

# Amazon Web Services

## Exam Questions DOP-C01

AWS Certified DevOps Engineer- Professional



#### NEW QUESTION 1

A company has a mission-critical application on AWS that uses automatic scaling. The company wants the deployment lifecycle to meet the following parameters

- The application must be deployed one instance at a time to ensure the remaining fleet continues to serve traffic.
- the application is CPU intensive and must be closely monitored
- the deployment must automatically roll back if the CPU utilization of the deployment instance exceeds 85% Which solution will meet these requirements'?

- A. Use AWS CloudFormation to create an AWS Step Functions state machine and Auto Scaling lifecycle hooks to move to one instance at a time into a wait state
- B. Use AWS Systems Manager automation to deploy the update to each instance and move it back into the Auto Scaling group using the heartbeat timeout
- C. Use AWS CodeDeploy with Amazon EC2 Auto Scaling Configure an alarm tied to the CPU utilization metric Use the CodeDeployDefault OneAtATime configuration as a deployment strategy Configure automatic rollbacks within the deployment group to roll back the deployment if the alarm thresholds are breached
- D. Use AWS Elastic Beanstalk for load balancing and AWS Auto Scaling Configure an alarm tied to the CPU utilization metric Configure rolling deployments with a fixed batch size of one instance Enable enhanced health to monitor the status of the deployment and roll back based on the alarm previously created
- E. Use AWS Systems Manager to perform a blue/green deployment with Amazon EC2 Auto Scaling Configure an alarm tied to the CPU utilization metric Deploy updates one at a time Configure automatic rollbacks within the Auto Scaling group to roll back the deployment if the alarm thresholds are breached.

**Answer: B**

#### NEW QUESTION 2

An ecommerce company is receiving reports that its order history page is experiencing delays in reflecting the processing status of orders. The order processing system consists of an AWS Lambda function using reserved concurrency. The Lambda function processes order messages from an Amazon SQS queue and inserts processed orders into an Amazon DynamoDB table. The DynamoDB table has Auto Scaling enabled for read and write capacity.

Which actions will diagnose and resolve the delay? (Select TWO.)

- A. Check the ApproximateAgeOfOldestMessage metric for the SQS queue and increase the Lambda function concurrency limit.
- B. Check the ApproximateAgeOfOldestMessage metric for the SQS queue and configure a redrive policy on the SQS queue.
- C. Check the NumberOfMessagesSent metric for the SQS queue and increase the SQS queue visibility timeout.
- D. Check the ThrottledWriteRequests metric for the DynamoDB table and increase the maximum write capacity units for the table's Auto Scaling policy.
- E. Check the Throttles metric for the Lambda function and increase the Lambda function timeout.

**Answer: AB**

#### NEW QUESTION 3

A new zero-day vulnerability was found in OpenSSL requiring the immediate patching of a production web fleet running on Amazon Linux. Currently, OS updates are performed manually on a monthly basis and deployed using updates to the production Auto Scaling Group's launch configuration.

Which method should a DevOps Engineer use to update packages in-place without downtime?

- A. Use AWS CodePipeline and AWS CodeBuild to generate new copies of these packages, and update the Auto Scaling group's launch configuration.
- B. Use AWS Inspector to run "yum upgrade" on all running production instances, and manually update the AMI for the next maintenance window.
- C. Use Amazon EC2 Run Command to issue a package update command to all running production instances, and update the AMI for future deployments.
- D. Define a new AWS OpsWorks layer to match the running production instances, and use a recipe to issue a package update command to all running production instances.

**Answer: C**

#### Explanation:

<https://aws.amazon.com/blogs/aws/ec2-run-command-is-now-a-cloudwatch-events-target/>

" EC2 Run Command is part of EC2 Systems Manager. It allows you to operate on collections of EC2 instances and on-premises servers reliably and at scale, in a controlled and selective fashion. You can run scripts, install software, collect metrics and log files, manage patches, and much more, on both Windows and Linux."

