

## Amazon

SOA-C01

AWS Certified SysOps Administrator – Associate



## Question: 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 that VPC to deny access from the IP address block

**Answer: B**

Explanation:

Reference:

[http://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/VPC\\_SecurityGroups.html](http://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/VPC_SecurityGroups.html)

## Question: 2

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 3 answers

- 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 your control objectives

**Answer: A, B, D**

## Question: 3

---

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 (ELB) 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

**Answer: B**

### Question: 4

You have been asked to leverage Amazon VPC, EC2, and SQS 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 configured to watch the SQS queue size

**Answer: D**

Explanation:

The question is about most “scalable solution for communicating” for SQS that is parallel processing of SQS messages.

See

also:

–

<https://aws.amazon.com/articles/1464>

– <http://docs.aws.amazon.com/AWSSimpleQueueService/latest/SQSDeveloperGuide/throughput.html>

### Question: 5

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

**Answer: B**

Explanation:

Reference:

[https://media.amazonwebservices.com/AWS\\_Amazon\\_EMR\\_Best\\_Practices.pdf](https://media.amazonwebservices.com/AWS_Amazon_EMR_Best_Practices.pdf)

## Question: 6

When attached to an Amazon VPC which two components provide connectivity with external networks?  
Choose 2 answers

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

**Answer: C, D**

## Question: 7

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 prior to the marketing campaign

**Answer: B**

Explanation:

As the EC2 limit per region is max 20. You will need to fill an Amazon EC2 instance request form to increase the EC2 instances to 175.

[http://aws.amazon.com/ec2/faqs/#How many instances can I run in Amazon EC2](http://aws.amazon.com/ec2/faqs/#How_many_instances_can_I_run_in_Amazon_EC2)

I don't think the answer can be D, as the question says "expects a steady ramp up in traffic to follow an upcoming campaign that will result in a 20x growth in traffic over 4 weeks". To pre-warm your ELB, you have to put in a request to AWS. You can't do it.

Q: How do I reserve capacity for an existing, running instance?

To reserve capacity for a running instance, you can purchase a Reserved Instance or modify an existing reservation so it matches your instance's specifications. You can purchase Reserved Instances via the Amazon EC2 Console or by using the `PurchaseReservedInstancesOffering` API. You can modify existing Reserved Instances via the Amazon EC2 Console or by using the `ModifyReservedInstances` API call.

In both cases, the reservation must match the following attributes of the running instance you want to cover:

Availability Zone (e.g., us-east-1a)

Instance type (e.g., m3.large)

Platform (e.g., Linux/UNIX (Amazon VPC))

Tenancy (e.g., default)

Q: How do I control which instances are billed at the lower rate?

The `RunInstances` API command does not distinguish between On-Demand instances and the reservations that can be applied to them. When computing your bill, our system will automatically optimize which instances are charged at the lower rate to ensure you always pay the lowest amount.

For information about hourly billing, and how it applies to Reserved Instances, see [Billing Benefits and Payment Options](#).

Q: How many Reserved Instances can I purchase?

You can purchase up to 20 Reserved Instances per Availability Zone each month. If you need additional Reserved Instances, complete the form found [here](#). Information about previous generation Reserved Instance types can be found [here](#).

Q: Can I reassign my Reserved Instance from one instance type (e.g., c1.xlarge) to another (e.g., m1.large)?

No. A Reserved Instance is associated with a specific instance type for the duration of its term; however, you can change from one instance size (e.g., c3.large) to another (e.g., c3.xlarge) in the same type, if it is a Linux/UNIX Reserved Instance.

Q: Can I move a Reserved Instance from one region to another?

No. A Reserved Instance is associated with a specific region, which is fixed for the duration of the reservation's term.

Q: Can I modify a Reserved Instance?

Yes. You can request to modify active reservations that you own in one of the following ways:

Move between Availability Zones within the same region.

