



SOA-C02^{Q&As}

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.pass4itsure.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 company has an application that runs only on Amazon EC2 Spot Instances. The instances run in an Amazon EC2 Auto Scaling group with scheduled scaling actions.

However, the capacity does not always increase at the scheduled times, and instances terminate many times a day. A Sysops administrator must ensure that the instances launch on time and have fewer interruptions.

Which action will meet these requirements?

- A. Specify the capacity-optimized allocation strategy for Spot Instances. Add more instance types to the Auto Scaling group.
- B. Specify the capacity-optimized allocation strategy for Spot Instances. Increase the size of the instances in the Auto Scaling group.
- C. Specify the lowest-price allocation strategy for Spot Instances. Add more instance types to the Auto Scaling group.
- D. Specify the lowest-price allocation strategy for Spot Instances. Increase the size of the instances in the Auto Scaling group.

Correct Answer: A

The correct answer is A. Specify the capacity-optimized allocation strategy for Spot Instances. Add more instance types to the Auto Scaling group.

By using the capacity-optimized allocation strategy for Spot Instances, Amazon EC2 Auto Scaling will launch instances on the most available Spot Instance pools with the lowest prices. This helps to improve the chances of getting the required capacity at the scheduled times.

Adding more instance types to the Auto Scaling group also increases the chances of finding available Spot Instances at any given time. It provides flexibility in selecting different instance types based on the availability and cost of Spot Instances in different pools.

Option B (Specify the capacity-optimized allocation strategy for Spot Instances and increase the size of the instances in the Auto Scaling group) might not be the most efficient approach, as simply increasing the instance size may not necessarily address the issue of instances terminating frequently.

QUESTION 2

A company hosts an internet web application on Amazon EC2 instances. The company is replacing the application with a new AWS Lambda function. During a transition period, the company must route some traffic to the legacy application and some traffic to the new Lambda function. The company needs to use the URL path of request to determine the routing.

Which solution will meet these requirements?

- A. Configure a Gateway Load Balancer to use the URL path to direct traffic to the legacy application and the new Lambda function.
- B. Configure a Network Load Balancer to use the URL path to direct traffic to the legacy application and the new Lambda function.



C. Configure a Network Load Balancer to use a regular expression to match the URL path to direct traffic to the new Lambda function.

D. Configure an Application Load Balancer to use the URL path to direct traffic to the legacy application and the new Lambda function.

Correct Answer: D

Application load balancer as ALB only supports targets like EC2 Instance, IP Address and Lambda

QUESTION 3

A company has deployed an application on AWS. The application runs on a fleet of Linux Amazon EC2 instances that are in an Auto Scaling group. The Auto Scaling group is configured to use launch templates. The launch templates launch

Amazon Elastic Block Store (Amazon EBS) backed EC2 instances that use General Purpose SSD (gp3) EBS volumes for primary storage.

A SysOps administrator needs to implement a solution to ensure that all the EC2 instances can share the same underlying files. The solution also must ensure that the data is consistent.

Which solution will meet these requirements?

A. Create an Amazon Elastic File System (Amazon EFS) file system. Create a new launch template version that includes user data that mounts the EFS file system. Update the Auto Scaling group to use the new launch template version to cycle in newer EC2 instances and to terminate the older EC2 instances.

B. Enable Multi-Attach on the EBS volumes. Create a new launch template version that includes user data that mounts the EBS volume. Update the Auto Scaling group to use the new template version to cycle in newer EC2 instances and to terminate the older EC2 instances.

C. Create a cron job that synchronizes the data between the EBS volumes for all the EC2 instances in the Auto Scaling group. Create a lifecycle hook during instance launch to configure the cron job on all the EC2 instances. Rotate out the older EC2 instances.

D. Create a new launch template version that creates an Amazon Elastic File System (Amazon EFS) file system. Update the Auto Scaling group to use the new template version to cycle in newer EC2 instances and to terminate the older EC2 instances.

Correct Answer: A

QUESTION 4

A SysOps administrator manages a company's Amazon S3 buckets. The SysOps administrator has identified 5 GB of incomplete multipart uploads in an S3 bucket in the company's AWS account. The SysOps administrator needs to reduce the number of incomplete multipart upload objects in the S3 bucket.

Which solution will meet this requirement?

