AWS Certified DevOps Engineer - Professional

1. A company wants to implement a CI/CD pipeline for an application that is deployed on AWS. The company also has a source-code analysis tool hosted on premises that checks for security flaws. The tool has not yet been migrated to AWS and can be accessed only on premises. The company wants to run checks against the source code as part of the pipeline before the code is compiled. The checks take anywhere from minutes to an hour to complete. How can a DevOps Engineer meet these requirements'?

A. A. Use AWS CodePipeline to create a pipeline. Add an action to the pipeline to invoke an AWS Lambda function after the source stage. Have the Lambda function invoke the source-code analysis tool on premises against the source input from CodePipeline. The function then waits for the execution to complete and places the output in a specified Amazon S3 location.

B. B. Use AWS CodePipeline to create a pipeline, then create a custom action type. Create a job worker for the custom action that runs on hardware hosted on premises. The job worker handles running security checks with the on-premises code analysis tool and then returns the job results to CodePipeline. Have the pipeline invoke the custom action after the source stage.

C. C. Use AWS CodePipeline to create a pipeline. Add a step after the source stage to make an HTTPS request to the on-premises hosted web service that invokes a test with the source code analysis tool. When the analysis is complete, the web service sends the results back by putting the results in an Amazon S3 output location provided by CodePipeline.

D. D. Use AWS CodePipeline to create a pipeline. Create a shell script that copies the input source code to a location on premises. Invoke the source code analysis tool and return the results to CodePipeline. Invoke the shell script by adding a custom script action after the source stage.

Answer(s): B

---

2. A company is adopting AWS CodeDeploy to automate its application deployments for a JavaApache Tomcat application with an Apache webserver. The Development team started with a proof of concept, created a deployment group for a developer environment, and performed functional tests within the application. After completion, the team will create additional deployment groups for staging and production The current log level is configured within the Apache settings, but the team wants to change this configuration dynamically when the deployment occurs, so that

they can set different log level configurations depending on the deployment group without having a different application revision for each group. How can these requirements be met with the LEAST management overhead and without requiring different script versions for each deployment group?

A. A. Tag the Amazon EC2 instances depending on the deployment group. Then place a script into the application revision that calls the metadata service and the EC2 API to identify which deployment group the instance is part of. Use this information to configure the log level settings. Reference the script as part of the AfterInstall lifecycle hook in the appspec yml file.

B. B. Create a script that uses the CodeDeploy environment variable DEPLOYMENT_GROUP_NAME to identify which deployment group the instances is part of. Use this information to configure the log level settings. Reference this script as part of the BeforeInstall lifecycle hook in the appspec.yml file

C. C. Create a CodeDeploy custom environment variable for each environment. Then place a script into the application revision that checks this environment variable to identify which deployment group the instance is part of. Use this information to configure the log level settings. Reference this script as part of the ValidateService lifecycle hook in the appspec yml file.

D. D. Create a script that uses the CodeDeploy environment variable DEPLOYMENT_GROUP_ID to identify which deployment group the instance is part of to configure the log level settings. Reference this script as part of the Install lifecycle hook in the appspec yml file

Answer(s): B

---

3. A company has an application that has predictable peak traffic times. The company wants the application instances to scale up only during the peak times. The application stores state in Amazon DynamoDB. The application environment uses a standard Node.js application stack and custom Chef recipes stored in a private Git repository. Which solution is MOST cost-effective and requires the LEAST amount of management overhead when performing rolling updates of the application environment?

A. A. Create a custom AMI with the Node.js environment and application stack using Chef recipes. Use the AMI in an Auto Scaling group and set up scheduled scaling for the required times, then set up an Amazon EC2 IAM role that provides permission to access DynamoDB.

B. B. Create a Docker file that uses the Chef recipes for the application environment based on an official Node.js Docker image. Create an Amazon ECS cluster and a service for the application environment, then create a task based on this Docker image. Use scheduled scaling to scale the containers at the appropriate times and attach a task-level IAM role that provides permission to access DynamoDB.

C. C. Configure AWS OpsWorks stacks and use custom Chef cookbooks. Add the Git repository information where the custom recipes are stored, and add a layer in OpsWorks for the Node.js application server. Then configure the custom recipe to deploy the application in the deploy step. Configure time-based instances and attach an Amazon EC2 IAM role that provides permission to access DynamoDB.

