



# SOA-C02<sup>Q&As</sup>

AWS Certified SysOps Administrator - Associate (SOA-C02)

## Pass Amazon SOA-C02 Exam with 100% Guarantee

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

<https://www.passapply.com/soa-c02.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 SysOps administrator is setting up a fleet of Amazon EC2 instances in an Auto Scaling group for an application. The fleet should have 50% CPU available at that times to accommodate bursts of traffic. The load will increase significantly between the hours of

09:00 and 17:00, 7 days a week.

How should the SysOps administrator configure the scaling of the EC2 instances to meet these requirements?

- A. Create a target tracking scaling policy that runs when the CPU utilization is higher than 90%
- B. Create a target tracking scaling policy that runs when the CPU utilization is higher than 50%. Create a scheduled scaling policy that ensures that the fleet is available at 09:00 Create a second scheduled scaling policy that scales in the fleet at 17:00
- C. Set the Auto Scaling group to start with 2 instances by setting the desired instances maximum instances, and minimum instances to 2 Create a scheduled scaling policy that ensures that the fleet is available at 09:00
- D. Create a scheduled scaling policy that ensures that the fleet is available at 09:00. Create a second scheduled scaling policy that scales in the fleet at 17:00

Correct Answer: B

Create a target tracking scaling policy that runs when the CPU utilization is higher than 50%: This ensures that the fleet will have 50% CPU available at all times to accommodate bursts of traffic.

Create a scheduled scaling policy that ensures that the fleet is available at 09:00: By scheduling the scaling action, the SysOps administrator can ensure that the number of instances is increased before the load increases significantly at

09:00.

Create a second scheduled scaling policy that scales in the fleet at 17:00: Similarly, this second scheduled scaling policy allows scaling in the fleet after the peak hours are over at 17:00, which helps in optimizing costs during low traffic periods.

---

### QUESTION 2

A company is managing many accounts by using a single organization in AWS Organizations. The organization has all features enabled. The company wants to turn on AWS Config in all the accounts of the organization and in all AWS Regions.

What should a Sysops administrator do to meet these requirements in the MOST operationally efficient way?

- A. Use AVVS CloudFormation StackSets to deploy stack instances that turn on AWS Config in all accounts and in all Regions.
- B. Use AWS CloudFormation StackSets to deploy stack policies that turn on AWS Config in all accounts and in all Regions.
- C. Use service control policies (SCPs) to configure AWS Config in all accounts and in all Regions.
- D. Create a script that uses the AWS CLI to turn on AWS Config in all accounts in the organization. Run the script from



the organization\\s management account.

Correct Answer: A

Option A (Use AWS CloudFormation Stack Sets to deploy stack instances that turn on AWS Config in all accounts and in all Regions) is the best option for achieving the goal in a highly efficient and scalable manner. AWS CloudFormation Stack Sets allow you to deploy CloudFormation stacks across multiple accounts and Regions in a single operation. By defining a stack that enables AWS Config and using Stack Sets, you can easily enable AWS Config in all accounts and Regions within the AWS organization.

Option B (Use AWS CloudFormation Stack Sets to deploy stack policies that turn on AWS Config in all accounts and in all Regions) is not the correct option. Stack policies are used to control what actions can be performed on a CloudFormation stack, but they are not suitable for enabling AWS Config across multiple accounts and Regions.

---

### QUESTION 3

A company has two VPC networks named VPC A and VPC B. The VPC A CIDR block is 10.0.0.0/16 and the VPC B CIDR block is 172.31.0.0/16. The company wants to establish a VPC peering connection named pcx-12345 between both VPCs.

Which rules should appear in the route table of VPC A after configuration? (Select TWO.)

- A. Destination: 10.0.0.0/16, Target: Local
- B. Destination: 172.31.0.0/16, Target: Local
- C. Destination: 10.0.0.0/16, Target: pcx-12345
- D. Destination: 172.31.0.0/16, Target: pcx-12345
- E. Destination: 10.0.0.0/16, Target: 172.31.0.0/16

Correct Answer: AD

<https://docs.aws.amazon.com/vpc/latest/peering/vpc-peering-routing.html>

---

### QUESTION 4

A company runs an application on Amazon EC2 instances. The EC2 instances are in an Auto Scaling group and run behind an Application Load Balancer (ALB). The application experiences errors when total requests exceed 100 requests per second. A SysOps administrator must collect information about total requests for a 2-week period to determine when requests exceeded this threshold.

What should the SysOps administrator do to collect this data?

- A. Use the ALB\\s RequestCount metric. Configure a time range of 2 weeks and a period of 1 minute. Examine the chart to determine peak traffic times and volumes.
- B. Use Amazon CloudWatch metric math to generate a sum of request counts for all the EC2 instances over a 2-week period. Sort by a 1-minute interval.
- C. Create Amazon CloudWatch custom metrics on the EC2 launch configuration templates to create aggregated request metrics across all the EC2 instances.



D. Create an Amazon EventBridge (Amazon CloudWatch Events) rule. Configure an EC2 event matching pattern that creates a metric that is based on EC2 requests. Display the data in a graph.

Correct Answer: A

The ALB's RequestCount metric provides the total number of requests processed by the ALB. By configuring a time range of 2 weeks with a 1-minute period, you can collect detailed data on request counts for each minute over the 2-week period. This will allow the SysOps administrator to visualize the traffic patterns and identify peak times when the requests exceeded the threshold of 100 requests per second.

---

## QUESTION 5

A company's architecture team must receive immediate email notification whenever new Amazon EC2 instances are launched in the company's main AWS production account.

What should a SysOps administrator do to meet this requirement?

A. Create a user data script that sends an email message through a smart host connector. Include the architecture team's email address in the user data script as the recipient. Ensure that all new EC2 instances include the user data script as part of a standardized build process.

B. Create an Amazon Simple Notification Service (Amazon SNS) topic and a subscription that uses the email protocol. Enter the architecture team's email address as the subscriber. Create an Amazon EventBridge rule that reacts when EC2 instances are launched. Specify the SNS topic as the rule's target.

C. Create an Amazon Simple Queue Service (Amazon SQS) queue and a subscription that uses the email protocol. Enter the architecture team's email address as the subscriber. Create an Amazon EventBridge rule that reacts when EC2 instances are launched. Specify the SQS queue as the rule's target.

D. Create an Amazon Simple Notification Service (Amazon SNS) topic. Configure AWS Systems Manager to publish EC2 events to the SNS topic. Create an AWS Lambda function to poll the SNS topic. Configure the Lambda function to send any messages to the architecture team's email address.

Correct Answer: B

Should use Eventbridge rule -> SNS -> send email

[Latest SOA-C02 Dumps](#)

[SOA-C02 PDF Dumps](#)

[SOA-C02 VCE Dumps](#)