#### NEW QUESTION 4

According to Information Security Policy, changes to the contents of objects inside production Amazon S3 bucket that contain encrypted secrets should only be made by a trusted group of administrators.

How should a DevOps Engineer create real-time, automated checks to meet this requirement?

- A. Create an AWS Lambda function that is triggered by Amazon S3 data events for object changes and that also checks the IAM user's membership in an administrator's IAM role.
- B. Create a periodic AWS Config rule to query Amazon S3 Logs for changes and to check the IAM user's membership in an administrator's IAM role.
- C. Create a metrics filter for Amazon CloudWatch logs to check for Amazon S3 bucket-level permission changes and to check the IAM user's membership in an administrator's IAM role.
- D. Create a periodic AWS Config rule to query AWS CloudTrail logs for changes to the Amazon S3 bucket-level permissions and to check the IAM user's membership in an administrator's IAM role.

**Answer: A**

#### NEW QUESTION 5

An Amazon EC2 instance with no internet access is running in a Virtual Private Cloud (VPC) and needs to download an object from a restricted Amazon S3 bucket. When the DevOps Engineer tries to gain access to the object, an Access Denied error is received.

What are the possible causes for this error? (Select THREE.)

- A. The S3 bucket default encryption is enabled.
- B. There is an error in the S3 bucket policy.
- C. There is an error in the VPC endpoint policy.
- D. The object has been moved to Amazon Glacier.
- E. There is an error in the IAM role configuration.
- F. S3 versioning is enabled.

Answer: BCE

#### NEW QUESTION 6

A company is creating a software solution that executes a specific parallel-processing mechanism. The software can scale to tens of servers in some special scenarios. This solution uses a proprietary library that is license-based, requiring that each individual server have a single, dedicated license installed. The company has 200 licenses and is planning to run 200 server nodes concurrently at most.

The company has requested the following features:

"ç A mechanism to automate the use of the licenses at scale. "ç Creation of a dashboard to use in the future to verify which licenses are available at any moment. What is the MOST effective way to accomplish these requirements?

- A. Upload the licenses to a private Amazon S3 bucket
- B. Create an AWS CloudFormation template with a Mappings section for the license
- C. In the template, create an Auto Scaling group to launch the server
- D. In the user data script, acquire an available license from the Mappings section
- E. Create an Auto Scaling lifecycle hook, then use it to update the mapping after the instance is terminated.
- F. Upload the licenses to an Amazon DynamoDB table
- G. Create an AWS CloudFormation template that uses an Auto Scaling group to launch the server
- H. In the user data script, acquire an available license from the DynamoDB table
- I. Create an Auto Scaling lifecycle hook, then use it to update the mapping after the instance is terminated.
- J. Upload the licenses to a private Amazon S3 bucket
- K. Populate an Amazon SQS queue with the list of licenses stored in S3. Create an AWS CloudFormation template that uses an Auto Scaling group to launch the server
- L. In the user data script acquire an available license from SQS
- M. Create an Auto Scaling lifecycle hook, then use it to put the license back in SQS after the instance is terminated.
- N. Upload the licenses to an Amazon DynamoDB table
- O. Create an AWS CLI script to launch the servers by using the parameter --count, with min:max instances to launch
- P. In the user data script, acquire an available license from the DynamoDB table
- Q. Monitor each instance and, in case of failure, replace the instance, then manually update the DynamoDB table.

Answer: D

#### NEW QUESTION 7

An application running on a set of Amazon EC2 instances in an Auto Scaling group requires a configuration file to operate. The instances are created and maintained with AWS CloudFormation. A DevOps engineer wants the instances to have the latest configuration file when launched, and wants changes to the configuration file to be reflected on all the instances with a minimal delay when the CloudFormation template is updated. Company policy requires that application configuration files be maintained along with AWS infrastructure configuration files in source control.

Which solution will accomplish this?

