

# Exam Questions AWS-Certified-DevOps-Engineer-Professional

Amazon AWS Certified DevOps Engineer Professional

<https://www.2passeasy.com/dumps/AWS-Certified-DevOps-Engineer-Professional/>



**NEW QUESTION 1**

A company is using AWS Organizations to create separate AWS accounts for each of its departments. It needs to automate the following tasks:  
Updating the Linux AMIs with new patches periodically and generating a golden image  
Installing a new version of Chef agents in the golden image, if available  
Enforcing the use of the newly generated golden AMIs in the department's account  
Which option requires the LEAST management overhead?

- A. Write a script to launch an Amazon EC2 instance from the previous golden AMI, apply the patch updates, install the new version of the Chef agent, generate a new golden AMI, and then modify the AMI permissions to share only the new image with the departments' accounts.
- B. Use an AWS Systems Manager Run Command to update the Chef agent first, use Amazon EC2 Systems Manager Automation to generate an updated AMI, and then assume an IAM role to copy the new golden AMI into the departments' accounts.
- C. Use AWS Systems Manager Automation to update the Linux AMI using the previous image, provide the URL for the script that will update the Chef agent, and then use AWS Organizations to replace the previous golden AMI into the departments' accounts.
- D. Use AWS Systems Manager Automation to update the Linux AMI from the previous golden image, provide the URL for the script that will update the Chef agent, and then share only the newly generated AMI with the departments' accounts.

**Answer: C**

**NEW QUESTION 2**

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 component
- B. Create a CloudWatch subscription filter for each component with Lambda as the subscription feed destination.
- C. Enable Amazon CloudWatch Logs to log the EKS component
- D. Create CloudWatch Logs Insights queries linked to Amazon CloudWatch Events events that trigger Lambda.
- E. Enable Amazon S3 logging for the EKS component
- F. Configure an Amazon CloudWatch subscription filter for each component with Lambda as the subscription feed destination.
- G. Enable Amazon S3 logging for the EKS component
- H. Configure S3 PUT Object event notifications with AWS Lambda as the destination.

**Answer: A**

**NEW QUESTION 3**

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. Use an Application Load Balancer and an in-place deployment
- B. Associate the Auto Scaling group with the deployment group
- C. Use the Automatically copy option, and use CodeDeployDefault.OneAtATime as the deployment configuration
- D. Instruct AWS CodeDeploy to terminate the original Auto Scaling group instances in the deployment group, and use the AllowTraffic hook within appspec.yml to delete the temporary files.
- E. Use an Application Load Balancer and a blue/green deployment
- F. Associate the Auto Scaling group and the Application Load Balancer target group with the deployment group
- G. 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
- H. Instruct AWS CodeDeploy to terminate the original instances in the deployment group, and use the BeforeBlockTraffic hook within appspec.yml to delete the temporary files.
- I. Use an Application Load Balancer and a blue/green deployment
- J. Associate the Auto Scaling group and the Application Load Balancer target group with the deployment group
- K. Use the Automatically copy Auto Scaling group option, and use CodeDeployDefault.HalfAtATime as the deployment configuration
- L. 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.
- M. Use an Application Load Balancer and an in-place deployment
- N. Associate the Auto Scaling group and Application Load Balancer target group with the deployment group
- O. Use the Automatically copy AutoScaling group option, and use CodeDeployDefault.AllatOnce as a deployment configuration
- P. 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: C**

**Explanation:**

<https://docs.aws.amazon.com/codedeploy/latest/userguide/deployment-configurations.html>

[https://docs.aws.amazon.com/codedeploy/latest/APIReference/API\\_BlueGreenDeploymentConfiguration.html](https://docs.aws.amazon.com/codedeploy/latest/APIReference/API_BlueGreenDeploymentConfiguration.html)

**NEW QUESTION 4**

You have an application consisting of a stateless web server tier running on Amazon EC2 instances behind load balancer, and are using Amazon RDS with read replicas. Which of the following methods should you use to implement a self-healing and cost-effective architecture? Choose 2 answers from the options given below

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

unhealthy Amazon EC2 instances.

- B. Set up scripts on each Amazon EC2 instance to frequently send ICMP pings to the load balancer in order to determine which instance is unhealthy and replace it.
- C. Set up an Auto Scaling group for the web server tier along with an Auto Scaling policy that uses the Amazon RDS DB CPU utilization Cloud Watch metric to scale the instances.
- D. Set up an Auto Scaling group for the web server tier along with an Auto Scaling policy that uses the Amazon EC2 CPU utilization CloudWatch metric to scale the instances.
- E. Use a larger Amazon EC2 instance type for the web server tier and a larger DB instance type for the data storage layer to ensure that they don't become unhealthy.
- F. Set up an Auto Scaling group for the database tier along with an Auto Scaling policy that uses the Amazon RDS read replica lag CloudWatch metric to scale out the Amazon RDS read replicas.
- G. Use an Amazon RDS Multi-AZ deployment.

**Answer:** DG

**Explanation:**

The scaling of EC2 Instances in the Autoscaling group is normally done with the metric of the CPU utilization of the current instances in the Autoscaling group. For more information on scaling in your Autoscaling Group, please refer to the below link:

- <http://docs.aws.amazon.com/autoscaling/latest/userguide/as-scaling-simple-step.html>

Amazon RDS Multi-AZ deployments provide enhanced availability and durability for Database (DB) Instances, making them a natural fit for production database workloads. When you provision a Multi-AZ DB Instance, Amazon RDS automatically creates a primary DB Instance and synchronously replicates the data to a standby instance in a different Availability Zone (AZ). Each AZ runs on its own physically distinct, independent infrastructure, and is engineered to be highly reliable. In case of an infrastructure failure, Amazon RDS performs an automatic failover to the standby or to a read replica in the case of Amazon Aurora, so that you can resume database operations as soon as the failover is complete. For more information on RDS

Multi-AZ please refer to the below link:

- <https://aws.amazon.com/rds/details/multi-az/>

Option A is invalid because if you already have in-built metrics from Cloudwatch, why would you want to spend more in using a third-party monitoring solution.

Option B is invalid because health checks are already a feature of AWS CLB

Option C is invalid because the database CPU usage should not be used to scale the web tier.

Option D is invalid because increasing the instance size does not always guarantee that the solution will not become unhealthy.

Option F is invalid because increasing Read-Replica's will not suffice for write operations if the primary DB fails.

**NEW QUESTION 5**

An n-tier application requires a table in an Amazon RDS MySQL DB instance to be dropped and repopulated at each deployment. This process can take several minutes and the web tier cannot come online until the process is complete. Currently, the web tier is configured in an Amazon EC2 Auto Scaling group, with instances being terminated and replaced at each deployment. The MySQL table is populated by running a SQL query through an AWS CodeBuild job.

What should be done to ensure that the web tier does not come online before the database is completely configured?

- A. Use Amazon Aurora as a drop-in replacement for RDS MySQL
- B. Use snapshots to populate the table with the correct data.
- C. Modify the launch configuration of the Auto Scaling group to pause user data execution for 600 seconds, allowing the table to be populated.
- D. Use AWS Step Functions to monitor and maintain the state of data population
- E. Mark the database in service before continuing with the deployment.
- F. Use an EC2 Auto Scaling lifecycle hook to pause the configuration of the web tier until the table is populated.

**Answer:** D

**NEW QUESTION 6**

A DevOps Engineer has a single Amazon DynamoDB table that received shipping orders and tracks inventory. The Engineer has three AWS Lambda functions reading from a DynamoDB stream on that table. The Lambda functions perform various functions such as doing an item count, moving items to Amazon Kinesis Data Firehose, monitoring inventory levels, and creating vendor orders when parts are low.

While reviewing logs, the Engineer notices the Lambda functions occasionally fail under increased load, receiving a stream throttling error.

Which is the MOST cost-effective solution that requires the LEAST amount of operational management?

- A. Use AWS Glue integration to ingest the DynamoDB stream, then migrate the Lambda code to an AWS Fargate task.
- B. Use Amazon Kinesis streams instead of DynamoDB streams, then use Kinesis analytics to trigger the Lambda functions.
- C. Create a fourth Lambda function and configure it to be the only Lambda reading from the stream
- D. Then use this Lambda function to pass the payload to the other three Lambda functions.
- E. Have the Lambda functions query the table directly and disable DynamoDB stream
- F. Then have the Lambda functions query from a global secondary index.

**Answer:** C

**NEW QUESTION 7**

A global company with distributed Development teams built a web application using a microservices architecture running on Amazon ECS. Each application service is independent and runs as a service in the ECS cluster. The container build files and source code reside in a private GitHub source code repository. Separate ECS clusters exist for development, testing, and production environments.

Developers are required to push features to branches in the GitHub repository and then merge the changes into an environment-specific branch (development, test, or production). This merge needs to trigger an automated pipeline to run a build and a deployment to the appropriate ECS cluster.

What should the DevOps Engineer recommend as an automated solution to these requirements?

- A. Create an AWS CloudFormation stack for the ECS cluster and AWS CodePipeline service
- B. Store the container build files in an Amazon S3 bucket
- C. Use a post-commit hook to trigger a CloudFormation stack update that deploys the ECS cluster
- D. Add a task in the ECS cluster to build and push images to Amazon ECR, based on the container build files in S3.
- E. Create a separate pipeline in AWS CodePipeline for each environment
- F. Trigger each pipeline based on commits to the corresponding environment branch in GitHub
- G. Add a build stage to launch AWS CodeBuild to create the container image from the build file and push it to Amazon ECR

- H. Then add another stage to update the Amazon ECS task and service definitions in the appropriate cluster for that environment.
- I. Create a pipeline in AWS CodePipel
- J. Configure it to be triggered by commits to the master branch in GitHu
- K. Add a stage to use the Git commit message to determine which environment the commit should be applied to, then call the create-image Amazon ECR command to build the image, passing it to the container build fil
- L. Then add a stage to update the ECS task and service definitions in the appropriate cluster for that environment.
- M. Create a new repository in AWS CodeCommi
- N. Configure a scheduled project in AWS CodeBuild to synchronize the GitHub repository to the new CodeCommit repositor
- O. Create a separate pipeline for each environment triggered by changes to the CodeCommit repositor
- P. Add a stage using AWS Lambda to build the container image and push to Amazon EC
- Q. Then add another stage to update the ECS task and service definitions in the appropriate cluster for that environment.

**Answer: B**

**Explanation:**

<https://docs.aws.amazon.com/AmazonECS/latest/developerguide/ecs-cd-pipeline.html>

**NEW QUESTION 8**

A consulting company was hired to assess security vulnerabilities within a client company's application and propose a plan to remediate all identified issues. The architecture is identified as follows: Amazon S3 storage for content, an Auto Scaling group of Amazon EC2 instances behind an Elastic Load Balancer with attached Amazon EBS storage, and an Amazon RDS MySQL database. There are also several AWS Lambda functions that communicate directly with the RDS database using connection string statements in the code.

The consultants identified the top security threat as follows: the application is not meeting its requirement to have encryption at rest.

What solution will address this issue with the LEAST operational overhead and will provide monitoring for potential future violations?

- A. Enable SSE encryption on the S3 buckets and RDS databas
- B. Enable OS-based encryption of data on EBS volume
- C. Configure Amazon Inspector agents on EC2 instances to report on insecure encryption cipher
- D. Set up AWS Config rules to periodically check for non-encrypted S3 objects.
- E. Configure the application to encrypt each file prior to storing on Amazon S3. Enable OS-based encryption of data on EBS volume
- F. Encrypt data on write to RD
- G. Run cron jobs on each instance to check for encrypted data and notify via Amazon SN
- H. Use S3 Events to call an AWS Lambda function and verify if the file is encrypted.
- I. Enable Secure Sockets Layer (SSL) on the load balancer, ensure that AWS Lambda is using SSL to communicate to the RDS database, and enable S3 encryptio
- J. Configure the application to force SSL for incoming connections and configure RDS to only grant access if the session is encrypte
- K. Configure Amazon Inspector agents on EC2 instances to report on insecure encryption ciphers.
- L. Enable SSE encryption on the S3 buckets, EBS volumes, and the RDS databas
- M. Store RDS credentials in EC2 Parameter Stor
- N. Enable a policy on the S3 bucket to deny unencrypted put
- O. Set up AWS Config rules to periodically check for non-encrypted S3 objects and EBS volumes, and to ensure that RDS storage is encrypted.

**Answer: D**

**NEW QUESTION 9**

A government agency is storing highly confidential files in an encrypted Amazon S3 bucket. The agency has configured federated access and has allowed only a particular on-premises Active Directory user group to access this bucket.

The agency wants to maintain audit records and automatically detect and revert any accidental changes administrators make to the IAM policies used for providing this restricted federated access.

Which of the following options provide the FASTEST way to meet these requirements?

