Free Questions for SOA-C02

Shared by Stephenson on 04-10-2024

For More Free Questions and Preparation Resources

Check the Links on Last Page

Question 1

Question Type: MultipleChoice

A SysOps administrator is investigating a company's web application for performance problems The application runs on Amazon EC2 instances that are in an Auto Scaling group. The application receives large traffic increases at random times throughout the day. During periods of rapid traffic increases, the Auto Scaling group is not adding capacity fast enough. As a result, users are experiencing poor performance.

The company wants to minimize costs without adversely affecting the user experience when web traffic surges quickly. The company needs a solution that adds more capacity to me Auto Scaling group for larger traffic increases than for smaller traffic increases.

How should the SysOps administrator configure the Auto Scaling group to meet these requirements?

Options:

A- Create a simple scaling policy with settings to make larger adjustments in capacity when the system is under heavy load

B- Create a step scaling policy with settings to make larger adjustments in capacity when the system is under heavy load.

C- Create a target tracking scaling policy with settings to make larger adjustments in capacity when the system is under heavy load

D- Use Amazon EC2 Auto Scaling lifecycle hooks Adjust the Auto Scaling group's maximum number of instances after every scaling event

Answer:

B

Explanation:

Step Scaling Policy:

Step scaling policies allow you to define scaling actions based on different levels of CloudWatch alarms.

Steps:

Go to the AWS Management Console.

Navigate to EC2 Auto Scaling.

Select your Auto Scaling group.

Create or edit a scaling policy and choose 'Step scaling.'

Define different steps based on CloudWatch alarm thresholds (e.g., CPU usage or request count).

Configure larger adjustments for higher thresholds and smaller adjustments for lower thresholds.

Example Configuration:

For CPU > 80%, increase capacity by 4 instances.

For CPU > 60%, increase capacity by 2 instances.

For CPU > 40%, increase capacity by 1 instance.

Question 2

Question Type: MultipleChoice

A SysOps administrator must ensure that all of a company's current and future Amazon S3 buckets have logging enabled If an S3 bucket does not have logging enabled an automated process must enable logging for the S3 bucket.

Which solution will meet these requirements?

Options:

A- Use AWS Trusted Advisor 10 perform a check for S3 buckets that do not have logging enabled Configure the check to enable logging for S3 buckets that do not have logging enabled.

B- Configure an S3 bucket policy that requires all current and future S3 buckets to have logging enabled

C- Use the s3-bucket-logging-enabled AWS Config managed rule. Add a remediation action that uses an AWS Lambda function to enable logging.

D- Use the s3-bucket-logging-enabled AWS Config managed rule. Add a remediation action that uses the AWS-ConfigureS3BucketLoggmg AWS Systems Manager Automation runbook to enable logging.

Answer:

C, D

Explanation:

AWS Config Managed Rule for S3 Logging:

The s3-bucket-logging-enabled AWS Config rule checks whether S3 buckets have logging enabled.

Steps:

Go to the AWS Management Console.

Navigate to AWS Config.

Create a rule using s3-bucket-logging-enabled.

Add a remediation action using an AWS Lambda function or Systems Manager Automation runbook.

Using AWS Lambda for Remediation:

Create a Lambda function that enables logging on S3 buckets.

Steps:

Write a Lambda function in Python or Node.js to enable logging.

Configure the function to trigger on non-compliant buckets.

Using AWS Systems Manager Automation:

The AWS-ConfigureS3BucketLogging runbook automates enabling logging.

Steps:

Go to the AWS Management Console.

Navigate to Systems Manager.

Create an Automation document or use the existing AWS-ConfigureS3BucketLogging runbook.

Configure the remediation action to use this runbook.

Question 3

Question Type: **MultipleChoice**

A SysOps administrator needs to create a report that shows how many bytes are sent to and received from each target group member for an Application Load Balancer (ALB).

Which combination of steps should the SysOps administrator take to meet these requirements? (Select TWO.)

Options:

A- Enable access logging for the ALB. Save the logs to an Amazon S3 bucket.

B- Install the Amazon CloudWatch agent on the Instances in the target group.

C- Use Amazon Athena to query the ALB logs Query the table Use the received_bytes and senl_byt.es fields to calculate the total bytes grouped by the target:port field.

D- Use Amazon Athena to query the ALB logs Query the table. Use the received_bytes and sent_byt.es fields to calculate the total bytes grouped by the clientport field

E- Create an Amazon CloudWatch dashboard that shows the Sum statistic of the ProcessedBytes metric for the ALB.

Answer:

A, C

Explanation:

Enable Access Logging for the ALB:

Access logging provides detailed information about requests sent to your load balancer.

Steps:

Go to the AWS Management Console.

Navigate to EC2 and select 'Load Balancers.'

Select your Application Load Balancer.

