

Demo Questions

Amazon SOA-C01 Exam

AWS Certified SysOps Administrator – Associate

Thank you for downloading SOA-C01 Exam PDF

Question #1 Topic 1

You are currently hosting multiple applications in a VPC and have logged numerous port scans coming in from a specific IP address block. Your security team has requested that all access from the offending IP address block be denied for the next 24 hours.

Which of the following is the best method to quickly and temporarily deny access from the specified IP address block?

- A. Create an AD policy to modify Windows Firewall settings on all hosts in the VPC to deny access from the IP address block
- B. Modify the Network ACLs associated with all public subnets in the VPC to deny access from the IP address block
- C. Add a rule to all of the VPC 5 Security Groups to deny access from the IP address block
- D. Modify the Windows Firewall settings on all Amazon Machine Images (AMIs) that your organization uses in

Correct Answer: B

http://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/VPC_SecurityGroups.html

Question #2 Topic 1

When preparing for a compliance assessment of your system built inside of AWS. what are three best-practices for you to prepare for an audit? (Choose three.)

- A. Gather evidence of your IT operational controls
- B. Request and obtain applicable third-party audited AWS compliance reports and certifications
- C. Request and obtain a compliance and security tour of an AWS data center for a pre-assessment security review
- D. Request and obtain approval from AWS to perform relevant network scans and in-depth penetration tests of your system's Instances and endpoints
- E. Schedule meetings with AWS's third-party auditors to provide evidence of AWS compliance that maps to

Correct Answer: *ABD*

Question #3*Topic 1*

You have started a new job and are reviewing your company's infrastructure on AWS You notice one web application where they have an Elastic Load Balancer (&B) in front of web instances in an Auto Scaling Group

When you check the metrics for the ELB in CloudWatch you see four healthy instances in Availability Zone

(AZ) A and zero in AZ B There are zero unhealthy instances.

What do you need to fix to balance the instances across AZs?

- A. Set the ELB to only be attached to another AZ
- B. Make sure Auto Scaling is configured to launch in both AZs
- C. Make sure your AMI is available in both AZs
- D. Make sure the maximum size of the Auto Scaling Group is greater than 4

Correct Answer: *B*

Question #4*Topic 1*

You have been asked to leverage Amazon VPC BC2 and SOS to implement an application that submits and receives millions of messages per second to a message queue. You want to ensure your application has sufficient bandwidth between your EC2 instances and SQS

Which option will provide the most scalable solution for communicating between the application and SQS?

- A. Ensure the application instances are properly configured with an Elastic Load Balancer

- B. Ensure the application instances are launched in private subnets with the EBS-optimized option enabled
- C. Ensure the application instances are launched in public subnets with the associate-public-IP-address=true option enabled
- D. Launch application instances in private subnets with an Auto Scaling group and Auto Scaling triggers

Correct Answer: D

Bandwidth literally means network not IO Bandwidth. Having alerts to scale the Autoscaling is most sophisticated option.

Question #5 Topic 1

You have identified network throughput as a bottleneck on your m1.small EC2 instance when uploading data into Amazon S3 in the same region. How do you remedy this situation?

- A. Add an additional ENI
- B. Change to a larger Instance
- C. Use DirectConnect between EC2 and S3
- D. Use EBS PIOPS on the local volume

Correct Answer: B

https://media.amazonwebservices.com/AWS_Amazon_EMR_Best_Practices.pdf

Question #6 Topic 1

When attached to an Amazon VPC, which two components provide connectivity with external networks? (Choose two.)

- A. Elastic IP (EIP)
- B. NAT Gateway (NAT)
- C. Internet Gateway (IGW)
- D. Virtual Private Gateway (VGW)

Correct Answer: *CD*

Question #7*Topic 1*

Your application currently leverages AWS Auto Scaling to grow and shrink as load increases/decreases and has been performing well. Your marketing team expects a steady ramp up in traffic to follow an upcoming campaign that will result in a 20x growth in traffic over 4 weeks. Your forecast for the approximate number of Amazon EC2 instances necessary to meet the peak demand is 175. What should you do to avoid potential service disruptions during the ramp up in traffic?