- A. Configure an Amazon CloudWatch Events Event Bus on an AWS CloudTrail API for triggering the AWS Lambda function that detects and reverts the change.
- B. Configure an AWS Config rule to detect the configuration change and execute an AWS Lambda function to revert the change.
- C. Schedule an AWS Lambda function that will scan the IAM policy attached to the federated access role for detecting and reverting any changes.
- D. Restrict administrators in the on-premises Active Directory from changing the IAM policies

**Answer: B**

**Explanation:**

<https://www.puresec.io/blog/aws-security-best-practices-config-rules-lambda-security> "Cloudwatch Event Bus" are used for -> "Sending and Receiving Events Between AWS Accounts"

<https://aws.amazon.com/about-aws/whats-new/2017/06/cloudwatch-events-adds-cross-account-event-delivery-s>

<https://docs.aws.amazon.com/config/latest/developerguide/evaluate-config-rules.html>

**NEW QUESTION 10**

A DevOps Engineer is leading the implementation for automating patching of Windows-based workstations in a hybrid cloud environment by using AWS Systems Manager (SSM).

What steps should the Engineer follow to set up Systems Manager to automate patching in this environment? (Select TWO.)

- A. Create multiple IAM service roles for Systems Manager so that the ssm.amazonaws.com service can execute the AssumeRole operation on every instanc
- B. Register the role on a per-resource level to enable the creation of a service toke
- C. Perform managed-instance activation with the newly created service role attached to each managed instance.
- D. Create an IAM service role for Systems Manager so that the ssm.amazonaws.com service can execute the AssumeRole operatio
- E. Register the role to enable the creation of a service toke
- F. Perform managed-instance activation with the newly created service role.
- G. Using previously obtained activation codes and activation IDs, download and install the SSM Agent on the hybrid servers, and register the servers or virtual machines on the Systems Manager servic
- H. Hybrid instances will show with an "mi-" prefix in the SSM console.
- I. Using previously obtained activation codes and activation IDs, download and install the SSM Agent on the hybrid servers, and register the servers or virtual

- machines on the Systems Manager service
- J. Hybrid instances will show with an "i-" prefix in the SSM console as if they were provisioned as a regular Amazon EC2 instance.
- K. Run AWS Config to create a list of instances that are unpatched and not compliant
- L. Create an instance scheduler job, and through an AWS Lambda function, perform the instance patching to bring them up to compliance.

**Answer:** BC

**Explanation:**

<https://docs.aws.amazon.com/systems-manager/latest/userguide/sysman-managed-instance-activation.html>  
<https://docs.aws.amazon.com/systems-manager/latest/userguide/sysman-install-managed-win.html>

**NEW QUESTION 10**

A company wants to migrate a legacy application to AWS and develop a deployment pipeline that uses AWS services only. A DevOps engineer is migrating all of the application code from a Git repository to AWS CodeCommit while preserving the history of the repository. The DevOps engineer has set all the permissions within CodeCommit, installed the Git client and the AWS CLI on a local computer, and is ready to migrate the repository. Which actions will follow?

- A. Create the CodeCommit repository using the AWS CL
- B. Clone the Git repository directly to CodeCommit using the AWS CL
- C. Validate that the files were migrated, and publish the CodeCommit repository.
- D. Create the CodeCommit repository using the AWS Management Console
- E. Clone both the Git and CodeCommit repositories to the local compute
- F. Copy the files from the Git repository to the CodeCommit repository on the local compute
- G. Commit the CodeCommit repository
- H. Validate that the files were migrated, and share the CodeCommit repository.
- I. Create the CodeCommit repository using the AWS Management Console
- J. Use the console to clone the Git repository into the CodeCommit repository
- K. Validate that the files were migrated, and publish the CodeCommit repository.
- L. Create the CodeCommit repository using the AWS Management Console or the AWS CL
- M. Clone the Git repository with a mirror argument to the local computer and push the repository to CodeCommit
- N. Validate that the files were migrated, and share the CodeCommit repository.

**Answer:** D

**NEW QUESTION 11**

A DevOps engineer wants to deploy a serverless web application based on AWS Lambda. The deployment must meet the following requirements:

- Provide staging and production environments.
- Restrict the developers from accessing the production environment.
- Avoid hard coding passwords in the Lambda functions
- Store source code in AWS CodeCommit.
- Use AWS CodePipeline to automate the deployment. Which solution will accomplish this?

- A. Create separate staging and production accounts to segregate deployment targets
- B. Use AWS KMS to store environment-specific values Use CodePipeline to automate deployments with AWS CodeDeploy.
- C. Create separate staging and production accounts to segregate deployment targets
- D. Use Lambda environment variables to store environment-specific values
- E. Use CodePipeline to automate deployments with AWS CodeDeploy.
- F. Define tagging conventions for staging and production environments to segregate deployment targets. Use AWS KMS to store environment-specific values Use CodePipeline to automate deployments with AWS CodeDeploy.
- G. Define naming conventions for staging and production environments to segregate deployment targets. Use Lambda environment variables to store environment-specific values
- H. Use CodePipeline to automate deployments with AWS CodeDeploy

**Answer:** A

**NEW QUESTION 16**

A Security team is concerned that a Developer can unintentionally attach an Elastic IP address to an Amazon EC2 instance in production. No Developer should be allowed to attach an Elastic IP address to an instance. The Security team must be notified if any production server has an Elastic IP address at any time. How can this task be automated?

- A. Use Amazon Athena to query AWS CloudTrail logs to check for any associate-address attempt
- B. Create an AWS Lambda function to dissociate the Elastic IP address from the instance, and alert the Security team.
- C. Attach an IAM policy to the Developer's IAM group to deny associate-address permission
- D. Create a custom AWS Config rule to check whether an Elastic IP address is associated with any instance tagged as production, and alert the Security team.
- E. Ensure that all IAM groups associated with Developers do not have associate-address permissions. Create a scheduled AWS Lambda function to check whether an Elastic IP address is associated with any instance tagged as production, and alert the Security team if an instance has an Elastic IP address associated with it.
- F. Create an AWS Config rule to check that all production instances have the EC2 IAM roles that include deny associate-address permission
- G. Verify whether there is an Elastic IP address associated with any instance, and alert the Security team if an instance has an Elastic IP address associated with it.

**Answer:** B

**Explanation:**

<https://docs.aws.amazon.com/vpc/latest/userguide/vpc-migrate-ipv6.html#vpc-migrate-ipv6-sg-rules>

**NEW QUESTION 20**

A production account has a requirement that any Amazon EC2 instance that has been logged into manually must be terminated within 24 hours. All applications in the production account are using Auto Scaling groups with Amazon CloudWatch Logs agent configured.

How can this process be automated?

- A. Create a CloudWatch Logs subscription to an AWS Step Functions applicatio
- B. Configure the function to add a tag to the EC2 instance that produced the login event and mark the instance to be decommissione
- C. Then create a CloudWatch Events rule to trigger a second AWS Lambda function once a day that will terminate all instances with this tag.
- D. Create a CloudWatch alarm that will trigger on the login even
- E. Send the notification to an Amazon SNS topic that the Operations team is subscribed to, and have them terminate the EC2 instance within 24 hours.
- F. Create a CloudWatch alarm that will trigger on the login even
- G. Configure the alarm to send to an Amazon SQS queu
- H. Use a group of worker instances to process messages from the queue, which then schedules the Amazon CloudWatch Events rule to trigger.
- I. Create a CloudWatch Logs subscription in an AWS Lambda functio
- J. Configure the function to add a tag to the EC2 instance that produced the login event and mark the instance to be decommissione
- K. Create a CloudWatch Events rule to trigger a daily Lambda function that terminates all instances withthis ta

**Answer:** D

**Explanation:**

<https://boto3.amazonaws.com/v1/documentation/api/latest/guide/cw-example-subscription-filters.html>

**NEW QUESTION 21**

A company wants to use Amazon DynamoDB for maintaining metadata on its forums. See the sample data set in the image below.

**Thread**

ForumName	Subject	LastPostDateTime	Thread
"S3"	"aaa"	"2015-03-15:17:24:31"	12
"S3"	"bbb"	"2015-01-22:23:18:01"	3
"S3"	"ccc"	"2015-02-31:13:14:21"	4
"S3"	"ddd"	"2015-01-03:09:21:11"	9
"EC2"	"yyy"	"2015-02-12:11:07:56"	18
"EC2"	"zzz"	"2015-01-18:07:33:42"	0
"RDS"	"rrr"	"2015-01-19:01:13:24"	3
"RDS"	"sss"	"2015-03-11:06:53:00"	11
"RDS"	"ttt"	"2015-10-22:12:19:44"	5

A DevOps Engineer is required to define the table schema with the partition key, the sort key, the local secondary index, projected attributes, and fetch operations. The schema should support the following example searches using the least provisioned read capacity units to minimize cost.

- Search within ForumName for items where the subject starts with "'a'.
- Search forums within the given LastPostDateTime time frame.
- Return the thread value where LastPostDateTime is within the last three months. Which schema meets the requirements?

- A. Use Subject as the primary key and ForumName as the sort ke
- B. Have LSI with LastPostDateTime as the sort key and fetch operations for thread.
- C. Use ForumName as the primary key and Subject as the sort ke
- D. Have LSI with LastPostDateTime as the sort key and the projected attribute thread.
- E. Use ForumName as the primary key and Subject as the sort ke
- F. Have LSI with Thread as the sort key and the projected attribute LastPostDateTime.
- G. Use Subject as the primary key and ForumName as the sort ke
- H. Have LSI with Thread as the sort key and fetch operations for LastPostDateTime.

**Answer:** B

**Explanation:**

<https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/LSI.html>

**NEW QUESTION 26**

A company is running an application on Amazon EC2 instances in an Auto Scaling group. Recently, an issue occurred that prevented EC2 instances from launching successfully, and it took several hours for the support team to discover the issue. The support team wants to be notified by email whenever an EC2 instance does not start successfully.

Which action will accomplish this?

- A. Add a health check to the Auto Scaling group to invoke an AWS Lambda function whenever an instance status is impaired.
- B. Configure the Auto Scaling group to send a notification to an Amazon SNS topic whenever a failed instance launch occurs.
- C. Create an Amazon CloudWatch alarm that invokes an AWS Lambda function when a failed AttachInstances Auto Scaling API call is made.
- D. Create a status check alarm on Amazon EC2 to send a notification to an Amazon SNS topic whenever a status check fail occurs.

**Answer:** B

#### NEW QUESTION 27

A DevOps engineer is architecting a continuous development strategy for a company's software as a service (SaaS) web application running on AWS. For application and security reasons, users subscribing to this application are distributed across multiple Application Load Balancers (ALBs), each of which has a dedicated Auto Scaling group and fleet of Amazon EC2 instances. The application does not require a build stage, and when it is committed to AWS CodeCommit, the application must trigger a simultaneous deployment to all ALBs. Auto Scaling groups, and EC2 fleets. Which architecture will meet these requirements with the LEAST amount of configuration?

- A. Create a single AWS CodePipeline pipeline that deploys the application in parallel using unique AWS CodeDeploy applications and deployment groups created for each ALB-Auto Scaling group pair.
- B. Create a single AWS CodePipeline pipeline that deploys the application using a single AWS CodeDeploy application and single deployment group.
- C. Create a single AWS CodePipeline pipeline that deploys the application in parallel using a single AWS CodeDeploy application and unique deployment group for each ALB-Auto Scaling group pair.
- D. Create an AWS CodePipeline pipeline for each ALB-Auto Scaling group pair that deploys the application using an AWS CodeDeploy application and deployment group created for the same ALB-Auto Scaling group pair.

**Answer: C**

#### NEW QUESTION 29

A DevOps engineer is assisting with a multi-Region disaster recovery solution for a new application. The application consists of Amazon EC2 instances running in an Auto Scaling group and an Amazon Aurora MySQL DB cluster. The application must be available with an RTO of 120 minutes and an RPO of 60 minutes. What is the MOST cost-effective way to meet these requirements?

- A. Launch an Aurora DB cluster as an Aurora Replica in a different Region
- B. Create an AWS CloudFormation template for all compute resources and create a stack in two Region
- C. Write a script that promotes the Aurora Replica to the primary instance in the event of a failure.
- D. Launch an Aurora DB cluster as an Aurora Replica in a different Region and configure automatic cross-Region failover
- E. Create an AWS CloudFormation template that includes an Auto Scaling group, and create a stack in two Region
- F. Write a script that updates the CloudFormation stack in the disaster recovery Region to increase the number of instances.
- G. Use AWS Lambda to create and copy a snapshot of the Aurora DB cluster to the destination Region hourly
- H. Create an AWS CloudFormation template that includes an Auto Scaling group, and create a stack in two Region
- I. Restore the Aurora DB cluster from a snapshot and update the Auto Scaling group to start launching instances.
- J. Configure Amazon DynamoDB cross-Region replication
- K. Create an AWS CloudFormation template that includes an Auto Scaling group, and create a stack in two Region
- L. Write a script that will update the CloudFormation stack in the disaster recovery Region and promote the DynamoDB replica to the primary instance in the event of a failure.

