

Amazon

DOP-C01

AWS Certified DevOps Engineer- Professional



QUESTION & ANSWERS

QUESTION 1

To run an application, a DevOps Engineer launches an Amazon EC2 instances 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 whitelist only outbound traffic to the artifact repository. Remove the security group rule once the install is complete.

Correct Answer: C

Explanation/Reference:

EC2 instances running in private subnets of a VPC can now have controlled access to S3 buckets, objects, and API functions that are in the same region as the VPC. You can use an S3 bucket policy to indicate which VPCs and which VPC Endpoints have access to your S3 buckets 1-

<https://aws.amazon.com/pt/blogs/aws/new-vpc-endpoint-for-amazon-s3/>

QUESTION 2

A DevOps Engineer is asked to implement a strategy for deploying updates to a web application with zero downtime. The application infrastructure is defined in AWS CloudFormation and is made up of an Amazon Route 53 record, an Application Load Balancer, Amazon EC2 instances in an EC2 Auto Scaling group, and Amazon DynamoDB tables. To avoid downtime, there must be an active instance serving the application at all times.

Which strategies will ensure the deployment happens with zero downtime? (Select TWO.)

- A. In the CloudFormation template, modify the `AWS::AutoScaling::AutoScalingGroup` resource and add an `UpdatePolicy` attribute to define the required elements for a deployment with zero downtime.
- B. In the CloudFormation template, modify the `AWS::AutoScaling::DeploymentUpdates` resource and add an `UpdatePolicy` attribute to define the required elements for a deployment with zero downtime.
- C. Add a new Application Load Balancer and Auto Scaling group to the CloudFormation template. Deploy new changes to the inactive Auto Scaling group. Use Route 53 to change the active Application Load Balancer.
- D. Add a new Application Load Balancer and Auto Scaling group to the CloudFormation template.

Modify the `AWS::AutoScaling::AutoScalingGroup` resource and add an `UpdatePolicy` attribute to perform rolling updates.

- E. In the CloudFormation template, modify the `UpdatePolicy` attribute for the CloudFormation stack and specify the Auto Scaling group that will be updated. Configure `MinSuccessfulInstancesPercent` and `PauseTime` to ensure the deployment happens with zero downtime.

Correct Answer: A,C

QUESTION 3

The Deployment team has grown substantially in recent months and so has the number of projects that use separate code repositories. The current process involves configuring AWS CodePipeline manually, and there have been service limit alerts for the count of Amazon S3 buckets.

Which pipeline option will reduce S3 bucket sprawl alerts?

- A. Combine the multiple separate code repositories into a single one, and deploy using a global AWS CodePipeline that has logic for each project.
- B. Create new pipelines by using the AWS API or AWS CLI, and configure them to use a single global S3 bucket with separate prefixes for each project.
- C. Create a new pipeline in a different region for each project to bypass the service limits for S3 buckets in a single region.
- D. Create a new pipeline and for S3 bucket for each project by using the AWS API or AWS CLI to bypass the service limits for S3 buckets in a single account

Correct Answer: C

QUESTION 4

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:

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

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: B

QUESTION 5

A DevOps Engineer must implement monitoring for a workload running on Amazon EC2 and Amazon RDS MySQL. The monitoring must include:

Application logs and operating system metrics for the Amazon EC2 instances

Database logs and operating system metrics for the Amazon RDS database

Which steps should the Engineer take?

- A. Install an Amazon CloudWatch agent on the EC2 and RDS instances. Configure the agent to send the operating system metrics and application and database logs to CloudWatch.
- B. Install an Amazon CloudWatch agent on the EC2 instance, and configure the agent to send the application logs and operating system metrics to CloudWatch. Enable RDS Enhanced Monitoring, and modify the RDS instance to publish database logs to CloudWatch Logs.
- C. Install an Amazon CloudWatch Logs agent on the EC2 instance and configure it to send application logs to CloudWatch.
- D. Set up scheduled tasks on the EC2 and RDS instances to put operating system metrics and application and database logs into an Amazon S3 bucket. Set up an event on the bucket to invoke an AWS Lambda function to monitor for errors each time an object is put into the bucket.

Correct Answer: B

QUESTION 6

A company wants to ensure that their EC2 instances are secure. They want to be notified if any new vulnerabilities are discovered on their instances, and they also want an audit trail of all login activities on the instances. Which solution will meet these requirements?