- A. Ensure that you have pre-allocated 175 Elastic IP addresses so that each server will be able to obtain one as it launches
- B. Check the service limits in Trusted Advisor and adjust as necessary so the forecasted count remains within limits.
- C. Change your Auto Scaling configuration to set a desired capacity of 175 prior to the launch of the marketing campaign
- D. Pre-warm your Elastic Load Balancer to match the requests per second anticipated during peak demand

Correct Answer: *D*

Amazon ELB is able to handle the vast majority of use cases for our customers without requiring "pre-warming" (configuring the load balancer to have the appropriate level of capacity based on expected traffic). <https://aws.amazon.com/articles/1636185810492479#pre-warming>

Question #8*Topic 1*

You have an Auto Scaling group associated with an Elastic Load Balancer (ELB). You have noticed that instances launched via the Auto Scaling group are being marked unhealthy due to an ELB health check, but these unhealthy instances are not being terminated. What do you need to do to ensure instances marked unhealthy by the ELB will be terminated and replaced?

- A. Change the thresholds set on the Auto Scaling group health check
- B. Add an Elastic Load Balancing health check to your Auto Scaling group
- C. Increase the value for the Health check interval set on the Elastic Load Balancer
- D. Change the health check set on the Elastic Load Balancer to use TCP rather than HTTP checks

Correct Answer: *B*

<http://docs.aws.amazon.com/AutoScaling/latest/DeveloperGuide/as-add-elb-healthcheck.html>

Add an Elastic Load Balancing Health Check to your Auto Scaling Group

By default, an Auto Scaling group periodically reviews the results of EC2 instance status to determine the health state of each instance. However, if you have associated your Auto Scaling group with an Elastic Load

Balancing load balancer, you can choose to use the Elastic Load Balancing health check. In this case, Auto

Scaling determines the health status of your instances by checking the results of both the EC2 instance status check and the Elastic Load Balancing instance health check.

For information about EC2 instance status checks, see

Monitor Instances With Status Checks

in the Amazon

EC2 User Guide for Linux Instances. For information about Elastic Load Balancing health checks, see

Health -

Check -

in the Elastic Load Balancing Developer Guide.

This topic shows you how to add an Elastic Load Balancing health check to your Auto Scaling group, assuming that you have created a load balancer and have registered the load balancer with your Auto Scaling group. If you have not registered the load balancer with your Auto Scaling group, see

Set Up a Scaled and -

Load-Balanced Application -

.

Auto Scaling marks an instance unhealthy if the calls to the Amazon EC2 action

DescribeInstanceStatus -

return any state other than running, the system status shows impaired, or the calls to Elastic Load Balancing action

DescribeInstanceHealth -

returns OutOfService in the instance state field.

If there are multiple load balancers associated with your Auto Scaling group, Auto Scaling checks the health state of your EC2 instances by making health check calls to each load balancer. For each call, if the Elastic

Load Balancing action returns any state other than InService, the instance is marked as unhealthy.

After Auto

Scaling marks an instance as unhealthy, it remains in that state, even if subsequent calls from other load balancers return an InService state for the same instance.

Question #9 Topic 1

Which two AWS services provide out-of-the-box user configurable automatic backup-as-a-service and backup rotation options? (Choose two.)

- A. Amazon S3
- B. Amazon RDS
- C. Amazon EBS
- D. Amazon Red shift

Correct Answer: BD

By default, and at no additional charge, Amazon RDS enables automated backups of your DB Instance with a 1-day retention period.

By default, Amazon Redshift enables automated backups of your data warehouse cluster with a 1-day retention period.

Question #10 Topic 1

An organization has configured a VPC with an Internet Gateway (IGW), pairs of public and private subnets (each with one subnet per Availability Zone), and an Elastic Load Balancer (ELB) configured to use the public subnets. The application's web tier leverages the ELB. Auto Scaling and a multi-AZ RDS database instance

The organization would like to eliminate any potential single points of failure in this design.

What step should you take to achieve this organization's objective?

- A. Nothing, there are no single points of failure in this architecture.
- B. Create and attach a second IGW to provide redundant internet connectivity.
- C. Create and configure a second Elastic Load Balancer to provide a redundant load balancer.
- D. Create a second multi-AZ RDS instance in another Availability Zone and configure replication to provide a

Correct Answer: A

You need multiple ELB if you want HA across regions.

"AWS Load Balancer – Cross Network

Many times it happens that after setting up your ELB, you experience significant drops in your performance.