- A. In the CloudFormation template, add an AWS Config rule
- B. Place the configuration file content in the rule's InputParameters property, and set the Scope property to the EC2 Auto Scaling group
- C. Add an AWS Systems Manager Resource Data Sync resource to the template to poll for updates to the configuration.
- D. In the CloudFormation template, add an EC2 launch template resource
- E. Place the configuration file content in the launch template
- F. Configure the cfn-init script to run when the instance is launched, and configure the cfn-hup script to poll for updates to the configuration.
- G. In the CloudFormation template, add an EC2 launch template resource
- H. Place the configuration file content in the launch template
- I. Add an AWS Systems Manager Resource Data Sync resource to the template to poll for updates to the configuration.
- J. In the CloudFormation template, add CloudFormation init metadata
- K. Place the configuration file content in the metadata
- L. Configure the cfn-init script to run when the instance is launched, and configure the cfn-hup script to poll for updates to the configuration.

Answer: B

#### NEW QUESTION 8

A DevOps engineer is tasked with migrating Docker containers used for a workload to AWS. The solution must allow for changes to be deployed into development and test environments automatically by updating each container and checking it into a container registry. Once the containers are pushed, they must be deployed automatically.

Which solution will meet these requirements?

- A. Store container images in Amazon S3. Run the containers in AWS Elastic Beanstalk using a multicontainer Docker environment
- B. Configure Elastic Beanstalk to redeploy the containers if it detects a new version in Amazon S3.
- C. Store container images in AWS Artifact. Use AWS CodePipeline to trigger a deployment if a new container version is created.
- D. Use AWS CodeDeploy to deploy new containers to Amazon EKS.
- E. Store container images in Amazon ECR. Use AWS CodePipeline to trigger a deployment if a new container version is created. Use AWS CodeDeploy to deploy the image to AWS Fargate.
- F. Store container images in Docker Hub. Install Docker on an Amazon EC2 instance and use AWS CodePipeline and AWS CodeDeploy to deploy any new containers.

Answer: C

#### NEW QUESTION 9

A DevOps Engineer is developing a deployment strategy that will allow for data-driven decisions before a feature is fully approved for general availability. The current deployment process uses AWS CloudFormation and blue/green-style deployments. The development team has decided that customers should be randomly assigned to groups, rather than using a set percentage, and redirects should be avoided.

What process should be followed to implement the new deployment strategy?

- A. Configure Amazon Route 53 weighted records for the blue and green stacks, with 50% of traffic configured to route to each stack.
- B. Configure Amazon CloudFront with an AWS Lambda@Edge function to set a cookie when CloudFront receives a request.
- C. Assign the user to a version A or B, and configure the web server to redirect to version A or B.

- D. Configure Amazon CloudFront with an AWS Lambda@Edge function to set a cookie when CloudFront receives a request.
- E. Assign the user to a version A or B, then return the corresponding version to the viewer.
- F. Configure Amazon Route 53 with an AWS Lambda function to set a cookie when Amazon CloudFront receives a request.
- G. Assign the user to version A or B, then return the corresponding version to the viewer.

**Answer: C**

**Explanation:**

[https://docs.aws.amazon.com/zh\\_cn/AmazonCloudFront/latest/DeveloperGuide/lambda-examples.html](https://docs.aws.amazon.com/zh_cn/AmazonCloudFront/latest/DeveloperGuide/lambda-examples.html)

**NEW QUESTION 10**

A Development team is adding a new country to an e-commerce application. This addition requires that new application features be added to the shipping component of the application. The team has not decided if all new features should be added, as some will take approximately six weeks to build. While the final decision on the shipping component features is being made, other team members are continuing to work on other features of the application. Based on this situation, how should the application feature deployments be managed?

- A. Add the code updates as commits to the release branch
- B. The team can delay the deployment until all features are ready.
- C. Add the code updates as commits to a feature branch
- D. Merge the commits to a release branch as features are ready.
- E. Add the code updates as a single commit when a feature is ready
- F. Tag this commit with "new-country."
- G. Create a new repository named "new-country". Commit all the code changes to the new repository.

