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

Your system automatically provisions EIPs to EC2 instances in a VPC on boot. The system provisions the whole VPC and stack at once. You have two of them per VPC. On your new AWS account, your attempt to create a Development

environment failed, after successfully creating Staging and Production environments in the same region.

What happened?

A. You didn\\'t choose the Development version of the AMI you are using.

B. You didn\\'t set the Development flag to true when deploying EC2 instances.

C. You hit the soft limit of 5 EIPs per region and requested a 6th.

D. You hit the soft limit of 2 VPCs per region and requested a 3rd.

Correct Answer: C

There is a soft limit of 5 EIPs per Region for VPC on new accounts. The third environment could not allocate the 6th EIP. Reference: http://docs.aws.amazon.com/general/latest/gr/aws_service_limits.html#limits_vpc

QUESTION 2

You have an application consisting of a stateless web server tier running on Amazon EC2 instances behind load balancer, and are using Amazon RDS with read replicas. Which of the following methods should you use to implement a self-healing and cost-effective architecture? (Choose two.)

A. Set up a third-party monitoring solution on a cluster of Amazon EC2 instances in order to emit custom CloudWatch metrics to trigger the termination of unhealthy Amazon EC2 instances.

B. Set up scripts on each Amazon EC2 instance to frequently send ICMP pings to the load balancer in order to determine which instance is unhealthy and replace it.

C. Set up an Auto Scaling group for the web server tier along with an Auto Scaling policy that uses the Amazon RDS DB CPU utilization CloudWatch metric to scale the instances.

D. Set up an Auto Scaling group for the web server tier along with an Auto Scaling policy that uses the Amazon EC2 CPU utilization CloudWatch metric to scale the instances.

E. Use a larger Amazon EC2 instance type for the web server tier and a larger DB instance type for the data storage layer to ensure that they don\\'t become unhealthy.

F. Set up an Auto Scaling group for the database tier along with an Auto Scaling policy that uses the Amazon RDS read replica lag CloudWatch metric to scale out the Amazon RDS read replicas.

G. Use an Amazon RDS Multi-AZ deployment.

Correct Answer: DG

QUESTION 3

A company needs to introduce automatic DNS failover for a distributed web application to a disaster recovery or standby installation. The DevOps Engineer plans to configure Amazon Route 53 to provide DNS routing to alternate endpoint in the event of an application failure. What steps should the Engineer take to accomplish this? (Choose two.)

A. Create Amazon Route 53 health checks for each endpoint that cannot be entered as alias records. Ensure firewall and routing rules allow Amazon Route 53 to send requests to the endpoints that are specified in the health checks.

B. Create alias records that route traffic to AWS resources and set the value of the Evaluate Target Health option to Yes, then create all the non-alias records.

C. Create a governing Amazon Route 53 record set, set it to failover, and associate it with the primary and secondary Amazon Route 53 record sets to distribute traffic to healthy DNS entries.

D. Create an Amazon CloudWatch alarm to monitor the primary Amazon Route 53 DNS entry. Then create an associated AWS Lambda function to execute the failover API call to Route 53 to the secondary DNS entry.

E. Map the primary and secondary Amazon Route 53 record sets to an Amazon CloudFront distribution using primary and secondary origins.

Correct Answer: AC

QUESTION 4

You need the absolute highest possible network performance for a cluster computing application. You already selected homogeneous instance types supporting 10 gigabit enhanced networking, made sure that your workload was network bound, and put the instances in a placement group. What is the last optimization you can make?

A. Use 9001 MTU instead of 1500 for Jumbo Frames, to raise packet body to packet overhead ratios.

B. Segregate the instances into different peered VPCs while keeping them all in a placement group, so each one has its own Internet Gateway.

C. Bake an AMI for the instances and relaunch, so the instances are fresh in the placement group and do not have noisy neighbors.

D. Turn off SYN/ACK on your TCP stack or begin using UDP for higher throughput.

Correct Answer: A

For instances that are collocated inside a placement group, jumbo frames help to achieve the maximum network throughput possible, and they are recommended in this case. For more information, see Placement Groups.