**Answer: D**

#### NEW QUESTION 33

A development team is using AWS CodeCommit to version control application code and AWS CodePipeline to orchestrate software deployments. The team has decided to use a remote master branch as the trigger (or the pipeline to integrate code changes. A developer has pushed code changes to the CodeCommit repository, but noticed that the pipeline had no reaction, even after 10 minutes. Which of the following actions should be taken to troubleshoot this issue?

- A. Check that an Amazon CloudWatch Events rule has been created for the master branch to trigger the pipeline.
- B. Check that the CodePipeline service role has permission to access the CodeCommit repository.
- C. Check that the developer's IAM role has permission to push to the CodeCommit repository.
- D. Check to see if the pipeline failed to start because of CodeCommit errors in Amazon CloudWatch Logs.

**Answer: C**

#### NEW QUESTION 37

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

- A. Create Amazon Route 53 health checks for each endpoint that cannot be entered as alias record
- B. Ensure firewall and routing rules allow Amazon Route 53 to send requests to the endpoints that are specified in the health checks.
- C. Create alias records that route traffic to AWS resources and set the value of the Evaluate Target Health option to Yes, then create all the non-alias records.
- D. Create a governing Amazon Route 53 record set, set it to failover, and associate it with the primary and secondary Amazon Route 53 record sets to distribute traffic to healthy DNS entries.
- E. Create an Amazon CloudWatch alarm to monitor the primary Amazon Route 53 DNS entry
- F. Then create an associated AWS Lambda function to execute the failover API call to Route 53 to the secondary DNS entry.

**Answer: AC**

#### NEW QUESTION 38

A DevOps Engineer administers an application that manages video files for a video production company. The application runs on Amazon EC2 instances behind an ELB Application Load Balancer. The instances run in an Auto Scaling group across multiple Availability Zones. Data is stored in an Amazon RDS PostgreSQL Multi-AZ DB instance, and the video files are stored in an Amazon S3 bucket. On a typical day, 50 GB of new video are added to the S3 bucket. The Engineer must implement a multi-region disaster recovery plan with the least data loss and the lowest recovery times. The current application infrastructure is already described using AWS CloudFormation.

Which deployment option should the Engineer choose to meet the uptime and recovery objectives for the system?

- A. Launch the application from the CloudFormation template in the second region, which sets the capacity of the Auto Scaling group to 1. Create an Amazon RDS read replica in the second region
- B. In the second region, enable cross-region replication between the original S3 bucket and a new S3 bucket
- C. To fail over, promote the read replica as master
- D. Update the CloudFormation stack and increase the capacity of the Auto Scaling group.

- E. Launch the application from the CloudFormation template in the second region, which sets the capacity of the Auto Scaling group to 1. Create a scheduled task to take daily Amazon RDS cross-region snapshots to the second region
- F. In the second region, enable cross-region replication between the original S3 bucket and Amazon Glacie
- G. In a disaster, launch a new application stack in the second region and restore the database from the most recent snapshot.
- H. Launch the application from the CloudFormation template in the second region which sets the capacity of the Auto Scaling group to 1. Use Amazon CloudWatch Events to schedule a nightly task to take a snapshot of the database, copy the snapshot to the second region, and replace the DB instance in the second region from the snapshot
- I. In the second region, enable cross-region replication between the original S3 bucket and a new S3 bucket
- J. To fail over, increase the capacity of the Auto Scaling group.
- K. Use Amazon CloudWatch Events to schedule a nightly task to take a snapshot of the database and copy the snapshot to the second region
- L. Create an AWS Lambda function that copies each object to a new S3 bucket in the second region in response to S3 event notification
- M. In the second region, launch the application from the CloudFormation template and restore the database from the most recent snapshot.

**Answer: A**

#### NEW QUESTION 41

A DevOps engineer is researching the least expensive way to implement an image batch processing cluster on AWS. The application cannot run in Docker containers and must run on Amazon EC2. The batch job stores checkpoint data on an NFS and can tolerate interruptions. Configuring the cluster software from a generic EC2 Linux image takes 30 minutes.

What is the MOST cost-effective solution?

- A. Use Amazon EFS for checkpoint data
- B. To complete the job
- C. Use an EC2 Auto Scaling group and an On-Demand pricing model to provision EC2 instances temporarily.
- D. Use GlusterFS on EC2 instances for checkpoint data
- E. To run the batch job
- F. Configure EC2 instances manually
- G. When the job completes, shut down the instances manually.
- H. Use Amazon EFS for checkpoint data
- I. Use EC2 Fleet to launch EC2 Spot Instances, and utilize user data to configure the EC2 Linux instance on startup.
- J. Use Amazon EFS for checkpoint data
- K. Use EC2 Fleet to launch EC2 Spot Instance
- L. Create a custom AMI for the cluster and use the latest AMI when creating instances.

**Answer: A**

#### NEW QUESTION 42

A company has multiple child accounts that are part of an organization in AWS Organizations. The security team needs to review every Amazon EC2 security group and their inbound and outbound rules. The security team wants to programmatically retrieve this information from the child accounts using an AWS Lambda function in the master account of the organization.

Which combination of access changes will meet these requirements? (Select THREE.)

- A. Create a trust relationship that allows users in the child accounts to assume the master account IAM role.
- B. Create a trust relationship that allows users in the master account to assume the IAM roles of the child accounts.
- C. Create an IAM role in each child account that has access to the AmazonEC2ReadOnlyAccess managed policy.
- D. Create an IAM role in each child account to allow the sts:AssumeRole action against the master account IAM role's ARN.
- E. Create an IAM role in the master account that allows the sts:AssumeRole action against the child account IAM role's ARN.
- F. Create an IAM role in the master account that has access to the AmazonEC2ReadOnlyAccess managed policy.

**Answer: ADF**

#### NEW QUESTION 46

A company is implementing an Amazon ECS cluster to run its workload. The company architecture will run multiple ECS services on the cluster, with an Application Load Balancer on the front end, using multiple target groups to route traffic. The Application Development team has been struggling to collect logs that must be collected and sent to an Amazon S3 bucket for near-real time analysis

What must the DevOps Engineer configure in the deployment to meet these requirements? (Select THREE)

- A. Install the Amazon CloudWatch Logs logging agent on the ECS instance
- B. Change the logging driver in the ECS task definition to 'awslogs'.
- C. Download the Amazon CloudWatch Logs container instance from AWS and configure it as a task. Update the application service definitions to include the logging task.
- D. Use Amazon CloudWatch Events to schedule an AWS Lambda function that will run every 60 seconds running the create-export -task CloudWatch Logs command, then point the output to the logging S3 bucket.
- E. Enable access logging on the Application Load Balancer, then point it directly to the S3 logging bucket.
- F. Enable access logging on the target groups that are used by the ECS services, then point it directly to the S3 logging bucket.
- G. Create an Amazon Kinesis Data Firehose with a destination of the S3 logging bucket, then create an Amazon CloudWatch Logs subscription filter for Kinesis

**Answer: BDF**

#### NEW QUESTION 47

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. Build an application that will create a list of all AWS access keys in the account and search each key on Git repository hosting service
- B. If a match is found, configure the application to disable the associated access key
- C. Then deploy the application to an AWS Elastic Beanstalk worker environment and define a periodic task to invoke the application every hour.
- D. Use Amazon Inspector to detect when a key has been exposed online
- E. Have Amazon Inspector send a notification to an Amazon SNS topic when a key has been exposed

- F. 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.
- G. Configure AWS Trusted Advisor and create an Amazon CloudWatch Events rule that uses Trusted Advisor as the event source.
- H. Configure the CloudWatch Events rule to invoke an AWS Lambda function as the target.
- I. If the Lambda function finds the exposed access keys, then have it disable the access key so that it cannot be used.
- J. Create an AWS Config rule to detect when a key is exposed online.
- K. Have AWS Config send change notifications to an SNS topic.
- L. 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: C**

**Explanation:**

<https://github.com/aws/Trusted-Advisor-Tools/tree/master/ExposedAccessKeys/stepbystep>

**NEW QUESTION 50**

A DevOps Engineer has been asked to recommend a tool to deploy the components of a three-tier web application. This application will use Amazon DynamoDB as a database.

Which deployment requires the LEAST amount of operational management?

- A. Use AWS CloudFormation to create a Classic Load Balancer and an Auto Scaling group.
- B. Use AWS OpsWorks to create the application and database resources. Deploy application updates with OpsWorks using lifecycle events.
- C. Use AWS OpsWorks to create a Classic Load Balancer, an Auto Scaling group, application, and database resources. Deploy application updates using OpsWorks lifecycle events.
- D. Use AWS OpsWorks to create a Classic Load Balancer, Auto Scaling, and application resources. Use AWS CloudFormation to create the database resources. Deploy application updates using CloudFormation rolling updates.
- E. Use AWS CloudFormation to create a Classic Load Balancer, an Auto Scaling group, and database resources. Deploy application updates using CloudFormation rolling updates.

**Answer: B**

**NEW QUESTION 51**

An Engineering team manages a Node.js e-commerce application. The current environment consists of the following components: "☐ Amazon S3 buckets for storing content" "☐ Amazon EC2 for the front-end web servers" "☐ AWS Lambda for executing image processing" "☐ Amazon DynamoDB for storing session-related data."

The team expects a significant increase in traffic to the site. The application should handle the additional load without interruption. The team ran initial tests by adding new servers to the EC2 front-end to handle the larger load, but the instances took up to 20 minutes to become fully configured. The team wants to reduce this configuration time.

What changes will the Engineering team need to implement to make the solution the MOST resilient and highly available while meeting the expected increase in demand?

- A. Use AWS OpsWorks to automatically configure each new EC2 instance as it is launched.
- B. Configure the EC2 instances by using an Auto Scaling group behind an Application Load Balancer across multiple Availability Zones.
- C. Implement Amazon DynamoDB Auto Scaling.
- D. Use Amazon Route 53 to point the application DNS record to the Application Load Balancer.
- E. Deploy a fleet of EC2 instances, doubling the current capacity, and place them behind an Application Load Balancer.
- F. Increase the Amazon DynamoDB read and write capacity units.
- G. Add an alias record that contains the Application Load Balancer endpoint to the existing Amazon Route 53 DNS record that points to the application.
- H. Configure Amazon CloudFront and have its origin point to Amazon S3 to host the web application. Implement Amazon DynamoDB Auto Scaling.
- I. Use Amazon Route 53 to point the application DNS record to the CloudFront DNS name.
- J. Use AWS Elastic Beanstalk with a custom AMI including all web components.
- K. Deploy the platform by using an Auto Scaling group behind an Application Load Balancer across multiple Availability Zones.
- L. Implement Amazon DynamoDB Auto Scaling.
- M. Use Amazon Route 53 to point the application DNS record to the Elastic Beanstalk load balancer.

**Answer: D**

**NEW QUESTION 56**

A company has a website in an AWS Elastic Beanstalk load balancing and automatic scaling environment. This environment has an Amazon RDS MySQL instance configured as its database resource. After a sudden increase in traffic, the website started dropping traffic. An administrator discovered that the application on some instances is not responding as the result of out-of-memory errors. Classic Load Balancer marked those instances as out of service, and the health status of Elastic Beanstalk enhanced health reporting is degraded. However, Elastic Beanstalk did not replace those instances. Because of the diminished capacity behind the Classic Load Balancer, the application response times are slower for the customers.

Which action will permanently fix this issue?

- A. Clone the Elastic Beanstalk environment.
- B. When the new environment is up, swap CNAME and terminate the earlier environment.
- C. Temporarily change the maximum number of instances in the Auto Scaling group to allow the group to support more traffic.
- D. Change the setting for the Auto Scaling group health check from Amazon EC2 to Elastic Load Balancing, and increase the capacity of the group.
- E. Write a cron script for restraining the web server process when memory is full, and deploy it with AWS Systems Manager.

**Answer: C**

**NEW QUESTION 60**

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. Using the UserData section of an AWS CloudFormation template, install tcpdump on every provisioned Amazon EC2 instance.