Change the network platform from EC2-Classic to EC2-VPC (for EC2-Classic-enabled customers).

Change the instance type of your Linux/UNIX Reserved Instances to a larger or smaller size in the same instance type (e.g., convert 8 m1.smalls into 4 m1.mediums, or vice versa).

Instance type modifications are only supported for Linux/UNIX platform reservations. However, due to licensing differences Linux Reserved Instances cannot be modified to RedHat or SUSE Linux Reserved Instances.

---

The reservations that you modify must have been purchased on the same day, be the same instance type, and in the same Availability Zone and region. It is not possible to combine reservations. However, if you have multiple instances in the same reservation (i.e., the reservation was purchased to apply to 10 instances), you can modify each of these instances either individually or as a whole.

Q: How do I request changes or modifications?

You can submit a modification request from the Amazon EC2 Console or by using the `ModifyReservedInstances` API. We process your requests as soon as possible, depending on available capacity. There is no additional cost for modifying your Reserved Instances.

To learn more about modification, see the [Amazon EC2 User Guide](#).

## Question: 8

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

**Answer: B**

Explanation:

Reference:

<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

Which two AWS services provide out-of-the-box user configurable automatic backup-as-a-service and backup rotation options? Choose 2 answers

- A. Amazon S3
- B. Amazon RDS
- C. Amazon EBS
- D. Amazon Redshift

**Answer: B, D**

Explanation:

By default: 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

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 redundant database.

**Answer: A**

### Question: 11

Which of the following are characteristics of Amazon VPC subnets? Choose 2 answers

- 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. Each subnet spans at least 2 Availability zones to provide a high-availability environment

**Answer: A, D**

Explanation:

You can create a VPC that spans multiple Availability Zones. For more information, see [Creating a VPC](#). After creating a VPC, you can add one or more subnets in each Availability Zone. Each subnet must reside entirely within one Availability Zone and cannot span zones. Availability Zones are distinct locations that are engineered to be isolated from failures in other Availability Zones. By launching instances in separate Availability Zones, you can protect your applications from the failure of a single location. AWS assigns a unique ID to each subnet.

- B is wrong: /28 is the smallest
- C is wrong: private subnet should go via NAT (EIP only in public subnet)
- E is wrong: subnet can only map to ONE AZ (not span multiple)

## Question: 12

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 Put MetricData permission and modify the Auto Scaling launch configuration to inject the users credentials 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 Put MetricData permission and put the credentials in a private repository and have applications on the server pull the credentials as needed

**Answer: A**

## Question: 13

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 ESS snapshot.
- C. Data is automatically deleted.
- D. Data is unavailable until the instance is restarted.

**Answer: C**

Explanation:

Reference:

<http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/ComponentsAMIs.html>

## Question: 14

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 causes the site to go down.

What could you do to help alleviate the pressure on the infrastructure while maintaining availability during these events?

Choose 3 answers

- 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.

**Answer: A, B, C**

Explanation:

The questions mention RDS so an answer that includes that as part of the solution makes sense. Also, Route53 does nothing to alleviate pressure on the infrastructure, it's for failover. E is counterproductive. It talks about failing over to an error page on S3.

## Question: 15

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 application?

- 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 between your new and existing domains and use the new domain for authentication

**Answer: C**

Explanation:

Since it requires no changes to the authentication infrastructure as requested in the question. Option D creates a new LDAP, trusts, etc.

### Question: 16

You need to design a VPC for a web-application consisting of an Elastic Load Balancer (ELB). A fleet of web/application servers, and an RDS database The Entire Infrastructure must be distributed over 2 availability zones.

Which VPC configuration works while assuring the database is not available from the Internet?