D. D. Configure AWS OpsWorks stacks and push the custom recipes to an Amazon S3 bucket and configure custom recipes to point to the S3 bucket. Then add an application layer type for a standard Node.js application server and configure the custom recipe to deploy the application in the deploy step from the S3 bucket. Configure time-based instances and attach an Amazon EC2 IAM role that provides permission to access DynamoDB.

Answer(s): C

---

4. The Development team at an online retailer has moved to Business support and want to take advantage of the AWS Health Dashboard and the AWS Health API to automate remediation actions for issues with the health of AWS resources. The first use case is to respond to AWS detecting an IAM access key that is listed on a public code repository site. The automated response will be to delete the IAM access key and send a notification to the Security team. How should this be achieved?

A. A. Create an AWS Lambda function to delete the IAM access key. Send AWS CloudTrail logs to AWS CloudWatch logs. Create a CloudWatch Logs metric filter for the AWS_RISK_CREDENTIALS_EXPOSED event with two actions: first, run the Lambda function; second, use Amazon SNS to send a notification to the Security team.

B. B. Create an AWS Lambda function to delete the IAM access key. Create an AWS Config rule for changes to aws.health and the AWS_RISK_CREDENTIALS_EXPOSED event with two actions: first, run the Lambda function; second, use Amazon SNS to send a notification to the Security team.

C. C. Use AWS Step Functions to create a function to delete the IAM access key, and then use Amazon SNS to send a notification to the Security team. Create an AWS Personal Health Dashboard rule for the AWS_RISK_CREDENTIALS_EXPOSED event; set the target of the Personal Health Dashboard rule to Step Functions.

D. D. Use AWS Step Functions to create a function to delete the IAM access key, and then use Amazon SNS to send a notification to the Security team. Create an Amazon CloudWatch Events rule with an aws.health event source and the AWS_RISK_CREDENTIALS_EXPOSED event, set the target of the CloudWatch Events rule to Step Functions.

Answer(s): D

---

5. The Security team depends on AWS CloudTrail to detect sensitive security issues in the company's AWS account. The DevOps Engineer needs a solution to auto-remediate CloudTrail being turned off in an AWS account. What solution ensures the LEAST amount of downtime for the CloudTrail log deliveries?

A. A. Create an Amazon CloudWatch Events rule for the CloudTrail StopLogging event. Create an AWS Lambda function that uses the AWS SDK to call StartLogging on the ARN of the resource in which StopLogging was called. Add the Lambda function ARN as a target to the CloudWatch Events rule.

B. B. Deploy the AWS-managed CloudTrail-enabled AWS Config rule, set with a periodic interval of 1 hour. Create an Amazon CloudWatch Events rule for AWS Config rules compliance change. Create an AWS Lambda function that uses the AWS SDK to call StartLogging on the ARN of the resource in which StopLogging was called. Add the Lambda function ARN as a target to the CloudWatch Events rule.

C. C. Create an Amazon CloudWatch Events rule for a scheduled event every 5 minutes. Create an AWS Lambda function that uses the AWS SDK to call StartLogging on an CloudTrail trail in the AWS account. Add the Lambda function ARN as a target to the CloudWatch Events rule.

D. D. Launch a t2.nano instance with a script running every 5 minutes that uses the AWS SDK to query CloudTrail in the current account. If the CloudTrail trail is disabled, have the script re-enable the trail.

Answer(s): A

---

6. A DevOps Engineer has been asked by the Security team to ensure that AWS CloudTrail files are not tampered with after being created. Currently, there is a process with multiple trails, using AWS IAM to restrict access to specific trails. The Security team wants to ensure they can trace the integrity of each file and make sure there has been no tampering. Which option will require the LEAST effort to implement and ensure the legitimacy of the file while allowing the Security team to prove the authenticity of the logs?

A. A. Create an Amazon CloudWatch Events rule that triggers an AWS Lambda function when a new file is delivered. Configure the Lambda function to perform an MD5 hash check on the file, store the name and location of the file, and post the returned hash to an Amazon DynamoDB table. The Security team can use the values stored in DynamoDB to verify the file authenticity.

B. B. Enable the CloudTrail file integrity feature on an Amazon S3 bucket. Create an IAM policy that grants the Security team access to the file integrity logs stored in the S3 bucket.