- B. The output of the tool is sent to Amazon EFS for aggregation and queryin
- C. 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.
- D. 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
- E. 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.
- F. Create a flow log for the production VP
- G. 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 VP
- H. If the VPC flow logs are not configured, return a "'NON\_COMPLIANT' status and notify the security organization.
- I. Configure a new trail using AWS CloudTrail service
- J. Using the UserData section of an AWS CloudFormation template, install tcpdump on every provisioned Amazon EC2 instance
- K. Connect Amazon Athena to the CloudTrail and write an AWS Lambda function that monitors for a flow log disabled event
- L. Once the CloudTrail entry has been spotted, alert the security organization

**Answer: C**

#### NEW QUESTION 63

A Development team is building more than 40 applications. Each app is a three-tiered web application based on an ELB Application Load Balancer, Amazon EC2, and Amazon RDS. Because the applications will be used internally, the Security team wants to allow access to the 40 applications only from the corporate network and block access from external IP addresses. The corporate network reaches the internet through proxy servers. The proxy servers have 12 proxy IP addresses that are being changed one or two times per month. The Network Infrastructure team manages the proxy servers; they upload the file that contains the latest proxy IP addresses into an Amazon S3 bucket. The DevOps Engineer must build a solution to ensure that the applications are accessible from the corporate network. Which solution achieves these requirements with MINIMAL impact to application development, MINIMAL operational effort, and the LOWEST infrastructure cost?

- A. Implement an AWS Lambda function to read the list of proxy IP addresses from the S3 object and to update the ELB security groups to allow HTTPS only from the given IP addresses
- B. Configure the S3 bucket to invoke the Lambda function when the object is updated
- C. Save the IP address list to the S3 bucket when they are changed.
- D. Ensure that all the applications are hosted in the same Virtual Private Cloud (VPC). Otherwise, consolidate the applications into a single VPC
- E. Establish an AWS Direct Connect connection with an active/standby configuration
- F. Change the ELB security groups to allow only inbound HTTPS connections from the corporate network IP addresses.
- G. Implement a Python script with the AWS SDK for Python (Boto), which downloads the S3 object that contains the proxy IP addresses, scans the ELB security groups, and updates them to allow only HTTPS inbound from the given IP addresses
- H. Launch an EC2 instance and store the script in the instance
- I. Use a cron job to execute the script daily.
- J. Enable ELB security groups to allow HTTPS inbound access from the Internet
- K. Use Amazon Cognito to integrate the company's Active Directory as the identity provider
- L. Change the 40 applications to integrate with Amazon Cognito so that only company employees can log into the applications
- M. Save the user access logs to Amazon CloudWatch Logs to record user access activities

**Answer: A**

#### NEW QUESTION 67

A healthcare provider has a hybrid architecture that includes 120 on-premises VMware servers running RedHat and 50 Amazon EC2 instances running Amazon Linux. The company is in the middle of an all-in migration to AWS and wants to implement a solution for collecting information from the on-premises virtual machines and the EC2 instances for data analysis. The information includes:

- Operating system type and version
- Data for installed applications
- Network configuration information, such as MAC and IP addresses
- Amazon EC2 instance AMI ID and IAM profile

How can these requirements be met with the LEAST amount of administration?

- A. Write a shell script to run as a cron job on EC2 instances to collect and push the data to Amazon S3. For on-premises resources, use VMware vSphere to collect the data and write it into a file gateway for storing the data in S3. Finally, use Amazon Athena on the S3 bucket for analytics.
- B. Use a script on the on-premises virtual machines as well as the EC2 instances to gather and push the data into Amazon S3, and then use Amazon Athena for analytics.
- C. Install AWS Systems Manager agents on both the on-premises virtual machines and the EC2 instances. Enable inventory collection and configure resource data sync to an Amazon S3 bucket to analyze the data with Amazon Athena.
- D. Use AWS Application Discovery Service for deploying Agentless Discovery Connector in the VMware environment and Discovery Agents on the EC2 instances for collecting the data
- E. Then use the AWS Migration Hub Dashboard for analytics.

**Answer: C**

#### NEW QUESTION 69

An AWS CodePipeline pipeline has implemented a code release process. The pipeline is integrated with AWS CodeDeploy to deploy versions of an application to multiple Amazon EC2 instances for each CodePipeline stage.

During a recent deployment, the pipeline failed due to a CodeDeploy issue. The DevOps team wants to improve monitoring and notifications during deployment to decrease resolution times.

What should the DevOps Engineer do to create notifications when issues are discovered?

- A. Implement AWS CloudWatch Logs for CodePipeline and CodeDeploy, create an AWS Config rule to evaluate code deployment issues, and create an Amazon SNS topic to notify stakeholders of deployment issues.
- B. Implement AWS CloudWatch Events for CodePipeline and CodeDeploy, create an AWS Lambda function to evaluate code deployment issues, and create an Amazon SNS topic to notify stakeholders of deployment issues.
- C. Implement AWS CloudTrail to record CodePipeline and CodeDeploy API call information, create an AWS Lambda function to evaluate code deployment issues, and create an Amazon SNS topic to notify stakeholders of deployment issues.
- D. Implement AWS CloudWatch Events for CodePipeline and CodeDeploy, create an Amazon Inspector assessment target to evaluate code deployment issues,

and create an Amazon SNS topic to notify stakeholders of deployment issues.

**Answer:** A

#### NEW QUESTION 72

A DevOps Engineer is deploying an Amazon API Gateway API with an AWS Lambda function providing the backend functionality. The Engineer needs to record the source IP address and response status of every API call.

Which combination of actions should the DevOps Engineer take to implement this functionality? (Choose three.)

- A. Configure AWS X-Ray to enable access logging for the API Gateway requests.
- B. Configure the API Gateway stage to enable access logging and choose a logging format.
- C. Create a new Amazon CloudWatch Logs log group or choose an existing log group to store the logs.
- D. Grant API Gateway permission to read and write logs to Amazon CloudWatch through an IAM role.
- E. Create a new Amazon S3 bucket or choose an existing S3 bucket to store the logs.
- F. Configure API Gateway to stream its log data to Amazon Kinesis.

**Answer:** BDE

#### NEW QUESTION 77

A DevOps Engineer is designing a deployment strategy for a web application. The application will use an Auto Scaling group to launch Amazon EC2 instances using an AMI. The same infrastructure will be deployed in multiple environments (development, test, and quality assurance). The deployment strategy should meet the following requirements: "Minimize the startup time for the instance" "Allow the same AMI to work in multiple environments" "Store secrets for multiple environments securely"

How should this be accomplished?

- A. Preconfigure the AMI using an AWS Lambda function that launches an Amazon EC2 instance, and then runs a script to install the software and create the AMI
- B. Configure an Auto Scaling lifecycle hook to determine which environment the instance is launched in, and, based on that finding, run a configuration script
- C. Save the secrets on an .ini file and store them in Amazon S3. Retrieve the secrets using a configuration script in EC2 user data.
- D. Preconfigure the AMI by installing all the software using AWS Systems Manager automation and configure Auto Scaling to tag the instances at launch with their specific environment
- E. Then use a bootstrap script in user data to read the tags and configure settings for the environment
- F. Use the AWS Systems Manager Parameter Store to store the secrets using AWS KMS.
- G. Use a standard AMI from the AWS Marketplace
- H. Configure Auto Scaling to detect the current environment
- I. Install the software using a script in Amazon EC2 user data
- J. Use AWS Secrets Manager to store the credentials for all environments.
- K. Preconfigure the AMI by installing all the software and configuration for all environments
- L. Configure Auto Scaling to tag the instances at launch with their environment
- M. Use the Amazon EC2 user data to trigger an AWS Lambda function that reads the instance ID and then reconfigures the setting for the proper environment
- N. Use the AWS Systems Manager Parameter Store to store the secrets using AWS KMS.

**Answer:** A

#### NEW QUESTION 81

A company is using Docker containers for an application deployment and wants to move its application to AWS. The company currently manages its own clusters on premises to manage the deployment of these containers. It wants to deploy its application to a managed service in AWS and wants the entire flow of the deployment process to be automated. In addition, the company has the following requirements:

Focus first on the development workload. The environment must be easy to manage.

Deployment should be repeatable and reusable for new environments. Store the code in a GitHub repository.

Which solution will meet these requirements?

- A. Set up an Amazon ECS environment
- B. Use AWS CodePipeline to create a pipeline that is triggered on a commit to the GitHub repository
- C. Use AWS CodeBuild to create the container images and AWS CodeDeploy to publish the container image to the ECS environment.
- D. Use AWS CodePipeline that triggers on a commit from the GitHub repository, build the container images with AWS CodeBuild, and publish the container images to Amazon EC2
- E. In the final stage, use AWS CloudFormation to create an Amazon ECS environment that gets the container images from the ECR repository.
- F. Create a Kubernetes Cluster on Amazon EC2. Use AWS CodePipeline to create a pipeline that is triggered when the code is committed to the repository
- G. Create the container images with a Jenkins server on EC2 and store them in the Docker Hub
- H. Use AWS Lambda from the pipeline to trigger the deployment to the Kubernetes Cluster.
- I. Set up an Amazon ECS environment
- J. Use AWS CodePipeline to create a pipeline that is triggered on a commit to the GitHub repository
- K. Use AWS CodeBuild to create the container and store it in the Docker Hub
- L. Use an AWS Lambda function to trigger a deployment and pull the new container image from the Docker Hub.

**Answer:** A

#### NEW QUESTION 86

A DevOps team needs to query information in application logs that are generated by an application running multiple Amazon EC2 instances deployed with AWS Elastic Beanstalk.

Instance log streaming to Amazon CloudWatch Logs was enabled on Elastic Beanstalk. Which approach would be the MOST cost-efficient?

- A. Use a CloudWatch Logs subscription to trigger an AWS Lambda function to send the log data to an Amazon Kinesis Data Firehose stream that has an Amazon S3 bucket destination
- B. Use Amazon Athena to query the log data from the bucket.
- C. Use a CloudWatch Logs subscription to trigger an AWS Lambda function to send the log data to an Amazon Kinesis Data Firehose stream that has an Amazon S3 bucket destination
- D. Use a new Amazon Redshift cluster and Amazon Redshift Spectrum to query the log data from the bucket.
- E. Use a CloudWatch Logs subscription to send the log data to an Amazon Kinesis Data Firehose stream that has an Amazon S3 bucket destination

- F. Use Amazon Athena to query the log data from the bucket.
- G. Use a CloudWatch Logs subscription to send the log data to an Amazon Kinesis Data Firehose stream that has an Amazon S3 bucket destination.
- H. Use a new Amazon Redshift cluster and Amazon Redshift Spectrum to query the log data from the bucket.

**Answer:** C

**Explanation:**

<https://docs.aws.amazon.com/AmazonCloudWatch/latest/logs/Subscriptions.html>

**NEW QUESTION 91**

After a recent audit, a company decided to implement a new disaster recovery strategy for its Amazon S3 data and its MySQL database running on Amazon EC2. Management wants the ability to recover to a secondary AWS Region with an RPO under 5 seconds and a RTO under 1 minute. Which actions will meet the requirements while MINIMIZING operational overhead? (Select TWO.)

- A. Modify the application to write to both Regions at the same time when uploading objects to Amazon S3
- B. Migrate the database to an Amazon Aurora multi-master in the primary and secondary Regions.
- C. Migrate the database to Amazon RDS with a read replica in the secondary Region
- D. Migrate to Amazon Aurora Global Database.
- E. Set up S3 cross-Region replication with a replication SLA for the S3 buckets where objects are being put.

**Answer:** AE

**NEW QUESTION 95**

A company has a legacy application running on AWS. The application can only run on one Amazon EC2 instance at a time. Application metadata is stored in Amazon S3 and must be retrieved if the instance is restarted. The instance should be automatically restarted or relaunched if performance degrades. Which solution will satisfy those requirements?

- A. Create an Amazon Cloud Watch alarm to monitor the EC2 instance
- B. When the StatusCheckFailed system alarm is triggered use the recover action to stop and start the instance
- C. Use a trigger in Amazon S3 to push the metadata to the instance when it is back up and running
- D. Use the auto healing feature in AWS OpsWorks to stop and start the EC2 instance
- E. Use a lifecycle event in OpsWorks to pull the data from Amazon S3 and update it on the instance.
- F. Use the Auto Recovery feature in Amazon EC2 to automatically stop and start the EC2 instance in case of a failure. Use a trigger in Amazon S3 to push the metadata to the instance when it is back up and running
- G. Use AWS CloudFormation to create an EC2 instance that includes the user-data property for the EC2 resource
- H. Add a command in user data to retrieve the application metadata from Amazon S3.

**Answer:** C

**NEW QUESTION 100**

A company's web application will be migrated to AWS. The application is designed so that there is no server-side code required. As part of the migration, the company would like to improve the security of the application by adding HTTP response headers, following the Open Web Application Security Project (OWASP) secure headers recommendations. How can this solution be implemented to meet the security requirements using best practices?