A. Create an S3 Lifecycle rule on the S3 bucket to delete expired markers or incomplete multipart uploads.

B. Require users that perform uploads of files into Amazon S3 to use the S3 TransferUtility.



- C. Enable S3 Versioning on the S3 bucket that contains the incomplete multipart uploads.
- D. Create an S3 Object Lambda Access Point to delete incomplete multipart uploads.

Correct Answer: A

S3 Lifecycle rules allow you to define actions that Amazon S3 should take on objects in the bucket over time. This includes transitioning objects between storage classes and deleting objects when they meet certain criteria. To reduce the number of incomplete multipart upload objects in the S3 bucket, you can create an S3 Lifecycle rule that targets incomplete multipart uploads and specifies a deletion action for them. This will help in automatically cleaning up the incomplete multipart uploads after a certain period.

QUESTION 5

A company needs to deploy instances of an application and associated infrastructure to multiple AWS Regions. The company wants to use a single AWS CloudFormation template to achieve this goal. The company uses AWS Organizations and wants to administer and run this template from a central administration account.

What should a SysOps administrator do to meet these requirements?

- A. Create a CloudFormation template that is stored in Amazon S3. Configure Cross-Region Replication (CRR) on the S3 bucket. Reference the required accounts and remote Regions in the input template parameters.
- B. In the central administration account, create a CloudFormation primary template that loads CloudFormation nested stacks from Amazon S3 buckets in the target Regions.
- C. Create CloudFormation nested stacks by using a primary template in the central administration account. Configure the required accounts and Regions for deployment of the nested stacks.
- D. Create a CloudFormation stack set that includes service-managed permissions. Deploy the stack set into the required accounts and Regions from the central administration account.

Correct Answer: D

Using CloudFormation stack sets with service-managed permissions is the recommended approach when you need to deploy CloudFormation stacks to multiple AWS accounts and/or regions from a central administration account. Stack sets simplify the process of deploying and managing infrastructure across multiple accounts and regions, ensuring consistency and ease of administration

QUESTION 6

A company has implemented a Kubernetes cluster on Amazon Elastic Kubernetes Service (Amazon EKS) to host a microservices-based application. The company expects application traffic to increase significantly for the next month and wants to prevent the application from crashing because of the high number of requests.

Which solution will meet these requirements with the LEAST administrative overhead?

- A. Create a second EKS cluster. Load balance the workload between the two clusters.
- B. Implement the Kubernetes Horizontal Pod Autoscaler. Set a target CPU utilization percentage.
- C. Migrate the application from Amazon EKS to Amazon EC2 for the next month. Migrate the application back to Amazon EKS when the month ends.



D. Implement the Kubernetes Vertical Pod Autoscaler. Set a target CPU utilization percentage.

Correct Answer: B

QUESTION 7

An organization with a large IT department has decided to migrate to AWS. With different job functions in the IT department, it is not desirable to give all users access to all AWS resources. Currently, the organization handles access via LDAP group membership.

What is the BEST method to allow access using current LDAP credentials?

- A. Create an AWS Directory Service Simple AD. Replicate the on-premises LDAP directory to Simple AD.
- B. Create a Lambda function to read LDAP groups and automate the creation of IAM users.
- C. Use AWS CloudFormation to create IAM roles. Deploy Direct Connect to allow access to the on-premises LDAP server.
- D. Federate the LDAP directory with IAM using SAML. Create different IAM roles to correspond to different LDAP groups to limit permissions.

Correct Answer: D

QUESTION 8

A company runs workloads on 90 Amazon EC2 instances in the eu-west-1 Region in an AWS account. In 2 months, the company will migrate the workloads from eu-west-1 to the eu-west-3 Region.

The company needs to reduce the cost of the EC2 instances. The company is willing to make a 1-year commitment that will begin next week. The company must choose an EC2 Instance purchasing option that will provide discounts for the 90 EC2 instances regardless of Region during the 1-year period.

Which solution will meet these requirements?

- A. Purchase EC2 Standard Reserved Instances.
- B. Purchase an EC2 Instance Savings Plan.
- C. Purchase EC2 Convertible Reserved Instances.
- D. Purchase a Compute Savings Plan.

