

DOP-C01^{Q&As}

AWS Certified DevOps Engineer - Professional (DOP-C01)

Pass Amazon DOP-C01 Exam with 100% Guarantee

Free Download Real Questions & Answers **PDF** and **VCE** file from:

<https://www.pass2lead.com/aws-devops-engineer-professional.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

When thinking of AWS OpsWorks, which of the following is true?

- A. Stacks have many layers, layers have many instances.
- B. Instances have many stacks, stacks have many layers.
- C. Layers have many stacks, stacks have many instances.
- D. Layers have many instances, instances have many stacks.

Correct Answer: A

The stack is the core AWS OpsWorks component. It is basically a container for AWS resources--Amazon EC2 instances, Amazon RDS database instances, and so on--that have a common purpose and should be logically managed together.

You define the stack's constituents by adding one or more layers. A layer represents a set of Amazon EC2 instances that serve a particular purpose, such as serving applications or hosting a database server. An instance represents a single

computing resource, such as an Amazon EC2 instance.

Reference: <http://docs.aws.amazon.com/opsworks/latest/userguide/welcome.html>

QUESTION 2

A DevOps engineer has been tasked with ensuring that all Amazon S3 buckets, except for those with the word "public" in the name, allow access only to authorized users utilizing S3 bucket policies. The security team wants to be notified when a bucket is created without the proper policy and for the policy to be automatically updated.

Which solutions will meet these requirements?

- A. Create a custom AWS Config rule that will trigger an AWS Lambda function when an S3 bucket is created or updated. Use the Lambda function to look for S3 buckets that should be private, but that do not have a bucket policy that enforces privacy. When such a bucket is found, invoke a remediation action and use Amazon SNS to notify the security team.
- B. Create an Amazon EventBridge (Amazon CloudWatch Events) rule that triggers when an S3 bucket is created. Use an AWS Lambda function to determine whether the bucket should be private. If the bucket should be private, update the PublicAccessBlock configuration. Configure a second EventBridge (CloudWatch Events) rule to notify the security team using Amazon SNS when PutBucketPolicy is called.
- C. Create an Amazon S3 event notification that triggers when an S3 bucket is created that does not have the word "public" in the name. Define an AWS Lambda function as a target for this notification and use the function to apply a new default policy to the S3 bucket. Create an additional notification with the same filter and use Amazon SNS to send an email to the security team.
- D. Create an Amazon EventBridge (Amazon CloudWatch Events) rule that triggers when a new object is created in a bucket that does not have the word "public" in the name. Target and use an AWS Lambda function to update the PublicAccessBlock configuration. Create an additional notification with the same filter and use Amazon SNS to send an email to the security team.

Correct Answer: D

QUESTION 3

An Engineering team manages a Node.js e-commerce application. The current environment consists of the following components:

1.
Amazon S3 buckets for storing content
2.
Amazon EC2 for the front-end web servers
3.
AWS Lambda for executing image processing
4.
Amazon DynamoDB for storing session-related data

The team expects a significant increase in traffic to the site. The application should handle the additional load without interruption. The team ran initial tests by adding new servers to the EC2 front-end to handle the larger load, but the instances took up to 20 minutes to become fully configured. The team wants to reduce this configuration time.

What changes will the Engineering team need to implement to make the solution the MOST resilient and highly available while meeting the expected increase in demand?

- A. Use AWS OpsWorks to automatically configure each new EC2 instance as it is launched. Configure the EC2 instances by using an Auto Scaling group behind an Application Load Balancer across multiple Availability Zones. Implement Amazon DynamoDB Auto Scaling. Use Amazon Route 53 to point the application DNS record to the Application Load Balancer.
- B. Deploy a fleet of EC2 instances, doubling the current capacity, and place them behind an Application Load Balancer. Increase the Amazon DynamoDB read and write capacity units. Add an alias record that contains the Application Load Balancer endpoint to the existing Amazon Route 53 DNS record that points to the application.
- C. Configure Amazon CloudFront and have its origin point to Amazon S3 to host the web application. Implement Amazon DynamoDB Auto Scaling. Use Amazon Route 53 to point the application DNS record to the CloudFront DNS name.
- D. Use AWS Elastic Beanstalk with a custom AMI including all web components. Deploy the platform by using an Auto Scaling group behind an Application Load Balancer across multiple Availability Zones. Implement Amazon DynamoDB Auto Scaling. Use Amazon Route 53 to point the application DNS record to the Elastic Beanstalk load balancer.

Correct Answer: D

QUESTION 4

Your application's Auto Scaling Group scales up too quickly, too much, and stays scaled when traffic decreases. What

should you do to fix this?

- A. Set a longer cooldown period on the Group, so the system stops overshooting the target capacity. The issue is that the scaling system doesn't allow enough time for new instances to begin servicing requests before measuring aggregate load again.
- B. Calculate the bottleneck or constraint on the compute layer, then select that as the new metric, and set the metric thresholds to the bounding values that begin to affect response latency.
- C. Raise the CloudWatch Alarms threshold associated with your autoscaling group, so the scaling takes more of an increase in demand before beginning.
- D. Use larger instances instead of lots of smaller ones, so the Group stops scaling out so much and wasting resources as the OS level, since the OS uses a higher proportion of resources on smaller instances.

Correct Answer: B

Systems will always over-scale unless you choose the metric that runs out first and becomes constrained first. You also need to set the thresholds of the metric based on whether or not latency is affected by the change, to justify adding capacity instead of wasting money. Reference: http://docs.aws.amazon.com/AutoScaling/latest/DeveloperGuide/policy_creating.html

QUESTION 5

An application is deployed on Amazon EC2 instances running in an Auto Scaling group. During the bootstrapping process, the instances register their private IP addresses with a monitoring system. The monitoring system performs health checks frequently by sending ping requests to those IP addresses and sending alerts if an instance becomes non-responsive.

The existing deployment strategy replaces the current EC2 instances with new ones. A DevOps Engineer has noticed that the monitoring system is sending false alarms during a deployment, and is tasked with stopping these false alarms.

Which solution will meet these requirements without affecting the current deployment method?

- A. Define an Amazon CloudWatch Events target, an AWS Lambda function, and a lifecycle hook attached to the Auto Scaling group. Configure CloudWatch Events to invoke Amazon SNS to send a message to the Systems Administrator group for remediation.
- B. Define an AWS Lambda function and a lifecycle hook attached to the Auto Scaling group. Configure the lifecycle hook to invoke the Lambda function, which removes the entry of the private IP from the monitoring system upon instance termination.
- C. Define an Amazon CloudWatch Events target, an AWS Lambda function, and a lifecycle hook attached to the Auto Scaling group. Configure CloudWatch Events to invoke the Lambda function, which removes the entry of the private IP from the monitoring system upon instance termination.
- D. Define an AWS Lambda function that will run a script when instance termination occurs in an Auto Scaling group. The script will remove the entry of the private IP from the monitoring system.

Correct Answer: C

Reference: <https://aws.amazon.com/blogs/compute/using-aws-lambda-with-auto-scaling-lifecycle-hooks/>

[DOP-C01 Practice Test](#)

[DOP-C01 Study Guide](#)

[DOP-C01 Braindumps](#)