- A. Use an Amazon S3 bucket configured for website hosting, then set up server access logging on the S3 bucket to track user activity
- B. Then configure the static website hosting and execute a scheduled AWS Lambda function to verify, and if missing, add security headers to the metadata.
- C. Use an Amazon S3 bucket configured for website hosting, then set up server access logging on the S3 bucket to track user activity
- D. Configure the static website hosting to return the required security headers.
- E. Use an Amazon S3 bucket configured for website hosting
- F. Create an Amazon CloudFront distribution that refers to this S3 bucket, with the origin response event set to trigger a Lambda@Edge Node.js function to add in the security headers.
- G. Set an Amazon S3 bucket configured for website hosting
- H. Create an Amazon CloudFront distribution that refers to this S3 bucket
- I. Set "Cache Based on Selected Request Headers" to "Whitelist," and add the security headers into the whitelist.

**Answer:** C

**Explanation:**

<https://aws.amazon.com/blogs/networking-and-content-delivery/adding-http-security-headers-using-lambdaedge>

**NEW QUESTION 101**

An ecommerce company uses a large number of Amazon EBS backed Amazon EC2 instances. To decrease manual work across all the instances, a DevOps engineer is tasked with automating restart actions when EC2 instance retirement events are scheduled. How can this be accomplished?

- A. Create a scheduled Amazon CloudWatch Events rule to execute an AWS Systems Manager automation document that checks if any EC2 instances are scheduled for retirement once a week
- B. If the instance is scheduled for retirement, the automation document will hibernate the instance.
- C. Enable EC2 Auto Recovery on all of the instances
- D. Create an AWS Config rule to limit the recovery to occur during a maintenance window only.
- E. Reboot all EC2 instances during an approved maintenance window that is outside of standard business hours
- F. Set up Amazon CloudWatch alarms to send a notification in case any instance is failing EC2 instance status checks.
- G. Set up an AWS Health Amazon CloudWatch Events rule to execute AWS Systems Manager automation documents that stop and start the EC2 instance when a retirement scheduled event occurs.

**Answer:** D

**NEW QUESTION 103**

A Developer is maintaining a fleet of 50 Amazon EC2 Linux servers. The servers are part of an Amazon EC2 Auto Scaling group, and also use Elastic Load Balancing for load balancing. Occasionally, some application servers are being terminated after failing ELB HTTP health checks. The Developer would like to perform a root cause analysis on the issue, but before being able to access application logs, the server is terminated. How can log collection be automated?

- A. Use Auto Scaling lifecycle hooks to put instances in a Pending:Wait stat
- B. Create an Amazon CloudWatch Alarm for EC2 Instance Terminate and trigger an AWS Lambda function that executes an SSM Run Command script to collect logs, push them to Amazon S3, and complete the Successful lifecycle action once logs are collected.
- C. Use Auto Scaling lifecycle hooks to put instances in a Terminating:Wait stat
- D. Create a Config rule for EC2 Instance-terminate Lifecycle and trigger a step function that executes a script to collect logs, push them to Amazon S3, and complete the lifecycle action once logs are collected
- E. Action
- F. Use Auto Scaling lifecycle hooks to put instances in a Terminating:Wait stat
- G. Create an Amazon CloudWatch subscription filter for EC2 Instance and trigger a CloudWatch agent that executes a script to collect logs, push them to Amazon S3, and complete the lifecycle action Terminate Successful once logs are collected.
- H. Use Auto Scaling lifecycle hooks to put instances in a Terminating:Wait stat
- I. Create an Amazon CloudWatch Events rule for EC2 Instance- and trigger an AWS Lambda function that executes a SSM Run Command script to collect logs, push them to Amazon S3, terminate Lifecycle Action and complete the lifecycle action once logs are collected.

**Answer:** D

**Explanation:**

<https://docs.aws.amazon.com/autoscaling/ec2/userguide/lifecycle-hooks.html>

**NEW QUESTION 104**

A company has multiple development teams sharing one AWS account. The development team's manager wants to be able to automatically stop Amazon EC2 instances and receive notifications if resources are idle and not tagged as production resources. Which solution will meet these requirements?

- A. Use a scheduled Amazon CloudWatch Events rule to filter for Amazon EC2 instance status checks and identify idle EC2 instances
- B. Use the CloudWatch Events rule to target an AWS Lambda function to stop non-production instances and send notifications.
- C. Use a scheduled Amazon CloudWatch Events rule to filter AWS Systems Manager events and identify idle EC2 instances and resource
- D. Use the CloudWatch Events rule to target an AWS Lambda function to stop non-production instances and send notifications.
- E. Use a scheduled Amazon CloudWatch Events rule to target a custom AWS Lambda function that runs AWS Trusted Advisor checks. Create a second CloudWatch Events rule to filter events from Trusted Advisor to trigger a Lambda function to stop idle non-production instances and send notifications
- F. Use a scheduled Amazon CloudWatch Events rule to target Amazon Inspector events for idle EC2 instances. Use the CloudWatch Events rule to target the AWS Lambda function to stop non-production instances and send notifications

**Answer:** A

**NEW QUESTION 108**

A retail company wants to use AWS Elastic Beanstalk to host its online sales website running on Java. Since this will be the production website, the CTO has the following requirements for the deployment strategy:

- \*Zero downtime. While the deployment is ongoing, the current Amazon EC2 instances in service should remain in service. No deployment or any other action should be performed on the EC2 instances because they serve production traffic.
- \*A new fleet of instances should be provisioned for deploying the new application version.
- \*Once the new application version is deployed successfully in the new fleet of instances, the new instances should be placed in service and the old ones should be removed.
- \*The rollback should be as easy as possible. If the new fleet of instances fail to deploy the new application version, they should be terminated and the current instances should continue serving traffic as normal.
- \*The resources within the environment (EC2 Auto Scaling group, Elastic Load Balancing, Elastic Beanstalk DNS CNAME) should remain the same and no DNS change should be made.

Which deployment strategy will meet the requirements?

- A. Use rolling deployments with a fixed amount of one instance at a time and set the healthy threshold to OK.
- B. Use rolling deployments with additional batch with a fixed amount of one instance at a time and set the healthy threshold to OK.
- C. launch a new environment and deploy the new application version there, then perform a CNAME swap between environments.
- D. Use immutable environment updates to meet all the necessary requirements.

**Answer:** D

**Explanation:**

<https://aws.amazon.com/about-aws/whats-new/2016/04/aws-elastic-beanstalk-adds-two-new-deployment-policies/>  
<https://docs.aws.amazon.com/elasticbeanstalk/latest/dg/environmentmgmt-updates-immutable.html>  
<https://docs.aws.amazon.com/elasticbeanstalk/latest/dg/command-options-general.html#command-options-general>

**NEW QUESTION 110**

A company is using AWS CodeBuild, AWS CodeDeploy, and AWS CodePipeline to deploy applications automatically to an Amazon EC2 instance. A DevOps Engineer needs to perform a security assessment scan of the operating system on every application deployment to the environment. How should this be automated?

- A. Use Amazon CloudWatch Events to monitor for Auto Scaling event notifications of new instances and configure CloudWatch Events to trigger an Amazon Inspector scan.
- B. Use Amazon CloudWatch Events to monitor for AWS CodeDeploy notifications of a successful code deployment and configure CloudWatch Events to trigger an Amazon Inspector scan.
- C. Use Amazon CloudWatch Events to monitor for CodePipeline notifications of a successful code deployment and configure CloudWatch Events to trigger an AWS X-Ray scan.
- D. Use Amazon Inspector as a CodePipeline task after the successful use of CodeDeploy to deploy the code to the systems.

Answer: A

#### NEW QUESTION 115

An application is deployed on Amazon EC2 instances running in an Auto Scaling group. During the bootstrapping process, the instances register their private IP addresses with a monitoring system. The monitoring system performs health checks frequently by sending ping requests to those IP addresses and sending alerts if an instance becomes non-responsive.

The existing deployment strategy replaces the current EC2 instances with new ones. A DevOps engineer has noticed that the monitoring system is sending false alarms during a deployment, and is tasked with stopping these false alarms.

Which solution will meet these requirements without affecting the current deployment method?

- A. Define an Amazon CloudWatch Events target, an AWS Lambda function, and a lifecycle hook attached to the Auto Scaling group
- B. Configure CloudWatch Events to invoke Amazon SNS to send a message to the systems administrator group for remediation.
- C. Define an AWS Lambda function and a lifecycle hook attached to the Auto Scaling group
- D. Configure the lifecycle hook to invoke the Lambda function, which removes the entry of the private IP from the monitoring system upon instance termination.
- E. Define an Amazon CloudWatch Events target, an AWS Lambda function, and a lifecycle hook attached to the Auto Scaling group
- F. Configure CloudWatch Events to invoke the Lambda function, which removes the entry of the private IP from the monitoring system upon instance termination.
- G. Define an AWS Lambda function that will run a script when instance termination occurs in an Auto Scaling group
- H. The script will remove the entry of the private IP from the monitoring system.

Answer: C

#### NEW QUESTION 116

A company has multiple development groups working in a single shared AWS account. The Senior Manager of the groups wants to be alerted via a third-party API call when the creation of resources approaches the service limits for the account.

Which solution will accomplish this with the LEAST amount of development effort?

- A. Create an Amazon CloudWatch Event rule that runs periodically and targets an AWS Lambda function. Within the Lambda function, evaluate the current state of the AWS environment and compare deployed resource values to resource limits on the account
- B. Notify the Senior Manager if the account is approaching a service limit.
- C. Deploy an AWS Lambda function that refreshes AWS Trusted Advisor checks, and configure an Amazon CloudWatch Events rule to run the Lambda function periodically
- D. Create another CloudWatch Events rule with an event pattern matching Trusted Advisor events and a target Lambda function
- E. In the target Lambda function, notify the Senior Manager.
- F. Deploy an AWS Lambda function that refreshes AWS Personal Health Dashboard checks, and configure an Amazon CloudWatch Events rule to run the Lambda function periodically
- G. Create another CloudWatch Events rule with an event pattern matching Personal Health Dashboard events and a target Lambda function
- H. In the target Lambda function, notify the Senior Manager.
- I. Add an AWS Config custom rule that runs periodically, checks the AWS service limit status, and streams notifications to an Amazon SNS topic
- J. Deploy an AWS Lambda function that notifies the Senior Manager, and subscribe the Lambda function to the SNS topic.

Answer: B

#### NEW QUESTION 118

A DevOps engineer has automated a web service deployment using AWS CodePipeline with the following steps:

- An AWS CodeBuild project compiles the deployment artifact and runs unit tests.
- An AWS CodeDeploy deployment group deploys the web service to Amazon EC2 instances in the staging environment.
- A CodeDeploy deployment group deploys the web service to EC2 instances in the production environment. The quality assurance (QA) team has asked for permission to inspect the build artifact before the deployment to the production environment occurs. The QA team wants to run an internal automated penetration testing tool (invoked using a REST API call) to run some manual tests.

Which combination of actions will fulfill this request? (Select TWO.)

- A. Insert a manual approval action between the test and deployment actions of the pipeline.
- B. Modify the buildspec.yml file for the compilation stage to require manual approval before completion.
- C. Update the CodeDeploy deployment group so it requires manual approval to proceed
- D. Update the pipeline to directly trigger the REST API for the automated penetration testing tool.
- E. Update the pipeline to invoke a Lambda function that triggers the REST API for the automated penetration testing tool.

Answer: BD

#### NEW QUESTION 122

A company gives its employees limited rights to AWS. DevOps engineers have the ability to assume an administrator role. For tracking purposes, the security team wants to receive a near-real-time notification when the administrator role is assumed.

How should this be accomplished?

- A. Configure AWS Config to publish logs to an Amazon S3 bucket
- B. Use Amazon Athena to query the logs and send a notification to the security team when the administrator role is assumed.
- C. Configure Amazon GuardDuty to monitor when the administrator role is assumed and send a notification to the security team.
- D. Create an Amazon EventBridge (Amazon CloudWatch Events) event rule using an AWS Management Console sign-in events event pattern that publishes a message to an Amazon SNS topic if the administrator role is assumed
- E. [^
- F. Create an Amazon EventBridge (Amazon CloudWatch Events) event rule using an AWS API call that uses an AWS CloudTrail event pattern to trigger an AWS Lambda function that publishes a message to an Amazon SNS topic if the administrator role is assumed.

Answer: D

#### NEW QUESTION 123

A company's legacy application uses IAM user credentials to access resources in the company's AWS Organizations organization. A DevOps engineer needs to ensure new IAM users cannot be created unless the employee creating the IAM user is on an exception list.

Which solution will meet these requirements?

- A. Attach an Organizations SCP with an explicit deny for all iam:CreateAccessKey actions with a condition that excludes StringNotEquals for aws:username with a value of the exception list.
- B. Attach an Organizations SCP with an explicit deny for all iam:CreateUser actions with a condition that includes StringEquals for aws:username with a value of the exception list.
- C. Create an Amazon EventBridge (Amazon CloudWatch Events) rule with a pattern that matches the iam:CreateAccessKey action with an AWS Lambda function target
- D. The function will check the user name account against an exception list
- E. If the user is not in the exception list, the function will delete the user.
- F. Create an Amazon EventBridge (Amazon CloudWatch Events) rule with a pattern that matches the iam:CreateUser action with an AWS Lambda function target
- G. The function will check the user name and account against an exception list
- H. If the user is not in the exception list, the function will delete the user.