- A. Use AWS Systems Manager to detect vulnerabilities on the EC2 instances. Install the Amazon Kinesis Agent to capture system logs and deliver them to Amazon S3.
- B. Use AWS Systems Manager to detect vulnerabilities on the EC2 instances. Install the Systems Manager Agent to capture system logs and view login activity in the CloudTrail console.
- C. Configure Amazon CloudWatch to detect vulnerabilities on the EC2 instances. Install the AWS Config daemon to capture system logs and view them in the AWS Config console.
- D. Configure Amazon Inspector to detect vulnerabilities on the EC2 instances. Install the Amazon CloudWatch Agent to capture system logs and record them via Amazon CloudWatch Logs.

Correct Answer: D

QUESTION 7

A DevOps engineer is tasked with creating a more stable deployment solution for a web application in AWS. Previous deployments have resulted in user-facing bugs, premature user traffic, and inconsistencies between web servers running behind an Application Load Balancer. The current strategy uses AWS CodeCommit to store the code for the application. When developers push to the master branch of the repository, CodeCommit triggers an AWS Lambda deploy function, which invokes an AWS Systems Manager run command to build and deploy the new code to all Amazon EC2 instances.

Which combination of actions should be taken to implement a more stable deployment solution? (Select TWO.)

- A. Create a pipeline in AWS CodePipeline with CodeCommit as a source provider. Create parallel pipeline stages to build and test the application. Pass the build artifact to AWS CodeDeploy.
- B. Create a pipeline in AWS CodePipeline with CodeCommit as a source provider. Create separate pipeline stages to build and then test the application. Pass the build artifact to AWS CodeDeploy.
- C. Create and use an AWS CodeDeploy application and deployment group to deploy code updates to the EC2 fleet. Select the Application Load Balancer for the deployment group.
- D. Create individual Lambda functions to run all build, test, and deploy actions using AWS CodeDeploy instead of AWS Systems Manager.
- E. Modify the Lambda function to build a single application package to be shared by all instances. Use AWS CodeDeploy instead of AWS Systems Manager to update the code on the EC2 fleet.

Correct Answer: C,E

QUESTION 8

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

How can the Engineer improve the performance of the website?

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

Correct Answer: C

Explanation/Reference:

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

QUESTION 9

A mobile application running on eight Amazon EC2 instances is relying on a third-party API endpoint. The third-party service has a high failure rate because of limited capacity which is expected to be resolved in a few weeks.

In the meantime the mobile application developers have added a retry mechanism and are logging failed API requests. A DevOps Engineer must automate the monitoring of application logs and count the specific error messages if there are more than 10 errors within a 1-minute window the system must issue an alert

How can the requirements be met with MINIMAL management overhead?

- A. Install AfterAllowTraffic hook to the AppSpec file that forces traffic not having fully propagated before the push the application logs to CloudWatch Logs Use metric filters to count the error messages every minute and trigger a CloudWatch alarm if the count exceeds errors.
- B. Install the Amazon CloudWatch Logs agent on all instances to push the access logs to CloudWatch Logs Create a CloudWatch Events rule to count the error messages every minute and trigger a CloudWatch alarm if the count exceeds 10 errors
- C. Install the Amazon CloudWatch Logs agent on all instances to push the application logs to CloudWatch Logs Use a metric filter to generate a custom CloudWatch metric that records the number of failures and triggers a CloudWatch alarm if the custom metric reaches 10 errors in a 1-minute period
- D. Deploy a custom script on all instances to check application logs regularly in a job Count the number of error messages every minute and push a data point to a custom CloudWatch metric Trigger a CloudWatch alarm if the custom metric reaches 10 errors in a 1-minute period

Correct Answer: C

QUESTION 10

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. 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. 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. 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. 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.

Correct Answer: B

Explanation/Reference:

Explanation: <https://docs.aws.amazon.com/codebuild/latest/userguide/sample-build-notifications.html>

QUESTION 11

An application runs on Amazon EC2 instances behind an Application Load Balancer. Amazon RDS MySQL is used on the backend. The instances run in an Auto Scaling group across multiple Availability Zones. The Application Load Balancer health check ensures the web servers are operating and able to make read/write SQL connections. Amazon Route 53 provides DNS functionality with a record pointing to the Application Load Balancer. A new policy requires a geographically isolated disaster recovery site with an RTO of 4 hours and an RPO of 15 minutes.

Which disaster recovery strategy will require the LEAST amount of changes to the application stack?

