amazon
web services

DOP-C02

AWS Certified

DevOps Engineer - Professional

Amazon Web Services

Exam DOP-C02

AWS Certified DevOps Engineer - Professional

Version: 4.0

[Total Questions: 81]

Question No : 1

A company is using AWS CodePipeline to automate its release pipeline. AWS CodeDeploy is being used in the pipeline to deploy an application to Amazon Elastic Container Service (Amazon ECS) using the blue/green deployment model. The company wants to implement scripts to test the green version of the application before shifting traffic. These scripts will complete in 5 minutes or less. If errors are discovered during these tests, the application must be rolled back.

Which strategy will meet these requirements?

A. Add a stage to the CodePipeline pipeline between the source and deploy stages. Use AWS CodeBuild to create a runtime environment and build commands in the buildspec file to invoke test scripts. If errors are found, use the aws deploy stop-deployment command to stop the deployment.
B. Add a stage to the CodePipeline pipeline between the source and deploy stages. Use this stage to invoke an AWS Lambda function that will run the test scripts. If errors are found, use the aws deploy stop-deployment command to stop the deployment.
C. Add a hooks section to the CodeDeploy AppSpec file. Use the AfterAllowTestTraffic lifecycle event to invoke an AWS Lambda function to run the test scripts. If errors are found, exit the Lambda function with an error to initiate rollback.
D. Add a hooks section to the CodeDeploy AppSpec file. Use the AfterAllowTraffic lifecycle event to invoke the test scripts. If errors are found, use the aws deploy stop-deployment CLI command to stop the deployment.

Answer: C

Question No : 2

A company has an on-premises application that is written in Go. A DevOps engineer must move the application to AWS. The company's development team wants to enable blue/green deployments and perform A/B testing.

Which solution will meet these requirements?

A. Deploy the application on an Amazon EC2 instance, and create an AMI of the instance. Use the AMI to create an automatic scaling launch configuration that is used in an Auto Scaling group. Use Elastic Load Balancing to distribute traffic. When changes are made to the application, a new AMI will be created, which will initiate an EC2 instance refresh.
B. Use Amazon Lightsail to deploy the application. Store the application in a zipped format in an Amazon S3 bucket. Use this zipped version to deploy new versions of the application to Lightsail. Use Lightsail deployment options to manage the deployment.
C. Use AWS CodeArtifact to store the application code. Use AWS CodeDeploy to deploy the application to a fleet of Amazon EC2 instances. Use Elastic Load Balancing to

distribute the traffic to the EC2 instances. When making changes to the application, upload a new version to CodeArtifact and create a new CodeDeploy deployment.
D. Use AWS Elastic Beanstalk to host the application. Store a zipped version of the application in Amazon S3. Use that location to deploy new versions of the application. Use Elastic Beanstalk to manage the deployment options.

Answer: D

Question No : 3

A company is hosting a static website from an Amazon S3 bucket. The website is available to customers at example.com. The company uses an Amazon Route 53 weighted routing policy with a TTL of 1 day. The company has decided to replace the existing static website with a dynamic web application. The dynamic web application uses an Application Load Balancer (ALB) in front of a fleet of Amazon EC2 instances.

On the day of production launch to customers, the company creates an additional Route 53 weighted DNS record entry that points to the ALB with a weight of 255 and a TTL of 1 hour. Two days later, a DevOps engineer notices that the previous static website is displayed sometimes when customers navigate to example.com.

How can the DevOps engineer ensure that the company serves only dynamic content for example.com?

A. Delete all objects, including previous versions, from the S3 bucket that contains the static website content.
B. Update the weighted DNS record entry that points to the S3 bucket. Apply a weight of 0. Specify the domain reset option to propagate changes immediately.
C. Configure webpage redirect requests on the S3 bucket with a hostname that redirects to the ALB.
D. Remove the weighted DNS record entry that points to the S3 bucket from the example.com hosted zone. Wait for DNS propagation to become complete.

Answer: B

Question No : 4

A company runs an application with an Amazon EC2 and on-premises configuration. A DevOps engineer needs to standardize patching across both environments. Company policy dictates that patching only happens during non-business hours.

Which combination of actions will meet these requirements? (Choose three.)