Under the 'Attributes' tab, enable 'Access logs.'

Specify an S3 bucket where the logs will be saved.

Use Amazon Athena to Query the ALB Logs:

Athena allows you to run SQL queries on data stored in S3.

Steps:

Go to the AWS Management Console.

Navigate to Athena.

Create a table for the ALB logs using the appropriate schema.

Run queries to calculate the total bytes sent and received, grouped by the target

field.

Example query:

SELECT target, SUM(received_bytes) as total_received, SUM(sent_bytes) as total_sent

FROM alb_logs

GROUP BY target, port

Question 4

Question Type: MultipleChoice

A company wants to prohibit its developers from using a particular family of Amazon EC2 instances The company uses AWS Organizations and wants to apply the restriction across multiple accounts

What is the MOST operationally efficient way for the company lo apply service control policies (SCPs) to meet these requirements?

Options:

A- Add the accounts to an organizational unit (OUf Apply the SCPs to the OU.

B- Add the accounts to resource groups in AWS Resource Groups. Apply the SCPs to the resource groups.

C- Apply the SCPs to each developer account.

D- Enroll the accounts with AWS Control Tower. Apply the SCPs to the AWS Control Tower management account.

Answer:

A

Explanation:

Applying SCPs to an Organizational Unit:

Service Control Policies (SCPs) allow you to manage permissions for multiple AWS accounts within an organization.

Steps:

Go to the AWS Management Console.

Navigate to AWS Organizations.

Create an Organizational Unit (OU) if not already created.

Move the target accounts into the OU.

Create an SCP that denies the use of the specific EC2 instance family.

Attach the SCP to the OU.

This approach ensures that the policy is applied consistently across all accounts in the OU.

Question 5

Question Type: MultipleChoice

An AWS Cloud Formation template creates an Amazon RDS instance This template is used to build up development environments as needed and then delete the stack when the environment is no longer required. The RDS-persisted data must be retained for further use. even after the Cloud Format ton stack is deleted

How can this be achieved in a reliable and efficient way?

Options:

A- Write a script to continue backing up the RDS instance every live minutes.

B- Create an AWS Lambda function to take a snapshot of the RDS instance, and manually invoke the function before deleting the stack.

C- Use the Snapshot Deletion Policy in the Cloud Formation template definition of the RDS instance.

D- Create a new CloudFormallon template to perform backups of the RDS instance, and run this template before deleting the stack.

Answer:

C

Explanation:

Snapshot Deletion Policy:

The Snapshot Deletion Policy ensures that a snapshot is created when an RDS instance is deleted as part of a CloudFormation stack deletion.

Steps:

Update your CloudFormation template to include the DeletionPolicy attribute for the RDS instance resource.

Example template snippet:

Resources:

MyDBInstance:

Type: AWS::RDS::DBInstance

Properties:

DB instance properties

DeletionPolicy: Snapshot

This configuration retains a snapshot of the RDS instance data when the stack is deleted.

Reference: AWS CloudFormation DeletionPolicy

Question 6

Question Type: MultipleChoice

A company has many accounts in an organization in AWS Organizations The company must automate resource provisioning from the organization's management account to the member accounts.

Which solution will meet this requirement?

Options:

A- Create an AWS CkHJdFormation change set Deploy the change set to all member accounts

B- Create an AWS CtoudFormation nested stack Deploy the nested stack to all member accounts.

C- Create an AWS CtoudFormation stack set Deploy the stack set to all member accounts.

D- Create an AWS Serverless Application Model (AWS SAM) template. Deploy the template to all member accounts.

Answer:

C

Explanation:

Using CloudFormation Stack Sets:

CloudFormation stack sets allow you to deploy CloudFormation stacks across multiple AWS accounts and regions.

Steps:

Go to the AWS Management Console.

Navigate to CloudFormation and select 'StackSets.'

Click on 'Create StackSet.'

Provide the template URL or upload a template file.

Configure the stack set options and specify the accounts and regions.

Deploy the stack set to the specified accounts and regions.

Question 7

Question Type: **MultipleChoice**

A company hosts an application on Amazon EC2 instances The instances are in an Amazon EC2 Auto Scaling group that uses a launch template The amount of application traffic changes throughout the day. Scaling events happen frequently.

A SysOps administrator needs to help developers troubleshoot the application. When a scaling event removes an instance. EC2 Auto Scaling terminates the instance before the developers can log in to the instance to diagnose issues.

Which solution will prevent termination of the instance so that the developers can log in to the instance?

Options:

A- Ensure that the Delete on termination setting is turned off in the UserData section of the launch template

B- Update the Auto Scaling group by enabling instance scale-in protection for newly launched instances.

C- Use Amazon Inspector to configure a rules package to protect the instances from termination.

D- Use Amazon GuardDuty to configure rules to protect the instances from termination.

