Free Questions for SOA-C02 by certsdeals

Shared by Serrano on 09-08-2024

For More Free Questions and Preparation Resources

Check the Links on Last Page

Question 1

Question Type: MultipleChoice

A company receives an alert from an Amazon CloudWatch alarm The alarm indicates that a web application that Is running on Amazon EC2 instances is not responding to requests The EC2 instances have a Red Hat Enterprise Linux operating system and are in an Auto Scaling group. The Auto Scaling group has a minimum capacity of 2 and a maximum capacity of 5.

An Investigation reveals that the web application is experiencing oul-of-memory errors. The company adds memory lo the web application and wants to track operating system memory utilization. A CloudWatch memory metric does not currently exist tor the EC2 Instances in the Auto Scaling group

What should a SysOps administrator do to provide a CloudWatch memory metric for the EC2 instances?

Options:

A- Use an Amazon Machine Image (AMI) that includes the CloudWatch agent.

B- Turn on CloudWatch detailed monitoring

C- Turn on Instance Metadata Service Version 2 (IMOSv2).

D- Use an Amazon Machine Image (AMI) that is based on Amazon Linux.

Answer:

A

Explanation:

Using an AMI with CloudWatch Agent:

The CloudWatch agent can collect memory utilization metrics and send them to CloudWatch.

Steps:

Create or use an existing AMI that includes the CloudWatch agent installed and configured.

Ensure the CloudWatch agent is configured to collect memory metrics.

Use this AMI for instances in the Auto Scaling group.

Question 2

Question Type: MultipleChoice

A company is using AWS to deploy a critical application on a fleet of Amazon EC2 instances The company is rewriting the application because the application failed a security review The application will take 12 months to rewrite While this rewrite happens, the company needs to rotate IAM access keys that the application uses.

A SysOps administrator must implement an automated solution that finds and rotates IAM access Keys that are at least 30 days old. The solution must then continue to rotate the IAM access Keys every 30 days.

Which solution will meet this requirement with the MOST operational efficiency?

Options:

A- Use an AWS Config rule to identify IAM access Keys that are at least 30 days old. Configure AWS Config to invoKe an AWS Systems Manager Automation runbook to rotate the identified IAM access keys.

B- Use AWS Trusted Advisor to identify IAM access Keys that are at least 30 days old. Configure Trusted Advisor to invoke an AWS Systems Manager Automation runbook to rotate the identified IAM access keys

C- Create a script that checks the age of IAM access Keys and rotates them if they are at least 30 days old. Launch an EC2 instance. Schedule the script to run as a cron expression on the EC2 instance every day.

D- Create an AWS Lambda function that checks the age of IAM access keys and rotates them if they are at least 30 days old Use an Amazon EventBridge rule to invoke the Lambda function every time a new IAM access key is created.

Answer:

D

Explanation:

Lambda Function to Rotate IAM Access Keys:

A Lambda function can be used to automate the rotation of IAM access keys based on their age.

Steps:

Write a Lambda function that checks the age of IAM access keys.

The function should rotate keys that are at least 30 days old.

Deploy the Lambda function.

Amazon EventBridge Rule:

EventBridge can trigger the Lambda function periodically and when a new key is created.

Steps:

Create an EventBridge rule that triggers the Lambda function on a schedule (e.g., daily) and on IAM key creation events.

Question 3

Question Type: **MultipleChoice**

Users of a company's internal web application recently experienced application performance issues for a brief period The application includes frontend web servers that run in an Amazon Elastic Kubernetes Service (Amazon EKS) cluster The application also includes a bacKend Amazon Aurora PostgreSQL DB cluster that includes one DB instance.

A SysOps administrator determines that the source of the performance issues was high utilization of the DB cluster. The single writer instance experienced more than 90% utilization for 11 minutes The cause of the high utilization was an automated report that is scheduled to run one time each week

What should the SysOps administrator do to ensure that users do not experience performance Issues each week when the report runs?

Options:

A- Increase the size of the DB instance. Monitor the performance during the next scheduled run of the report

B- Add a reader instance. Change the database connection string of the report application to use the newly created reader instance.

C- Add another writer instance Change the database connection string of the report application to use the newly created writer instance.

D- Configure auto scaling for the DB cluster Set the minimum capacity units, maximum capacity units, and target utilization

Answer:

A

Explanation:

Increasing DB Instance Size:

Increasing the instance size provides more CPU and memory resources, which can help handle higher loads.

Steps:

Go to the AWS Management Console.

Navigate to RDS and select the DB instance.

Modify the instance to increase its size.

Apply the changes during the next maintenance window or immediately if it is a critical issue.

Monitoring Performance:

After resizing, monitor the instance during the next report run to ensure that it handles the load effectively.

Question 4

Question Type: MultipleChoice

A company has a list of pre-appf oved Amazon Machine Images (AMIs) for developers lo use to launch Amazon EC2 instances However, developers are still launching EC2 instances from unapproved AMIs.

A SysOps administrator must implement a solution that automatically terminates any instances that are launched from unapproved AMIs.

Which solution will meet mis requirement?