A. Add the physical machines into AWS Systems Manager using Systems Manager Hybrid Activations.
B. Attach an IAM role to the EC2 instances, allowing them to be managed by AWS Systems Manager.
C. Create IAM access keys for the on-premises machines to interact with AWS Systems Manager.
D. Run an AWS Systems Manager Automation document to patch the systems every hour.
E. Use Amazon EventBridge scheduled events to schedule a patch window.
F. Use AWS Systems Manager Maintenance Windows to schedule a patch window.

Answer: A,B,F

Question No : 5

An ecommerce company has chosen AWS to host its new platform. The company's DevOps team has started building an AWS Control Tower landing zone. The DevOps team has set the identity store within AWS IAM Identity Center (AWS Single Sign-On) to external identity provider (IdP) and has configured SAML 2.0.

The DevOps team wants a robust permission model that applies the principle of least privilege. The model must allow the team to build and manage only the team's own resources.

Which combination of steps will meet these requirements? (Choose three.)

A. Create IAM policies that include the required permissions. Include the aws:PrincipalTag condition key.
B. Create permission sets. Attach an inline policy that includes the required permissions and uses the aws:PrincipalTag condition key to scope the permissions.
C. Create a group in the IdP. Place users in the group. Assign the group to accounts and the permission sets in IAM Identity Center.
D. Create a group in the IdP. Place users in the group. Assign the group to OUs and IAM policies.
E. Enable attributes for access control in IAM Identity Center. Apply tags to users. Map the tags as key-value pairs.
F. Enable attributes for access control in IAM Identity Center. Map attributes from the IdP as key-value pairs.

Answer: A,B,C

Question No : 6

A company wants to set up a continuous delivery pipeline. The company stores application code in a private GitHub repository. The company needs to deploy the application components to Amazon Elastic Container Service (Amazon ECS). Amazon EC2, and AWS Lambda. The pipeline must support manual approval actions.

Which solution will meet these requirements?

A. Use AWS CodePipeline with Amazon ECS. Amazon EC2, and Lambda as deploy providers.
B. Use AWS CodePipeline with AWS CodeDeploy as the deploy provider.
C. Use AWS CodePipeline with AWS Elastic Beanstalk as the deploy provider.
D. Use AWS CodeDeploy with GitHub integration to deploy the application.

Answer: B

Question No : 7

A company provides an application to customers. The application has an Amazon API Gateway REST API that invokes an AWS Lambda function. On initialization, the Lambda function loads a large amount of data from an Amazon DynamoDB table. The data load process results in long cold-start times of 8-10 seconds. The DynamoDB table has DynamoDB Accelerator (DAX) configured.

Customers report that the application intermittently takes a long time to respond to requests. The application receives thousands of requests throughout the day. In the middle of the day, the application experiences 10 times more requests than at any other time of the day. Near the end of the day, the application's request volume decreases to 10% of its normal total.

A DevOps engineer needs to reduce the latency of the Lambda function at all times of the day.

Which solution will meet these requirements?

A. Configure provisioned concurrency on the Lambda function with a concurrency value of 1. Delete the DAX cluster for the DynamoDB table.
B. Configure reserved concurrency on the Lambda function with a concurrency value of 0.
C. Configure provisioned concurrency on the Lambda function. Configure AWS Application Auto Scaling on the Lambda function with provisioned concurrency values set to a minimum of 1 and a maximum of 100.
D. Configure reserved concurrency on the Lambda function. Configure AWS Application Auto Scaling on the API Gateway API with a reserved concurrency maximum value of 100.

Answer: C

Question No : 8

A company runs an application on one Amazon EC2 instance. Application metadata is stored in Amazon S3 and must be retrieved if the instance is restarted. The instance must restart or relaunch automatically if the instance becomes unresponsive.

Which solution will meet these requirements?

A. Create an Amazon CloudWatch alarm for the StatusCheckFailed metric. Use the recover action to stop and start the instance. Use an S3 event notification to push the metadata to the instance when the instance is back up and running.
B. Configure AWS OpsWorks, and use the auto healing feature to stop and start the instance. Use a lifecycle event in OpsWorks to pull the metadata from Amazon S3 and update it on the instance.
C. Use EC2 Auto Recovery to automatically stop and start the instance in case of a failure. Use an S3 event notification to push the metadata to the instance when the instance is back up and running.
D. Use AWS CloudFormation to create an EC2 instance that includes the UserData property for the EC2 resource. Add a command in UserData to retrieve the application metadata from Amazon S3.

