

Exam Questions SOA-C02

AWS Certified SysOps Administrator - Associate (SOA-C02)

<https://www.2passeasy.com/dumps/SOA-C02/>



NEW QUESTION 1

A SysOps administrator is creating two AWS CloudFormation templates. The first template will create a VPC with associated resources, such as subnets, route tables, and an internet gateway. The second template will deploy application resources within the VPC that was created by the first template. The second template should refer to the resources created by the first template.

How can this be accomplished with the LEAST amount of administrative effort?

- A. Add an export field to the outputs of the first template and import the values in the second template.
- B. Create a custom resource that queries the stack created by the first template and retrieves the required values.
- C. Create a mapping in the first template that is referenced by the second template.
- D. Input the names of resources in the first template and refer to those names in the second template as a parameter.

Answer: C

NEW QUESTION 2

A company is running a website on Amazon EC2 instances that are in an Auto Scaling group. When the website traffic increases, additional instances take several minutes to become available because of a long-running user data script that installs software. A SysOps administrator must decrease the time that is required for new instances to become available.

Which action should the SysOps administrator take to meet this requirement?

- A. Reduce the scaling thresholds so that instances are added before traffic increases.
- B. Purchase Reserved Instances to cover 100% of the maximum capacity of the Auto Scaling group.
- C. Update the Auto Scaling group to launch instances that have a storage optimized instance type.
- D. Use EC2 Image Builder to prepare an Amazon Machine Image (AMI) that has pre-installed software.

Answer: C

NEW QUESTION 3

A company is running a flash sale on its website. The website is hosted on burstable performance Amazon EC2 instances in an Auto Scaling group. The Auto Scaling group is configured to launch instances when the CPU utilization is above 70%.

A couple of hours into the sale, users report slow load times and error messages for refused connections. A SysOps administrator reviews Amazon CloudWatch metrics and notices that the CPU utilization is at 20% across the entire fleet of instances.

The SysOps administrator must restore the website's functionality without making changes to the network infrastructure. Which solution will meet these requirements?

- A. Activate unlimited mode for the instances in the Auto Scaling group.
- B. Implement an Amazon CloudFront distribution to offload the traffic from the Auto Scaling group.
- C. Move the website to a different AWS Region that is closer to the users.
- D. Reduce the desired size of the Auto Scaling group to artificially increase CPU average utilization.

Answer: C

NEW QUESTION 4

A company has launched a social media website that gives users the ability to upload images directly to a centralized Amazon S3 bucket. The website is popular in areas that are geographically distant from the AWS Region where the S3 bucket is located. Users are reporting that uploads are slow. A SysOps administrator must improve the upload speed.

What should the SysOps administrator do to meet these requirements?

- A. Create S3 access points in Regions that are closer to the users.
- B. Create an accelerator in AWS Global Accelerator for the S3 bucket.
- C. Enable S3 Transfer Acceleration on the S3 bucket.
- D. Enable cross-origin resource sharing (CORS) on the S3 bucket.

Answer: A

NEW QUESTION 5

An organization is running multiple applications for their customers. Each application is deployed by running a base AWS CloudFormation template that configures a new VPC. All applications are run in the same AWS account and AWS Region. A SysOps administrator has noticed that when trying to deploy the same AWS CloudFormation stack, it fails to deploy.

What is likely to be the problem?

- A. The Amazon Machine image used is not available in that region.
- B. The AWS CloudFormation template needs to be updated to the latest version.
- C. The VPC configuration parameters have changed and must be updated in the template.
- D. The account has reached the default limit for VPCs allowed.

Answer: D

NEW QUESTION 6

A company hosts a web application on an Amazon EC2 instance in a production VPC. Client connections to the application are failing. A SysOps administrator inspects the VPC flow logs and finds the following entry:

```
2 111122223333 eni-####> 192.0.2.15 203.0.113.56 40711 443 6 1 40 1418530010 1418530070 REJECT OK
```

What is a possible cause of these failed connections?

