Amazon-Web-Services

Exam Questions SOA-C02

AWS Certified SysOps Administrator - Associate (SOA-C02)

NEW QUESTION 1
A company asks a SysOps administrator to ensure that AWS CloudTrail files are not tampered with after they are created. Currently, the company uses AWS Identity and Access Management (IAM) to restrict access to specific trails. The company's security team needs the ability to trace the integrity of each file.
What is the MOST operationally efficient solution that meets these requirements?

A. Create an Amazon EventBridge (Amazon CloudWatch Events) rule that invokes an AWS Lambda function when a new file is delivere
B. Configure the Lambda function to compute an MD5 hash check on the file and store the result in anAmazon DynamoDB tabl
C. The security team can use the values that are stored in DynamoDB to verify the integrity of the delivered files.
D. Create an AWS Lambda function that is invoked each time a new file is delivered to the CloudTrail bucke
E. Configure the Lambda function to compute an MD5 hash check on the file and store the result as a tag in an Amazon 53 object.The security team can use the information in the tag to verify the integrity of the delivered files.
F. Enable the CloudTrail file integrity feature on an Amazon S3 bucke
G. Create an IAM policy that grants the security team access to the file integrity logs that are stored in the S3 bucket.
H. Enable the CloudTrail file integrity feature on the trai
I. The security team can use the digest file that is created by CloudTrail to verify the integrity of the delivered files.

Answer: C

NEW QUESTION 2
A company has an Amazon RDS DB instance. The company wants to implement a caching service while maintaining high availability. Which combination of actions will meet these requirements? (Choose two.)

A. Add Auto Discovery to the data store.
B. Create an Amazon ElastiCache for Memcached data store.
C. Create an Amazon ElastiCache for Redis data store.
D. Enable Multi-AZ for the data store.
E. Enable Multi-threading for the data store.

Answer: AD

NEW QUESTION 3
A SysOps administrator notices a scale-up event for an Amazon EC2 Auto Scaling group. Amazon CloudWatch shows a spike in the RequestCount metric for the associated Application Load Balancer. The administrator would like to know the IP addresses for the source of the requests.
Where can the administrator find this information?

A. Auto Scaling logs
B. AWS CloudTrail logs
C. EC2 instance logs
D. Elastic Load Balancer access logs

Answer: A

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
A SysOps administrator has created a VPC that contains a public subnet and a private subnet. Amazon EC2 instances that were launched in the private subnet cannot access the internet. The default network ACL is active on all subnets in the VPC, and all security groups allow all outbound traffic:
Which solution will provide the EC2 instances in the private subnet with access to the internet?

A. Create a NAT gateway in the public subne
B. Create a route from the private subnet to the NAT gateway.
C. Create a NAT gateway in the public subne
D. Create a route from the public subnet to the NAT gateway.
E. Create a NAT gateway in the private subne
F. Create a route from the public subnet to the NAT gateway.
G. Create a NAT gateway in the private subne
H. Create a route from the private subnet to the NAT gateway.

Answer: A

NEW QUESTION 6
A company uses an Amazon Elastic File System (Amazon EFS) file system to share files across many Linux Amazon EC2 instances. A SysOps administrator notices that the file system's PercentIOLimit metric is consistently at 100% for 15 minutes or longer. The SysOps administrator also notices that the application that reads and writes to that file system is performing poorly. They application requires high throughput and IOPS while accessing the file system.
What should the SysOps administrator do to remediate the consistently high PercentIOLimit metric?

A. Create a new EFS file system that uses Max I/O performance mod
B. Use AWS DataSync to migrate data to the new EFS file system.
C. Create an EFS lifecycle policy to transition future files to the Infrequent Access (IA) storage class to improve performanc
D. Use AWS DataSync to migrate existing data to IA storage.
E. Modify the existing EFS file system and activate Max I/O performance mode.
F. Modify the existing EFS file system and activate Provisioned Throughput mode.