Answer:

B

Explanation:

Enabling Instance Scale-In Protection:

Instance scale-in protection prevents Auto Scaling from terminating specific instances.

Steps:

Go to the AWS Management Console.

Navigate to EC2 and select 'Auto Scaling Groups.'

Select your Auto Scaling group.

Go to the 'Instance management' tab.

Select the instances you want to protect and click 'Actions.'

Choose 'Enable scale-in protection.'

This ensures that instances are not terminated during troubleshooting.

Question 8

Question Type: **MultipleChoice**

A company needs to monitor the disk utilization of Amazon Elastic Block Store (Amazon EBS) volumes The EBS volumes are attached to Amazon EC2 Linux Instances A SysOps administrator must set up an Amazon CloudWatch alarm that provides an alert when disk utilization increases to more than 80%.

Which combination of steps must the SysOps administrator lake lo meet these requirements? (Select THREE.)

Options:

A- Create an 1AM role that includes the Cloud Watch AgentServerPol icy AWS managed policy Attach me role to the instances

B- Create an 1AM role that includes the CloudWatchApplicationInsightsReadOnlyAccess AWS managed policy. Attach the role to the instances

C- Install and start the CloudWatch agent by using AWS Systems Manager or the command line

D- Install and start the CloudWatch agent by using an 1AM role. Attach the Cloud Watch AgentServerPolicy AWS managed policy to the role.

E- Configure a CloudWatch alarm to enter ALARM state when the disk_used_percent CloudWatch metric is greater than 80%.

F- Configure a CloudWatch alarm to enter ALARM state when the disk_used CloudWatch metric is greater than 80% or when the disk_free CloudWatch metric is less than 20%.

Answer:

A, C, E

Explanation:

Create an IAM role with the CloudWatchAgentServerPolicy:

This policy grants the necessary permissions for the CloudWatch agent to collect and send metrics.

Steps:

Go to the AWS Management Console.

Navigate to IAM and create a new role.

Choose 'EC2' as the trusted entity.

Attach the 'CloudWatchAgentServerPolicy' managed policy to the role.

Attach this IAM role to your EC2 instances.

Install and start the CloudWatch agent:

The CloudWatch agent must be installed and configured to collect disk utilization metrics.

Steps:

Use AWS Systems Manager or SSH to connect to your instances.

Install the CloudWatch agent using the following commands:

sudo yum install amazon-cloudwatch-agent

sudo /opt/aws/amazon-cloudwatch-agent/bin/amazon-cloudwatch-agent-config-wizard

sudo /opt/aws/amazon-cloudwatch-agent/bin/amazon-cloudwatch-agent-ctl -a fetch-config -m ec2 -c file:/path/to/your-config-file.json -s

Start the agent:

sudo systemctl start amazon-cloudwatch-agent

Configure a CloudWatch alarm:

Create an alarm based on the disk_used_percent metric.

Steps:

Go to the AWS Management Console.

Navigate to CloudWatch and select 'Alarms' from the left-hand menu.

Click on 'Create alarm.'

Select the disk_used_percent metric.

Set the threshold to 80% and configure the alarm actions (e.g., sending a notification).

Question 9

Question Type: **MultipleChoice**

A company runs a single-page web application on AWS The application uses Amazon CloudFront lo deliver static content from an Amazon S3 bucket origin The application also uses an Amazon Elastic Kubemetes Service (Amazon EKS) duster to serve API calls

Users sometimes report that the website is not operational, even when monitoring shows that the index page is reachable and that the EKS cluster is healthy. A SysOps administrator must Implement additional monitoring that can delect when the website is not operational before users report the problem.

Which solution will meet these requirements?

Options:

A- Create an Amazon CloudWatch Synthetics heartbeat monitor canary that points to the fully qualified domain name (FQDN) of the website.

B- Create an Amazon CloudWatch Synthetics API canary that monitors the availability of API endpoints from the EKS cluster.

C- Create an Amazon CloudWatch RUM app monitor that points to the fully qualified domain name (FQDN) of the website. Configure the app monitor to collect performance telemetry and JavaScript errors

D- Create an Amazon CloudWatch RUM app monitor that uses the API endpoints from the EKS cluster

Answer:

A

Explanation:

Amazon CloudWatch Synthetics:

CloudWatch Synthetics allows you to create canaries to monitor your endpoints and API calls, simulating user behavior to detect issues before users do.

Steps:

Go to the AWS Management Console.

Navigate to CloudWatch and select 'Synthetics.'

Click on 'Create canary.'

Choose 'Heartbeat monitoring' as the blueprint.

Configure the canary to point to the FQDN of the website.

Set the frequency and retention settings as per your requirement.

Create the canary.

This setup continuously checks the operational status of your website, alerting you if it becomes unreachable or has issues.