster placement groups work and their benefits.

2 describes the characteristics and use cases of compute optimized EC2 instances.

3 provides an overview of Amazon FSx for NetApp ONTAP and its features.

質問 # 660

A company is developing a new application that uses a relational database to store user data and application configurations. The company expects the application to have steady user growth. The company expects the database usage to be variable and read-heavy, with occasional writes.

The company wants to cost-optimize the database solution. The company wants to use an AWS managed database solution that will

provide the necessary performance.

Which solution will meet these requirements MOST cost-effectively?

- A. Deploy the database on Amazon RDS Use magnetic storage and use read replicas to accommodate the workload
- B. Deploy the database on Amazon Aurora Serverless to automatically scale the database capacity based on actual usage to accommodate the workload.
- C. Deploy the database on Amazon DynamoDB. Use on-demand capacity mode to automatically scale throughput to accommodate the workload.
- D. Deploy the database on Amazon RDS. Use Provisioned IOPS SSD storage to ensure consistent performance for read and write operations.

正解： B

解説：

Amazon Aurora Serverless is a cost-effective, on-demand, autoscaling configuration for Amazon Aurora. It automatically adjusts the database's capacity based on the current demand, which is ideal for workloads with variable and unpredictable usage patterns. Since the application is expected to be read-heavy with occasional writes and steady growth, Aurora Serverless can provide the necessary performance without requiring the management of database instances.

Cost-Optimization: Aurora Serverless only charges for the database capacity you use, making it a more cost-effective solution compared to always running provisioned database instances, especially for workloads with fluctuating demand.

Scalability: It automatically scales database capacity up or down based on actual usage, ensuring that you always have the right amount of resources available.

Performance: Aurora Serverless is built on the same underlying storage as Amazon Aurora, providing high performance and availability.

Why Not Other Options?:

Option A (RDS with Provisioned IOPS SSD): While Provisioned IOPS SSD ensures consistent performance, it is generally more expensive and less flexible compared to the autoscaling nature of Aurora Serverless.

Option C (DynamoDB with On-Demand Capacity): DynamoDB is a NoSQL database and may not be the best fit for applications requiring relational database features.

Option D (RDS with Magnetic Storage and Read Replicas): Magnetic storage is outdated and generally slower. While read replicas help with read-heavy workloads, the overall performance might not be optimal, and magnetic storage doesn't provide the necessary performance.

AWS Reference:

[Amazon Aurora Serverless](#) - Information on how Aurora Serverless works and its use cases.

[Amazon Aurora Pricing](#) - Details on the cost-effectiveness of Aurora Serverless.

質問 # 661

A company is migrating some workloads to AWS. However, many workloads will remain on premises. The on-premises workloads require secure and reliable connectivity to AWS with consistent, low-latency performance.

The company has deployed the AWS workloads across multiple AWS accounts and multiple VPCs. The company plans to scale to hundreds of VPCs within the next year.

The company must establish connectivity between each of the VPCs and from the on-premises environment to each VPC.

Which solution will meet these requirements?

- A. Use multiple AWS Site-to-Site VPN connections to connect the on-premises environment to AWS. Create a transit gateway to establish connectivity between VPCs.
- B. Use an AWS Direct Connect connection with a Direct Connect gateway to connect the on-premises environment to AWS. Create a transit gateway to establish connectivity between VPCs. Associate the transit gateway with the Direct Connect gateway.
- C. Use an AWS Direct Connect connection to connect the on-premises environment to AWS. Configure VPC peering to establish connectivity between VPCs.
- D. Use an AWS Site-to-Site VPN connection to connect the on-premises environment to AWS. Configure VPC peering to establish connectivity between VPCs.

正解： B

解説：

The optimal solution for scalable and resilient hybrid networking is to use AWS Direct Connect with a Direct Connect gateway for secure, low-latency access to AWS, and an AWS Transit Gateway to manage connectivity among hundreds of VPCs.

By associating the Transit Gateway with the Direct Connect gateway, you enable transitive routing between on-premises and all

VPCs, while minimizing network complexity and maintaining high performance.

VPC peering does not scale well, and VPNs don't offer the same performance or consistency.