Answer: A

NEW QUESTION 7
A company using AWS Organizations requires that no Amazon S3 buckets in its production accounts should ever be deleted.
What is the SIMPLEST approach the SysOps administrator can take to ensure S3 buckets in those accounts can never be deleted?

A. Set up MFA Delete on all the S3 buckets to prevent the buckets from being deleted.
B. Use service control policies to deny the s3:DeleteBucket action on all buckets in production accounts.
C. Create an IAM group that has an IAM policy to deny the s3:DeleteBucket action on all buckets in production accounts.
D. Use AWS Shield to deny the s3:DeleteBucket action on the AWS account instead of all S3 buckets.

Answer: B

NEW QUESTION 8
A company uses Amazon Route 53 to manage the public DNS records for the domain example.com. The company deploys an Amazon CloudFront distribution to deliver static assets for a new corporate website. The company wants to create a subdomain that is named "static" and must route traffic for the subdomain to the CloudFront distribution.
How should a SysOps administrator create a new record for the subdomain in Route 53?

A. Create a CNAME recor
B. Enter static.cloudfront.net as the record nam
C. Enter the CloudFront distribution's public IP address as the value.
D. Create a CNAME recor
E. Enter static.example.com as the record nam
F. Enter the CloudFront distribution's private IP address as the value.
G. Create an A recor
H. Enter static.cloudfront.net as the record nam
I. Enter the CloudFront distribution's ID as an alias target.
J. Create an A recor
K. Enter static.example.com as the record nam
L. Enter the CloudFront distribution's domain name as an alias target.

Answer: D

NEW QUESTION 9
A SysOps administrator is maintaining a web application using an Amazon CloudFront web distribution, an Application Load Balancer (ALB), Amazon RDS, and Amazon EC2 in a VPC. All services have logging enabled. The administrator needs to investigate HTTP Layer 7 status codes from the web application.
Which log sources contain the status codes? (Choose two.)

A. VPC Flow Logs
B. AWS CloudTrail logs
C. ALB access logs
D. CloudFront access logs
E. RDS logs

Answer: CD

NEW QUESTION 10
A SysOps administrator is deploying a test site running on Amazon EC2 instances. The application requires both incoming and outgoing connectivity to the internet. Which combination of steps are required to provide internet connectivity to the EC2 instances? (Choose two.)

A. Add a NAT gateway to a public subnet.
B. Attach a private address to the elastic network interface on the EC2 instance.
C. Attach an Elastic IP address to the internet gateway.
D. Add an entry to the route table for the subnet that points to an internet gateway.
E. Create an internet gateway and attach it to a VPC.

Answer: DE

NEW QUESTION 10
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 15
A SysOps administrator is notified that an Amazon EC2 instance has stopped responding. The AWS Management Console indicates that the system checks are failing.
What should the administrator do first to resolve this issue?

A. Reboot the EC2 instance so it can be launched on a new host.
B. Stop and then start the EC2 instance so that it can be launched on a new host.
C. Terminate the EC2 instance and relaunch it.
D. View the AWS CloudTrail log to investigate what changed on the EC2 instance.

Answer: B

NEW QUESTION 17
A company needs to create a daily Amazon Machine Image (AMI) of an existing Amazon Linux EC2 instance that hosts the operating system, application, and database on multiple attached Amazon Elastic Block Store (Amazon EBS) volumes. File system integrity must be maintained.
Which solution will meet these requirements?

A. Create an AWS Lambda function to call the CreateImage API operation with the EC2 instance ID and the no-reboot parameter enable
B. Create a daily scheduled Amazon EventBridge (Amazon CloudWatch Events) rule that invokes thefunction.
C. Create an AWS Lambda function to call the CreateImage API operation with the EC2 instance ID and the reboot parameter enable
D. Create a daily scheduled Amazon EventBridge (Amazon CloudWatch Events) rule that invokes thefunction.
E. Use AWS Backup to create a backup plan with a backup rule that runs dail
F. Assign the resource ID of the EC2 instance with the no-reboot parameter enabled.
G. Use AWS Backup to create a backup plan with a backup rule that runs dail
H. Assign the resource ID of the EC2 instance with the reboot parameter enabled.