Correct Answer: D

Compute Savings Plans provide the most flexibility and help to reduce your costs by up to 66% (just like Convertible RIs). These plans automatically apply to EC2 instance usage regardless of instance family, size, AZ, Region, operating system, or tenancy, and also apply to Fargate and Lambda usage. For example, with Compute Savings Plans, you can change from C4 to M5 instances, shift a workload from EU (Ireland) to Europe (London), or move a workload from Amazon EC2 to Fargate or Lambda at any time and automatically continue to pay the Savings Plans price.

**QUESTION 9**

A company runs several workloads on AWS. The company identifies five AWS Trusted Advisor service quota metrics to monitor in a specific AWS Region. The company wants to receive email notification each time resource usage exceeds 60% of one of the service quotas.

Which solution will meet these requirements?

- A. Create five Amazon CloudWatch alarms, one for each Trusted Advisor service quota metric. Configure an Amazon Simple Notification Service (Amazon SNS) topic for email notification each time that usage exceeds 60% of one of the service quotas.
- B. Create five Amazon CloudWatch alarms, one for each Trusted Advisor service quota metric. Configure an Amazon Simple Queue Service (Amazon SQS) queue for email notification each time that usage exceeds 60% of one of the service quotas.
- C. Use the AWS Service Health Dashboard to monitor each Trusted Advisor service quota metric. Configure an Amazon Simple Queue Service (Amazon SQS) queue for email notification each time that usage exceeds 60% of one of the service quotas.
- D. Use the AWS Service Health Dashboard to monitor each Trusted Advisor service quota metric. Configure an Amazon Simple Notification Service (Amazon SNS) topic for email notification each time that usage exceeds 60% of one of the service quotas.

Correct Answer: A

Amazon CloudWatch allows you to set up alarms on various metrics, including service quota metrics from AWS Trusted Advisor. In this case, the company wants to monitor five Trusted Advisor service quota metrics in a specific AWS Region and receive email notifications when resource usage exceeds 60% of any of these quotas.

Option D is incorrect because it mentions using the AWS Service Health Dashboard to monitor service quota metrics. While the Service Health Dashboard provides information about the overall health of AWS services, it does not provide specific monitoring capabilities for individual service quota metrics as required in this scenario.

QUESTION 10

A SysOps administrator wants to monitor the free disk space that is available on a set of Amazon EC2 instances that have Amazon Elastic Block Store (Amazon EBS) volumes attached. The SysOps administrator wants to receive a notification when the used disk space of the EBS volumes exceeds a threshold value, but only when the DiskReadOps metric also exceeds a threshold value. The SysOps administrator has set up an Amazon Simple Notification Service (Amazon SNS) topic.

How can the SysOps administrator receive notification only when both metrics exceed their threshold values?

- A. Install the Amazon CloudWatch agent on the EC2 instances. Create a metric alarm for the disk space and a metric alarm for the DiskReadOps metric. Create a composite alarm that includes the two metric alarms to publish a notification to the SNS topic.
- B. Install the Amazon CloudWatch agent on the EC2 instances. Create a metric alarm for the disk space and a metric alarm for the DiskReadOps metric. Configure each alarm to publish a notification to the SNS topic.
- C. Create a metric alarm for the EBSByteBalance% metric and a metric alarm for the DiskReadOps metric. Create a composite alarm that includes the two metric alarms to publish a notification to the SNS topic.
- D. Configure detailed monitoring for the EC2 instances. Create a metric alarm for the disk space and a metric alarm for



the DiskReadOps metric. Create a composite alarm that includes the two metric alarms to publish a notification to the SNS topic.

Correct Answer: A

To receive a notification only when both metrics exceed their threshold values, the SysOps administrator should follow these steps:

Install the Amazon CloudWatch agent on the EC2 instances.

Configure detailed monitoring for the EC2 instances.

Create a metric alarm for the DiskReadOps metric and set a threshold value for it.

Create a metric alarm for the free disk space and set a threshold value for it.

Create a composite alarm that includes the two metric alarms, and set a threshold value for it.

Configure the composite alarm to publish a notification to the SNS topic when it is triggered.

QUESTION 11

An application runs on Amazon EC2 instances in an Auto Scaling group. Following the deployment of a new feature on the EC2 instances, some instances were marked as unhealthy and then replaced by the Auto Scaling group. The EC2 instances terminated before a SysOps administrator could determine the cause of the health status changes. To troubleshoot this issue, the SysOps administrator wants to ensure that an AWS Lambda function is invoked in this situation. How should the SysOps administrator meet these requirements?