Answer: B

Question No : 9

A company has an application that runs on Amazon EC2 instances that are in an Auto Scaling group. When the application starts up. the application needs to process data from an Amazon S3 bucket before the application can start to serve requests.

The size of the data that is stored in the S3 bucket is growing. When the Auto Scaling group adds new instances, the application now takes several minutes to download and process the data before the application can serve requests. The company must reduce the time that elapses before new EC2 instances are ready to serve requests.

Which solution is the MOST cost-effective way to reduce the application startup time?

A. Configure a warm pool for the Auto Scaling group with warmed EC2 instances in the Stopped state. Configure an autoscaling:EC2_INSTANCE_LAUNCHING lifecycle hook on the Auto Scaling group. Modify the application to complete the lifecycle hook when the application is ready to serve requests.
B. Increase the maximum instance count of the Auto Scaling group. Configure an autoscaling:EC2_INSTANCE_LAUNCHING lifecycle hook on the Auto Scaling group. Modify the application to complete the lifecycle hook when the application is ready to serve requests.

C. Configure a warm pool for the Auto Scaling group with warmed EC2 instances in the Running state. Configure an autoscaling:EC2_INSTANCE_LAUNCHING lifecycle hook on the Auto Scaling group. Modify the application to complete the lifecycle hook when the application is ready to serve requests.
D. Increase the maximum instance count of the Auto Scaling group. Configure an autoscaling:EC2_INSTANCE_LAUNCHING lifecycle hook on the Auto Scaling group. Modify the application to complete the lifecycle hook and to place the new instance in the Standby state when the application is ready to serve requests.

Answer: C

Question No : 10

A company has multiple accounts in an organization in AWS Organizations. The company's SecOps team needs to receive an Amazon Simple Notification Service (Amazon SNS) notification if any account in the organization turns off the Block Public Access feature on an Amazon S3 bucket. A DevOps engineer must implement this change without affecting the operation of any AWS accounts. The implementation must ensure that individual member accounts in the organization cannot turn off the notification.

Which solution will meet these requirements?

A. Designate an account to be the delegated Amazon GuardDuty administrator account. Turn on GuardDuty for all accounts across the organization. In the GuardDuty administrator account, create an SNS topic. Subscribe the SecOps team's email address to the SNS topic. In the same account, create an Amazon EventBridge rule that uses an event pattern for GuardDuty findings and a target of the SNS topic.
B. Create an AWS CloudFormation template that creates an SNS topic and subscribes the SecOps team's email address to the SNS topic. In the template, include an Amazon EventBridge rule that uses an event pattern of CloudTrail activity for s3:PutBucketPublicAccessBlock and a target of the SNS topic. Deploy the stack to every account in the organization by using CloudFormation StackSets.
C. Turn on AWS Config across the organization. In the delegated administrator account, create an SNS topic. Subscribe the SecOps team's email address to the SNS topic. Deploy a conformance pack that uses the s3-bucket-level-public-access-prohibited AWS Config managed rule in each account and uses an AWS Systems Manager document to publish an event to the SNS topic to notify the SecOps team.
D. Turn on Amazon Inspector across the organization. In the Amazon Inspector delegated administrator account, create an SNS topic. Subscribe the SecOps team's email address to the SNS topic. In the same account, create an Amazon EventBridge rule that uses an event pattern for public network exposure of the S3 bucket and publishes an event to the SNS topic to notify the SecOps team.

Answer: B

Question No : 11

A company wants to migrate its content sharing web application hosted on Amazon EC2 to a serverless architecture. The company currently deploys changes to its application by creating a new Auto Scaling group of EC2 instances and a new Elastic Load Balancer, and then shifting the traffic away using an Amazon Route 53 weighted routing policy.

For its new serverless application, the company is planning to use Amazon API Gateway and AWS Lambda. The company will need to update its deployment processes to work with the new application. It will also need to retain the ability to test new features on a small number of users before rolling the features out to the entire user base.