- A. Launch a replica stack of everything except RDS in a different Availability Zone. Create an RDS read-only replica in a new Availability Zone and configure the new stack to point to the local RDS instance. Add the new stack to the Route 53 record set with a failover routing policy.
- B. Launch a replica stack of everything except RDS in a different region. Create an RDS read-only replica in a new region and configure the new stack to point to the local RDS instance. Add the new stack to the Route 53 record set with a latency routing policy.
- C. Launch a replica stack of everything except RDS in a different region. Upon failure, copy the snapshot over from the primary region to the disaster recovery region. Adjust the Amazon Route 53 record set to point to the disaster recovery region's Application Load Balancer.
- D. Launch a replica stack of everything except RDS in a different region. Create an RDS read-only replica in a new region and configure the new stack to point to the local RDS instance. Add the new stack to the Amazon Route 53 record set with a failover routing policy.

Correct Answer: D

QUESTION 12

A company has 100 GB of log data in an Amazon S3 bucket stored in .csv format. SQL developers want to query this data and generate graphs to visualize it. They also need an efficient, automated way to store metadata from the .csv file.

Which combination of steps should be taken to meet these requirements with the LEAST amount of effort? (Select THREE.)

- A. Filter the data through AWS X-Ray to visualize the data.
- B. Filter the data through Amazon QuickSight to visualize the data.
- C. Query the data with Amazon Athena.
- D. Query the data with Amazon Redshift.
- E. Use AWS Glue as the persistent metadata store.
- F. Use Amazon S3 as the persistent metadata store.

Correct Answer: B,C,E

QUESTION 13

A company wants to use AWS development tools to replace its current bash deployment scripts. The company currently deploys a LAMP application to a group of Amazon EC2 instances behind an Application Load Balancer (ALB). During the deployments, the company unit tests the committed application, stops and starts services, unregisters and re-registers instances with the load balancer, and updates file permissions. The company wants to maintain the same deployment functionality through the shift to using AWS services.

Which solution will meet these requirements?

- A. Use AWS CodeBuild to test the application. Use bash scripts invoked by AWS CodeDeploy's appspec.yml file to restart services, and deregister and register instances with the ALB. Use the appspec.yml file to update file permissions without a custom script.
- B. Use AWS CodePipeline to move the application from the AWS CodeCommit repository to AWS CodeDeploy. Use CodeDeploy's deployment group to test the application, unregister and reregister instances with the ALB, and restart services. Use the appspec.yml file to update file permissions without a custom script.
- C. Use AWS CodePipeline to move the application source code from the AWS CodeCommit repository to AWS CodeDeploy. Use CodeDeploy to test the application. Use CodeDeploy's appspec.yml file to restart services and update permissions without a custom script. Use AWS CodeBuild to unregister and re-register instances with the ALB.
- D. Use AWS CodePipeline to trigger AWS CodeBuild to test the application. Use bash scripts invoked by AWS CodeDeploy's appspec.yml file to restart services. Unregister and re-register the instances in the AWS CodeDeploy deployment group with the ALB. Update the appspec.yml file to update file permissions without a custom script.

Correct Answer: D

QUESTION 14

A web application with multiple services runs on Amazon EC2 instances behind an Application Load Balancer. The application stores data in an Amazon RDS Multi-AZ DB instance. The instance health check used by the load balancer returns PASS if at least one service is running on the instance. The company uses AWS CodePipeline with AWS CodeBuild and AWS CodeDeploy steps to deploy code to test and production environments. Recently, a new version was unable to connect to the database server in the test environment. One process was running, so the health checks reported healthy and the application was promoted to production, causing a production outage. The company wants to ensure that test builds are fully functional before a promotion to production.

Which changes should a DevOps Engineer make to the test and deployment process? (Choose two.)

- A. Add an automated functional test to the pipeline that ensures solid test cases are performed.
- B. Add a manual approval action to the CodeDeploy deployment pipeline that requires a Testing Engineer to validate the testing environment.
- C. Refactor the health check endpoint the Elastic Load Balancer is checking to better validate actual application functionality.
- D. Refactor the health check endpoint the Elastic Load Balancer is checking to return a text-based status result and configure the load balancer to check for a valid response.
- E. Add a dependency checking step to the existing testing framework to ensure compatibility.

