Reliable Amazon AWS-Solutions-Architect-Associate Real Exam - AWS-Solutions-Architect-Associate Valid Exam Practice

P.S. Free & New AWS-Solutions-Architect-Associate dumps are available on Google Drive shared by Actualtests4sure: https://drive.google.com/open?id=1LqenwxVfo8PhZ0DXYeykdB2pINXUToJ7

The study system of our company will provide all customers with the best study materials. If you buy the AWS-Solutions-Architect-Associate latest questions of our company, you will have the right to enjoy all the AWS-Solutions-Architect-Associate certification training materials from our company. By updating the study system of the AWS-Solutions-Architect-Associate Training Materials, we can guarantee that our company can provide the newest information about the exam for all people. We believe that getting the newest information about the exam will help all customers pass the AWS-Solutions-Architect-Associate exam easily.

The AWS Certified Solutions Architect - Associate (SAA-C03) certification exam is designed for individuals who have at least one year of experience in designing and deploying scalable, cost-effective, and fault-tolerant systems on AWS. It covers a range of topics, including AWS architecture, deployment, security, and troubleshooting. AWS-Solutions-Architect-Associate exam is designed to test the candidate's knowledge and ability to apply AWS best practices and design principles to real-world scenarios.

The AWS Certified Solutions Architect - Associate (SAA-C02) exam is a rigorous and challenging certification exam that tests the candidate’s ability to design and deploy scalable and highly available systems on the AWS platform. Passing AWS-Solutions-Architect-Associate Exam demonstrates the candidate’s expertise in AWS services and their functionalities, making them a valuable asset to any organization looking to adopt cloud technology.

>> Reliable Amazon AWS-Solutions-Architect-Associate Real Exam <<

AWS-Solutions-Architect-Associate Valid Exam Practice - AWS-Solutions-Architect-Associate Practice Test

Even the fierce competition cannot stop demanding needs from exam candidates. To get more specific information about our AWS-Solutions-Architect-Associate learning quiz, we are here to satisfy your wish with following details. So you can get detailed information with traits and information about our AWS-Solutions-Architect-Associate Real Exam requested on the website. You can free download the demos of our AWS-Solutions-Architect-Associate exam questions and click on every detail that you are interested.

Amazon AWS Certified Solutions Architect - Associate (SAA-C03) Sample Questions (Q209-Q214):

NEW QUESTION # 209
A company has an ecommerce application running In a single VPC. The application stack has a single web server and an Amazon RDS Multi-AZ DB instance The company launches new products twice a month. This Increases website traffic by approximately 400% for a minimum of 72 hours During product launches, users experience slow response times and frequent timeout errors in their browsers What should a solutions architect do to mitigate the slow response times and timeout errors while minimizing operational

overhead?

- A. Add an Application Load Balancer and an additional web server.
- B. Increase the instance size of the web server.
- C. Add Amazon EC2 Auto Scaling and an Application Load Balancer
- D. Deploy an Amazon ElastiCache cluster to store frequently accessed data.

Answer: B

NEW QUESTION # 210
A company owns an asynchronous API that is used to ingest user requests and, based on the request type, dispatch requests to the appropriate microservice for processing. The company is using Amazon API Gateway to deploy the API front end, and an AWS Lambda function that invokes Amazon DynamoDB to store user requests before dispatching them to the processing microservices. The company provisioned as much DynamoDB throughput as its budget allows, but the company is still experiencing availability issues and is losing user requests.
What should a solutions architect do to address this issue without impacting existing users?

- A. Create a secondary index in DynamoDB for the table with the user requests.
- B. Use DynamoDB Accelerator (DAX) and Lambda to buffer writes to DynamoDB.
- C. Use the Amazon Simple Queue Service (Amazon SQS) queue and Lambda to buffer writes to DynamoDB.
- D. Add throttling on the API Gateway with server-side throttling limits.

Answer: C

Explanation:
Explanation
By using an SQS queue and Lambda, the solutions architect can decouple the API front end from the processing microservices and improve the overall scalability and availability of the system. The SQS queue acts as a buffer, allowing the API front end to continue accepting user requests even if the processing microservices are experiencing high workloads or are temporarily unavailable. The Lambda function can then retrieve requests from the SQS queue and write them to DynamoDB, ensuring that all user requests are stored and processed. This approach allows the company to scale the processing microservices independently from the API front end, ensuring that the API remains available to users even during periods of high demand.