C. C. Enable the CloudTrail file integrity feature on the trail. Use the digest file created by CloudTrail to verify the integrity of the delivered CloudTrail files.

D. D. Create an AWS Lambda function that is triggered each time a new file is delivered to the CloudTrail bucket. Configure the Lambda function to execute an MD5 hash check on the file, and store the result on a tag in an Amazon S3 object. The Security team can use the information on the tag to verify the integrity of the file.

Answer(s): C

---

7. A company is building a web and mobile application that uses a serverless architecture powered by AWS Lambda and Amazon API Gateway. The company wants to fully automate the backend Lambda deployment based on code that is pushed to the appropriate environment branch in an AWS CodeCommit repository. The deployment must have the following: Separate environment pipelines for testing and production. Automatic deployment that occurs for test environments only. Which steps should be taken to meet these requirements?

A. A. Configure a new AWS CodePipeline service. Create a CodeCommit repository for each environment. Set up CodePipeline to retrieve the source code from the appropriate repository. Set up a deployment step to deploy the Lambda functions with AWS CloudFormation.

B. B. Create two AWS CodePipeline configurations for test and production environments. Configure the production pipeline to have a manual approval step. Create a CodeCommit repository for each environment. Set up each CodePipeline to retrieve the source code from the appropriate repository. Set up the deployment step to deploy the Lambda functions with AWS CloudFormation.

C. C. Create two AWS CodePipeline configurations for test and production environments. Configure the production pipeline to have a manual approval step. Create one CodeCommit repository with a branch for each environment. Set up each CodePipeline to retrieve the source code from the appropriate branch in the repository. Set up the deployment step to deploy the Lambda functions with AWS CloudFormation.

D. D. Create an AWS CodeBuild configuration for test and production environments. Configure the production pipeline to have a manual approval step. Create one CodeCommit repository with a branch for each environment. Push the Lambda function code to an Amazon S3 bucket. Set up the deployment step to deploy the Lambda functions from the S3 bucket.

Answer(s): C

---

8. A company is using AWS for an application. The Development team must automate its deployments. The team has set up an AWS CodePipeline to deploy the application to Amazon EC2 instances by using AWS CodeDeploy after it has been built using the AWS CodeBuild service. The team would like to add automated testing to the pipeline to confirm that the application is healthy before deploying it to the next stage of the pipeline using the same code. The team requires a manual approval action before the application is deployed, even if the test is

successful. The testing and approval must be accomplished at the lowest costs, using the simplest management solution. Which solution will meet these requirements?

A. A. Add a manual approval action after the last deploy action of the pipeline. Use Amazon SNS to inform the team of the stage being triggered. Next, add a test action using CodeBuild to do the required tests. At the end of the pipeline, add a deploy action to deploy the application to the next stage.

B. B. Add a test action after the last deploy action of the pipeline. Configure the action to use CodeBuild to perform the required tests. If these tests are successful, mark the action as successful. Add a manual approval action that uses Amazon SNS to notify the team, and add a deploy action to deploy the application to the next stage.

C. C. Create a new pipeline that uses a source action that gets the code from the same repository as the first pipeline. Add a deploy action to deploy the code to a test environment. Use a test action using AWS Lambda to test the deployment. Add a manual approval action by using Amazon SNS to notify the team, and add a deploy action to deploy the application to the next stage.

D. D. Add a test action after the last deployment action. Use a Jenkins server on Amazon EC2 to do the required tests and mark the action as successful if the tests pass. Create a manual approval action that uses Amazon SQS to notify the team and add a deploy action to deploy the application to the next stage

Answer(s): B

---

9. A company is building a solution for storing files containing Personally Identifiable Information (PII) on AWS. Requirements state: All data must be encrypted at rest and in transit. Al data must be replicated in at least two locations that are at least 500 miles apart. Which solution meets these requirements?

A. A. Create primary and secondary Amazon S3 buckets in two separate Availability Zones that are at least 500 miles apart. Use a bucket policy to enforce access to the buckets only through HTTPS. Use a bucket policy to enforce Amazon S3 SSE-C on all objects uploaded to the bucket. Configure cross-region replication between the two buckets.