Reference:

http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/network_mtu.html#jumbo_frame_instances

QUESTION 5

A company uses AWS CodePipeline to manage and deploy infrastructure as code. The infrastructure is defined in AWS CloudFormation templates and is primarily comprised of multiple Amazon EC2 instances and Amazon RDS databases. The Security team has observed many operators creating inbound security group rules with a source CIDR of 0 0 0 0/0 and would like to proactively stop the deployment of rules with open CIDRs The DevOps Engineer will implement a predeptoyment step that runs some security checks over the CloudFormation template before the pipeline processes it.

This check should allow only inbound security group rules with a source CIDR of 0.0.0.0/0 if the rule has the description "Security Approval Ref XXXXX (where XXXXX is a preallocated reference). The pipeline step should fail if this condition is not met and the deployment should be blocked. How should this be accomplished?

A. Enable a SCP in AWS Organizations. The policy should deny access to the API call Create Security GroupRule if the rule specifies 0.0.0.0/0 without a description referencing a security approval.

B. Add an initial stage to CodePipeline called Security Check. This stage should call an AWS Lambda function that scans the CloudFormation template and fails the pipeline if it finds 0.0.0.0/0 in a security group without a description referencing a security approval.

C. Create an AWS Config rule that is triggered on creation or edit of resource type EC2 SecurityGroup. This rule should call an AWS Lambda function to send a failure notification if the security group has any rules with a source CIDR of 0.0.0.0/0 without a description referencing a security approval.

D. Modify the IAM role used by CodePipeline. The IAM policy should deny access.

Correct Answer: B

QUESTION 6

Your development team wants account-level access to production instances in order to do live debugging of a highly secure environment. Which of the following should you do?

A. Place the credentials provided by Amazon Elastic Compute Cloud (EC2) into a secure Amazon Sample Storage Service (S3) bucket with encryption enabled. Assign AWS Identity and Access Management (IAM) users to each developer so they can download the credentials file.

B. Place an internally created private key into a secure S3 bucket with server-side encryption using customer keys and configuration management, create a service account on all the instances using this private key, and assign IAM users to each developer so they can download the file.

C. Place each developer\\'s own public key into a private S3 bucket, use instance profiles and configuration management to create a user account for each developer on all instances, and place the user\\'s public keys into the appropriate account.

D. Place the credentials provided by Amazon EC2 onto an MFA encrypted USB drive, and physically share it with each developer so that the private key never leaves the office.

Correct Answer: C

QUESTION 7

A DevOps engineer is creating a CI/CD pipeline for an Amazon ECS service. The ECS container instances run behind an Application Load Balancer as the web tier of a three-tier application. An acceptance criterion for a successful deployment is the verification that the web tier can communicate with the database and middleware tiers of the application upon deployment.

How can this be accomplished in an automated fashion?

A. Create a health check endpoint in the web application that tests connectivity to the data and middleware tiers. Use this endpoint as the health check URL for the load balancer.

B. Create an approval step for the quality assurance team to validate connectivity. Reject changes in the pipeline if there is an issue with connecting to the dependent tiers.

C. Use an Amazon RDS active connection count and an Amazon CloudWatch ELB metric to alarm on a significant change to the number of open connections.

D. Use Amazon Route 53 health checks to detect issues with the web service and roll back the CI/CD pipeline if there is an error.

Correct Answer: A

QUESTION 8

You run a SIP-based telephony application that uses Amazon EC2 for its web tier and uses MySQL on Amazon RDS as its database. The application stores only the authentication profile data for its existing users in the database and

therefore is read-intensive. Your monitoring system shows that your web instances and the database have high CPU utilization.

Which of the following steps should you take in order to ensure the continual availability of your application? (Choose two.)

A. Use a CloudFront RTMP download distribution with the application tier as the origin for the distribution.

B. Set up an Auto Scaling group for the application tier and a policy that scales based on the Amazon EC2 CloudWatch CPU utilization metric.

