

1. A company hosts a website on multiple Amazon EC2 instances that run in an Auto Scaling group. Users are reporting slow responses during peak times between 6 PM and 11 PM every weekend. A SysOps administrator must implement a solution to improve performance during these peak times. What is the MOST operationally efficient solution that meets these requirements?

A. Create a scheduled Amazon EventBridge (Amazon CloudWatch Events) rule to invoke an AWS Lambda function to increase the desired capacity before peak times.

B. Configure a scheduled scaling action with a recurrence option to change the desired capacity before and after peak times. Most Voted

C. Create a target tracking scaling policy to add more instances when memory utilization is above 70%.

D. Configure the cooldown period for the Auto Scaling group to modify desired capacity before and after peak times.

2. A company hosts its website on Amazon EC2 instances behind an Application Load Balancer. The company manages its DNS with Amazon Route 53, and wants to point its domain's zone apex to the website. Which type of record should be used to meet these requirements?

A. An AAAA record for the domain's zone apex

B. An A record for the domain's zone apex

C. A CNAME record for the domain's zone apex

D. An alias record for the domain's zone apex Most Voted

3. A company is running distributed computing software to manage a fleet of 20 Amazon EC2 instances for calculations. The fleet includes 2 control nodes and 18 task nodes to run the calculations. Control nodes can automatically start the task nodes. Currently, all the nodes run on demand. The control nodes must be available 24 hours a day, 7 days a week. The task nodes run for 4 hours each day. A SysOps administrator needs to optimize the cost of this solution. Which combination of actions will meet these requirements? (Choose two.)

A. Purchase EC2 Instance Savings Plans for the control nodes. Most Voted

B. Use Dedicated Hosts for the control nodes.

C. Use Reserved Instances for the task nodes.

D. Use Spot Instances for the control nodes. Use On-Demand Instances if there is no Spot availability.

E. Use Spot Instances for the task nodes. Use On-Demand Instances if there is no Spot availability.
Most Voted

4. A company has an Auto Scaling group of Amazon EC2 instances that scale based on average CPU utilization. The Auto Scaling group events log indicates an `InsufficientInstanceCapacity` error. Which actions should a SysOps administrator take to remediate this issue? (Choose two.)

A. Change the instance type that the company is using. Most Voted

B. Configure the Auto Scaling group in different Availability Zones. Most Voted

C. Configure the Auto Scaling group to use different Amazon Elastic Block Store (Amazon EBS) volume sizes.

D. Increase the maximum size of the Auto Scaling group.

E. Request an increase in the instance service quota.

5. An Amazon CloudFront distribution has a single Amazon S3 bucket as its origin. A SysOps administrator must ensure that users can access the S3 bucket only through requests from the CloudFront endpoint. Which solution will meet these requirements?

A. Configure S3 Block Public Access on the S3 bucket. Update the S3 bucket policy to allow the `GetObject` action from only the CloudFront distribution.

B. Configure Origin Shield in the CloudFront distribution. Update the CloudFront origin to include a custom `Origin_Shield` header.

C. Create an origin access identity (OAI). Assign the OAI to the CloudFront distribution. Update the S3 bucket policy to restrict access to the OAI.

D. Create an origin access identity (OAI). Assign the OAI to the S3 bucket. Update the CloudFront origin to include a custom Origin header with the OAI value.