B. B. Create primary and secondary Amazon S3 buckets in two separate AWS Regions that are at least 500 miles apart. Use a bucket policy to enforce access to the buckets only through HTTPS. Use a bucket policy to enforce S3-Managed Keys (SSE-S3) on all objects uploaded to the bucket. Configure cross-region replication between the two buckets

C. C. Create primary and secondary Amazon S3 buckets in two separate AWS Regions that are at least 500 miles apart. Use an IAM role to enforce access to the buckets only through HTTPS. Use a bucket policy to enforce Amazon S3-Managed Keys (SSE-S3) on all objects uploaded to the bucket. Configure cross-region replication between the two buckets.

D. D. Create primary and secondary Amazon S3 buckets in two separate Availability Zones that are at least 500 miles apart. Use a bucket policy to enforce access to the buckets only through HTTPS. Use a bucket policy to enforce AWS KMS encryption on all objects uploaded to the bucket. Configure cross-region replication between the two buckets. Create a KMS Customer Master Key (CMK) in the primary region for encrypting objects.

Answer(s): B

---

10. A company is using AWS CodeDeploy to automate software deployment. The deployment must meet these requirements: A number of instances must be available to serve traffic during the deployment. Traffic must be balanced across those instances, and the instances must automatically heal in the event of failure. A new fleet of instances must be launched for deploying a new revision automatically, with no manual provisioning. Traffic must be rerouted to the new environment to half of the new instances at a time. The deployment should succeed if traffic is rerouted to at least half of the instances; otherwise, it should fail. Before routing traffic to the new fleet of instances, the temporary files generated during the deployment process must be deleted. At the end of a successful deployment, the original instances in the deployment group must be deleted immediately to reduce costs. How can a DevOps Engineer meet these requirements?

A. A. Use an Application Load Balancer and an in-place deployment. Associate the Auto Scaling group with the deployment group. Use the Automatically copy Auto Scaling group option, and use CodeDeployDefault.OneAtAtime as the deployment configuration. Instruct AWS CodeDeploy to terminate the original instances in the deployment group, and use the AllowTraffic hook within appspec.yml to delete the temporary files.

B. B. Use an Application Load Balancer and a blue/green deployment. Associate the Auto Scaling group and the Application Load Balancer target group with the deployment group. Use the Automatically copy

Auto Scaling group option, create a custom deployment configuration with minimum healthy hosts defined as 50%, and assign the configuration to the deployment group. Instruct AWS CodeDeploy to terminate the original instances in the deployment group, and use the BeforeBlock Traffic hook within appspec.yml to delete the temporary files.

C. C. Use an Application Load Balancer and a blue/green deployment. Associate the Auto Scaling group and the Application Load Balancer target group with the deployment group. Use the Automatically copy Auto Scaling group option, and use CodeDeployDefault HalfAtAtime as the deployment configuration. Instruct AWS CodeDeploy to terminate the original instances in the deployment group, and use the BeforeAllowTraffic hook within appspec.yml to delete the temporary files.

D. D. Use an Application Load Balancer and an in-place deployment. Associate the Auto Scaling group and Application Load Balancer target group with the deployment group. Use the Automatically copy Auto Scaling group option, and use CodeDeployDefault AllatOnce as a deployment configuration. Instruct AWS CodeDeploy to terminate the original instances in the deployment group, and use the BlockTraffic hook within appspec.yml to delete the temporary files.

Answer(s): C

---

11. A DevOps Engineer is working with an application deployed to 12 Amazon EC2 instances across 3 Availability Zones. New instances can be started from an AMI image. On a typical day, each EC2 instance has 30% utilization during business hours and 10% utilization after business hours. The CPU utilization has an immediate spike in the first few minutes of business hours. Other increases in CPU utilization rise gradually. The Engineer has been asked to reduce costs while retaining the same or higher reliability. Which solution meets these requirements?

A. A. Create two Amazon CloudWatch Events rules with schedules before and after business hours begin and end. Create two AWS Lambda functions, one invoked by each rule. The first function should stop nine instances after business hours end, the second function should restart the nine instances before the business day begins.

B. B. Create an Amazon EC2 Auto Scaling group using the AMI image, with a scaling action based on the Auto Scaling group’s CPU Utilization average with a target of 75%. Create a scheduled action for the group to adjust the minimum number of instances to three after business hours end and reset to six before business hours begin.

C. C. Create two Amazon CloudWatch Events rules with schedules before and after business hours begin and end. Create an AWS CloudFormation stack, which creates an EC2 Auto Scaling group, with a parameter for the number of instances. Invoke the stack from each rule, passing a parameter value of three in the morning, and six in the evening.