C. Vertically scale up the Amazon EC2 instances manually.

D. Set up an Auto Scaling group for the application tier and a policy that scales based on the Amazon RDS CloudWatch CPU utilization metric.

E. Switch to General Purpose (SSD) Storage from Provisioned IOPS Storage (PIOPS) for the Amazon RDS database.

F. Use multiple Amazon RDS read replicas.

Correct Answer: BF

QUESTION 9

A DevOps Engineer is using AWS CodeDeploy across a fleet of Amazon EC2 instances in an EC2 Auto Scaling group. The associated CodeDeploy deployment group, which is integrated with EC2 Auto Scaling, is configured to perform in-

place deployments with CodeDeployDefault.OneAtATime. During an ongoing new deployment, the Engineer discovers that, although the overall deployment finished successfully, two out of five instances have the previous application revision

deployed. The other three instances have the newest application revision.

What is likely causing this issue?

A. The two affected instances failed to fetch the new deployment.

B. A failed AfterInstall lifecycle event hook caused the CodeDeploy agent to roll back to the previous version on the affected instances.

C. The CodeDeploy agent was not installed in two affected instances.

D. EC2 Auto Scaling launched two new instances while the new deployment had not yet finished, causing the previous version to be deployed on the affected instances.

Correct Answer: D

QUESTION 10

A company requires that its internally facing web application be highly available. The architecture is made up of one Amazon EC2 web server instance and one NAT instance that provides outbound internet access for updates and accessing public data.

Which combination of architecture adjustments should the company implement to achieve high availability? (Choose two.)

A. Add the NAT instance to an EC2 Auto Scaling group that spans multiple Availability Zones. Update the route tables.

B. Create additional EC2 instances spanning multiple Availability Zones. Add an Application Load Balancer to split the load between them.

C. Configure an Application Load Balancer in front of the EC2 instance. Configure Amazon CloudWatch alarms to recover the EC2 instance upon host failure.

D. Replace the NAT instance with a NAT gateway in each Availability Zone. Update the route tables.

E. Replace the NAT instances with a NAT gateway that spans multiple Availability Zones. Update the route tables.

Correct Answer: BD

QUESTION 11

A government agency has multiple AWS accounts, many of which store sensitive citizen information. A Security team wants to detect anomalous account and network activities (such as SSH brute force attacks) in any account and centralize that information in a dedicated security account. Event information should be stored in an Amazon S3 bucket in the security account, which is monitored by the department\\'s Security Information and Event Management (SIEM) system. How can this be accomplished?

A. Enable Amazon Macie in every account. Configure the security account as the Macie Administrator for every member account using invitation/acceptance. Create an Amazon CloudWatch Events rule in the security account to send all findings to Amazon Kinesis Data Firehose, which should push the findings to the S3 bucket.

B. Enable Amazon Macie in the security account only. Configure the security account as the Macie Administrator for every member account using invitation/acceptance. Create an Amazon CloudWatch Events rule in the security account to send all findings to Amazon Kinesis Data Streams. Write an application using KCL to read data from the Kinesis Data Streams and write to the S3 bucket.

C. Enable Amazon GuardDuty in every account. Configure the security account as the GuardDuty Administrator for

every member account using invitation/acceptance. Create an Amazon CloudWatch rule in the security account to send all findings to Amazon Kinesis Data Firehose, which will push the findings to the S3 bucket.

D. Enable Amazon GuardDuty in the security account only. Configure the security account as the GuardDuty Administrator for every member account using invitation/acceptance. Create an Amazon CloudWatch rule in the security account to send all findings to Amazon Kinesis Data Streams. Write an application using KCL to read data from Kinesis Data Streams and write to the S3 bucket.

Correct Answer: C

QUESTION 12

A company that runs many workloads on AWS has an Amazon EBS spend that has increased over time. The DevOps team notices there are many unattached EBS volumes. Although there are workloads where volumes are detached, volumes over 14 days old are stale and no longer needed. A DevOps engineer has been tasked with creating automation that deletes unattached EBS volumes that have been unattached for 14 days.

