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.leads4pass.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

You set up a scalable continuous integration platform on AWS. The platform consists of a master node that can delegate project build jobs to multiple slave nodes, all running on Amazon EC2. The build output will be stored in Amazon S3.

You always have five slave nodes deployed. Each slave node can handle 10 build jobs simultaneously. Your master node publishes a custom Amazon CloudWatch metric with the name "RunningBuildiobs" that Slows you to programmatically

track how many build jobs are running across your platform.

Which two configuration options will allow you to flexibly scale your platform to support more than 50 simultaneous build jobs while minimizing costs? (Choose two.)

A. Place your fleet of slave nodes in an Auto Scaling group. Configure a CloudWatch alarm that triggers an Auto Scaling policy to launch Amazon EC2 Instances when "RunningBuildJobs" is greater than 45 for more than five minutes.

B. Configure a CloudWatch alarm that sends an alert when "RunningBuildJobs" is greater than 45 for more than five minutes. Use Amazon Simple Queue Service to process additional build jobs when the CloudWatch alarm is triggered.

C. Configure your fleet of slave nodes to fully utilize all of your purchased Amazon EC2 Heavy Utilization Reserved Instances. Configure a CloudWatch alarm that launches new Amazon EC2 instances when "RunningBuildJobs" is less than 40 for more than five minutes.

D. Run your fleet of slave nodes in an Auto Scaling group. Configure a Cloudwatch alarm that launches new Amazon EC2 Dedicated Instances when "RunningBuildJobs" is less than 40 for more than five minutes.

E. Place your fleet of slave nodes in an Auto Scaling group. Configure a CloudWatch alarm that triggers an Auto Scaling policy to terminate Amazon EC2 instances when "RunningBuildJobs" is less than 40 for more than five minutes.

Correct Answer: AE

QUESTION 2

You have an Auto Sealing group of Instances that processes messages from an Amazon Simple Queue Service (SQS) queue. The group scales on the size of the queue. Processing Involves calling a third-party web service. The web service

is complaining about the number of failed and repeated calls it is receiving from you. You have noticed that when the group scales in, instances are being terminated while they are processing.

What cost-effective solution can you use to reduce the number of incomplete process attempts?

A. Create a new Auto Scaling group with minimum and maximum of 2 and instances running web proxy software. Configure the VPC route table to route HTTP traffic to these web proxies.

B. Modify the application running on the instances to enable termination protection while it processes a task and disable it when the processing is complete.

C. Increase the minimum and maximum size for the Auto Scaling group, and change the scaling policies so they scale less dynamically.

D. Modify the application running on the instances to put itself into an Auto Scaling Standby state while it processes a task and return itself to InService when the processing is complete.

Correct Answer: B

QUESTION 3

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 services. Retrieve the database credentials from a Systems Manager SecureString parameter.

B. Launch the EC2 instances with an EC2 IAM role to access AWS services. Retrieve the database credentials from AWS Secrets Manager.

C. 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.

D. Launch the EC2 instances with an EC2 IAM role to access AWS services. Store the database passwords in an encrypted config file with the application artifacts.

Correct Answer: B

QUESTION 4

A company is setting up a centralized logging solution on AWS and has several requirements. The company wants its Amazon CloudWatch Logs and VPC Flow logs to come from different sub accounts and to be delivered to a single auditing account. However, the number of sub accounts keeps changing. The company also needs to index the logs in the auditing account to gather actionable insight. How should a DevOps Engineer implement the solution to meet all of the company\\'s requirements?

A. Use AWS Lambda to write logs to Amazon ES in the auditing account. Create an Amazon CloudWatch subscription filter and use Amazon Kinesis Data Streams in the sub accounts to stream the logs to the Lambda function deployed in the auditing account.

B. Use Amazon Kinesis Streams to write logs to Amazon ES in the auditing account. Create a CloudWatch subscription filter and use Kinesis Data Streams in the sub accounts to stream the logs to the Kinesis stream in the auditing account.

C. Use Amazon Kinesis Firehose with Kinesis Data Streams to write logs to Amazon ES in the auditing account. Create a CloudWatch subscription filter and stream logs from sub accounts to the Kinesis stream in the auditing account.

D. Use AWS Lambda to write logs to Amazon ES in the auditing account. Create a CloudWatch subscription filter and use Lambda in the sub accounts to stream the logs to the Lambda function deployed in the auditing account.

Correct Answer: C

QUESTION 5

A defect was discovered in production and a new sprint item has been created for deploying a hotfix. However, any code change must go through the following steps before going into production:

1.

Scan the code for security breaches, such as password and access key leaks.

2.

Run the code through extensive, long-running unit tests.

Which source control strategy should a DevOps Engineer use in combination with AWS CodePipeline to complete this process?

A. Create a hotfix tag on the last commit of the master branch. Trigger the development pipeline from the hotfix tag. Use AWS CodeDeploy with Amazon ECS to do a content scan and run unit tests. Add a manual approval stage that merges the hotfix tag into the master branch.

B. Create a hotfix branch from the master branch. Trigger the development pipeline from the hotfix branch. Use AWS CodeBuild to do a content scan and run unit tests. Add a manual approval stage that merges the hotfix branch into the master branch.

C. Create a hotfix branch from the master branch. Trigger the development pipeline from the hotfix branch. Use AWS Lambda to do a content scan and run unit tests. Add a manual approval stage that merges the hotfix branch into the master branch.

D. Create a hotfix branch from the master branch. Create a separate source stage for the hotfix branch in the production pipeline. Trigger the pipeline from the hotfix branch. Use AWS Lambda to do a content scan and use AWS CodeBuild to run unit tests. Add a manual approval stage that merges the hotfix branch into the master branch.

Correct Answer: D