Options:

A- Set up an AWS Config managed rule to check if instances are running from AMIs that are on the list of pre-approved AMIs. Configure an automatic remediation action so that an AWS Systems Manager Automation runbook terminates any instances that are noncompliant with the rule

B- Store the list of pre-approved AMIs in an Amazon DynamoDB global table that is replicated to all AWS Regions that the developers use. Create Regional EC2 launch templates. Configure the launch templates to check AMIs against the list and to terminate any instances that are not on the list

C- Select the Amazon CloudWatch metric that shows all running instances and the AMIs that the instances were launched from Create a CloudWatch alarm that terminates an instance if the metric shows the use of an unapproved AMI.

D- Create a custom Amazon Inspector finding to compare a running instance's AMI against the list of pre-approved AMIs Create an AWS Lambda function that

terminates instances. Configure Amazon Inspector to report findings of unapproved AMIs to an Amazon Simple Queue Service (Amazon SQS) queue to invoke the Lambda function.

Answer:

A

Explanation:

AWS Config Managed Rule:

AWS Config can be used to assess, audit, and evaluate the configurations of AWS resources. The managed rule can check if instances are launched from approved AMIs.

Steps:

Go to the AWS Management Console.

Navigate to AWS Config.

Create a managed rule that checks for EC2 instances running approved AMIs.

Configure the rule to use a list of approved AMIs.

Automatic Remediation with Systems Manager Automation:

AWS Systems Manager Automation runbooks can automate the process of remediating non-compliant resources.

Steps:

Create a Systems Manager Automation runbook that terminates instances not running approved AMIs.

Attach the runbook to the AWS Config rule for automatic remediation.

Question 5

Question Type: MultipleChoice

A SysOps administrator needs to ensure that an Amazon RDS for PostgreSQL DB instance has available backups The DB instance has automated backups turned on with a backup retention period of 7 days. However, no automated backups for the DB instance have been created in the past month.

What could be the cause of the lack of automated backups?

Options:

A- The Amazon S3 bucket that stores the backups is full

B- The DB instance is in the STORAGE_FULL state

C- The DB instance is not configured for Multi-AZ.

D- The backup retention period must be 30 days.

Answer:

B

Explanation:

STORAGE_FULL State:

When an RDS instance is in the STORAGE_FULL state, automated backups cannot be performed because there is insufficient storage available.

Steps to Resolve:

Go to the AWS Management Console.

Navigate to RDS and select the DB instance.

Check the storage metrics to confirm the STORAGE_FULL state.

Increase the allocated storage for the DB instance to provide sufficient space for automated backups.

Question 6

Question Type: MultipleChoice

A company is uploading important files as objects to Amazon S3 The company needs to be informed if an object is corrupted during the upload

What should a SysOps administrator do to meet this requirement?

Options:

A- Pass the Content-Disposition value as a request body during the object upload.

B- Pass the Content-MD5 value as a request header during the object upload.

C- Pass x-amz-objecWock-mode as a request header during the object upload

D- Pass x-amz-server-side-encryption-customer-algorithm as a request body during the object upload.

Answer:

B

Explanation:

Content-MD5 Header:

The Content-MD5 header provides an MD5 checksum of the object being uploaded. Amazon S3 uses this checksum to verify the integrity of the object.

Steps:

When uploading an object to S3, calculate the MD5 checksum of the object.

Include the Content-MD5 header with the base64-encoded MD5 checksum value in the upload request.

This ensures that S3 can detect if the object is corrupted during the upload process.

Question 7

Question Type: **MultipleChoice**

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 wilt 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?

Options:

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.

Answer:

A

Explanation:

Cluster Placement Group:

Cluster placement groups are used to ensure low-latency networking between EC2 instances. They place instances physically close to each other within the same Availability Zone.

Steps:

Go to the AWS Management Console.

Navigate to EC2 and select 'Placement Groups.'

Create a new cluster placement group.

Back up the existing EC2 instance to an AMI.

Launch new EC2 instances from the AMI into the cluster placement group.

Ensure all instances are in the same Availability Zone.

Question 8

Question Type: MultipleChoice

A company runs an application on hundreds of Amazon EC2 instances in three Availability Zones The application calls a third-parly API over the public internet A SysOps administrator must provide the third party with a list of static IP addresses so that the third party can allow traffic from the application

Which solution will meet these requirements?

Options:

A- Add a NAT gateway in the public subnet of each Availability Zone. Make the NAT gateway the default route of all private subnets In those Availability Zones.

B- Allocate one Elastic IP address in each Availability Zone. Associate the Elastic IP address with all the instances in the Availability Zone

C- Place the instances behind a Network Load Balancer (NLB). Send the traffic to the interne! through the private IP address of the NLB

D- Update the main route table to send the traffic to the internet through an Elastic IP address that is assigned to each instance.

Answer:

A

Explanation:

NAT Gateway Setup:

A NAT gateway allows instances in a private subnet to connect to the internet or other AWS services, but prevents the internet from initiating a connection with those instances.

Steps:

Go to the AWS Management Console.

Navigate to VPC and select 'NAT Gateways.'

Create a NAT gateway in the public subnet of each Availability Zone.

Allocate an Elastic IP address to each NAT gateway.

Update the route tables for the private subnets to route internet-bound traffic to the NAT gateways.