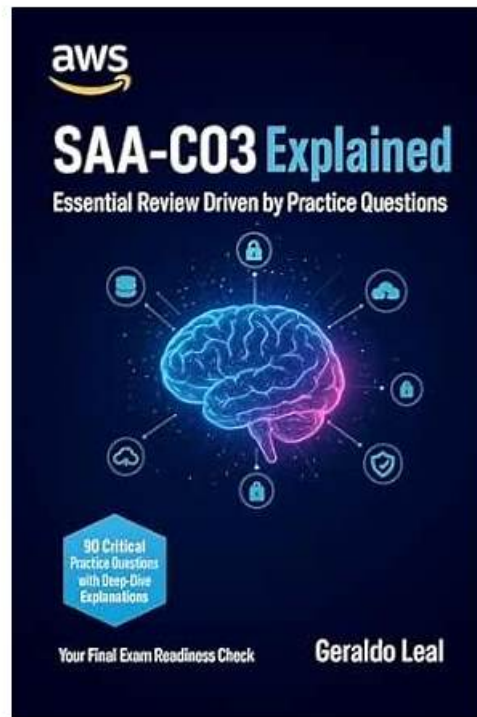


Amazon SAA-C03 Review Guide & Reliable SAA-C03 Exam Blueprint



BONUS!!! Download part of DumpsMaterials SAA-C03 dumps for free: <https://drive.google.com/open?id=1UyKOyiAXKVpJLtAs50JXc6tJUDTxmlyZ>

These people who used our products have thought highly of our SAA-C03 study materials. If you decide to buy our products and take it seriously consideration, we can make sure that it will be very easy for you to simply pass your exam and get the SAA-C03 certification in a short time. We are also willing to help you achieve your dream. Now give me a chance to show you our SAA-C03 Study Materials. You will have no regret spending your valuable time on our introduction. Besides, our SAA-C03 study quiz is priced reasonably, so we do not overcharge you at all.

The SAA-C03 exam covers a wide range of topics, including AWS services, security, networking, databases, and deployment. SAA-C03 exam consists of multiple-choice and multiple-response questions, and candidates are given 130 minutes to complete it. The passing score for the exam is 720 out of 1000, and the exam fee is \$150.

Amazon SAA-C03 (Amazon AWS Certified Solutions Architect - Associate) exam is a certification exam designed for IT professionals who want to validate their skills and knowledge in designing and deploying scalable, highly available, and fault-tolerant systems on Amazon Web Services (AWS) platforms. SAA-C03 Exam is intended for individuals who have some experience with AWS and are familiar with its core services and features.

Quiz SAA-C03 - Professional AWS Certified Solutions Architect - Associate Review Guide

With so many online resources, knowing where to start when preparing for an Amazon SAA-C03 exam can be tough. But with Amazon SAA-C03 practice test, you can be confident you're getting the best possible Amazon SAA-C03 Exam Dumps. Amazon exam mirrors the Amazon SAA-C03 exam-taking experience, so you know what to expect on Amazon SAA-C03 exam day.

Amazon AWS Certified Solutions Architect - Associate Sample Questions (Q121-Q126):

NEW QUESTION # 121

[Design Secure Architectures]

A company uses on-premises servers to host its applications. The company is running out of storage capacity. The applications use both block storage and NFS storage. The company needs a high-performing solution that supports local caching without re-architecting its existing applications.

Which combination of actions should a solutions architect take to meet these requirements? (Select TWO.)

- A. Mount Amazon S3 as a file system to the on-premises servers.
- B. Deploy AWS Snowball Edge to provision NFS mounts to on-premises servers.
- C. Deploy an AWS Storage Gateway file gateway to replace NFS storage.
- D. Deploy Amazon Elastic File System (Amazon EFS) volumes and mount them to on-premises servers.
- E. Deploy an AWS Storage Gateway volume gateway to replace the block storage

Answer: C,E

Explanation:

<https://aws.amazon.com/storagegateway/file/>

File Gateway provides a seamless way to connect to the cloud in order to store application data files and backup images as durable objects in Amazon S3 cloud storage. File Gateway offers SMB or NFS-based access to data in Amazon S3 with local caching. It can be used for on-premises applications, and for Amazon EC2-based applications that need file protocol access to S3 object storage.