D. D. Create an EC2 Auto Scaling group using the AMI image, with a scaling action based on the Auto Scaling group’s CPU Utilization average with a target of 75%. Create a scheduled action to terminate

nine instances each evening after the close of business.

Answer(s): B

---

12. A DevOps Engineer must improve the monitoring of a Finance team payments microservice that handles transactions for an e-commerce platform. The microservice runs on multiple Amazon EC2 instances. The Finance team would like to know the number of payments per minute, and the team would like to be notified when this metric falls below a specified threshold. How can this be cost-effectively automated?

A. A. Have the Development team log successful transactions to an application log. Set up Logstash on each instance, which sends logs to an Amazon ES cluster. Create a Kibana dashboard for the Finance team that graphs the metric.

B. B. Have the Development team post the number of successful transactions to Amazon CloudWatch as a custom metric. Create a CloudWatch alarm when the threshold is breached, and use Amazon SNS to notify the Finance team.

C. C. Have the Development team log successful transactions to an application log. On each instance, set up the Amazon CloudWatch Logs agent to send application logs to CloudWatch Logs. Use an EC2 instance to monitor a metric filter, and send notifications to the Finance team.

D. D. Have the Development team log successful transactions to an application log. Set up the Amazon CloudWatch agent on each instance. Create a CloudWatch alarm when the threshold is breached, and use Amazon SNS to notify the Finance team.

Answer(s): B

---

13. A company is migrating an application to AWS that runs on a single Amazon EC2 instance. Because of licensing limitations, the application does not support horizontal scaling. The application will be using Amazon Aurora for its database. How can the DevOps Engineer architect automated healing to automatically recover from EC2 and Aurora failures, in addition to recovering across Availability Zones (AZs), in the MOST cost-effective manner?

A. A. Create an EC2 Auto Scaling group with a minimum and maximum instance count of 1, and have it span across AZs. Use a single-node Aurora instance.

B. B. Create an EC2 instance and enable instance recovery. Create an Aurora database with a read replica in a second AZ, and promote it to a primary database instance if the primary database instance fails.

C. C. Create an Amazon CloudWatch Events rule to trigger an AWS Lambda function to start a new EC2 instance in an available AZ when the instance status reaches a failure state. Create an Aurora database with a read replica in a second AZ, and promote it to a primary database instance when the primary database instance fails.

D. D. Assign an Elastic IP address on the instance. Create a second EC2 instance in a second AZ. Create an Amazon CloudWatch Events rule to trigger an AWS Lambda function to move the Elastic IP address to the second instance when the first instance fails. Use a single-node Aurora instance

Answer(s): C

---

14. An Application team has three environments for their application: development, pre-production, and production. The team recently adopted AWS CodePipeline. However, the team has had several deployments of misconfigured or nonfunctional development code into the production environment, resulting in user disruption and downtime. The DevOps Engineer must review the pipeline and add steps to identify problems with the application before it is deployed. What should the Engineer do to identify functional issues during the deployment process? (Choose two.)

A. A. Use Amazon Inspector to add a test action to the pipeline. Use the Amazon Inspector Runtime Behavior Analysis Inspector rules package to check that the deployed code complies with company security standards before deploying it to production.

B. B. Using AWS CodeBuild to add a test action to the pipeline to replicate common user activities and ensure that the results are as expected before progressing to production deployment.

C. C. Create an AWS CodeDeploy action in the pipeline with a deployment configuration that automatically deploys the application code to a limited number of instances. The action then pauses the deployment so that the QA team can review the application functionality. When the review is complete, CodeDeploy resumes and deploys the application to the remaining production Amazon EC2 instances.

D. D. After the deployment process is complete, run a testing activity on an Amazon EC2 instance in a different region that accesses the application to simulate user behavior if unexpected results occur, the testing activity sends a warning to an Amazon SNS topic. Subscribe to the topic to get updates.

E. E. Add an AWS CodeDeploy action in the pipeline to deploy the latest version of the development code to pre-production. Add a manual approval action in the pipeline so that the QA team can test and confirm the expected functionality. After the manual approval action, add a second CodeDeploy action that deploys the approved code to the production environment.