- A. Activate the instance scale-in protection setting for the Auto Scaling group. Invoke the Lambda function through Amazon EventBridge (Amazon CloudWatch Events).
- B. Activate the instance scale-in protection setting for the Auto Scaling group. Invoke the Lambda function through Amazon Route 53.
- C. Add a lifecycle hook to the Auto Scaling group to invoke the Lambda function through Amazon EventBridge (Amazon CloudWatch Events).
- D. Add a lifecycle hook to the Auto Scaling group to invoke the Lambda function through Amazon Route 53.

Correct Answer: C

A is Wrong because Instance scale-in protection does not protect Auto Scaling instances from the following: Health check replacement if the instance fails health checks C is the Correct Answers -- When a scale-in event occurs, a lifecycle hook pauses the instance before it is terminated and sends you a notification using Amazon EventBridge. While the instance is in the wait state, you can invoke an AWS Lambda function or connect to the instance to download logs or other data before the instance is fully terminated. <https://docs.aws.amazon.com/autoscaling/ec2/userguide/lifecycle-hooks.html>

QUESTION 12

A SysOps administrator wants to manage a web server application with AWS Elastic Beanstalk. The Elastic Beanstalk service must maintain full capacity for new deployments at all times. Which deployment policies satisfy this requirement? (Select TWO.)



- A. All at once
- B. Immutable
- C. Rebuild
- D. Rolling
- E. Rolling with additional batch

Correct Answer: BE

Immutable deployments perform an immutable update to launch a full set of new instances running the new version of the application in a separate Auto Scaling group, alongside the instances running the old version. Immutable deployments can prevent issues caused by partially completed rolling deployments. If the new instances don't pass health checks, Elastic Beanstalk terminates them, leaving the original instances untouched.

To maintain full capacity during deployments, you can configure your environment to launch a new batch of instances before taking any instances out of service. This option is known as a rolling deployment with an additional batch. When the deployment completes, Elastic Beanstalk terminates the additional batch of instances.

<https://docs.aws.amazon.com/elasticbeanstalk/latest/dg/using-features.rolling-version-deploy.html>

QUESTION 13

A company runs a high performance computing (HPC) application on an Amazon EC2 instance. The company needs to scale this architecture to two or more EC2 instances. The EC2 instances will need to communicate with each other at high speeds with low latency to support the application.

The company wants to ensure that the network performance can support the required communication between the EC2 instances

What should a SysOps administrator do to meet these requirements?

- A. Create a cluster placement group. Back up the existing EC2 instance to an Amazon Machine Image (AMI). Restore the EC2 instance from the AMI into the placement group. Launch the additional EC2 instances into the placement group.
- B. Back up the existing EC2 instance to an Amazon Machine Image (AMI). Create a launch template from the existing EC2 instance by specifying the AMI. Create an Auto Scaling group and configure the desired instance count.
- C. Create a Network Load Balancer (NLB) and a target group. Launch the new EC2 instances and register them with the target group. Register the existing EC2 instance with the target group. Pass all application traffic through the NLB.
- D. Back up the existing EC2 instance to an Amazon Machine Image (AMI). Create additional clones of the EC2 instance from the AMI in the same Availability Zone where the existing EC2 instance is located.

Correct Answer: A

low latency = Cluster Placement Group

QUESTION 14



A SysOps administrator has been able to consolidate multiple, secure websites onto a single server, and each site is running on a different port. The administrator now wants to start a duplicate server in a second Availability Zone and put both behind a load balancer for high availability.

What would be the command line necessary to deploy one of the sites' certificates to the load balancer?

- A. `aws kms modify-listener ?load-balancer-name my-load-balancer -?ertificates CertificateArn=arn:aws:iam::123456789012:server-certifiante/my-new-server-cert`
- B. `aws elb set-load-balancer-listener-ssl-certificate --load-balancer-name my-load-balancer ?load-balancer-port 443 ?ssl-certificate-id arn:aws:iam::123456789012:server-certificate/new-server-cert`
- C. `aws ec2 put-ssl-certificate ?load-balancer-name my-load-balancer ?load-balancer-port 443 ?ssl-certificate-id arn:aws:iam::123456789012:server-certificate/new-server-cert`
- D. `aws acm put-ssl-certificate ?load-balancer-name my-load-balancer ?load-balancer-port 443 ?ssl-certificate-id arn:aws:iam::123456789012:server-certificate/new-server-cert`

Correct Answer: B

<https://docs.aws.amazon.com/cli/latest/reference/elb/set-load-balancer-listener-ssl-certificate.html>

QUESTION 15

CORRECT TEXT

A webpage is stored in an Amazon S3 bucket behind an Application Load Balancer (ALB). Configure the SS bucket to serve a static error page in the event of a failure at the primary site.