<https://aws.amazon.com/storagegateway/volume/>

Volume Gateway presents cloud-backed iSCSI block storage volumes to your on-premises applications. Volume Gateway stores and manages on-premises data in Amazon S3 on your behalf and operates in either cache mode or stored mode. In the cached Volume Gateway mode, your primary data is stored in Amazon S3, while retaining your frequently accessed data locally in the cache for low latency access.

NEW QUESTION # 122

An online gaming company hosts its platform on Amazon EC2 instances behind Network Load Balancers (NLBs) across multiple AWS Regions. The NLBs can route requests to targets over the internet. The company wants to improve the customer playing experience by reducing end-to-end load time for its global customer base. Which solution will meet these requirements?

- A. Create a standard accelerator in AWS Global Accelerator. Configure the existing NLBs as target endpoints.
- B. Configure Amazon Route 53 to route equally weighted traffic to the NLBs in each Region.
- C. Create additional NLBs and EC2 instances in other Regions where the company has large customer bases.
- D. Create Application Load Balancers (ALBs) in each Region to replace the existing NLBs. Register the existing EC2 instances as targets for the ALBs in each Region.

Answer: A

Explanation:

The company wants to reduce end-to-end load time for its global customer base. AWS Global Accelerator provides a network optimization service that reduces latency by routing traffic to the nearest AWS edge locations, improving the user experience for globally distributed customers.

AWS Global Accelerator:

Global Accelerator improves the performance of your applications by routing traffic through AWS's global network infrastructure.

This reduces the number of hops and latency compared to using the public internet.

By creating a standard accelerator and configuring the existing NLBs as target endpoints, Global Accelerator ensures that traffic from users around the world is routed to the nearest AWS edge location and then through optimized paths to the NLBs in each region. This significantly improves end-to-end load time for global customers.

Why Not the Other Options?:

Option A (ALBs instead of NLBs): ALBs are designed for HTTP/HTTPS traffic and provide layer 7 features, but they wouldn't solve the latency issue for a global customer base. The key problem here is latency, and Global Accelerator is specifically designed to address that.

Option B (Route 53 weighted routing): Route 53 can route traffic to different regions, but it doesn't optimize network performance. It simply balances traffic between endpoints without improving latency.

Option C (Additional NLBs in more regions): This could potentially improve latency but would require setting up infrastructure in multiple regions. Global Accelerator is a simpler and more efficient solution that leverages AWS's existing global network.

NEW QUESTION # 123

A media company hosts its website on AWS. The website application's architecture includes a fleet of Amazon EC2 instances behind an Application Load Balancer (ALB) and a database that is hosted on Amazon Aurora. The company's cyber security team reports that the application is vulnerable to SQL injection.

How should the company resolve this issue?

- A. Subscribe to AWS Shield Advanced to block all SQL injection attempts automatically.
- **B. Use AWS WAF in front of the ALB. Associate the appropriate web ACLs with AWS WAF.**
- C. Create an ALB listener rule to reply to SQL injection with a fixed response.
- D. Set up Amazon Inspector to block all SQL injection attempts automatically.

Answer: B

Explanation:

<https://aws.amazon.com/premiumsupport/knowledge-center/waf-block-common-attacks/#:~:text=To%20protect%20your%20applications%20against,%2C%20query%20string%2C%20or%20URI.> -----

----- Protect against SQL injection and cross-site scripting To protect your applications against SQL injection and cross-site scripting (XSS) attacks, use the built-in SQL injection and cross-site scripting engines. Remember that attacks can be performed on different parts of the HTTP request, such as the HTTP header, query string, or URI. Configure the AWS WAF rules to inspect different parts of the HTTP request against the built-in mitigation engines.

NEW QUESTION # 124

A company wants to rearchitect a large-scale web application to a serverless microservices architecture. The application uses Amazon EC2 instances and is written in Python.

The company selected one component of the web application to test as a microservice. The component supports hundreds of requests each second. The company wants to create and test the microservice on an AWS solution that supports Python. The solution must also scale automatically and require minimal infrastructure and minimal operational support.

Which solution will meet these requirements?