Which solution will accomplish this?

A. Configure the AWS Config ec2-volume-inuse-check managed rule with a configuration changes trigger type and an Amazon EC2 volume resource target. Create a new Amazon CloudWatch Events rule scheduled to execute an AWS Lambda function in 14 days to delete the specified EBS volume.

B. Use Amazon EC2 and Amazon Data Lifecycle Manager to configure a volume lifecycle policy. Set the interval period for unattached EBS volumes to 14 days and set the retention rule to delete. Set the policy target volumes as *.

C. Create an Amazon CloudWatch Events rule to execute an AWS Lambda function daily. The Lambda function should find unattached EBS volumes and tag them with the current date, and delete unattached volumes that have tags with dates that are more than 14 days old.

D. Use AWS Trusted Advisor to detect EBS volumes that have been detached for more than 14 days. Execute an AWS Lambda function that creates a snapshot and then deletes the EBS volume.

Correct Answer: B

QUESTION 13

When thinking of AWS Elastic Beanstalk, the \\'Swap Environment URLs\\' feature most directly aids in what?

A. Immutable Rolling Deployments

B. Mutable Rolling Deployments

C. Canary Deployments

D. Blue-Green Deployments

Correct Answer: D

Simply upload the new version of your application and let your deployment service (AWS Elastic Beanstalk, AWS CloudFormation, or AWS OpsWorks) deploy a new version (green). To cut over to the new version, you simply replace the ELB URLs in your DNS records. Elastic Beanstalk has a Swap Environment URLs feature to facilitate a simpler cutover process. Reference: https://d0.awsstatic.com/ whitepapers/overview-of-deployment-options-on-aws.pdf

QUESTION 14

A Development team creates a build project in AWS CodeBuild. The build project invokes automated tests of modules that access AWS services. Which of the following will enable the tests to run the MOST securely?

A. Generate credentials for an IAM user with a policy attached to allow the actions on AWS services. Store credentials as encrypted environment variables for the build project. As part of the build script, obtain the credentials to run the integration tests.

B. Have CodeBuild run only the integration tests as a build job on a Jenkins server. Create a role that has a policy attached to allow the actions on AWS services. Generate credentials for an IAM user that is allowed to assume the role. Configure the credentials as secrets in Jenkins, and allow the build job to use them to run the integration tests.

C. Create a service role in IAM to be assumed by CodeBuild with a policy attached to allow the actions on AWS services. Configure the build project to use the role created.

D. Use AWS managed credentials. Encrypt the credentials with AWS KMS. As part of the build script, decrypt with AWS KMS and use these credentials to run the integration tests.

Correct Answer: C

QUESTION 15

You recently encountered a major bug in your Windows-based web application during a deployment cycle. During this failed deployment, it took the team four hours to roll back to a previously working state, which left customers with a poor user experience. During the post-mortem, your team discussed the need to provide a quicker way to roll back failed deployments. You currently run your web application on Amazon EC2 using Windows 2012R2 and use Elastic Load Balancing for your load balancing needs. Which technique should you use to solve this problem?

A. Create deployable versioned bundles of your application. Store the bundles on Amazon S3. Re-deploy your web application on Elastic Beanstalk, and enable the Elastic Beanstalk autorollback feature tied to CloudWatch metrics that define failure.

B. Re-deploy your web application using an AWS OpsWorks stack, and use the AWS OpsWorks auto- rollback feature to initiate a rollback during failures.

C. Create deployable versioned bundles of your application. Store the bundle on Amazon S3. Re-deploy your web application using an AWS OpsWorks stack, and use AWS OpsWorks application versioning to initiate a rollback during failures.

D. Re-deploy your web application using Elastic Beanstalk, and use the Elastic Beanstalk application versions when deploying. During failures, re-deploy the previous version to the Elastic Beanstalk environment.

E. Re-deploy your web application using Elastic Beanstalk, and use the Elastic Beanstalk API to trigger a FailedDeployment API call to initiate a rollback to the previous version.

Correct Answer: B