Correct Answer: D,E

QUESTION 15

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

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

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

Correct Answer: C

Explanation/Reference:

Explanation:

https://docs.aws.amazon.com/zh_cn/AmazonCloudFront/latest/DeveloperGuide/lambdaexamples.html

QUESTION 16

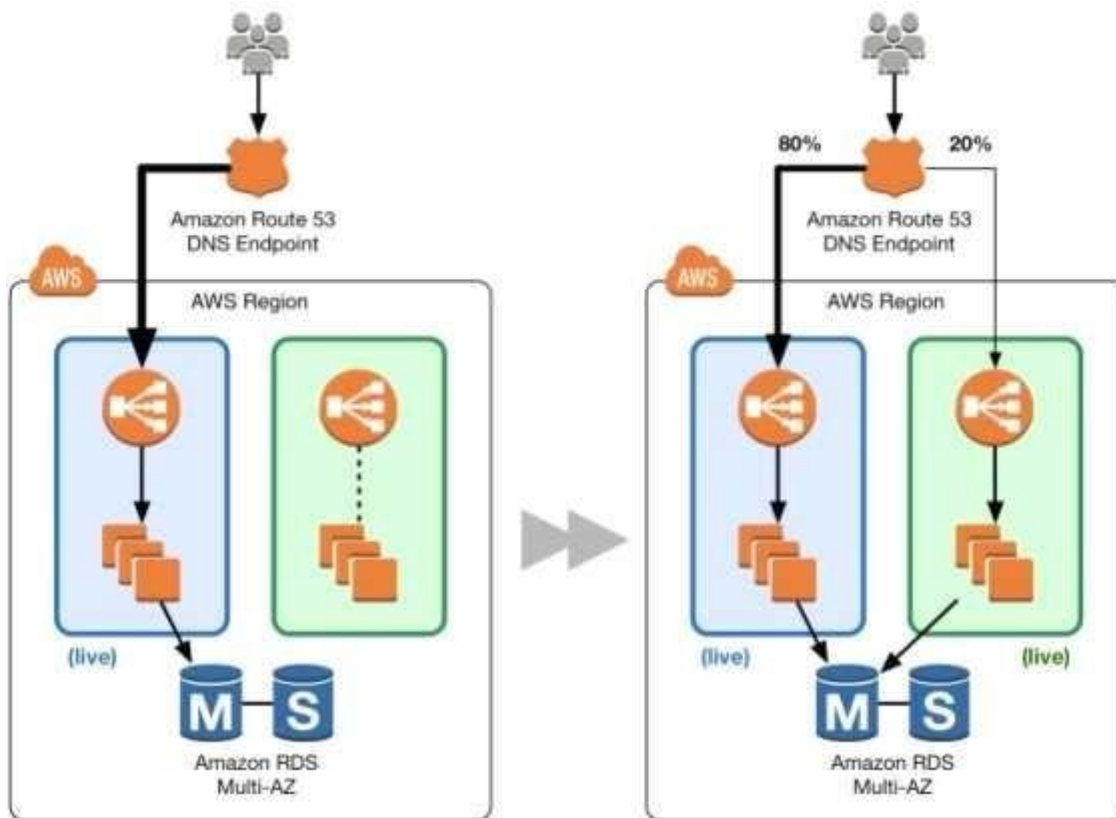
Your application is currently running on Amazon EC2 instances behind a load balancer. Your management has decided to use a Blue/Green deployment strategy. How should you implement this for each deployment?

- A. Set up Amazon Route 53 health checks to fail over from any Amazon EC2 instance that is currently being deployed to.
- B. Using AWS CloudFormation, create a test stack for validating the code, and then deploy the code to each production Amazon EC2 instance.
- C. Create a new load balancer with new Amazon EC2 instances, carry out the deployment, and then switch DNS over to the new load balancer using Amazon Route 53 after testing.
- D. Launch more Amazon EC2 instances to ensure high availability, de-register each Amazon EC2 instance from the load balancer, upgrade it, and test it, and then register it again with the load balancer.

Correct Answer: C

Explanation/Reference:

The below diagram shows how this can be done



- 1) First create a new ELB which will be used to point to the new production changes.
 - 2) Use the Weighted Route policy for Route53 to distribute the traffic to the 2 ELB's based on a 80-20% traffic scenario. This is the normal case, the ?n be changed based on the requirement.
 - 3) Finally when all changes have been tested, Route53 can be set to 100% for the new ELB.
- Option A is incorrect because this is a failover scenario and cannot be used for Blue green deployments. In Blue Green deployments, you need to have 2 environments running side by side.
- Option B is incorrect, because you need to a have a production stack with the changes which will run side by side.
- Option D is incorrect because this is not a blue green deployment scenario. You cannot control which users will go the new EC2 instances.
- For more information on blue green deployments, please refer to the below document link: from AWS
https://d0awsstatic.com/whitepapers/AWS_Blue_Green_Deployments.pdf