NEW QUESTION # 211
A company has Amazon EC2 instances that run nightly batch jobs to process data. The EC2 instances run in an Auto Scaling group that uses On-Demand billing. If a job fails on one instance: another instance will reprocess the job. The batch jobs run between 12:00 AM and 06 00 AM local time every day.
Which solution will provide EC2 instances to meet these requirements MOST cost-effectively'?

- A. Purchase a 1-year Savings Plan for Amazon EC2 that covers the instance family of the Auto Scaling group that the batch job uses.
- B. Create a new launch template for the Auto Scaling group Increase the instance size Set a policy to scale out based on CPU usage.
- C. Create a new launch template for the Auto Scaling group Set the instances to Spot Instances Set a policy to scale out based on CPU usage.
- D. Purchase a 1-year Reserved Instance for the specific instance type and operating system of the instances in the Auto Scaling group that the batch job uses.

Answer: C

Explanation:
This option is the most cost-effective solution because it leverages the Spot Instances, which are unused EC2 instances that are available at up to 90% discount compared to On-Demand prices. Spot Instances can be interrupted by AWS when the demand for On-Demand instances increases, but since the batch jobs are fault-tolerant and can be reprocessed by another instance, this is not a major issue. By using a launch template, the company can specify the configuration of the Spot Instances, such as the instance type, the operating system, and the user data. By using an Auto Scaling group, the company can automatically scale the number of Spot Instances based on the CPU usage, which reflects the load of the batch jobs. This way, the company can optimize the performance and the cost of the EC2 instances for the nightly batch jobs.
A: Purchase a 1-year Savings Plan for Amazon EC2 that covers the instance family of the Auto Scaling group that the batch job

uses. This option is not optimal because it requires a commitment to a consistent amount of compute usage per hour for a one-year term, regardless of the instance type, size, region, or operating system.
This can limit the flexibility and scalability of the Auto Scaling group and result in overpaying for unused compute capacity.
Moreover, Savings Plans do not provide a capacity reservation, which means the company still needs to reserve capacity with On-Demand Capacity Reservations and pay lower prices with Savings Plans.
B: Purchase a 1-year Reserved Instance for the specific instance type and operating system of the instances in the Auto Scaling group that the batch job uses. This option is not ideal because it requires a commitment to a specific instance configuration for a one-year term, which can reduce the flexibility and scalability of the Auto Scaling group and result in overpaying for unused compute capacity. Moreover, Reserved Instances do not provide a capacity reservation, which means the company still needs to reserve capacity with On-Demand Capacity Reservations and pay lower prices with Reserved Instances.
D: Create a new launch template for the Auto Scaling group Increase the instance size Set a policy to scale out based on CPU usage. This option is not cost-effective because it does not take advantage of the lower prices of Spot Instances. Increasing the instance size can improve the performance of the batch jobs, but it can also increase the cost of the On-Demand instances. Moreover, scaling out based on CPU usage can result in launching more instances than needed, which can also increase the cost of the system.
References:
1 Spot Instances - Amazon Elastic Compute Cloud
2 Launch templates - Amazon Elastic Compute Cloud
3 Auto Scaling groups - Amazon EC2 Auto Scaling
[4] Savings Plans - Amazon EC2 Reserved Instances and Other AWS Reservation Models

NEW QUESTION # 212
A company is deploying an application that processes streaming data in near-real time The company plans to use Amazon EC2 instances for the workload The network architecture must be configurable to provide the lowest possible latency between nodes Which combination of network solutions will meet these requirements? (Select TWO)

- A. Group the EC2 instances in separate accounts
- B. Run the EC2 instances in a cluster placement group
- C. Attach multiple elastic network interfaces to each EC2 instance
- D. Use Amazon Elastic Block Store (Amazon EBS) optimized instance types.
- E. Enable and configure enhanced networking on each EC2 instance