**Answer: B**

**NEW QUESTION 11**

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
- B. Use the Lambda function to look for S3 buckets that should be private, but that do not have a bucket policy that enforces privacy
- C. When such a bucket is found, invoke a remediation action and use Amazon SNS to notify the security team.
- D. Create an Amazon EventBridge (Amazon CloudWatch Events) rule that triggers when an S3 bucket is created
- E. Use an AWS Lambda function to determine whether the bucket should be private
- F. If the bucket should be private, update the PublicAccessBlock configuration
- G. Configure a second EventBridge (CloudWatch Events) rule to notify the security team using Amazon SNS when PutBucketPolicy is called.
- H. Create an Amazon S3 event notification that triggers when an S3 bucket is created that does not have the word "public" in the name
- I. 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
- J. Create an additional notification with the same filter and use Amazon SNS to send an email to the security team.
- K. 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
- L. Target and use an AWS Lambda function to update the PublicAccessBlock configuration
- M. Create an additional notification with the same filter and use Amazon SNS to send an email to the security team.

**Answer: A**

**NEW QUESTION 12**

A healthcare provider has a hybrid architecture that includes 120 on-premises VMware servers running RedHat and 50 Amazon EC2 instances running Amazon Linux. The company is in the middle of an all-in migration to AWS and wants to implement a solution for collecting information from the on-premises virtual machines and the EC2 instances for data analysis. The information includes:

- Operating system type and version
- Data for installed applications
- Network configuration information, such as MAC and IP addresses
- Amazon EC2 instance AMI ID and IAM profile

How can these requirements be met with the LEAST amount of administration?

- A. Write a shell script to run as a cron job on EC2 instances to collect and push the data to Amazon S3. For on-premises resources, use VMware vSphere to collect the data and write it into a file gateway for storing the data in S3. Finally, use Amazon Athena on the S3 bucket for analytics.
- B. Use a script on the on-premises virtual machines as well as the EC2 instances to gather and push the data into Amazon S3, and then use Amazon Athena for analytics.
- C. Install AWS Systems Manager agents on both the on-premises virtual machines and the EC2 instances. Enable inventory collection and configure resource data sync to an Amazon S3 bucket to analyze the data with Amazon Athena.
- D. Use AWS Application Discovery Service for deploying Agentless Discovery Connector in the VMware environment and Discovery Agents on the EC2 instances for collecting the data
- E. Then use the AWS Migration Hub Dashboard for analytics.

**Answer: C**

**NEW QUESTION 13**

A company is using an AWS CodeBuild project to build and package an application. The packages are copied to a shared Amazon S3 bucket before being deployed across multiple AWS accounts.

The buildspec.yml file contains the following:

```
version: 0.2
phases:
  build:
    commands:
      - go build -o myapp
  post_build:
    commands:
      - aws s3 cp --acl authenticated-read myapp s3://artifacts/
```

The DevOps Engineer has noticed that anybody with an AWS account is able to download the artifacts. What steps should the DevOps Engineer take to stop this?

- A. Modify the post\_build to command to use ""-acl public-read and configure a bucket policy that grants read access to the relevant AWS accounts only.
- B. Configure a default ACL for the S3 bucket that defines the set of authenticated users as the relevant AWS accounts only and grants read-only access.
- C. Create an S3 bucket policy that grants read access to the relevant AWS accounts and denies read access to the principal ""
- D. Modify the post\_build command to remove ""-acl authenticated-read and configure a bucket policy that allows read access to the relevant AWS accounts only.

**Answer: D**

#### NEW QUESTION 14

A company is testing a web application that runs on Amazon EC2 instances behind an Application Load Balancer. The instances run in an Auto Scaling group across multiple Availability Zones. The company uses a blue/green deployment process with immutable instances when deploying new software. During testing, users are being automatically logged out of the application at random times. Testers also report that, when a new version of the application is deployed, all users are logged out. The Development team needs a solution to ensure users remain logged in across scaling events and application deployments. What is the MOST efficient way to ensure users remain logged in?