QUESTION 17

A company is adopting serverless computing and is migrating some of its existing applications to AWS Lambda A DevOps engineer must come up with an automated deployment strategy using AWS CodePipeline that should include proper version controls, branching strategies, and rollback methods Which combination of steps should the DevOps engineer follow when setting up the pipeline? (Select THREE)

- A. Use Amazon S3 as the source code repository
- B. Use AWS CodeCommit as the source code repository
- C. Use AWS CloudFormation to create an AWS Serverless Application Model (AWS SAM) template for deployment.
- D. Use AWS CodeBuild to create an AWS Serverless Application Model (AWS SAM) template for deployment
- E. Use AWS CloudFormation to deploy the application

F. Use AWS CodeDeploy to deploy the application.

Correct Answer: A,B,C

QUESTION 18

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. 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. 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. 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. 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.

Correct Answer: C

Explanation/Reference:

Explanation: <https://docs.aws.amazon.com/awsccloudtrail/latest/userguide/cloudtrail-log-file-validationintro.html>

QUESTION 19

Which Auto Scaling process would be helpful when testing new instances before sending traffic to them, while still keeping them in your Auto Scaling Group?

- A. Suspend the process AZ Rebalance
- B. Suspend the process Health Check
- C. Suspend the process Replace Unhealthy
- D. Suspend the process AddToLoadBalancer

Correct Answer: D

Explanation/Reference:

If you suspend AddTo Load Balancer, Auto Scaling launches the instances but does not add them to the load balancer or target group. If you resume the AddTo Load Balancer process, Auto Scaling resumes adding instances to the load balancer or target group when they are launched. However, Auto Scaling does not add the instances that were launched while this process was suspended. You must register those instances manually. Option A is invalid because this just balances the number of EC2 instances in the group across the Availability Zones in the region. Option B is invalid because this just checks the health of the instances. Auto Scaling marks an instance as unhealthy if Amazon EC2 or Elastic Load Balancing tells Auto Scaling that the instance is unhealthy. Option C is invalid because this process just terminates instances that are marked as unhealthy and later creates new instances to replace them. For more information on process suspension, please refer to the below document link: from AWS

<http://docs.aws.amazon.com/autoscaling/latest/userguide/as-suspend-resume-processes.html>

QUESTION 20

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

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

Correct Answer: B,C,E

QUESTION 21

A company has established tagging and configuration standards for its infrastructure resources running on AWS. A DevOps Engineer is developing a design that will provide a near-real-time dashboard of the compliance posture with the ability to highlight violations. Which approach meets the stated requirements?

- A. Define the resource configurations in AWS Service Catalog, and monitor the AWS Service Catalog compliance and violations in Amazon CloudWatch. Then, set up and share a live CloudWatch dashboard. Set up Amazon SNS notifications for violations and corrections.
- B. Use AWS Config to record configuration changes and output the data to an Amazon S3 bucket. Create an Amazon QuickSight analysis of the dataset, and use the information on dashboards and

mobile devices.

- C. Create a resource group that displays resources with the specified tags and those without tags. Use the AWS Management Console to view compliant and non-compliant resources.
- D. Define the compliance and tagging requirements in Amazon inspector. Output the results to Amazon CloudWatch Logs. Build a metric filter to isolate the monitored elements of interest and present the data in a CloudWatch dashboard.

Correct Answer: C

Explanation/Reference:

Explanation:

<https://aws.amazon.com/about-aws/whats-new/2019/03/aws-config-now-supports-tagging-of-awsconfig-resources/>

QUESTION 22

A DevOps Engineer is implementing a mechanism for canary testing an application on AWS. The application was recently modified and went through security, unit, and functional testing. The application needs to be deployed on an AutoScaling group and must use a Classic Load Balancer. Which design meets the requirement for canary testing?

- A. Create a different Classic Load Balancer and Auto Scaling group for blue/green environments. Use Amazon Route 53 and create weighted A records on Classic Load Balancer.
- B. Create a single Classic Load Balancer and an Auto Scaling group for blue/green environments. Use Amazon Route 53 and create A records for Classic Load Balancer IPs. Adjust traffic using A records.
- C. Create a single Classic Load Balancer and an Auto Scaling group for blue/green environments. Create an Amazon CloudFront distribution with the Classic Load Balancer as the origin. Adjust traffic using CloudFront.
- D. Create a different Classic Load Balancer and Auto Scaling group for blue/green environments. Create an Amazon API Gateway with a separate stage for the Classic Load Balancer. Adjust traffic by giving weights to this stage.