**Answer: B**

#### NEW QUESTION 125

A media customer has several thousand Amazon EC2 instances in an AWS account. The customer is using a Slack channel for team communications and important updates. A DevOps Engineer was told to send all AWS-scheduled EC2 maintenance notifications to the company Slack channel. Which method should the Engineer use to implement this process in the LEAST amount of steps?

- A. Integrate AWS Trusted Advisor with AWS Config
- B. Based on the AWS Config rules created, the AWS Config event can invoke an AWS Lambda function to send notifications to the Slack channel.
- C. Integrate AWS Personal Health Dashboard with Amazon CloudWatch Event
- D. Based on the CloudWatch Events created, the event can invoke an AWS Lambda function to send notifications to the Slack channel.
- E. Integrate EC2 events with Amazon CloudWatch monitor
- F. Based on the CloudWatch Alarm created, the alarm can invoke an AWS Lambda function to send EC2 maintenance notifications to the Slack channel.
- G. Integrate AWS Support with AWS CloudTrail
- H. Based on the CloudTrail lookup event created, the event can invoke an AWS Lambda function to pass EC2 maintenance notifications to the Slack channel.

**Answer: B**

#### Explanation:

<https://docs.aws.amazon.com/health/latest/ug/cloudwatch-events-health.html>

#### NEW QUESTION 126

A company wants to adopt a methodology for handling security threats from leaked and compromised IAM access keys. The DevOps Engineer has been asked to automate the process of acting upon compromised access keys, which includes identifying users, revoking their permissions, and sending a notification to the Security team.

Which of the following would achieve this goal?

- A. Use the AWS Trusted Advisor generated security report for access key
- B. Use Amazon EMR to run analytics on the report
- C. Identify compromised IAM access keys and delete the
- D. Use Amazon CloudWatch with an EMR Cluster State Change event to notify the Security team.
- E. Use AWS Trusted Advisor to identify compromised access key
- F. Create an Amazon CloudWatch Events rule with Trusted Advisor as the event source, and AWS Lambda and Amazon SNS as target
- G. Use AWS Lambda to delete compromised IAM access keys and Amazon SNS to notify the Security team.
- H. Use the AWS Trusted Advisor generated security report for access key
- I. Use AWS Lambda to scan through the report
- J. Use scan result inside AWS Lambda and delete compromised IAM access key
- K. Use Amazon SNS to notify the Security team.
- L. Use AWS Lambda with a third-party library to scan for compromised access key
- M. Use scan result inside AWS Lambda and delete compromised IAM access key
- N. Create Amazon CloudWatch custom metrics for compromised key
- O. Create a CloudWatch alarm on the metrics to notify the Security team.

**Answer: B**

#### NEW QUESTION 127

A company runs a three-tier web application in its production environment, which is built on a single AWS CloudFormation template made up of Amazon EC2 instances behind an ELB Application Load Balancer. The instances run in an EC2 Auto Scaling group across multiple Availability Zones. Data is stored in an Amazon RDS Multi-AZ DB instance with read replicas. Amazon Route 53 manages the application's public DNS record.

A DevOps Engineer must create a workflow to mitigate a failed software deployment by rolling back changes in the production environment when a software cutover occurs for new application software.

What steps should the Engineer perform to meet these requirements with the LEAST amount of downtime?

- A. Use CloudFormation to deploy an additional staging environment and configure the Route 53 DNS with weighted record
- B. During cutover, change the Route 53 A record weights to achieve an even traffic distribution between the two environments
- C. Validate the traffic in the new environment and immediately terminate the old environment if tests are successful.
- D. Use a single AWS Elastic Beanstalk environment to deploy the staging and production environments. Update the environment by uploading the ZIP file with the new application code
- E. Swap the Elastic Beanstalk environment CNAME
- F. Validate the traffic in the new environment and immediately terminate the old environment if tests are successful.
- G. Use a single AWS Elastic Beanstalk environment and an AWS OpsWorks environment to deploy the staging and production environment
- H. Update the environment by uploading the ZIP file with the new application code into the Elastic Beanstalk environment deployed with the OpsWorks stack
- I. Validate the traffic in the new environment and immediately terminate the old environment if tests are successful.
- J. Use AWS CloudFormation to deploy an additional staging environment, and configure the Route 53 DNS with weighted record
- K. During cutover, increase the weight distribution to have more traffic directed to the new staging environment as workloads are successfully validated
- L. Keep the old production environment in place until the new staging environment handles all traffic.

Answer: D

#### NEW QUESTION 130

A DevOps Engineer needs to design and implement a backup mechanism for Amazon EFS. The Engineer is given the following requirements:

- \*The backup should run on schedule.
  - \*The backup should be stopped if the backup window expires.
  - \*The backup should be stopped if the backup completes before the backup window.
  - \*The backup logs should be retained for further analysis.
- The design should support highly available and fault-tolerant paradigms.
- \*Administrators should be notified with backup metadata. Which design will meet these requirements?

- A. Use AWS Lambda with an Amazon CloudWatch Events rule for scheduling the start/stop of backup activit
- B. Run backup scripts on Amazon EC2 in an Auto Scaling grou
- C. Use Auto Scaling lifecycle hooks and the SSM Run Command on EC2 for uploading backup logs to Amazon S3. Use Amazon SNS to notify administrators with backup activity metadata.
- D. Use Amazon SWF with an Amazon CloudWatch Events rule for scheduling the start/stop of backup activit
- E. Run backup scripts on Amazon EC2 in an Auto Scaling grou
- F. Use Auto Scaling lifecycle hooks and the SSM Run Command on EC2 for uploading backup logs to Amazon Redshif
- G. Use CloudWatch Alarms to notify administrators with backup activity metadata.
- H. Use AWS Data Pipeline with an Amazon CloudWatch Events rule for scheduling the start/stop of backup activit
- I. Run backup scripts on Amazon EC2 in a single Availability Zon
- J. Use Auto Scaling lifecycle hooks and the SSM Run Command on EC2 for uploading the backup logs to Amazon RD
- K. Use Amazon SNS to notify administrators with backup activity metadata.
- L. Use AWS CodePipeline with an Amazon CloudWatch Events rule for scheduling the start/stop of backup activit
- M. Run backup scripts on Amazon EC2 in a single Availability Zon
- N. Use Auto Scaling lifecycle hooks and the SSM Run Command on Amazon EC2 for uploading backup logs to Amazon S3. Use Amazon SES to notify admins with backup activity metadata.

Answer: A

#### Explanation:

<https://docs.aws.amazon.com/efs/latest/ug/alternative-efs-backup.html>

#### NEW QUESTION 133

A company is hosting a web application in an AWS Region. For disaster recovery purposes, a second region is being used as a standby. Disaster recovery requirements state that session data must be replicated between regions in near-real time and 1% of requests should route to the secondary region to continuously verify system functionality. Additionally, if there is a disruption in service in the main region, traffic should be automatically routed to the secondary region, and the secondary region must be able to scale up to handle all traffic.

How should a DevOps Engineer meet these requirements?

- A. In both regions, deploy the application on AWS Elastic Beanstalk and use Amazon DynamoDB global tables for session dat
- B. Use an Amazon Route 53 weighted routing policy with health checks to distribute the traffic across the regions.
- C. In both regions, launch the application in Auto Scaling groups and use DynamoDB for session dat
- D. Use a Route 53 failover routing policy with health checks to distribute the traffic across the regions.
- E. In both regions, deploy the application in AWS Lambda, exposed by Amazon API Gateway, and use Amazon RDS PostgreSQL with cross-region replication for session dat
- F. Deploy the web application with client-side logic to call the API Gateway directly.
- G. In both regions, launch the application in Auto Scaling groups and use DynamoDB global tables for session dat
- H. Enable an Amazon CloudFront weighted distribution across region
- I. Point the Amazon Route 53 DNS record at the CloudFront distribution.

Answer: B

#### NEW QUESTION 135

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. Create an AWS CodeCommit repository for each project, use the master branch for production code, and create a testing branch for code deployed to testin
- B. Use feature branches to develop new features and pull requests to merge code to testing and master branches.
- C. Create another S3 bucket for each project for testing code, and use an AWS Lambda function to promote code changes between testing and production bucket
- D. Enable versioning on all buckets to prevent code conflicts.
- E. Create an AWS CodeCommit repository for each project, and use the master branch for production and test code with different deployment pipelines for each environmen
- F. Use feature branches to develop new features.
- G. Enable versioning and branching on each S3 bucket, use the master branch for production code, and create a testing branch for code deployed to testin
- H. Have Developers use each branch for developing in each environment.

Answer: A

#### NEW QUESTION 137

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. Use AWS CodePipeline to create a pipelin

- B. Add an action to the pipeline to invoke an AWS Lambda function after the source stag
- C. Have the Lambda function invoke the source-code analysis tool on premises against the source input from CodePipel
- D. The function then waits for the execution to complete and places the output in a specified Amazon S3 location.
- E. Use AWS CodePipeline to create a pipeline, then create a custom action typ
- F. Create a job worker for the custom action that runs on hardware hosted on premise
- G. The job worker handles running security checks with the on-premises code analysis tool and then returns the job results to CodePipeline
- H. Have the pipeline invoke the custom action after the source stage.
- I. Use AWS CodePipeline to create a pipelin
- J. 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 too
- K. 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.
- L. Use AWS CodePipeline to create a pipelin
- M. Create a shell script that copies the input source code to a location on premise
- N. Invoke the source code analysis tool and return the results to CodePipeline
- O. Invoke the shell script by adding a custom script action after the source stage.

**Answer: B**

#### NEW QUESTION 139

A DevOps Engineer manages a web application that runs on Amazon EC2 instances behind an Application Load Balancer (ALB). The instances run in an EC2 Auto Scaling group across multiple Availability Zones. The Engineer needs to implement a deployment strategy that:  
Launches a second fleet of instances with the same capacity as the original fleet. Maintains the original fleet unchanged while the second fleet is launched.  
Transitions traffic to the second fleet when the second fleet is fully deployed. Terminates the original fleet automatically 1 hour after transition.  
Which solution will satisfy these requirements?

- A. Use an AWS CloudFormation template with a retention policy for the ALB set to 1 hou
- B. Update the Amazon Route 53 record to reflect the new ALB.
- C. Use two AWS Elastic Beanstalk environments to perform a blue/green deployment from the original environment to the new on
- D. Create an application version lifecycle policy to terminate the original environment in 1 hour.
- E. Use AWS CodeDeploy with a deployment group configured with a blue/green deployment configuratio
- F. Select the option Terminate the original instances in the deployment group with a waiting period of 1 hour.
- G. Use AWS Elastic Beanstalk with the configuration set to Immutabl
- H. Create an .ebextension using the Resources key that sets the deletion policy of the ALB to 1 hour, and deploy the application.

**Answer: B**

#### NEW QUESTION 142

For auditing, analytics, and troubleshooting purposes, a DevOps Engineer for a data analytics application needs to collect all of the application and Linux system logs from the Amazon EC2 instances before termination. The company, on average, runs 10,000 instances in an Auto Scaling group. The company requires the ability to quickly find logs based on instance IDs and date ranges.  
Which is the MOST cost-effective solution?

- A. Create an EC2 Instance-terminate Lifecycle Action on the group, write a termination script for pushing logs into Amazon S3, and trigger an AWS Lambda function based on S3 PUT to create a catalog of log files in an Amazon DynamoDB table with the primary key being Instance ID and sort key being Instance Termination Date.
- B. Create an EC2 Instance-terminate Lifecycle Action on the group, write a termination script for pushing logs into Amazon CloudWatch Logs, create a CloudWatch Events rule to trigger an AWS Lambda function to create a catalog of log files in an Amazon DynamoDB table with the primary key being Instance ID and sort key being Instance Termination Date.
- C. Create an EC2 Instance-terminate Lifecycle Action on the group, create an Amazon CloudWatch Events rule based on it to trigger an AWS Lambda function for storing the logs in Amazon S3, and create a catalog of log files in an Amazon DynamoDB table with the primary key being Instance ID and sort key being Instance Termination Date.
- D. Create an EC2 Instance-terminate Lifecycle Action on the group, push the logs into Amazon Kinesis Data Firehouse, and select Amazon ES as the destination for providing storage and search capability.