Which deployment strategy will meet these requirements?

A. Use AWS CDK to deploy API Gateway and Lambda functions. When code needs to be changed, update the AWS CloudFormation stack and deploy the new version of the APIs and Lambda functions. Use a Route 53 failover routing policy for the canary release strategy.
B. Use AWS CloudFormation to deploy API Gateway and Lambda functions using Lambda function versions. When code needs to be changed, update the CloudFormation stack with the new Lambda code and update the API versions using a canary release strategy. Promote the new version when testing is complete.
C. Use AWS Elastic Beanstalk to deploy API Gateway and Lambda functions. When code needs to be changed, deploy a new version of the API and Lambda functions. Shift traffic gradually using an Elastic Beanstalk blue/green deployment.
D. Use AWS OpsWorks to deploy API Gateway in the service layer and Lambda functions in a custom layer. When code needs to be changed, use OpsWorks to perform a blue/green deployment and shift traffic gradually.

Answer: B

Question No : 12

A company has migrated its container-based applications to Amazon EKS and want to establish automated email notifications. The notifications sent to each email address are for specific activities related to EKS components. The solution will include Amazon SNS topics and an AWS Lambda function to evaluate incoming log events and publish messages to the correct SNS topic.

Which logging solution will support these requirements?

A. Enable Amazon CloudWatch Logs to log the EKS components. Create a CloudWatch subscription filter for each component with Lambda as the subscription feed destination.
B. Enable Amazon CloudWatch Logs to log the EKS components. Create CloudWatch Logs Insights queries linked to Amazon EventBridge events that invoke Lambda.
C. Enable Amazon S3 logging for the EKS components. Configure an Amazon CloudWatch subscription filter for each component with Lambda as the subscription feed destination.
D. Enable Amazon S3 logging for the EKS components. Configure S3 PUT Object event notifications with AWS Lambda as the destination.

Answer: C

Question No : 13

A company builds a container image in an AWS CodeBuild project by running Docker commands. After the container image is built, the CodeBuild project uploads the container image to an Amazon S3 bucket. The CodeBuild project has an 1AM service role that has permissions to access the S3 bucket.

A DevOps engineer needs to replace the S3 bucket with an Amazon Elastic Container Registry (Amazon ECR) repository to store the container images. The DevOps engineer creates an ECR private image repository in the same AWS Region of the CodeBuild project. The DevOps engineer adjusts the 1AM service role with the permissions that are necessary to work with the new ECR repository. The DevOps engineer also places new repository information into the docker build command and the docker push command that are used in the buildspec.yml file.

When the CodeBuild project runs a build job, the job fails when the job tries to access the ECR repository.

Which solution will resolve the issue of failed access to the ECR repository?

A. Update the buildspec.yml file to log in to the ECR repository by using the aws ecr get-login-password AWS CLI command to obtain an authentication token. Update the docker login command to use the authentication token to access the ECR repository.
B. Add an environment variable of type SECRETS_MANAGER to the CodeBuild project. In the environment variable, include the ARN of the CodeBuild project's IAM service role. Update the buildspec.yml file to use the new environment variable to log in with the docker login command to access the ECR repository.
C. Update the ECR repository to be a public image repository. Add an ECR repository policy that allows the 1AM service role to have access.
D. Update the buildspec.yml file to use the AWS CLI to assume the 1AM service role for ECR operations. Add an ECR repository policy that allows the 1AM service role to have access.

Answer: A

Question No : 14

To run an application, a DevOps engineer launches an Amazon EC2 instance with public IP addresses in a public subnet. A user data script obtains the application artifacts and installs them on the instances upon launch. A change to the security classification of the application now requires the instances to run with no access to the internet. While the instances launch successfully and show as healthy, the application does not seem to be installed.

Which of the following should successfully install the application while complying with the new rule?

A. Launch the instances in a public subnet with Elastic IP addresses attached. Once the application is installed and running, run a script to disassociate the Elastic IP addresses afterwards.
B. Set up a NAT gateway. Deploy the EC2 instances to a private subnet. Update the private subnet's route table to use the NAT gateway as the default route.
C. Publish the application artifacts to an Amazon S3 bucket and create a VPC endpoint for S3. Assign an IAM instance profile to the EC2 instances so they can read the application artifacts from the S3 bucket.
D. Create a security group for the application instances and allow only outbound traffic to the artifact repository. Remove the security group rule once the install is complete.