Answer: B,E

Explanation:
These options are the most suitable ways to configure the network architecture to provide the lowest possible latency between nodes. Option A enables and configures enhanced networking on each EC2 instance, which is a feature that improves the network performance of the instance by providing higher bandwidth, lower latency, and lower jitter. Enhanced networking uses single root I/O virtualization (SR-IOV) or Elastic Fabric Adapter (EFA) to provide direct access to the network hardware. You can enable and configure enhanced networking by choosing a supported instance type and a compatible operating system, and installing the required drivers. Option C runs the EC2 instances in a cluster placement group, which is a logical grouping of instanceswithin a single Availability Zone that are placed close together on the same underlying hardware.
Cluster placement groups provide the lowest network latency and the highest network throughput among the placement group options. You can run the EC2 instances in a cluster placement group by creating a placement group and launching the instances into it.
Option B is not suitable because grouping the EC2 instances in separate accounts does not provide the lowest possible latency between nodes. Separate accounts are used to isolate and organize resources for different purposes, such as security, billing, or compliance. However, they do not affect the network performance or proximity of the instances. Moreover, grouping the EC2 instances in separate accounts would incur additional costs and complexity, and it would require setting up cross-account networking and permissions.
Option D is not suitable because attaching multiple elastic network interfaces to each EC2 instance does not provide the lowest possible latency between nodes. Elastic network interfaces are virtual network interfaces that can be attached to EC2 instances to provide additional network capabilities, such as multiple IP addresses, multiple subnets, or enhanced security. However, they do not affect the network performance or proximity of the instances. Moreover, attaching multiple elastic network interfaces to each EC2 instance would consume additional resources and limit the instance type choices.
Option E is not suitable because using Amazon EBS optimized instance types does not provide the lowest possible latency between nodes. Amazon EBS optimized instance types are instances that provide dedicated bandwidth for Amazon EBS volumes, which are block storage volumes that can be attached to EC2 instances.
EBS optimized instance types improve the performance and consistency of the EBS volumes, but they do not affect the network performance or proximity of the instances. Moreover, using EBS optimized instance types would incur additional costs and may not

be necessary for the streaming data workload. References:
Enhanced networking on Linux
Placement groups
Elastic network interfaces
Amazon EBS-optimized instances

NEW QUESTION # 213
A solutions architect is designing a solution where users will De directed to a backup static error page it the primary website is unavailable The primary website's DNS records are hosted in Amazon Route 53 where their domain is pointing to an Application Load Balancer (ALB) Which configuration should the solutions architect use to meet the company's needs while minimizing changes and infrastructure overhead?

- A. Update the Route 53 record to use a latency-based routing policy Add the backup static error page hosted within an Amazon S3 bucket to the record so the traffic is sent to the most responsive endpoints
- B. Set up a Route 53 active-active configuration with the ALB and an Amazon EC2 instance hosting a static error page as endpoints Route 53 will only send requests to the instance if the health checks fail for the ALB
- C. Set up a Route 53 active-passive failover configuration Direct traffic to a static error page hosted within an Amazon S3 bucket when Route 53 health checks determine that the ALB endpoint is unhealthy
- D. Point a Route 53 alias record to an Amazon CloudFront distribution with the ALB as one of its origins Then, create custom error pages for the distribution

Answer: D

Explanation:
Explanation
Active-passive failover
Use an active-passive failover configuration when you want a primary resource or group of resources to be available the majority of the time and you want a secondary resource or group of resources to be on standby in case all the primary resources become unavailable. When responding to queries, Route 53 includes only the healthy primary resources. If all the primary resources are unhealthy, Route 53 begins to include only the healthy secondary resources in response to DNS queries.
To create an active-passive failover configuration with one primary record and one secondary record, you just create the records and specify Failover for the routing policy. When the primary resource is healthy, Route 53 responds to DNS queries using the primary record. When the primary resource is unhealthy, Route 53 responds to DNS queries using the secondary record.
How Amazon Route 53 averts cascading failures
As a first defense against cascading failures, each request routing algorithm (such as weighted and failover) has a mode of last resort. In this special mode, when all records are considered unhealthy, the Route 53 algorithm reverts to considering all records healthy. For example, if all instances of an application, on several hosts, are rejecting health check requests, Route 53 DNS servers will choose an answer anyway and return it rather than returning no DNS answer or returning an NXDOMAIN (non-existent domain) response. An application can respond to users but still fail health checks, so this provides some protection against misconfiguration. Similarly, if an application is overloaded, and one out of three endpoints fails its health checks, so that it's excluded from Route 53 DNS responses, Route 53 distributes responses between the two remaining endpoints.
If the remaining endpoints are unable to handle the additional load and they fail, Route 53 reverts to distributing requests to all three endpoints.
https://docs.aws.amazon.com/Route53/latest/DeveloperGuide/dns-failover-types.html
https://docs.aws.amazon.com/Route53/latest/DeveloperGuide/dns-failover-problems.html

NEW QUESTION # 214
.....