- A. A security group is denying traffic on port 443.
- B. The EC2 instance is shut down.
- C. The network ACL is blocking HTTPS traffic.

D. The VPC has no internet gateway attached.

Answer: A

NEW QUESTION 7

A company has a stateful web application that is hosted on Amazon EC2 instances in an Auto Scaling group. The instances run behind an Application Load Balancer (ALB) that has a single target group. The ALB is configured as the origin in an Amazon CloudFront distribution. Users are reporting random logouts from the web application. Which combination of actions should a SysOps administrator take to resolve this problem? (Choose two.)

- A. Change to the least outstanding requests algorithm on the ALB target group.
- B. Configure cookie forwarding in the CloudFront distribution cache behavior.
- C. Configure header forwarding in the CloudFront distribution cache behavior.
- D. Enable group-level stickiness on the ALB listener rule.
- E. Enable sticky sessions on the ALB target group.

Answer: CE

NEW QUESTION 8

A SysOps administrator is setting up an automated process to recover an Amazon EC2 instance in the event of an underlying hardware failure. The recovered instance must have the same private IP address and the same Elastic IP address that the original instance had. The SysOps team must receive an email notification when the recovery process is initiated. Which solution will meet these requirements?

- A. Create an Amazon CloudWatch alarm for the EC2 instance, and specify the StatusCheckFailed_Instance metri
- B. Add an EC2 action to the alarm to recover the instanc
- C. Add an alarm notification to publish a message to an AmazonSimple Notification Service (Amazon SNS) topi
- D. Subscribe the SysOps team email address to the SNS topic.
- E. Create an Amazon CloudWatch alarm for the EC2 instance, and specify the StatusCheckFailed_System metri
- F. Add an EC2 action to the alarm to recover the instanc
- G. Add an alarm notification to publish a message to an AmazonSimple Notification Service (Amazon SNS) topi
- H. Subscribe the SysOps team email address to the SNS topic.
- I. Create an Auto Scaling group across three different subnets in the same Availability Zone with a minimum, maximum, and desired size of 1. Configure the Auto Scaling group to use a launch template that specifies the private IP addressand the Elastic IP address
- J. Add an activity notification for the Auto Scaling group to send an email message to the SysOps team through Amazon Simple Email Service (Amazon SES).
- K. Create an Auto Scaling group across three Availability Zones with a minimum, maximum, and desired size of 1. Configure the Auto Scaling group to use a launch template that specifies the private IP address and the Elastic IP address.Add an activity notification for the Auto Scaling group to publish a message to an Amazon Simple Notification Service (Amazon SNS) topi
- L. Subscribe the SysOps team email address to the SNS topic.

Answer: A

NEW QUESTION 9

A company manages an application that uses Amazon ElastiCache for Redis with two extra-large nodes spread across two different Availability Zones. The company's IT team discovers that the ElastiCache for Redis cluster has 75% freeable memory. The application must maintain high availability. What is the MOST cost-effective way to resize the cluster?

- A. Decrease the number of nodes in the ElastiCache for Redis cluster from 2 to 1.
- B. Deploy a new ElastiCache for Redis cluster that uses large node type
- C. Migrate the data from the original cluster to the new cluste
- D. After the process is complete, shut down the original cluster.
- E. Deploy a new ElastiCache for Redis cluster that uses large node type
- F. Take a backup from the original cluster, and restore the backup in the new cluste
- G. After the process is complete, shut down the original cluster.
- H. Perform an online resizing for the ElastiCache for Redis cluste
- I. Change the node types from extra-large nodes to large nodes.

Answer: B

NEW QUESTION 10

A company uses AWS Organizations to manage multiple AWS accounts with consolidated billing enabled. Organization member account owners want the benefits of Reserved Instances (RIs) but do not want to share RIs with other accounts. Which solution will meet these requirements?

- A. Purchase RIs in individual member account
- B. Disable RI discount sharing in the management account.
- C. Purchase RIs in individual member account
- D. Disable RI discount sharing in the member accounts.
- E. Purchase RIs in the management accoun
- F. Disable RI discount sharing in the management account.
- G. Purchase RIs in the management accoun
- H. Disable RI discount sharing in the member accounts.

