SAA-C03^{Q&As}

AWS Certified Solutions Architect - Associate (SAA-C03)

Pass Amazon SAA-C03 Exam with 100% Guarantee

Free Download Real Questions & Answers **PDF** and **VCE** file from:

https://www.passapply.com/saa-c03.html

100% Passing Guarantee
100% Money Back Assurance

Following Questions and Answers are all new published by Amazon
Official Exam Center

⚙ **Instant Download** After Purchase

⚙ **100% Money Back** Guarantee

⚙ **365 Days** Free Update

⚙ **800,000+** Satisfied Customers

QUESTION 1

A company runs an internal browser-based application The application runs on Amazon EC2 instances behind an Application Load Balancer. The instances run in an Amazon EC2 Auto Scaling group across multiple Availability Zones. The Auto Scaling group scales up to 20 instances during work hours but scales down to 2 instances overnight Staff are complaining that the application is very slow when the day begins although it runs well by mid-morning.

How should the scaling be changed to address the staff complaints and keep costs to a minimum\\'?

A. Implement a scheduled action that sets the desired capacity to 20 shortly before the office opens

B. Implement a step scaling action triggered at a lower CPU threshold, and decrease the cooldown period.

C. Implement a target tracking action triggered at a lower CPU threshold, and decrease the cooldown period.

D. Implement a scheduled action that sets the minimum and maximum capacity to 20 shortly before the office opens

Correct Answer: C

This option will scale up capacity faster in the morning to improve performance, but will still allow capacity to scale down during off hours. It achieves this as follows:

A target tracking action scales based on a CPU utilization target. By triggering at a lower CPU threshold in the morning, the Auto Scaling group will start scaling up sooner as traffic ramps up, launching instances before utilization gets too high

and impacts performance.

Decreasing the cooldown period allows Auto Scaling to scale more aggressively, launching more instances faster until the target is reached. This speeds up the ramp-up of capacity. However, unlike a scheduled action to set a fixed minimum/

maximum capacity, with target tracking the group can still scale down during off hours based on demand. This helps minimize costs.

QUESTION 2

A company hosts multiple production applications. One of the applications consists of resources from Amazon EC2, AWS Lambda, Amazon RDS, Amazon Simple Notification Service (Amazon SNS), and Amazon Simple Queue Service (Amazon SQS) across multiple AWS Regions. All company resources are tagged with a tag name of "application" and a value that corresponds to each application. A solutions architect must provide the quickest solution for identifying all of the tagged components.

Which solution meets these requirements?

A. Use AWS CloudTrail to generate a list of resources with the application tag.

B. Use the AWS CLI to query each service across all Regions to report the tagged components.

C. Run a query in Amazon CloudWatch Logs Insights to report on the components with the application tag.

D. Run a query with the AWS Resource Groups Tag Editor to report on the resources globally with the application tag.

Correct Answer: D

https://docs.aws.amazon.com/tag-editor/latest/userguide/tagging.html

QUESTION 3

A company hosts rts sialic website by using Amazon S3 The company wants to add a contact form to its webpage The contact form will have dynamic server-sKle components for users to input their name, email address, phone number and user message The company anticipates that there will be fewer than 100 site visits each month

Which solution will meet these requirements MOST cost-effectively?

A. Host a dynamic contact form page in Amazon Elastic Container Service (Amazon ECS) Set up Amazon Simple Email Service (Amazon SES) to connect to any third-party email provider.

B. Create an Amazon API Gateway endpoinl with an AWS Lambda backend that makes a call to Amazon Simple Email Service (Amazon SES)

C. Convert the static webpage to dynamic by deploying Amazon Ughtsail Use client-side scnpting to build the contact form Integrate the form with Amazon WorkMail

D. Create a Q micro Amazon EC2 instance Deploy a LAMP (Linux Apache MySQL. PHP/Perl/Python) stack to host the webpage Use client-side scripting to buiW the contact form Integrate the form with Amazon WorkMail

Correct Answer: B

QUESTION 4

A development team runs monthly resource-intensive tests on its general purpose Amazon RDS for MySQL DB instance with Performance Insights enabled. The testing lasts for 48 hours once a month and is the only process that uses the database. The team wants to reduce the cost of running the tests without reducing the compute and memory attributes of the DB instance.

Which solution meets these requirements MOST cost-effectively?

A. Stop the DB instance when tests are completed. Restart the DB instance when required.

B. Use an Auto Scaling policy with the DB instance to automatically scale when tests are completed.

C. Create a snapshot when tests are completed. Terminate the DB instance and restore the snapshot when required.

D. Modify the DB instance to a low-capacity instance when tests are completed. Modify the DB instance again when required.

Correct Answer: C

QUESTION 5

A company wants to use an Amazon RDS for PostgreSQL DB cluster to simplify time-consuming database

administrative tasks for production database workloads. The company wants to ensure that its database is highly available and will provide automatic failover support in most scenarios in less than 40 seconds. The company wants to offload reads off of the primary instance and keep costs as low as possible.

Which solution will meet these requirements?

A. Use an Amazon RDS Multi-AZ DB instance deployment. Create one read replica and point the read workload to the read replica.

B. Use an Amazon RDS Multi-AZ DB duster deployment Create two read replicas and point the read workload to the read replicas.

C. Use an Amazon RDS Multi-AZ DB instance deployment. Point the read workload to the secondary instances in the Multi-AZ pair.

D. Use an Amazon RDS Multi-AZ DB cluster deployment Point the read workload to the reader endpoint.

Correct Answer: D

The company wants high availability, automatic failover support in less than 40 seconds, read offloading from the primary instance, and cost-effectiveness. Answer D is the best choice for several reasons:

1.

 Amazon RDS Multi-AZ deployments provide high availability and automatic failover support.

2.

 In a Multi-AZ DB cluster, Amazon RDS automatically provisions and maintains a standby in a different Availability Zone. If a failure occurs, Amazon RDS performs an automatic failover to the standby, minimizing downtime.

3.

 The "Reader endpoint" for an Amazon RDS DB cluster provides load-balancing support for read-only connections to the DB cluster. Directing read traffic to the reader endpoint helps in offloading read operations from the primary instance.

[SAA-C03 PDF Dumps](link)   [SAA-C03 VCE Dumps](link)   [SAA-C03 Practice Test](link)