- A. Enable smart sessions on the load balancer and modify the application to check for an existing session.
- B. Enable session sharing on the load balancer and modify the application to read from the session store.
- C. Store user session information in an Amazon S3 bucket and modify the application to read session information from the bucket.
- D. Modify the application to store user session information in an Amazon ElastiCache cluster.

**Answer: D**

#### NEW QUESTION 15

A company is required to collect user consent to a privacy agreement. An application is deployed in six AWS Regions with two in North America, two in Europe, and two in Asia with a user base of 20-30 million users. The company needs to read and write data related to each user's response, and ensure the responses are available in all six Regions.

What solution will satisfy these requirements while MINIMIZING latency?

- A. Implement Amazon Aurora Global Database in each of the six Regions.
- B. Implement Amazon DocumentDB (with MongoDB compatibility) in each of the six Regions.
- C. Implement Amazon DynamoDB global tables in each of the six Regions.
- D. Implement Amazon ElastiCache for Redis replication group in each of the six Regions.

**Answer: C**

#### NEW QUESTION 16

A publishing company used AWS Elastic Beanstalk, Amazon S3, and Amazon DynamoDB to develop a web application. The web application has increased dramatically in popularity, resulting in unpredictable spikes in traffic. A DevOps Engineer has noted that 90% of the requests are duplicate read requests. How can the Engineer improve the performance of the website?

- A. Use Amazon ElastiCache for Redis to cache repeated read requests to DynamoDB and AWS Elemental MediaStore to cache images stored in S3.
- B. Use Amazon ElastiCache for Memcached to cache repeated read requests to DynamoDB and Varnish to cache images stored in S3.
- C. Use DynamoDB Accelerator to cache repeated read requests to DynamoDB and Amazon CloudFront to cache images stored in S3.
- D. Use DynamoDB Streams to cache repeated read requests to DynamoDB and API Gateway to cache images stored in S3.

**Answer: C**

#### Explanation:

<https://aws.amazon.com/blogs/aws/amazon-dynamodb-accelerator-dax-in-memory-caching-for-read-intensive-w>  
<https://aws.amazon.com/dynamodb/dax/>

#### NEW QUESTION 17

A large enterprise is deploying a web application on AWS. The application runs on Amazon EC2 instances behind an Application Load Balancer. The instances run in an Auto Scaling group across multiple Availability Zones. The application stores data in an Amazon RDS Oracle DB instance and Amazon DynamoDB. There are separate environments for development, testing, and production.

What is the MOST secure and flexible way to obtain password credentials during deployment?

- A. Retrieve an access key from an AWS Systems Manager SecureString parameter to access AWS service
- B. Retrieve the database credentials from a Systems Manager SecureString parameter.
- C. Launch the EC2 instances with an EC2 IAM role to access AWS service
- D. Retrieve the database credentials from AWS Secrets Manager.
- E. Retrieve an access key from an AWS Systems Manager plaintext parameter to access AWS services. Retrieve the database credentials from a Systems Manager SecureString parameter.
- F. Launch the EC2 instances with an EC2 IAM role to access AWS service
- G. Store the database passwords in an encrypted config file with the application artifacts.

**Answer: B**

#### Explanation:

<https://www.1strategy.com/blog/2019/02/28/aws-parameter-store-vs-aws-secrets-manager/>

#### NEW QUESTION 18

A company is developing a web application's infrastructure using AWS CloudFormation. The database engineering team maintains the database resources in a CloudFormation template, and the software development team maintains the web application resources in a separate CloudFormation template. As the scope of the application grows, the software development team needs to use resources maintained by the database engineering team. However, both teams have their own review and lifecycle management processes that they want to keep. Both teams also require resource-level change-set reviews. The software development team would like to deploy changes to this template using their CI/CD pipeline. Which solution will meet these requirements?