Answer: B

NEW QUESTION 10

A company's IT department noticed an increase in the spend of their developer AWS account. There are over 50 developers using the account, and the finance team wants to determine the service costs incurred by each developer. What should a SysOps administrator do to collect this information? (Choose two.)

- A. Activate the createdBy tag in the account.
- B. Analyze the usage with Amazon CloudWatch dashboards.
- C. Analyze the usage with Cost Explorer.
- D. Configure AWS Trusted Advisor to track resource usage.
- E. Create a billing alarm in AWS Budgets.

Answer: AC

NEW QUESTION 12

A SysOps Administrator is troubleshooting Amazon EC2 connectivity issues to the internet. The EC2 instance is in a private subnet. Below is the route table that is applied to the subnet of the EC2 instance.

Destination – 10.2.0.0/16

Target – local

Status – Active

Propagated – No

Destination – 0.0.0.0/0

Target – nat-xxxxxxx

Status – Blackhole

Propagated – No

What has caused the connectivity issue?

- A. The NAT gateway no longer exists
- B. There is no route to the internet gateway.
- C. The routes are no longer propagating.
- D. There is no route rule with a destination for the internet.

Answer: A

NEW QUESTION 17

Your new application is hosted in an Auto Scaling group of EC2 instances. To improve the monitoring process, you have to configure it to keep the average aggregate CPU utilization of your Auto Scaling group at 50 percent. This should be done by specifying the scaling metrics and threshold values for the CloudWatch alarms that trigger the scaling process.

Which of the following scaling policy type you should use?

- A. Simple scaling
- B. Target tracking scaling
- C. Step scaling
- D. Threshold scaling

Answer: B

Explanation:

Target tracking scaling is the correct answer. With target tracking scaling policies, you select a scaling metric and set a target value. Amazon EC2 Auto Scaling creates and manages the CloudWatch alarms that trigger the scaling policy and calculates the scaling adjustment based on the metric and the target value. The scaling policy adds or removes capacity as required to keep the metric at, or close to, the specified target value. In addition to keeping the metric close to the target value, a target tracking scaling policy also adjusts to changes in the metric due to a changing load pattern.

For example, you can use target tracking scaling to:

Configure a target tracking scaling policy to keep the average aggregate CPU utilization of your Auto Scaling group at 50 percent.

Configure a target tracking scaling policy to keep the request count per target of your Application Load Balancer target group at 1000 for your Auto Scaling group.

Step scaling policies and simple scaling policies are incorrect. Step scaling policies and simple scaling policies are two of the dynamic scaling options available for you to use. Both require you to create CloudWatch alarms for the scaling policies. Both require you to specify the high and low thresholds for the alarms. Both require you to define whether to add or remove instances, and how many, or set the group to an exact size.

The main difference between the policy types is the step adjustments that you get with step scaling policies. When step adjustments are applied, and they increase or decrease the current capacity of your Auto Scaling group, the adjustments vary based on the size of the alarm breach.

In most cases, step scaling policies are a better choice than simple scaling policies, even if you have only a single scaling adjustment.

Threshold scaling is incorrect as it is a fictitious scaling policy type.

NEW QUESTION 21

Which of the following AWS service is a security management service which allows you to centrally configure and manage firewall rules across your accounts and applications in AWS Organization?

- A. AWS Shield
- B. AWS Secrets Manager
- C. AWS WAF
- D. AWS Firewall Manager

Answer: D

Explanation:

AWS Firewall Manager is the correct answer. AWS Firewall Manager is a security management service that allows you to centrally configure and manage firewall rules across your accounts and applications

in AWS Organization. As new applications are created, Firewall Manager makes it easy to bring new applications and resources into compliance by enforcing a common set of security rules. Now you have a single service to build firewall rules, create security policies, and enforce them in a consistent, hierarchical manner across your entire infrastructure.