**Answer: C**

#### Explanation:

<https://docs.aws.amazon.com/autoscaling/ec2/userguide/lifecycle-hooks.html>

#### NEW QUESTION 146

A web application for healthcare services runs on Amazon EC2 instances behind an ELB Application Load Balancer. The instances run in an Amazon EC2 Auto Scaling group across multiple Availability Zones. A DevOps Engineer must create a mechanism in which an EC2 instance can be taken out of production so its system logs can be analyzed for issues to quickly troubleshoot problems on the web tier.  
How can the Engineer accomplish this task while ensuring availability and minimizing downtime?

- A. Implement EC2 Auto Scaling groups cooldown period
- B. Use EC2 instance metadata to determine the instance state, and an AWS Lambda function to snapshot Amazon EBS volumes to preserve system logs.
- C. Implement Amazon CloudWatch Events rule
- D. Create an AWS Lambda function that can react to an instance termination to deploy the CloudWatch Logs agent to upload the system and access logs to Amazon S3 for analysis.
- E. Terminate the EC2 instances manuall
- F. The Auto Scaling service will upload all log information to CloudWatch Logs for analysis prior to instance termination.
- G. Implement EC2 Auto Scaling groups with lifecycle hook
- H. Create an AWS Lambda function that can modify an EC2 instance lifecycle hook into a standby state, extract logs from the instance through a remote script execution, and place them in an Amazon S3 bucket for analysis.

**Answer: D**

#### NEW QUESTION 150

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.

**Answer:** BCE

#### NEW QUESTION 154

A company requires an RPO of 2 hours and an RTO of 10 minutes for its data and application at all times. An application uses a MySQL database and Amazon EC2 web servers. The development team needs a strategy for failover and disaster recovery. Which combination of deployment strategies will meet these requirements? (Select TWO)

- A. Create an Amazon Aurora cluster in one Availability Zone across multiple Regions as the data store. Use Aurora's automatic recovery capabilities in the event of a discluster.
- B. Create an Amazon Aurora global database in two Regions as the data store. In the event of a failure, promote the secondary Region as the master for the application.
- C. Create an Amazon Aurora multi-master cluster across multiple Regions as the data store. Use an Amazon Network Load Balancer to balance the database traffic in different Regions.
- D. Set up the application in two Regions and use Amazon Route 53 failover-based routing that points to the Application Load Balancers in both Regions. Use health checks to determine the availability in a given Region.
- E. Use Auto Scaling groups in each Region to adjust capacity based on demand.
- F. Set up the application in two Regions and use a multi-Region Auto Scaling group behind Application Load Balancers to manage the capacity based on demand. In the event of a disaster, adjust the Auto Scaling group's desired instance count to increase baseline capacity in the failover Region.

**Answer:** BE

#### NEW QUESTION 157

A company is developing a web application's infrastructure using AWS CloudFormation. The database engineering team maintains the database resources in a CloudFormation template, and the software development team maintains the web application resources in a separate CloudFormation template. As the scope of the application grows, the software development team needs to use resources maintained by the database engineering team. However, both teams have their own review and lifecycle management processes that they want to keep. Both teams also require resource-level change-set reviews. The software development team would like to deploy changes to this template using their CI/CD pipeline. Which solution will meet these requirements?

- A. Create a stack export from the database CloudFormation template and import those references into the web application CloudFormation template.
- B. Create a CloudFormation nested stack to make cross-stack resource references and parameters available in both stacks.
- C. Create a CloudFormation stack set to make cross-stack resource references and parameters available in both stacks.
- D. Create input parameters in the web application CloudFormation template and pass resource names and IDs from the database stack.

**Answer:** A

#### NEW QUESTION 158

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.

**Answer:** BD

#### Explanation:

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>

#### NEW QUESTION 160

A large enterprise is deploying a web application on AWS. The application runs on Amazon EC2 instances behind an Application Load Balancer. The instances run in an Auto Scaling group across multiple Availability Zones. The application stores data in an Amazon RDS Oracle DB instance and Amazon DynamoDB. There are separate environments for development, testing, and production. What is the MOST secure and flexible way to obtain password credentials during deployment?

- A. Retrieve an access key from an AWS Systems Manager SecureString parameter to access AWS service.
- B. Retrieve the database credentials from a Systems Manager SecureString parameter.
- C. Launch the EC2 instances with an EC2 IAM role to access AWS service.
- D. Retrieve the database credentials from AWS Secrets Manager.
- E. Retrieve an access key from an AWS Systems Manager plaintext parameter to access AWS services. Retrieve the database credentials from a Systems

- Manager SecureString parameter.
- F. Launch the EC2 instances with an EC2 IAM role to access AWS service
- G. Store the database passwords in an encrypted config file with the application artifacts.

**Answer:** B

**Explanation:**

<https://www.1strategy.com/blog/2019/02/28/aws-parameter-store-vs-aws-secrets-manager/>

**NEW QUESTION 162**

A devops team uses AWS CloudFormation to build their infrastructure. The security team is concerned about sensitive parameters, such as passwords, being exposed.

Which combination of steps will enhance the security of AWS CloudFormation? (Select THREE.)

- A. Create a secure string with AWS KMS and choose a KMS encryption ke
- B. Reference the ARN of the secure string, and give AWS CloudFormation permission to the KMS key for decryption.
- C. Create secrets using the AWS Secrets Manager AWS::SecretsManager::Secret resource typ
- D. Reference the secret resource return attributes in resources that need a password, such as an Amazon RDS database.
- E. Store sensitive static data as secure strings in the AWS Systems Manager Parameter Stor
- F. Use dynamic references in the resources that need access to the data.
- G. Store sensitive static data in the AWS Systems Manager Parameter Store as string
- H. Reference the stored value using types of Systems Manager parameters.
- I. Use AWS KMS to encrypt the CloudFormation template.
- J. Use the CloudFormation NoEcho parameter property to mask the parameter value.

**Answer:** ABD

**NEW QUESTION 164**

A company is deploying a container-based application using AWS CodeBuild. The security team mandates that all containers are scanned for vulnerabilities prior to deployment using a password-protected endpoint. All sensitive information must be stored securely.

Which solution should be used to meet these requirements?

- A. Encrypt the password using AWS KM
- B. Store the encrypted password in the buildspec.yml file as an environment variable under the variables mappin
- C. Reference the environment variable to initiate scanning.
- D. Import the password into an AWS CloudHSM ke
- E. Reference the CloudHSM key in the buildpec.yml file as an environment variable under the variables mappin
- F. Reference the environment variable to initiate scanning.
- G. Store the password in the AWS Systems Manager Parameter Store as a secure strin
- H. Add the Parameter Store key to the buildspec.yml file as an environment variable under the parameter-store mappin
- I. Reference the environment variable to initiate scanning.
- J. Use the AWS Encryption SDK to encrypt the password and embed in the buildspec.yml file as a variable under the secrets mappin
- K. Attach a policy to CodeBuild to enable access to the required decryption key.

**Answer:** C

**NEW QUESTION 168**

A DevOps engineer is currently running a container-based workload on-premises The engineer wants to move the application to AWS, but needs to keep the on-premises solution active because not all APIs will move at the same time. The traffic between AWS and the on-premises network should be secure and encrypted at all times. Low management overload is also a requirement.

Which combination of actions will meet these criteria? (Select THREE.)

- A. Create a Network Load Balancer an
- B. for each service, create a listener that points to the correct set of containers either in AWS or on-premises.
- C. Create an Application Load Balancer and, for each service, create a listener that points to the correct set of containers either in AWS or on-premises.
- D. Host the AWS containers in Amazon ECS with an EC2 launch type.
- E. Host the AWS containers in Amazon ECS with a Fargate launch type
- F. Use Amazon API Gateway to front the workload, and create a VPC link so API Gateway can forward API calls to the on-premises network through a VPN connection.
- G. Use Amazon API Gateway to front the workload, and set up public endpoints for the on-premises APIs so API Gateway can access them.

**Answer:** BDF

**NEW QUESTION 172**

A DevOps Engineer is using AWS CodeDeploy across a fleet of Amazon EC2 instances in an EC2 Auto Scaling group. The associated CodeDeploy deployment group, which is integrated with EC2 Auto Scaling, is configured to perform in-place deployments with CodeDeployDefault.OneAtATime. During an ongoing new deployment, the Engineer discovers that, although the overall deployment finished successfully, two out of five instances have the previous application revision deployed. The other three instances have the newest application revision.

What is likely causing this issue?

- A. The two affected instances failed to fetch the new deployment.
- B. A failed AfterInstall lifecycle event hook caused the CodeDeploy agent to roll back to the previous version on the affected instances.
- C. The CodeDeploy agent was not installed in two affected instances.
- D. EC2 Auto Scaling launched two new instances while the new deployment had not yet finished, causing the previous version to be deployed on the affected instances.

**Answer:** D

### NEW QUESTION 176

A DevOps Engineer is working on a project that is hosted on Amazon Linux and has failed a security review. The DevOps Manager has been asked to review the company buildspec.yaml file for an AWS CodeBuild project and provide recommendations. The buildspec.yaml file is configured as follows:

```
env:
variables:
  AWS_ACCESS_KEY_ID: AKIAJF7BRFWJBA4GHXNA
  AWS_SECRET_ACCESS_KEY: ORjJns3At2mlh4O4tm0+zHxZqz7cNAvMLYRehcl
  AWS_DEFAULT_REGION: us-east-1
  DB_PASSWORD: cuj5RptFa3va
phases:
build:
commands:
  -aws s3 cp s3://db-deploy-bucket/my.cnf.template/tmp/my.cnf
  -sed-i '' s/DB_PW/S{DB_PASSWORD}/ /tmp/my.cnf
  -aws s3 cp s3://db-deploy-bucket/instance.key/tmp/instance.key
  -chmod 600/tmp/instance.key
  -scp-i /tmp/instance.key/tmp/my.cnf root@10.25.15.23 :/etc/my.cnf
  -ssh- i /tmp/instance.key root@10.25.15.23 /etc/init.d/mysqld restart
```

What changes should be recommended to comply with AWS security best practices? (Select THREE.)

- A. Add a post-build command to remove the temporary files from the container before termination to ensure they cannot be seen by other CodeBuild users.
- B. Update the CodeBuild project role with the necessary permissions and then remove the AWS credentials from the environment variable.
- C. Store the DB\_PASSWORD as a SecureString value in AWS Systems Manager Parameter Store and then remove the DB\_PASSWORD from the environment variables.
- D. Move the environment variables to the "db-deploy-bucket" Amazon S3 bucket, add a prebuild stage to download, then export the variables.
- E. Use AWS Systems Manager run command versus scp and ssh commands directly to the instance.
- F. Scramble the environment variables using XOR followed by Base64, add a section to install, and then run XOR and Base64 to the build phase.

**Answer:** BCE

#### Explanation:

<https://aws.amazon.com/codebuild/faqs/>

### NEW QUESTION 179

A Development team is working on a serverless application in AWS. To quickly identify and remediate potential production issues, the team decides to roll out changes to a small number of users as a test before the full release. The DevOps Engineer must develop a solution to minimize downtime and impact. Which of the following solutions should be used to meet the requirements? (Select TWO.)

- A. Create an Application Load Balancer with two target group
- B. Set up the Application Load Balancer for Amazon API Gateway private integration
- C. Associate one target group to the current version and the other target group to the new version
- D. Configure API Gateway to route 10% of incoming traffic to the new version
- E. As the new version becomes stable, configure API Gateway to send all traffic to the new version and detach the old version from the load balancer.
- F. Create an alias for an AWS Lambda function pointing to both the current and new version
- G. Configure the alias to route 10% of incoming traffic to the new version
- H. As the new version is considered stable, update the alias to route all traffic to the new version.
- I. Create a failover record set in AWS Route 53 pointing to the AWS Lambda endpoints for the old and new version
- J. Configure Route 53 to route 10% of incoming traffic to the new version
- K. As the new version becomes stable, update the DNS record to route all traffic to the new version.
- L. Create an ELB Network Load Balancer with two target group
- M. Set up the Network Load Balancer for Amazon API Gateway private integration Associate one target group with the current version and the other target group with the new version
- N. Configure the load balancer to route 10% of incoming traffic to the new version
- O. As the new version becomes stable, detach the old version from the load balancer.
- P. In Amazon API Gateway, create a canary release deployment by adding canary settings to the stage of a regular deployment
- Q. Configure API Gateway to route 10% of the incoming traffic to the canary release
- R. As the canary release is considered stable, promote it to a production release.

**Answer:** BE

#### Explanation:

<https://docs.aws.amazon.com/lambda/latest/dg/configuration-aliases.html> <https://docs.aws.amazon.com/apigateway/latest/developerguide/canary-release.html>  
<https://aws.amazon.com/blogs/compute/implementing-canary-deployments-of-aws-lambda-functions-with-alias-traffic-shifting/>

### NEW QUESTION 180

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? (Select 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. Execute an AWS Systems Manager Automation document to patch the systems every hour.
- E. Use Amazon CloudWatch Events scheduled events to schedule a patch window.
- F. Use AWS Systems Manager Maintenance Windows to schedule a patch window.