The best way to handle this situation is to start with identifying whether your ELB is single AZ or multiple AZ, as single AZ ELB is also considered as one of the Single Points of Failures on AWS Cloud. Once you identify your ELB, it is necessary to make sure ELB loads are kept cross regions."

<https://www.botmetric.com/blog/eliminating-single-points-of-failures-on-aws-cloud/>

Question #11 Topic 1

Which of the following are characteristics of Amazon VPC subnets? (Choose two.)

- A. Each subnet maps to a single Availability Zone
- B. A CIDR block mask of /25 is the smallest range supported
- C. Instances in a private subnet can communicate with the internet only if they have an Elastic IP.
- D. By default, all subnets can route between each other, whether they are private or public
- E. V Each subnet spans at least 2 Availability zones to provide a high-availability environment

Correct Answer: AD

"Each subnet must reside entirely within one Availability Zone and cannot span zones."

"Every subnet that you create is automatically associated with the main route table for the VPC."

http://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/VPC_Subnets.html

Question #12 Topic 1

You are creating an Auto Scaling group whose Instances need to insert a custom metric into CloudWatch.

Which method would be the best way to authenticate your CloudWatch PUT request?

- A. Create an IAM role with the Put MetricData permission and modify the Auto Scaling launch configuration to launch instances in that role
- B. Create an IAM user with the PutMetricData permission and modify the Auto Scaling launch configuration to inject the userscredentials into the instance User Data
- C. Modify the appropriate Cloud Watch metric policies to allow the Put MetricData permission to instances from the Auto Scaling group
- D. Create an IAM user with the PutMetricData permission and put the credentials in a private repository and

Correct Answer: A

Creates an IAM role is always the best practice to give permissions to EC2 instances in order to interact with other AWS services

Question #13 Topic 1

When an EC2 instance that is backed by an S3-based AMI is terminated, what happens to the data on the root volume?

- A. Data is automatically saved as an EBS volume.
- B. Data is automatically saved as an EBS snapshot.
- C. Data is automatically deleted.
- D. Data is unavailable until the instance is restarted.

Correct Answer: C

We recommend that you use AMIs backed by Amazon EBS, because they launch faster and use persistent storage.

<http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/RootDeviceStorage.html#choose-an-ami-by-root-device>

Question #14 Topic 1

You have a web application leveraging an Elastic Load Balancer (ELB) in front of the web servers deployed using an Auto Scaling Group. Your database is running on Relational Database Service (RDS). The application serves out technical articles and responses to them. In general, there are more views of an article than there are responses to the article. On occasion, an article on the site becomes extremely popular, resulting in significant traffic increases that cause the site to go down.

What could you do to help alleviate the pressure on the infrastructure while maintaining availability during these events? (Choose three.)

- A. Leverage CloudFront for the delivery of the articles.
- B. Add RDS read-replicas for the read traffic going to your relational database.
- C. Leverage ElastiCache for caching the most frequently used data.
- D. Use SQS to queue up the requests for the technical posts and deliver them out of the queue.
- E. Use Route53 health checks to fail over to an S3 bucket for an error page.

Correct Answer: ABC

E is wrong, we are not using multiple regions, Route53 ELB-failover is not required.

D is wrong, SQS will only add to the strain on the server, and is not useful in this situation.

A is correct, as you can have users read from edge locations.

B is correct, as you can service reads with read-replicas.

C is correct, ElastiCache would assist in this situation.

Question #15 Topic 1

The majority of your Infrastructure is on premises and you have a small footprint on AWS Your company has decided to roll out a new application that is heavily dependent on low latency connectivity to LOAP for authentication Your security policy requires minimal changes to the company's existing application user management processes.

What option would you implement to successfully launch this application1?

- A. Create a second, independent LOAP server in AWS for your application to use for authentication
- B. Establish a VPN connection so your applications can authenticate against your existing on-premises LDAP servers
- C. Establish a VPN connection between your data center and AWS create a LDAP replica on AWS and configure your application to use the LDAP replica for authentication
- D. Create a second LDAP domain on AWS establish a VPN connection to establish a trust relationship

Correct Answer: C

Create read replica(RODC) of main LDAP server so that LDAP read replica or RODC can authenticate with application locally.

Creating new domain and trust relationship would require lot of work and changes in exiting ldap configuration so D cannot be answer here.