Using AWS Firewall Manager, you can easily roll out AWS WAF rules for your Application Load Balancers, API Gateways, and Amazon CloudFront distributions. Similarly, you can create AWS Shield Advanced protections for your Application Load Balancers, ELB Classic Load Balancers, Elastic IP Addresses and CloudFront distributions. Finally, with AWS Firewall Manager, you can enable security groups for your Amazon EC2 and ENI resource types in Amazon VPCs.

Benefits

- * 1. Simplify management of firewall rules across your accounts
- * 2. Ensure compliance of existing and new applications
- * 3. Easily deploy managed rules across accounts
- * 4. Enable rapid response to internet attacks

AWS Secrets Manager is incorrect. AWS Secrets Manager helps you to securely encrypt, store, and retrieve credentials for your databases and other services. Instead of hardcoding credentials in your apps, you can make calls to Secrets Manager to retrieve your credentials whenever needed. Secrets Manager helps you protect access to your IT resources and data by enabling you to rotate and manage access to your secrets.

AWS Shield is incorrect. AWS provides two levels of protection against DDoS attacks: AWS Shield Standard and AWS Shield Advanced. AWS Shield Standard is automatically included at no extra cost beyond what you already pay for AWS WAF and your other AWS services. For added protection against DDoS attacks, AWS offers AWS Shield Advanced.

AWS WAF is incorrect. AWS WAF is a web application firewall that lets you monitor web requests that are forwarded to Amazon CloudFront distributions or an Application Load Balancer. You can also use AWS WAF to block or allow requests based on conditions that you specify, such as the IP addresses that requests originate from or values in the requests.

NEW QUESTION 26

A company for compliance purposes needs to assess how well its resource configurations comply with internal practices, industry guidelines, and regulations. Which tool should a SysOps administrator use to meet these requirements?

- A. AWS Security Hub
- B. AWS Shield
- C. AWS Health
- D. AWS Config

Answer: D

Explanation:

AWS Config is the correct answer. AWS Config can be used to assess how well your resource configurations comply with internal practices, industry guidelines, and regulations.

AWS Security Hub is incorrect. AWS Security Hub provides you with a comprehensive view of your security state in AWS and helps you check your environment against security industry standards and best practices.

Security Hub collects security data from across AWS accounts, services, and supported third-party partner products and helps you analyze your security trends and identify the highest priority security issues.

AWS Shield is incorrect. AWS provides two levels of protection against DDoS attacks: AWS Shield Standard and AWS Shield Advanced. AWS Shield Standard is automatically included at no extra cost beyond what you already pay for AWS WAF and your other AWS services.

For added protection against DDoS attacks, AWS offers AWS Shield Advanced. AWS Shield Advanced provides expanded DDoS attack protection for your Amazon EC2 instances, Elastic Load Balancing load balancers, Amazon CloudFront distributions, and Amazon Route 53 hosted zones.

AWS Health is incorrect. AWS Health provides personalized information about events that can affect your AWS infrastructure, guides you through scheduled changes, and accelerates the troubleshooting of issues that affect your AWS resources and accounts.

NEW QUESTION 30

Which of the following recommendations is NOT considered a best practice for using AWS CloudFormation more effectively and securely throughout its entire workflow?

- A. Reuse templates to replicate stacks in multiple environments
- B. Use nested stacks to reuse common template patterns
- C. Embed credentials in your templates
- D. Use IAM to control access

Answer: C

Explanation:

Embed credentials in your templates is the correct answer as it is not considered a best practice for using AWS CloudFormation effectively.

Best practices are recommendations that can help you use AWS CloudFormation more effectively and securely throughout its entire workflow. Learn how to plan and organize your stacks, create templates that describe your resources and the software applications that run on them, and manage your stacks and their resources. The following best practices are based on real-world experience from current AWS CloudFormation customers.

- * 1. Organize your stacks by lifecycle and ownership

Use the lifecycle and ownership of your AWS resources to help you decide what resources should go in each stack. Initially, you might put all your resources in one stack, but as your stack grows in scale and broadens in scope, managing a single stack can be cumbersome and time-consuming.