Answer(s): B E

---

15. A DevOps Engineer is responsible for the deployment of a PHP application. The Engineer is working in a hybrid deployment, with the application running on both on-premises servers and Amazon EC2 instances. The application needs access to a database containing highly confidential information. Application instances need access to database credentials, which must be encrypted at rest and in transit before reaching the instances. How should the Engineer automate the deployment process while also meeting the security requirements?

A. A. Use AWS Elastic Beanstalk with a PHP platform configuration to deploy application packages to the instances. Store database credentials on AWS Systems Manager Parameter Store using the Secure String data type. Define an IAM role for Amazon EC2 allowing access, and decrypt only the database credentials. Associate this role to all the instances.

B. B. Use AWS CodeDeploy to deploy application packages to the instances. Store database credentials on AWS Systems Manager Parameter Store using the Secure String data type. Define an IAM policy for allowing access, and decrypt only the database credentials. Attach the IAM policy to the role associated to the instance profile for CodeDeploy-managed instances, and to the role used for on-premises instances registration on CodeDeploy.

C. C. Use AWS CodeDeploy to deploy application packages to the instances. Store database credentials on AWS Systems Manager Parameter Store using the Secure String data type. Define an IAM role with an attached policy that allows decryption of the database credentials. Associate this role to all the instances and on-premises servers.

D. D. Use AWS CodeDeploy to deploy application packages to the instances. Store database credentials in the AppSpec file. Define an IAM policy for allowing access to only the database credentials. Attach the IAM policy to the role associated to the instance profile for CodeDeploymanaged instances and the role used for on-premises instances registration on CodeDeploy.

Answer(s): B

---

16. A company has a single Developer writing code for an automated deployment pipeline. The Developer is storing source code in an Amazon S3 bucket for each project. The company wants to add more Developers to the team but is concerned about code conflicts and lost work. The company also wants to build a test environment to deploy newer versions of code for testing and

allow Developers to automatically deploy to both environments when code is changed in the repository. What is the MOST efficient way to meet these requirements?

A. A. Create an AWS CodeCommit repository for each project, use the master branch for production code, and create a testing branch for code deployed to testing. Use feature branches to develop new features and pull requests to merge code to testing and master branches.

B. B. Create another S3 bucket for each project for testing code, and use an AWS Lambda function to promote code changes between testing and production buckets. Enable versioning on all buckets to prevent code conflicts.

C. C. Create an AWS CodeCommit repository for each project, and use the master branch for production and test code with different deployment pipelines for each environment. Use feature branches to develop new features.

D. D. Enable versioning and branching on each S3 bucket, use the master branch for production code, and create a testing branch for code deployed to testing. Have Developers use each branch for developing in each environment.

Answer(s): A

---

17. After presenting a working proof of concept for a new application that uses AWS API Gateway, a Developer must set up a team development environment for the project. Due to a tight timeline, the Developer wants to minimize time spent on infrastructure setup, and would like to reuse the code repository created for the proof of concept. Currently, all source code is stored in AWS CodeCommit. Company policy mandates having alpha, beta, and production stages with separate Jenkins servers to build code and run tests for every stage. The Development Manager must have the ability to block code propagation between admins at any time. The Security team wants to make sure that users will not be able to modify the environment without permission. How can this be accomplished?

A. A. Create API Gateway alpha, beta, and production stages. Create a CodeCommit trigger to deploy code to the different stages using an AWS Lambda function.

B. B. Create API Gateway alpha, beta, and production stages. Create an AWS CodePipeline that pulls code from the CodeCommit repository. Create CodePipeline actions to deploy code to the API Gateway stages.

C. C. Create Jenkins servers for the alpha, beta, and production stages on Amazon EC2 instances. Create multiple CodeCommit triggers to deploy code to different stages using an AWS Lambda function.

D. D. Create an AWS CodePipeline pipeline that pulls code from the CodeCommit repository. Create alpha, beta, and production stages with Jenkins servers on CodePipeline.

Answer(s): D

---

18. An online company uses Amazon EC2 Auto Scaling extensively to provide an excellent customer experience while minimizing the number of running EC2 instances. The company's self-hosted Puppet environment in the application layer manages the configuration of the instances. The IT manager wants the lowest licensing costs and wants to ensure that whenever the EC2 Auto Scaling group scales down, removed EC2 instances are deregistered from the Puppet master as soon as possible. How can the requirement be met?