- A. One public subnet for ELB one public subnet for the web-servers, and one private subnet for the database
- B. One public subnet for ELB two private subnets for the web-servers, two private subnets for RDS
- C. Two public subnets for ELB two private subnets for the web-servers and two private subnets for RDS
- D. Two public subnets for ELB two public subnets for the web-servers, and two public subnets for RDS

**Answer: C**

### Question: 17

An application that you are managing has EC2 instances & Dynamo DB tables deployed to several AWS Regions. In order to monitor the performance of the application globally, you would like to see two graphs: 1) Avg CPU Utilization across all EC2 instances and 2) Number of Throttled Requests for all DynamoDB tables.

How can you accomplish this?

- A. Tag your resources with the application name, and select the tag name as the dimension in the CloudWatch Management console to view the respective graphs
- B. Use the Cloud Watch CLI tools to pull the respective metrics from each regional endpoint Aggregate the data offline & store it for graphing in CloudWatch.
- C. Add SNMP traps to each instance and DynamoDB table Leverage a central monitoring server to capture data from each instance and table Put the aggregate data into Cloud Watch for graphing.

D. Add a CloudWatch agent to each instance and attach one to each DynamoDB table. When configuring the agent set the appropriate application name & view the graphs in CloudWatch.

**Answer: A**

Explanation:

Correct answer should be A. When you turn on detailed monitoring in CloudWatch, you can get 1) Avg CPU Utilization across all EC2 instances and 2) Number of Throttled Requests for all DynamoDB tables

Reference: <http://docs.aws.amazon.com/AmazonCloudWatch/latest/DeveloperGuide/GetSingleMetricAllDimensions.html>

## Question: 18

When assessing an organization's use of AWS API access credentials which of the following three credentials should be evaluated? Choose 3 answers

- A. Key pairs
- B. Console passwords
- C. Access keys
- D. Signing certificates
- E. Security Group memberships

**Answer: A, C, D**

Explanation:

Reference:

[http://media.amazonwebservices.com/AWS\\_Operational\\_Checklists.pdf](http://media.amazonwebservices.com/AWS_Operational_Checklists.pdf)

## Question: 19

You have a Linux EC2 web server instance running inside a VPC. The instance is in a public subnet and has an EIP associated with it so you can connect to it over the Internet via HTTP or SSH. The instance was also fully accessible when you last logged in via SSH, and was also serving web requests on port 80.

Now you are not able to SSH into the host nor does it respond to web requests on port 80 that were working fine last time you checked. You have double-checked that all networking configuration parameters (security groups, route tables, IGW/EIP, NACLs etc) are properly configured (and you haven't made any changes to those anyway since you were last able to reach the instance). You look at the EC2 console and notice that system status check shows "impaired."

Which should be your next step in troubleshooting and attempting to get the instance back to a healthy state so that you can log in again?

- A. Stop and start the instance so that it will be able to be redeployed on a healthy host system that most likely will fix the "impaired" system status

- 
- B. Reboot your instance so that the operating system will have a chance to boot in a clean healthy state that most likely will fix the 'impaired' system status
  - C. Add another dynamic private IP address to the instance and try to connect via that new path, since the networking stack of the OS may be locked up causing the "impaired" system status.
  - D. Add another Elastic Network Interface to the instance and try to connect via that new path since the networking stack of the OS may be locked up causing the "impaired" system status
  - E. un-map and then re-map the EIP to the instance, since the IGW/NAT gateway may not be working properly, causing the "impaired" system status

**Answer: A**

### Question: 20

What is a placement group?

- A. A collection of Auto Scaling groups in the same Region
- B. Feature that enables EC2 instances to interact with each other via high bandwidth, low latency connections
- C. A collection of Elastic Load Balancers in the same Region or Availability Zone
- D. A collection of authorized Cloud Front edge locations for a distribution

**Answer: B**

Explanation:

Reference:

<http://aws.amazon.com/ec2/faqs/>

For More Information – **Visit link below:**  
<https://www.testsexpert.com/>

# Features:

■ Money Back Guarantee.....



■ 100% Course Coverage.....



■ 90 Days Free Updates.....



■ Instant Email Delivery after Order.....