- * 2. Use IAM to control access

IAM is an AWS service that you can use to manage users and their permissions in AWS. You can use IAM with AWS CloudFormation to specify what AWS CloudFormation actions users can perform, such as viewing stack templates, creating stacks, or deleting stacks.

* 3. Verify quotas for all resource types

Before launching a stack, ensure that you can create all the resources that you want without hitting your AWS account limits. If you hit a limit, AWS CloudFormation won't create your stack successfully until you increase your quota or delete extra resources.

* 4. Reuse templates to replicate stacks in multiple environments

After you have your stacks and resources set up, you can reuse your templates to replicate your infrastructure in multiple environments. For example, you can create environments for development, testing, and production so that you can test changes before implementing them into production.

* 5. Do not embed credentials in your templates

Rather than embedding sensitive information in your AWS CloudFormation templates, we recommend you use dynamic references in your stack template. Dynamic references provide a compact, powerful way for you to reference external values that are stored and managed in other services, such as the AWS Systems Manager Parameter Store or AWS Secrets Manager.

NEW QUESTION 35

Suppose you have ELB load balancers in the US West (Oregon) Region and in the Asia Pacific (Singapore) Region and you created a latency record for each load balancer. What will happen when a user in London enters the name of your domain in a browser? (Choose all that apply.)

- A. If latency is lower between the London and Oregon regions, Route 53 responds to the query with the IP address for the Singapore load balancer
- B. If latency is lower between the London and Oregon regions, Route 53 responds to the query with the IP address for the Oregon load balancer
- C. DNS routes the query to a Route 53 name server
- D. Route 53 refers to its data on latency ONLY between London and the Singapore region
- E. Route 53 refers to its data on latency between London and the Singapore region and between London and the Oregon region

Answer: BCE

Explanation:

Explanation/Reference:

The correct answers are:

* 1. DNS routes the query to a Route 53 name server

* 2. Route 53 refers to its data on latency between London and the Singapore region and between London and the Oregon region

* 3. If latency is lower between the London and Oregon regions, Route 53 responds to the query with the IP address for the Oregon load balancer

If your application is hosted in multiple AWS Regions, you can improve performance for your users by serving their requests from the AWS Region that provides the lowest latency.

To use latency-based routing, you create latency records for your resources in multiple AWS Regions. When Route 53 receives a DNS query for your domain or subdomain (example.com or acme.example.com), it determines which AWS Regions you've created latency records for, determines which region gives the user the lowest latency, and then selects a latency record for that region. Route 53 responds with the value from the selected record, such as the IP address for a web server.

For example, suppose you have ELB load balancers in the US West (Oregon) Region and in the Asia Pacific (Singapore) Region. You created a latency record for each load balancer. Here's what happens when a user in London enters the name of your domain in a browser:

* 1. DNS routes the query to a Route 53 name server.

* 2. Route 53 refers to its data on latency between London and the Singapore region and between London and the Oregon region.

* 3. If latency is lower between the London and Oregon regions, Route 53 responds to the query with the IP address for the Oregon load balancer. If latency is lower between London and the Singapore region, Route 53 responds with the IP address for the Singapore load balancer.

NEW QUESTION 37

- A. Mastered
- B. Not Mastered

Answer: A

NEW QUESTION 40

.....

THANKS FOR TRYING THE DEMO OF OUR PRODUCT

Visit Our Site to Purchase the Full Set of Actual SOA-C02 Exam Questions With Answers.

We Also Provide Practice Exam Software That Simulates Real Exam Environment And Has Many Self-Assessment Features. Order the SOA-C02 Product From:

<https://www.2passeasy.com/dumps/SOA-C02/>

Money Back Guarantee

SOA-C02 Practice Exam Features:

- * SOA-C02 Questions and Answers Updated Frequently
- * SOA-C02 Practice Questions Verified by Expert Senior Certified Staff
- * SOA-C02 Most Realistic Questions that Guarantee you a Pass on Your First Try
- * SOA-C02 Practice Test Questions in Multiple Choice Formats and Updates for 1 Year