1.

Use the us-east-2 Region for all resources.

2.

Unless specified below, use the default configuration settings.

3.

There is an existing hosted zone named lab-751906329398-26023898.com that contains an A record with a simple routing policy that routes traffic to an existing ALB.

4.

Configure the existing S3 bucket named lab-751906329398-26023898.com as a static hosted website using the object named index.html as the index document

5.

For the index-html object, configure the S3 ACL to allow for public read access. Ensure public access to the S3 bucket is allowed.

6.

In Amazon Route 53, change the A record for domain lab-751906329398-26023898.com to a primary record for a



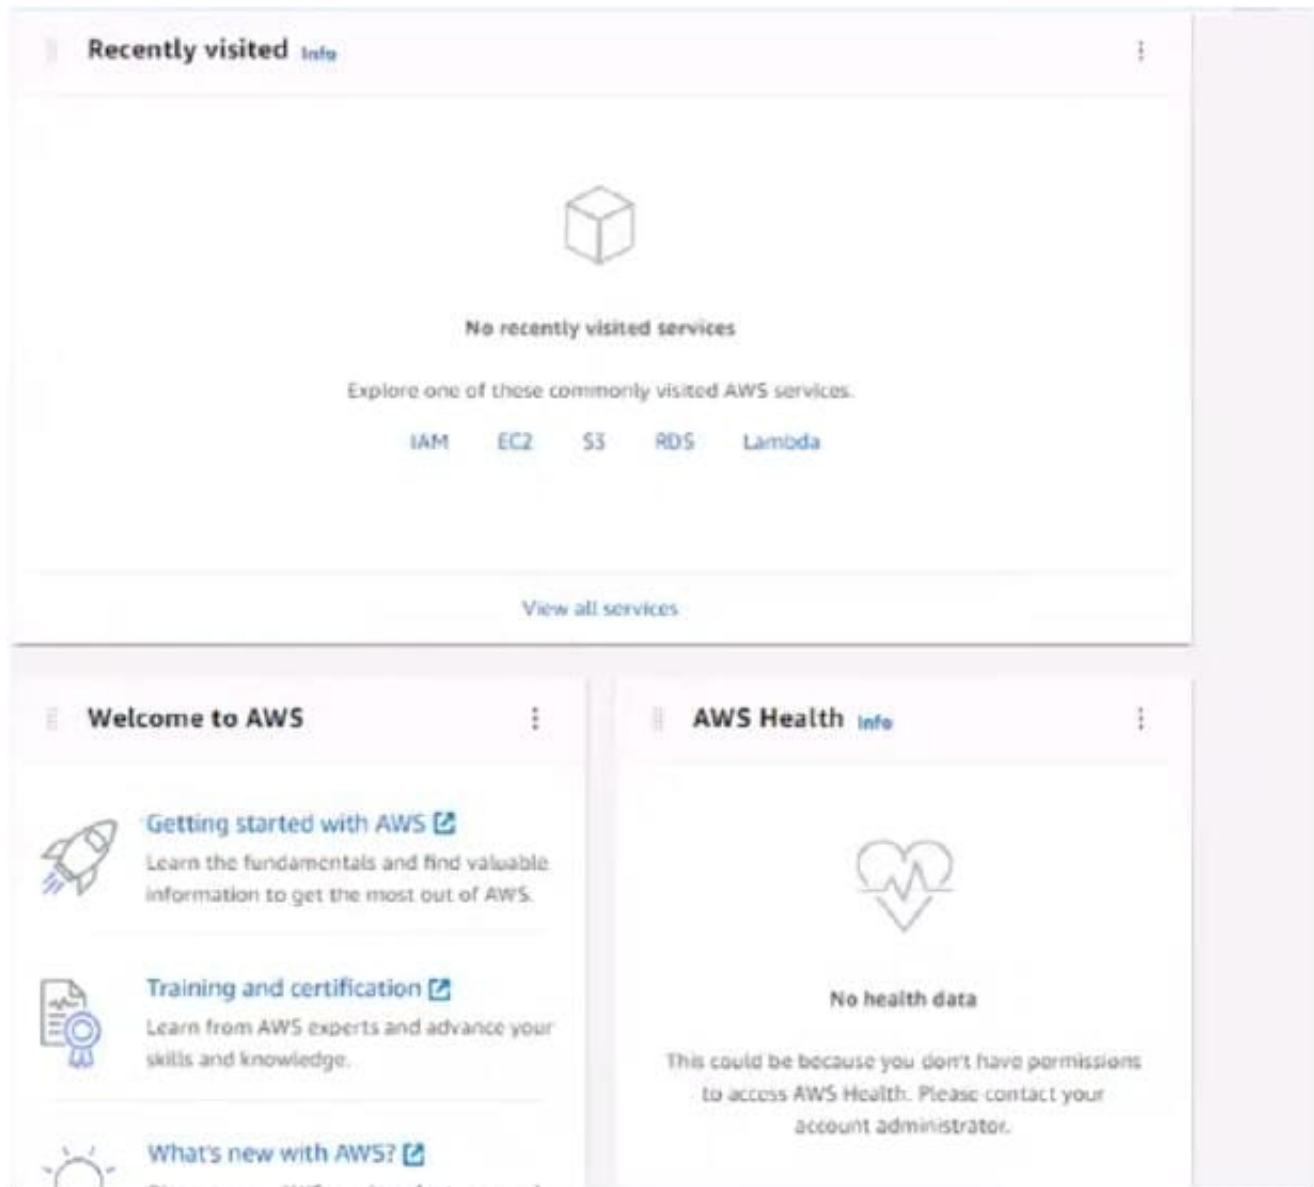
failover routing policy. Configure the record so that it evaluates the health of the ALB to determine failover.