Associate □ to download for free □AWS-Solutions-Architect-Associate Latest Test Online

- Reliable AWS-Solutions-Architect-Associate Exam Sims □ Reliable AWS-Solutions-Architect-Associate Study Guide □ □ Study AWS-Solutions-Architect-Associate Demo □ Search for ➠ AWS-Solutions-Architect-Associate □ on 【 www.pdfvce.com 】 immediately to obtain a free download □Vce AWS-Solutions-Architect-Associate Download
- Valid Test AWS-Solutions-Architect-Associate Bootcamp □ Vce AWS-Solutions-Architect-Associate Download □ Study AWS-Solutions-Architect-Associate Demo □ Download 【 AWS-Solutions-Architect-Associate 】 for free by simply entering ➤ www.prepawaypdf.com □ website □AWS-Solutions-Architect-Associate Exam Overviews
- Vce AWS-Solutions-Architect-Associate Download □ Guaranteed AWS-Solutions-Architect-Associate Success □ Vce AWS-Solutions-Architect-Associate Download □ Open ⇛ www.pdfvce.com ⇚ enter 【 AWS-Solutions-Architect-Associate 】 and obtain a free download □Study AWS-Solutions-Architect-Associate Demo
- AWS-Solutions-Architect-Associate Certified Questions □ AWS-Solutions-Architect-Associate Discount □ Training AWS-Solutions-Architect-Associate Tools □ The page for free download of ▹ AWS-Solutions-Architect-Associate ◃ on { www.practicevce.com } will open immediately □New AWS-Solutions-Architect-Associate Test Tutorial
- Study AWS-Solutions-Architect-Associate Demo □ AWS-Solutions-Architect-Associate Exam Tutorial □ AWS-Solutions-Architect-Associate Exam Quizzes □ The page for free download of （ AWS-Solutions-Architect-Associate ） on □ www.pdfvce.com □ will open immediately □Valid Test AWS-Solutions-Architect-Associate Bootcamp
- Latest AWS-Solutions-Architect-Associate Exam Fee □ AWS-Solutions-Architect-Associate Exam Tutorial □ Composite Test AWS-Solutions-Architect-Associate Price □ Download [AWS-Solutions-Architect-Associate] for free by simply searching on ▸ www.practicevce.com ◂ □AWS-Solutions-Architect-Associate Certified Questions
- Reliable AWS-Solutions-Architect-Associate Study Guide □ New AWS-Solutions-Architect-Associate Test Tutorial □ New AWS-Solutions-Architect-Associate Test Blueprint □ Enter 【 www.pdfvce.com 】 and search for [AWS-Solutions-Architect-Associate] to download for free □AWS-Solutions-Architect-Associate Discount
- Quiz Amazon - Pass-Sure Reliable AWS-Solutions-Architect-Associate Real Exam □ Simply search for [AWS-Solutions-Architect-Associate] for free download on 【 www.dumpsquestion.com 】 □New AWS-Solutions-Architect-Associate Test Braindumps
- Hot Reliable AWS-Solutions-Architect-Associate Real Exam | High Pass-Rate AWS-Solutions-Architect-Associate: AWS Certified Solutions Architect - Associate (SAA-C03) 100% Pass □ “ www.pdfvce.com ” is best website to obtain ➤ AWS-Solutions-Architect-Associate □ for free download □New AWS-Solutions-Architect-Associate Test Tutorial
- Hot Reliable AWS-Solutions-Architect-Associate Real Exam | High Pass-Rate AWS-Solutions-Architect-Associate: AWS Certified Solutions Architect - Associate (SAA-C03) 100% Pass □ Easily obtain 《 AWS-Solutions-Architect-Associate 》 for free download through [www.practicevce.com] □Study AWS-Solutions-Architect-Associate Demo
- www.stes.tyc.edu.tw, www.stes.tyc.edu.tw, 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, bbs.t-firefly.com, 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, www.stes.tyc.edu.tw, www.stes.tyc.edu.tw, Disposable vapes

2026 Latest Actualtests4sure AWS-Solutions-Architect-Associate PDF Dumps and AWS-Solutions-Architect-Associate Exam Engine Free Share: https://drive.google.com/open?id=1LqenwxVfo8PhZ0DXYeykdB2pINXUToJ7