質問 # 662

[Design Resilient Architectures]

A company has a regional subscription-based streaming service that runs in a single AWS Region. The architecture consists of web servers and application servers on Amazon EC2 instances. The EC2 instances are in Auto Scaling groups behind Elastic Load Balancers. The architecture includes an Amazon Aurora database cluster that extends across multiple Availability Zones.

The company wants to expand globally and to ensure that its application has minimal downtime.

- A. Deploy the web tier and the application tier to a second Region. Add an Aurora PostgreSQL cross-Region Aurora Replica in the second Region. Use Amazon Route 53 health checks with a failovers routing policy to the second Region, Promote the secondary to primary as needed.
- B. Deploy the web tier and the application tier to a second Region. Use an Amazon Aurora global database to deploy the database in the primary Region and the second Region. Use Amazon Route 53 health checks with a failover routing policy to the second Region. Promote the secondary to primary as needed.
- C. Deploy the web tier and the applicatin tier to a second Region. Create an Aurora PostgreSQL database in the second Region. Use AWS Database Migration Service (AWS DMS) to replicate the primary database to the second Region. Use Amazon Route 53 health checks with a failover routing policy to the second Region.
- D. Extend the Auto Scaling groups for the web tier and the application tier to deploy instances in Availability Zones in a second Region. Use an Aurora global database to deploy the database in the primary Region and the second Region. Use Amazon Route 53 health checks with a failover routing policy to the second Region.

正解: B

解説:

This option is the most efficient because it deploys the web tier and the application tier to a second Region, which provides high availability and redundancy for the application. It also uses an Amazon Aurora global database, which is a feature that allows a single Aurora database to span multiple AWS Regions¹. It also deploys the database in the primary Region and the second Region, which provides low latency global reads and fast recovery from a Regional outage. It also uses Amazon Route 53 health checks with a failover routing policy to the second Region, which provides data protection by routing traffic to healthy endpoints in different Regions². It also promotes the secondary to primary as needed, which provides data consistency by allowing write operations in one of the Regions at a time³. This solution meets the requirement of expanding globally and ensuring that its application has minimal downtime. Option A is less efficient because it extends the Auto Scaling groups for the web tier and the application tier to deploy instances in Availability Zones in a second Region, which could incur higher costs and complexity than deploying them separately. It also uses an Aurora global database to deploy the database in the primary Region and the second Region, which is correct. However, it does not use Amazon Route 53 health checks with a failover routing policy to the second Region, which could result in traffic being routed to unhealthy endpoints. Option B is less efficient because it deploys the web tier and the application tier to a second Region, which is correct. It also adds an Aurora PostgreSQL cross-Region Aurora Replica in the second Region, which provides read scalability across Regions. However, it does not use an Aurora global database, which provides faster replication and recovery than cross-Region replicas. It also uses Amazon Route 53 health checks with a failover routing policy to the second Region, which is correct. However, it does not promote the secondary to primary as needed, which could result in data inconsistency or loss. Option C is less efficient because it deploys the web tier and the application tier to a second Region, which is correct. It also creates an Aurora PostgreSQL database in the second Region, which provides data redundancy across Regions. However, it does not use an Aurora global database or cross-Region replicas, which provide faster replication and recovery than creating separate databases. It also uses AWS Database Migration Service (AWS DMS) to replicate the primary database to the second Region, which provides data migration between different sources and targets. However, it does not use an Aurora global database or cross-Region replicas, which provide faster replication and recovery than using AWS DMS. It also uses Amazon Route 53 health checks with a failover routing policy to the second Region, which is correct.

質問 # 663

弊社Tech4ExamのSAA-C03試験準備では、学習習慣を身に付けるのに役立ちます。SAA-C03学習教材を購入して使用すると、学習の良い習慣を身に付けることができます。さらに重要なことは、良い習慣は科学的な小道具の学習方法を見つけ、学習効率を高めるのに役立ちます。そして、短時間でSAA-C03試験に合格するのに役立ちます。弊社からSAA-C03テストガイドを急いで購入すると、多くのメリットが得られます。

SAA-C03受験記対策: <https://www.tech4exam.com/SAA-C03-pass-shiken.html>

P.S. Tech4ExamがGoogle Driveで共有している無料かつ新しいSAA-C03ダンプ: <https://drive.google.com/open?id=1jx39xbGLGZAmmqg6SuD7DZLdx6mdsoJJ>