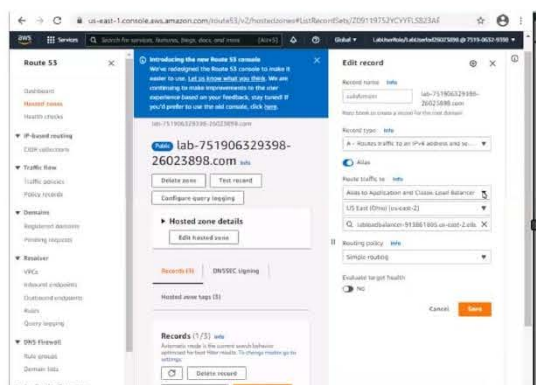
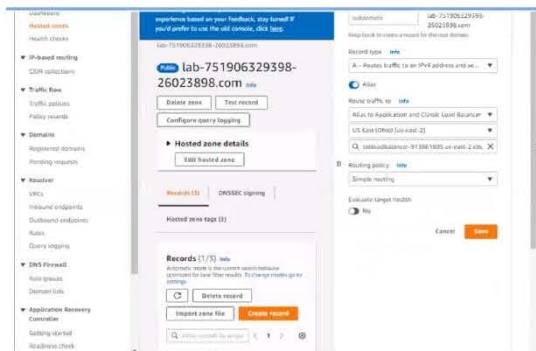
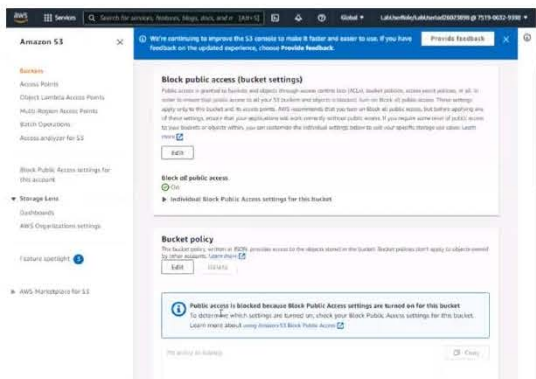
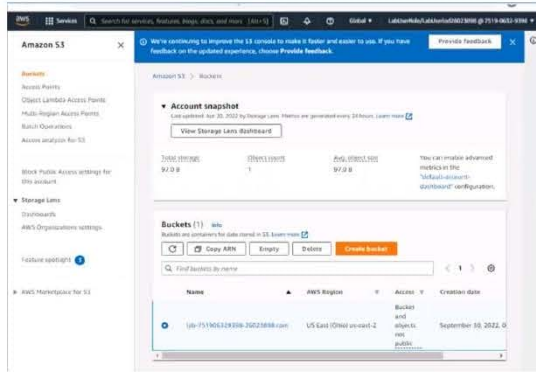
7.