- A. Use Amazon Elastic Kubernetes Service (Amazon EKS). Launch Auto Scaling groups of self-managed EC2 instances.
- B. Use a Spot Fleet with auto scaling of EC2 instances that run the most recent Amazon Linux operating system.
- **C. Use an AWS Lambda function that runs custom developed code.**
- D. Use an AWS Elastic Beanstalk web server environment that has high availability configured.

Answer: C

Explanation:

AWS Lambda is a serverless compute service that lets you run code without provisioning or managing servers. You can use Lambda to create and test microservices that are written in Python or other supported languages. Lambda scales automatically to handle the number of requests per second. You only pay for the compute time you consume. Lambda also integrates with other AWS services, such as Amazon API Gateway, Amazon S3, Amazon DynamoDB, and Amazon SQS, to enable event-driven architectures. Lambda has minimal infrastructure and operational overhead, as you do not need to manage servers, operating systems, patches, or scaling policies.

The other options are not serverless solutions and require more infrastructure and operational support. They also do not scale automatically to handle the number of requests per second. A Spot Fleet is a collection of EC2 instances that run on spare capacity

at low prices. However, Spot Instances can be interrupted by AWS at any time, which can affect the availability and performance of your microservice. AWS Elastic Beanstalk is a service that automates the deployment and management of web applications on EC2 instances. However, you still need to provision, configure, and monitor the underlying EC2 instances and load balancers. Amazon EKS is a service that runs Kubernetes on AWS. However, you still need to create, configure, and manage the EC2 instances that form the Kubernetes cluster and nodes. You also need to install and update the Kubernetes software and tools.

References:

- * What is AWS Lambda?
- * Building Lambda functions with Python
- * Create a layer for a Lambda Python function
- * AWS Lambda - Function in Python
- * How do I call my AWS Lambda function from a local python script?

NEW QUESTION # 125

A company stores several petabytes of data across multiple AWS accounts. The company uses AWS Lake Formation to manage its data lake. The company's data science team wants to securely share selective data from its accounts with the company's engineering team for analytical purposes.

Which solution will meet these requirements with the LEAST operational overhead?

- A. Use AWS Data Exchange to privately publish the required data to the required engineering team accounts
- **B. Use Lake Formation tag-based access control to authorize and grant cross-account permissions for the required data to the engineering team accounts**
- C. Use the Lake Formation permissions Grant command in each account where the data is stored to allow the required engineering team users to access the data.
- D. Copy the required data to a common account. Create an IAM access role in that account. Grant access by specifying a permission policy that includes users from the engineering team accounts as trusted entities.

Answer: B

Explanation:

Understanding the Requirement: The data science team needs to securely share selective data with the engineering team across multiple AWS accounts with minimal operational overhead.

Analysis of Options:

Copy data to a common account: Involves data duplication and increased storage costs, and requires managing additional permissions.

Lake Formation permissions Grant command: This method can be effective but may involve significant operational overhead if managing permissions across multiple accounts and datasets manually.

AWS Data Exchange: Designed for sharing data externally or between organizations, which adds unnecessary complexity for internal sharing within the same organization.

Lake Formation tag-based access control: Provides a scalable and efficient way to manage access permissions based on tags, allowing fine-grained control and simplified management across accounts.

Best Solution:

Lake Formation tag-based access control: This solution meets the requirements with the least operational overhead, allowing efficient management of cross-account permissions and secure data sharing.

References:

AWS Lake Formation

Tag-based access control

NEW QUESTION # 126

.....

DumpsMaterials is intent on keeping up with the latest technologies and applying them to the exam questions and answers not only on the content but also on the displays. That is why our pass rate is high as 98% to 100%. The data are unique-particular in this career. With our SAA-C03 study torrent, you can enjoy the leisure study experience as well as pass the SAA-C03 Exam with success ensured. For the content of our SAA-C03 preparation materials is simplified by our professional experts and the displays are designed effectually. Just try and enjoy it!

Reliable SAA-C03 Exam Blueprint: <https://www.dumpsmaterials.com/SAA-C03-real-torrent.html>

- Trusted SAA-C03 Review Guide - Leader in Qualification Exams - Valid Reliable SAA-C03 Exam Blueprint Search for