Correct Answer: A

QUESTION 23

Two teams are working together on different portions of an architecture and are using AWS CloudFormation to manage their resources. One team administers operating system-level updates and patches, while the other team manages application-level dependencies and updates. The Application team must take the most recent AMI when creating new instances and deploying the application.

What is the MOST scalable method for linking these two teams and processes?

- A. The Operating System team uses CloudFormation to create new versions of their AMIs and lists the Amazon Resource names (ARNs) of the AMIs in an encrypted Amazon S3 object as part of the stack output section. The Application team uses a cross-stack reference to load the encrypted S3 object and obtain the most recent AMI ARNs.
- B. The Operating System team uses CloudFormation stack to create an AWS CodePipeline pipeline that builds new AMIs, then places the latest AMI ARNs in an encrypted Amazon S3 object as part of the pipeline output. The Application team uses a cross-stack reference within their own CloudFormation template to get that S3 object location and obtain the most recent AMI ARNs to use when deploying their application.
- C. The Operating System team uses CloudFormation stack to create an AWS CodePipeline pipeline that builds new AMIs. The team then places the AMI ARNs as parameters in AWS Systems Manager Parameter Store as part of the pipeline output. The Application team specifies a parameter of type ssm in their CloudFormation stack to obtain the most recent AMI ARN from the Parameter Store.
- D. The Operating System team maintains a nested stack that includes both the operating system and Application team templates. The Operating System team uses a stack update to deploy updates to the application stack whenever the Application team changes the application code.

Correct Answer: B

QUESTION 24

During metric analysis, your team has determined that the company's website during peak hours is experiencing response times higher than anticipated. You currently rely on Auto Scaling to make sure that you are scaling your environment during peak windows. How can you improve your Auto Scaling policy to reduce this high response time? Choose 2 answers.

- A. Push custom metrics to CloudWatch to monitor your CPU and network bandwidth from your servers, which will allow your Auto Scaling policy to have better fine-grain insight.
- B. Increase your Auto Scaling group's number of max servers.
- C. Create a script that runs and monitors your servers; when it detects an anomaly in load, it posts to an Amazon SNS topic that triggers Elastic Load Balancing to add more servers to the load balancer.
- D. Push custom metrics to CloudWatch for your application that include more detailed information about your web application, such as how many requests it is handling and how many are waiting to be processed.

Correct Answer: B,D

Explanation/Reference:

Option B makes sense because maybe the max servers is low hence the application cannot handle the peak load.

Option D helps in ensuring Autoscaling can scale the group on the right metrics.

For more information on Autoscaling health checks, please refer to the below document link: from AWS

<http://docs.aws.amazon.com/autoscaling/latest/userguide/healthcheck.html>

QUESTION 25

A financial institution provides security-hardened AMIs of Red Hat Enterprise Linux 7.4 and Windows Server 2016 for its application teams to use in deployments. A DevOps Engineer needs to implement an automated daily check of each AMI to monitor for the latest CVE. How should the Engineer implement these checks using Amazon Inspector?

- A. Install the Amazon Inspector agent in each AMI. Configure AWS Step Functions to launch an Amazon EC2 instance for each operating system from the hardened AMI, and tag the instance with SecurityCheck: True. Once EC2 instances have booted up, Step Functions will trigger an Amazon Inspector assessment for all instances with the tag SecurityCheck: True. Implement a scheduled Amazon CloudWatch Events rule that triggers Step Functions once each day.
- B. Tag each AMI with SecurityCheck: True. Configure AWS Step Functions to first compose an Amazon Inspector assessment template for all AMIs that have the tag SecurityCheck: True and second to make a call to the Amazon Inspector API action StartAssessmentRun. Implement a scheduled Amazon CloudWatch Events rule that triggers Step Functions once each day.
- C. Tag each AMI with SecurityCheck: True. Implement a scheduled Amazon Inspector assessment to run once each day for all AMIs with the tag SecurityCheck: True. Amazon Inspector should automatically launch an Amazon EC2 instance for each AMI and perform a security assessment.
- D. Tag each instance with SecurityCheck: True. Implement a scheduled Amazon Inspector assessment to run once each day for all instances with the tag SecurityCheck: True. Amazon Inspector should automatically perform an in-place security assessment for each AMI.

Correct Answer: A

Explanation/Reference:

Explanation: <https://aws.amazon.com/pt/blogs/security/how-to-set-up-continuous-golden-ami-vulnerability-assessments-with-amazon-inspector/>