A. A. At instance launch time, use EC2 user data to deploy the AWS CodeDeploy agent. Use CodeDeploy to install the Puppet agent. When the Auto Scaling group scales out, run a script to register the newly deployed instances to the Puppet master. When the Auto Scaling group scales in, use the EC2 Auto Scaling EC2_INSTANCE_TERMINATING lifecycle hook to trigger de-registration from the Puppet master.

B. B. Bake the AWS CodeDeploy agent into the base AMI. When the Auto Scaling group scales out, use CodeDeploy to install the Puppet agent, and execute a script to register the newly deployed instances to the Puppet master. When the Auto Scaling group scales in, use the CodeDeploy ApplicationStop lifecycle hook to run a script to de-register the instance from the Puppet master.

C. C. At instance launch time, use EC2 user data to deploy the AWS CodeDeploy agent. When the Auto Scaling group scales out, use CodeDeploy to install the Puppet agent, and run a script to register the newly deployed instances to the Puppet master. When the Auto Scaling group scales in, use the EC2 user data instance stop script to run a script to de-register the instance from the Puppet master.

D. D. Bake the AWS Systems Manager agent into the base AMI. When the Auto Scaling group scales out, use the AWS Systems Manager to install the Puppet agent, and run a script to register the newly deployed instances to the Puppet master. When the Auto Scaling group scales in, use the Systems Manager instance stop lifecycle hook to run a script to de-register the instance from the Puppet master.

Answer(s): D

---

19. A company discovers that some IAM users have been storing their AWS access keys in configuration files that have been pushed to a Git repository hosting service. Which solution will require the LEAST amount of management overhead while preventing the exposed AWS access keys from being used?

A. A. Build an application that will create a list of all AWS access keys in the account and search each key on Git repository hosting services. If a match is found, configure the application to disable the associated access key. Then deploy the application to an AWS Elastic Beanstalk worker environment and define a periodic task to invoke the application every hour.

B. B. Use Amazon Inspector to detect when a key has been exposed online. Have Amazon Inspector send a notification to an Amazon SNS topic when a key has been exposed. Create an AWS Lambda function subscribed to the SNS topic to disable the IAM user to whom the key belongs, and then delete the key so that it cannot be used.

C. C. Configure AWS Trusted Advisor and create an Amazon CloudWatch Events rule that uses Trusted Advisor as the event source. Configure the CloudWatch Events rule to invoke an AWS Lambda function as the target. If the Lambda function finds the exposed access keys, then have it disable the access key so that it cannot be used.

D. D. Create an AWS Config rule to detect when a key is exposed online. Haw AWS Config send change notifications to an SNS topic. Configure an AWS Lambda function that is subscribed to the SNS topic to check the notification sent by AWS Config, and then disable the access key so it cannot be used.

Answer(s): C

---

20. Company policies require that information about IP traffic going between instances in the production Amazon VPC is captured. The capturing mechanism must always be enabled and the Security team must be notified when any changes in configuration occur. What should be done to ensure that these requirements are met?

A. A. Using the UserData section of an AWS CloudFormation template, install tcpdump on every provisioned Amazon EC2 instance. The output of the tool is sent to Amazon EFS for aggregation and querying. In addition, scheduling an Amazon CloudWatch Events rule calls an AWS Lambda function to check whether tcpdump is up and running and sends an email to the security organization when there is an exception.

B. B. Create a flow log for the production VPC and assign an Amazon S3 bucket as a destination for delivery. Using Amazon S3 Event Notification, set up an AWS Lambda function that is triggered when a new log file gets delivered. This Lambda function updates an entry in Amazon DynamoDB, which is periodically checked by scheduling an Amazon CloudWatch Events rule to notify security when logs have not arrived.

C. C. Create a flow log for the production VPC. Create a new rule using AWS Config that is triggered by configuration changes of resources of type ‘EC2:VPC’. As part of configuring the rule, create an AWS Lambda function that looks up flow logs for a given VPC. If the VPC flow logs are not configured, return a ‘NON_COMPLIANT’ status and notify the security organization.

D. D. Configure a new trail using AWS CloudTrail service. Using the UserData section of an AWS CloudFormation template, install tcpdump on every provisioned Amazon EC2 instance. Connect Amazon Athena to the CloudTrail and write an AWS Lambda function that monitors for a flow log disable event. Once the CloudTrail entry has been spotted, alert the security organization

Answer(s): C