Create a new secondary failover alias record for the domain lab-751906329398-26023898.com that routes traffic to the existing 53 bucket.

Correct Answer: Check the answer in explanation.

Solution as given below.







Introducing the new Route 53 console
We've redesigned the Route 53 console to make it easier for you. Let us know what you think. We are continuing to make improvements to the user experience based on your feedback, stay tuned if you'd prefer to use the old console, click here.

Quick create (recommended for expert users)
Choose this method if you are confident in the process of creating records and know which options you need.

Wizard (recommended for new users)
Choose this method if you need more explanations as you create your record.

Quick create record info

Record 1

Record name: info
subdomain lab-751906129398-20022898.com
Record type: info
A - Routes traffic to an IPv4 address and some AWS resources

Keep blank to create a record for the root domain.

Alias
Value: info
192.0.2.238

Enter multiple values on separate lines.

TTL (seconds): info
300
Routing policy: info
Simple routing

Recommended values: 60 to 172800 (max: 3600)

Add another record

Introducing the new Route 53 console
We've redesigned the Route 53 console to make it easier for you. Let us know what you think. We are continuing to make improvements to the user experience based on your feedback, stay tuned if you'd prefer to use the old console, click here.

Quick create record info

Record 1

Record name: info
subdomain lab-751906129398-20022898.com
Record type: info
A - Routes traffic to an IPv4 address and some AWS resources

Keep blank to create a record for the root domain.

Alias
Value: info
192.0.2.238

Enter multiple values on separate lines.

TTL (seconds): info
300
Routing policy: info
Simple routing

Recommended values: 60 to 172800 (max: 3600)

Add another record

View existing records
The following table lists the existing records in lab-751906129398-20022898.com.

Quick create record info

Record 1

Record name: info
subdomain lab-751906129398-20022898.com
Record type: info
A - Routes traffic to an IPv4 address and some AWS resources

Keep blank to create a record for the root domain.

Alias
Route traffic to: info
Alias to another record in this hosted zone
US East (N. Virginia)
An alias to a CloudFront distribution and another record in the same hosted zone are global and available only in US East (N. Virginia).
lab-751906129398-20022898.com
Alias hosted zone ID: Z0111752YVFL32234F
Routing policy: info
Fallover
Fallover record type: info
Secondary
Health check ID: optional info
Choose health check
Record ID: info
US West (and Oregon)

Add another record

Introducing the new Route 53 console
We've redesigned the Route 53 console to make it easier for you. Let us know what you think. We are continuing to make improvements to the user experience based on your feedback, stay tuned if you'd prefer to use the old console, click here.

Record creation method

Quick create (recommended for expert users)
Choose this method if you are confident in the process of creating records and know which options you need.

Wizard (recommended for new users)
Choose this method if you need more explanations as you create your record.

Learn more
Working with records

Quick create record info

Record 1

Record name: info
subdomain lab-751906129398-20022898.com
Record type: info
A - Routes traffic to an IPv4 address and some AWS resources

Keep blank to create a record for the root domain.

Alias
Route traffic to: info
Alias to another record in this hosted zone
US East (N. Virginia)
An alias to a CloudFront distribution and another record in the same hosted zone are global and available only in US East (N. Virginia).
lab-751906129398-20022898.com
Alias hosted zone ID: Z0111752YVFL32234F



Route 53 > Hosted zones > lab-751906329398-26023898.com > Create record

Quick create record [Info](#) [Switch to wizard](#)

Record 1 [Delete](#)

Record name [Info](#) Keep blank to create a record for the root domain

Record type [Info](#)

☒ Alias

Route traffic to [Info](#)

Alias hosted zone ID: Z3MCH0K5KTT1J

Routing policy [Info](#)

Failover record type

Health check ID - optional [Info](#)

Evaluate target health ☒ Yes

Record ID [Info](#)

[Add another record](#)

[SOA-C02 PDF Dumps](#)

[SOA-C02 Practice Test](#)

[SOA-C02 Braindumps](#)