- A. Create a stack export from the database CloudFormation template and import those references into the web application CloudFormation template.
- B. Create a CloudFormation nested stack to make cross-stack resource references and parameters available in both stacks.
- C. Create a CloudFormation stack set to make cross-stack resource references and parameters available in both stacks.
- D. Create input parameters in the web application CloudFormation template and pass resource names and IDs from the database stack.

**Answer:** A

#### NEW QUESTION 19

A web application for healthcare services runs on Amazon EC2 instances behind an ELB Application Load Balancer. The instances run in an Amazon EC2 Auto Scaling group across multiple Availability Zones. A DevOps Engineer must create a mechanism in which an EC2 instance can be taken out of production so its system logs can be analyzed for issues to quickly troubleshoot problems on the web tier.

How can the Engineer accomplish this task while ensuring availability and minimizing downtime?

- A. Implement EC2 Auto Scaling groups cooldown period.
- B. Use EC2 instance metadata to determine the instance state, and an AWS Lambda function to snapshot Amazon EBS volumes to preserve system logs.
- C. Implement Amazon CloudWatch Events rule.
- D. Create an AWS Lambda function that can react to an instance termination to deploy the CloudWatch Logs agent to upload the system and access logs to Amazon S3 for analysis.
- E. Terminate the EC2 instances manually.
- F. The Auto Scaling service will upload all log information to CloudWatch Logs for analysis prior to instance termination.
- G. Implement EC2 Auto Scaling groups with lifecycle hook.
- H. Create an AWS Lambda function that can modify an EC2 instance lifecycle hook into a standby state, extract logs from the instance through a remote script execution, and place them in an Amazon S3 bucket for analysis.

**Answer:** D

#### NEW QUESTION 20

The management team at a company with a large on-premises OpenStack environment wants to move non-production workloads to AWS. An AWS Direct Connect connection has been provisioned and configured to connect the environments. Due to contractual obligations, the production workloads must remain on-premises, and will be moved to AWS after the next contract negotiation. The company follows Center for Internet Security (CIS) standards for hardening images; this configuration was developed using the company's configuration management system.

Which solution will automatically create an identical image in the AWS environment without significant overhead?

- A. Write an AWS CloudFormation template that will create an Amazon EC2 instance.
- B. Use cloud-init to install the configuration management agent, use cfn-wait to wait for configuration management to successfully apply, and use an AWS Lambda-backed custom resource to create the AMI.
- C. Log in to the console, launch an Amazon EC2 instance, and install the configuration management agent. When changes are applied through the configuration management system, log in to the console and create a new AMI from the instance.
- D. Create a new AWS OpsWorks layer and mirror the image hardening standard.
- E. Use this layer as the baseline for all AWS workloads.
- F. When a change is made in the configuration management system, a job in Jenkins is triggered to use the VM Import command to create an Amazon EC2 instance in the Amazon VPC.
- G. Use lifecycle hooks to launch an AWS Lambda function to create the AMI.

**Answer:** D

#### Explanation:

<https://www.brad-x.com/2015/10/01/importing-an-openstack-vm-into-amazon-ec2/> <https://aws.amazon.com/ec2/vm-import/>

#### NEW QUESTION 21

.....

## **Thank You for Trying Our Product**

### **We offer two products:**

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

2nd - Questions and Answers in PDF Format

### **DOP-C01 Practice Exam Features:**

- \* DOP-C01 Questions and Answers Updated Frequently
- \* DOP-C01 Practice Questions Verified by Expert Senior Certified Staff
- \* DOP-C01 Most Realistic Questions that Guarantee you a Pass on Your First Try
- \* DOP-C01 Practice Test Questions in Multiple Choice Formats and Updates for 1 Year

**100% Actual & Verified — Instant Download, Please Click**  
**[Order The DOP-C01 Practice Test Here](#)**