Answer: C

Question No : 15

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 that uses reserved concurrency. The Lambda function processes order messages from an Amazon Simple Queue Service (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 should a DevOps engineer take to resolve this delay? (Choose two.)

A. Check the ApproximateAgeOfOldestMessage metric for the SQS queue. Increase the Lambda function concurrency limit.
B. Check the ApproximateAgeOfOldestMessage metnc for the SQS queue Configure a redrive policy on the SQS queue.

C. Check the NumberOfMessagesSent metric for the SQS queue. Increase the SQS queue visibility timeout.
D. Check the WriteThrottleEvents metric for the DynamoDB table. Increase the maximum write capacity units (WCUs) for the table's scaling policy.
E. Check the Throttles metric for the Lambda function. Increase the Lambda function timeout.

Answer: A,D

Question No : 16

A company must encrypt all AMIs that the company shares across accounts. A DevOps engineer has access to a source account where an unencrypted custom AMI has been built. The DevOps engineer also has access to a target account where an Amazon EC2 Auto Scaling group will launch EC2 instances from the AMI. The DevOps engineer must share the AMI with the target account.

The company has created an AWS Key Management Service (AWS KMS) key in the source account.

Which additional steps should the DevOps engineer perform to meet the requirements? (Choose three.)

A. In the source account, copy the unencrypted AMI to an encrypted AMI. Specify the KMS key in the copy action.
B. In the source account, copy the unencrypted AMI to an encrypted AMI. Specify the default Amazon Elastic Block Store (Amazon EBS) encryption key in the copy action.
C. In the source account, create a KMS grant that delegates permissions to the Auto Scaling group service-linked role in the target account.
D. In the source account, modify the key policy to give the target account permissions to create a grant. In the target account, create a KMS grant that delegates permissions to the Auto Scaling group service-linked role.
E. In the source account, share the unencrypted AMI with the target account.
F. In the source account, share the encrypted AMI with the target account.

Answer: A,C,D

Question No : 17

A global company manages multiple AWS accounts by using AWS Control Tower. The company hosts internal applications and public applications.

Each application team in the company has its own AWS account for application hosting. The accounts are consolidated in an organization in AWS Organizations. One of the AWS Control Tower member accounts serves as a centralized DevOps account with CI/CD pipelines that application teams use to deploy applications to their respective target AWS accounts. An 1AM role for deployment exists in the centralized DevOps account.

An application team is attempting to deploy its application to an Amazon Elastic Kubernetes Service (Amazon EKS) cluster in an application AWS account. An 1AM role for deployment exists in the application AWS account. The deployment is through an AWS CodeBuild project that is set up in the centralized DevOps account. The CodeBuild project uses an 1AM service role for CodeBuild. The deployment is failing with an Unauthorized error during attempts to connect to the cross-account EKS cluster from CodeBuild.

Which solution will resolve this error?

A. Configure the application account's deployment 1AM role to have a trust relationship with the centralized DevOps account. Configure the trust relationship to allow the sts:AssumeRole action. Configure the application account's deployment 1AM role to have the required access to the EKS cluster. Configure the EKS cluster aws-auth ConfigMap to map the role to the appropriate system permissions.

B. Configure the centralized DevOps account's deployment I AM role to have a trust relationship with the application account. Configure the trust relationship to allow the sts:AssumeRole action. Configure the centralized DevOps account's deployment 1AM role to allow the required access to CodeBuild.

C. Configure the centralized DevOps account's deployment 1AM role to have a trust relationship with the application account. Configure the trust relationship to allow the sts:AssumeRoleWithSAML action. Configure the centralized DevOps account's deployment 1AM role to allow the required access to CodeBuild.

D. Configure the application account's deployment 1AM role to have a trust relationship with the AWS Control Tower management account. Configure the trust relationship to allow the sts:AssumeRole action. Configure the application account's deployment 1AM role to have the required access to the EKS cluster. Configure the EKS cluster aws-auth ConfigMap to map the role to the appropriate system permissions.