**Answer:** ABF

#### NEW QUESTION 181

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
- B. Use Amazon SNS to inform the team of the stage being triggered
- C. Next, add a test action using CodeBuild to do the required test
- D. At the end of the pipeline, add a deploy action to deploy the application to the next stage.
- E. Add a test action after the last deploy action of the pipeline
- F. Configure the action to use CodeBuild to perform the required test
- G. If these tests are successful, mark the action as successful
- H. 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.
- I. Create a new pipeline that uses a source action that gets the code from the same repository as the first pipeline
- J. Add a deploy action to deploy the code to a test environment
- K. Use a test action using AWS Lambda to test the deployment
- L. 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.
- M. Add a test action after the last deployment action
- N. Use a Jenkins server on Amazon EC2 to do the required tests and mark the action as successful if the tests pass
- O. 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:** B

#### Explanation:

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

#### NEW QUESTION 183

A DevOps Engineer uses Docker container technology to build an image-analysis application. The application often sees spikes in traffic. The Engineer must automatically scale the application in response to customer demand while maintaining cost effectiveness and minimizing any impact on availability. What will allow the FASTEST response to spikes in traffic while fulfilling the other requirements?

- A. Create an Amazon ECS cluster with the container instances in an Auto Scaling group
- B. Configure the ECS service to use Service Auto Scaling
- C. Set up Amazon CloudWatch alarms to scale the ECS service and cluster.
- D. Deploy containers on an AWS Elastic Beanstalk Multicontainer Docker environment
- E. Configure Elastic Beanstalk to automatically scale the environment based on Amazon CloudWatch metrics.
- F. Create an Amazon ECS cluster using Spot instances
- G. Configure the ECS service to use Service Auto Scaling
- H. Set up Amazon CloudWatch alarms to scale the ECS service and cluster.
- I. Deploy containers on Amazon EC2 instances
- J. Deploy a container scheduler to schedule containers onto EC2 instances
- K. Configure EC2 Auto Scaling for EC2 instances based on available Amazon CloudWatch metrics.

**Answer:** A

#### Explanation:

<https://aws.amazon.com/blogs/compute/automatic-scaling-with-amazon-ecs/>

#### NEW QUESTION 188

A DevOps engineer is tasked with migrating Docker containers used for a workload to AWS. The solution must allow for changes to be deployed into development and test environments automatically by updating each container and checking it into a container registry. Once the containers are pushed, they must be deployed automatically. Which solution will meet these requirements?

- A. Store container images in Amazon S3. Run the containers in AWS Elastic Beanstalk using a multicontainer Docker environment
- B. Configure Elastic Beanstalk to redeploy the containers if it detects a new version in Amazon S3.
- C. Store container images in AWS Artifact. Use AWS CodePipeline to trigger a deployment if a new container version is created
- D. Use AWS CodeDeploy to deploy new containers to Amazon EKS.
- E. Store container images in Amazon ECR. Use AWS CodePipeline to trigger a deployment if a new container version is created. Use AWS CodeDeploy to deploy the image to AWS Fargate.
- F. Store container images in Docker Hub. Install Docker on an Amazon EC2 instance and use AWS CodePipeline and AWS CodeDeploy to deploy any new containers

**Answer:** C

#### NEW QUESTION 193

A company is creating a software solution that executes a specific parallel-processing mechanism. The software can scale to tens of servers in some special scenarios. This solution uses a proprietary library that is license-based, requiring that each individual server have a single, dedicated license installed. The

company has 200 licenses and is planning to run 200 server nodes concurrently at most.

The company has requested the following features:

• A mechanism to automate the use of the licenses at scale. • Creation of a dashboard to use in the future to verify which licenses are available at any moment. What is the MOST effective way to accomplish these requirements?

- A. Upload the licenses to a private Amazon S3 bucket
- B. Create an AWS CloudFormation template with a Mappings section for the license
- C. In the template, create an Auto Scaling group to launch the server
- D. In the user data script, acquire an available license from the Mappings section
- E. Create an Auto Scaling lifecycle hook, then use it to update the mapping after the instance is terminated.
- F. Upload the licenses to an Amazon DynamoDB table
- G. Create an AWS CloudFormation template that uses an Auto Scaling group to launch the server
- H. In the user data script, acquire an available license from the DynamoDB table
- I. Create an Auto Scaling lifecycle hook, then use it to update the mapping after the instance is terminated.
- J. Upload the licenses to a private Amazon S3 bucket
- K. Populate an Amazon SQS queue with the list of licenses stored in S3. Create an AWS CloudFormation template that uses an Auto Scaling group to launch the server
- L. In the user data script acquire an available license from SQS
- M. Create an Auto Scaling lifecycle hook, then use it to put the license back in SQS after the instance is terminated.
- N. Upload the licenses to an Amazon DynamoDB table
- O. Create an AWS CLI script to launch the servers by using the parameter --count, with min:max instances to launch
- P. In the user data script, acquire an available license from the DynamoDB table
- Q. Monitor each instance and, in case of failure, replace the instance, then manually update the DynamoDB table.

**Answer: D**

#### NEW QUESTION 198

A company runs a database on a single Amazon EC2 instance in a development environment. The data is stored on separate Amazon EBS volumes that are attached to the EC2 instance. An Amazon Route 53 A record has been created and configured to point to the EC2 instance. The company would like to automate the recovery of the database instance when an instance or Availability Zone (AZ) fails. The company also wants to keep its costs low. The RTO is 4 hours and RPO is 12 hours.

Which solution should a DevOps Engineer implement to meet these requirements?

- A. Run the database in an Auto Scaling group with a minimum and maximum instance count of 1 in multiple AZ
- B. Add a lifecycle hook to the Auto Scaling group and define an Amazon CloudWatch Events rule that is triggered when a lifecycle event occurs
- C. Have the CloudWatch Events rule invoke an AWS Lambda function to detach or attach the Amazon EBS data volumes from the EC2 instance based on the event
- D. Configure the EC2 instance UserData to mount the data volumes (retry on failure with a short delay), then start the database and update the Route 53 record.
- E. Run the database on two separate EC2 instances in different AZs with one active and the other as a standby
- F. Attach the data volumes to the active instance
- G. Configure an Amazon CloudWatch Events rule to invoke an AWS Lambda function on EC2 instance termination
- H. The Lambda function launches a replacement EC2 instance
- I. If the terminated instance was the active node, then the function attaches the data volumes to the standby node
- J. Start the database and update the Route 53 record.
- K. Run the database in an Auto Scaling group with a minimum and maximum instance count of 1 in multiple AZ
- L. Create an AWS Lambda function that is triggered by a scheduled Amazon CloudWatch Events rule every 4 hours to take a snapshot of the data volume and apply it to a standby
- M. Have the instance UserData get the latest snapshot, create a new volume from it, and attach and mount the volume
- N. Then start the database and update the Route 53 record.
- O. Run the database on two separate EC2 instances in different AZ
- P. Configure one of the instances as a master and the other as a standby
- Q. Set up replication between the master and standby instance
- R. Point the Route 53 record to the master
- S. Configure an Amazon CloudWatch Events rule to invoke an AWS Lambda function upon the EC2 instance termination
- T. The Lambda function launches a replacement EC2 instance
- . If the terminated instance was the active node, the function promotes the standby to master and points the Route 53 record to it.

**Answer: D**

#### NEW QUESTION 203

A company needs to introduce automatic DNS failover for a distributed web application to a disaster recovery or standby installation. The DevOps Engineer plans to configure Amazon Route 53 to provide DNS routing to alternate endpoint in the event of an application failure.

What steps should the Engineer take to accomplish this? (Select TWO.)

- A. Create Amazon Route 53 health checks for each endpoint that cannot be entered as alias record
- B. Ensure firewall and routing rules allow Amazon Route 53 to send requests to the endpoints that are specified in the health checks.
- C. Create alias records that route traffic to AWS resources and set the value of the Evaluate Target Health option to Yes, then create all the non-alias records.
- D. Create a governing Amazon Route 53 record set, set it to failover, and associate it with the primary and secondary Amazon Route 53 record sets to distribute traffic to healthy DNS entries.
- E. Create an Amazon CloudWatch alarm to monitor the primary Amazon Route 53 DNS entries
- F. Then create an associated AWS Lambda function to execute the failover API call to Route 53 to the secondary DNS entry.

**Answer: AC**

#### NEW QUESTION 206

A company needs to introduce automatic DNS failover for a distributed web application to a disaster recovery or standby installation. The DevOps Engineer plans to configure Amazon Route 53 to provide DNS routing to alternate endpoint in the event of an application failure.

What steps should the Engineer take to accomplish this? (Select TWO.)

- A. Create Amazon Route 53 health checks for each endpoint that cannot be entered as alias record

- B. Ensure firewall and routing rules allow Amazon Route 53 to send requests to the endpoints that are specified in the health checks.
- C. Create alias records that route traffic to AWS resources and set the value of the Evaluate Target Health option to Yes, then create all the non-alias records.
- D. Create a governing Amazon Route 53 record set, set it to failover, and associate it with the primary and secondary Amazon Route 53 record sets to distribute traffic to healthy DNS entries.
- E. Create an Amazon CloudWatch alarm to monitor the primary Amazon Route 53 DNS entr
- F. Then create an associated AWS Lambda function to execute the failover API call to Route 53 to the secondary DNS entry.
- G. Map the primary and secondary Amazon Route 53 record sets to an Amazon CloudFront distribution using primary and secondary origins.

**Answer:** AC

#### NEW QUESTION 207

A DevOps Engineer discovered a sudden spike in a website's page load times and found that a recent deployment occurred. A brief diff of the related commit shows that the URL for an external API call was altered and the connecting port changed from 80 to 443. The external API has been verified and works outside the application. The application logs show that the connection is now timing out, resulting in multiple retries and eventual failure of the call. Which debug steps should the Engineer take to determine the root cause of the issue?

- A. Check the VPC Flow Logs looking for denials originating from Amazon EC2 instances that are part of the web Auto Scaling group
- B. Check the ingress security group rules and routing rules for the VPC.
- C. Check the existing egress security group rules and network ACLs for the VP
- D. Also check the application logs being written to Amazon CloudWatch Logs for debug information.
- E. Check the egress security group rules and network ACLs for the VP
- F. Also check the VPC flow logs looking for accepts originating from the web Auto Scaling group.
- G. Check the application logs being written to Amazon CloudWatch Logs for debug informatio
- H. Check the ingress security group rules and routing rules for the VPC.

**Answer:** C

#### NEW QUESTION 211

An IT team has built an AWS CloudFormation template so others in the company can quickly and reliably deploy and terminate an application. The template creates an Amazon EC2 instance with a user data script to install the application and an Amazon S3 bucket that the application uses to serve static webpages while it is running.

All resources should be removed when the CloudFormation stack is deleted. However, the team observes that CloudFormation reports an error during stack deletion, and the S3 bucket created by the stack is not deleted.

How can the team resolve the error in the MOST efficient manner to ensure that all resources are deleted without errors?

- A. Add Deletion Policy attribute to the S3 bucket resource, with the value Delete forcing the bucket to be removed when the stack is deleted.
- B. Add a custom resource when an AWS Lambda function with the DependsOn attribute specifying the S3 bucket, and an IAM rol
- C. Writhe the Lambda function to delete all objects from the bucket when the RequestType is Delete.
- D. Identify the resource that was not delete
- E. From the S3 console, empty the S3 bucket and then delete it.
- F. Replace the EC2 and S3 bucket resources with a single AWS OpsWorks Stacks resourc
- G. Define a custom recipe for the stack to create and delete the EC2 instance and the S3 bucket.

**Answer:** C

#### NEW QUESTION 213

.....

## THANKS FOR TRYING THE DEMO OF OUR PRODUCT

Visit Our Site to Purchase the Full Set of Actual AWS-Certified-DevOps-Engineer-Professional Exam Questions With Answers.

We Also Provide Practice Exam Software That Simulates Real Exam Environment And Has Many Self-Assessment Features. Order the AWS-Certified-DevOps-Engineer-Professional Product From:

<https://www.2passeasy.com/dumps/AWS-Certified-DevOps-Engineer-Professional/>

## Money Back Guarantee

### **AWS-Certified-DevOps-Engineer-Professional Practice Exam Features:**

- \* AWS-Certified-DevOps-Engineer-Professional Questions and Answers Updated Frequently
- \* AWS-Certified-DevOps-Engineer-Professional Practice Questions Verified by Expert Senior Certified Staff
- \* AWS-Certified-DevOps-Engineer-Professional Most Realistic Questions that Guarantee you a Pass on Your FirstTry
- \* AWS-Certified-DevOps-Engineer-Professional Practice Test Questions in Multiple Choice Formats and Updatesfor 1 Year