Answer: C

NEW QUESTION 20
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 22
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 27
A company has multiple Amazon EC2 instances that run a resource-intensive application in a development environment. A SysOps administrator is implementing a solution to stop these EC2 instances when they are not in use.
Which solution will meet this requirement?

A. Assess AWS CloudTrail logs to verify that there is no EC2 API activit
B. Invoke an AWS lambda function to stop the EC2 instances.
C. Create an Amazon CloudWatch alarm to stop the EC2 instances when the average CPU utilization is lower than 5% for a 30-minute period.
D. Create an Amazon CloudWatch metric to stop the EC2 instances when the VolumeReadBytes metric is lower than 500 for a 30-minute period.
E. Use AWS Config to invoke an AWS Lambda function to stop the EC2 instances based on resource configuration changes.

Answer: A

NEW QUESTION 30
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 inorrect as it is a fictitious scaling policy type.

NEW QUESTION 34
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 37
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 withthe IP address for the Singapore load balancer
B. If latency is lower between the London and Oregon regions, Route 53 responds to the query withthe 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 betweenLondon 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 40
Which of the following AWS feature helps you ensure that you have the correct number of Amazon EC2 instances available to handle the load for your application?

A. Amazon EC2 Auto Scaling
B. AWS CloudFormation
C. AWS OpsWorks
D. AWS Elastic Beanstalk

Answer: A

Explanation:
Amazon EC2 Auto Scaling is the correct answer. Amazon EC2 Auto Scaling helps you ensure that you have the correct number of Amazon EC2 instances available to handle the load for your application.
You create collections of EC2 instances, called Auto Scaling groups. You can specify the minimum number of instances in each Auto Scaling group, and Amazon EC2 Auto Scaling ensures that your group never goes below this size.
You can specify the maximum number of instances in each Auto Scaling group, and Amazon EC2 Auto Scaling ensures that your group never goes above this size. If you specify the desired capacity, either when you create the group or at any time thereafter, Amazon EC2 Auto Scaling ensures that your group has this many instances. If you specify scaling policies, then Amazon EC2 Auto Scaling can
launch or terminate instances as demand on your application increases or decreases.
AWS CloudFormation is incorrect. AWS CloudFormation is a service that helps you model and set up your Amazon Web Services resources so that you can spend less time managing those resources and more time focusing on your applications that run in AWS. You create a template that describes all the AWS resources that you want (like Amazon EC2 instances or Amazon RDS DB instances), and AWS CloudFormation takes care of provisioning and configuring those resources for you. You don't need to individually create and configure AWS resources and figure out what's dependent on what; AWS CloudFormation handles all of that.
AWS Elastic Beanstalk is incorrect. With Elastic Beanstalk, you can quickly deploy and manage applications in the AWS Cloud without having to learn about the infrastructure that runs those applications. Elastic Beanstalk reduces management complexity without restricting choice or control. You simply upload your application, and Elastic Beanstalk automatically handles the details of capacity provisioning, load balancing, scaling, and application health monitoring.
AWS OpsWorks is incorrect. AWS OpsWorks is a configuration management service that helps you configure and operate applications in a cloud enterprise by using Puppet or Chef. AWS OpsWorks Stacks and AWS OpsWorks for Chef Automate let you use Chef cookbooks and solutions for configuration management, while OpsWorks for Puppet Enterprise lets you configure a Puppet Enterprise master server in AWS. Puppet offers a set of tools for enforcing the desired state of your infrastructure and automating on-demand tasks.

NEW QUESTION 45

A. Mastered
B. Not Mastered

Answer: A

NEW QUESTION 50
.....

Thank You for Trying Our Product

We offer two products:

1st - We have Practice Tests Software with Actual Exam Questions

2nd - Questons and Answers in PDF Format

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 FirstTry

* SOA-C02 Practice Test Questions in Multiple Choice Formats and Updatesfor 1 Year