Answer: D

Question No : 18

A company has an AWS CodePipeline pipeline that is configured with an Amazon S3 bucket in the eu-west-1 Region. The pipeline deploys an AWS Lambda application to the same Region. The pipeline consists of an AWS CodeBuild project build action and an AWS CloudFormation deploy action.

The CodeBuild project uses the aws cloudformation package AWS CLI command to build

---

an artifact that contains the Lambda function code's .zip file and the CloudFormation template. The CloudFormation deploy action references the CloudFormation template from the output artifact of the CodeBuild project's build action.

The company wants to also deploy the Lambda application to the us-east-1 Region by using the pipeline in eu-west-1. A DevOps engineer has already updated the CodeBuild project to use the aws cloudformation package command to produce an additional output artifact for us-east-1.

Which combination of additional steps should the DevOps engineer take to meet these requirements? (Choose two.)

A. Modify the CloudFormation template to include a parameter for the Lambda function code's zip file location. Create a new CloudFormation deploy action for us-east-1 in the pipeline. Configure the new deploy action to pass in the us-east-1 artifact location as a parameter override.
B. Create a new CloudFormation deploy action for us-east-1 in the pipeline. Configure the new deploy action to use the CloudFormation template from the us-east-1 output artifact.
C. Create an S3 bucket in us-east-1. Configure the S3 bucket policy to allow CodePipeline to have read and write access.
D. Create an S3 bucket in us-east-1. Configure S3 Cross-Region Replication (CRR) from the S3 bucket in eu-west-1 to the S3 bucket in us-east-1.
E. Modify the pipeline to include the S3 bucket for us-east-1 as an artifact store. Create a new CloudFormation deploy action for us-east-1 in the pipeline. Configure the new deploy action to use the CloudFormation template from the us-east-1 output artifact.

Answer: A,B

Question No : 19

A company is implementing a well-architected design for its globally accessible API stack. The design needs to ensure both high reliability and fast response times for users located in North America and Europe.

The API stack contains the following three tiers:

Amazon API Gateway

AWS Lambda

Amazon DynamoDB

Which solution will meet the requirements?

A. Configure Amazon Route 53 to point to API Gateway APIs in North America and Europe

---

using health checks. Configure the APIs to forward requests to a Lambda function in that Region. Configure the Lambda functions to retrieve and update the data in a DynamoDB table in the same Region as the Lambda function.

B. Configure Amazon Route 53 to point to API Gateway APIs in North America and Europe using latency-based routing and health checks. Configure the APIs to forward requests to a Lambda function in that Region. Configure the Lambda functions to retrieve and update the data in a DynamoDB global table.

C. Configure Amazon Route 53 to point to API Gateway in North America, create a disaster recovery API in Europe, and configure both APIs to forward requests to the Lambda functions in that Region. Retrieve the data from a DynamoDB global table. Deploy a Lambda function to check the North America API health every 5 minutes. In the event of a failure, update Route 53 to point to the disaster recovery API.

D. Configure Amazon Route 53 to point to API Gateway API in North America using latency-based routing. Configure the API to forward requests to the Lambda function in the Region nearest to the user. Configure the Lambda function to retrieve and update the data in a DynamoDB table.

Answer: B

Question No : 20

A company's DevOps engineer uses AWS Systems Manager to perform maintenance tasks during maintenance windows. The company has a few Amazon EC2 instances that require a restart after notifications from AWS Health. The DevOps engineer needs to implement an automated solution to remediate these notifications. The DevOps engineer creates an Amazon EventBridge rule.

How should the DevOps engineer configure the EventBridge rule to meet these requirements?

A. Configure an event source of AWS Health, a service of EC2. and an event type that indicates instance maintenance. Target a Systems Manager document to restart the EC2 instance.

B. Configure an event source of Systems Manager and an event type that indicates a maintenance window. Target a Systems Manager document to restart the EC2 instance.

C. Configure an event source of AWS Health, a service of EC2, and an event type that indicates instance maintenance. Target a newly created AWS Lambda function that registers an automation task to restart the EC2 instance during a maintenance window.

D. Configure an event source of EC2 and an event type that indicates instance maintenance. Target a newly created AWS Lambda function that registers an automation task to restart the EC2 instance during a maintenance window.

Answer: A