

## DOP-C01 Dumps

### AWS Certified DevOps Engineer- Professional

<https://www.certleader.com/DOP-C01-dumps.html>



**NEW QUESTION 1**

After conducting a disaster recovery exercise, an Enterprise Architect discovers that a large team of Database and Storage Administrators need more than seven hours of manual effort to make a flagship application's database functional in a different AWS Region. The Architect also discovers that the recovered database is often missing as much as two hours of data transactions.

Which solution provides improved RTO and RPO in a cross-region failover scenario?

- A. Deploy an Amazon RDS Multi-AZ instance backed by a multi-region Amazon EF
- B. Configure the RDS option group to enable multi-region availability for native automation of cross-region recovery and continuous data replicatio
- C. Create an Amazon SNS topic subscribed to RDS-impacted events to send emails to the Database Administration team when significant query Latency is detected in a single Availability Zone.
- D. Use Amazon SNS topics to receive published messages from Amazon RDS availability and backup event
- E. Use AWS Lambda for three separate functions with calls to Amazon RDS to snapshot a database instance, create a cross-region snapshot copy, and restore an instance from a snapsho
- F. Use a scheduled Amazon CloudWatch Events rule at a frequency matching the RPO to trigger the Lambda function to snapshot a database instanc
- G. Trigger the Lambda function to create a cross-region snapshot copy when the SNS topic for backup events receives a new messag
- H. Configure the Lambda function to restore an instance from a snapshot to trigger sending new messages published to the availability SNS topic.
- I. Create a scheduled Amazon CloudWatch Events rule to make a call to Amazon RDS to create a snapshot from a database instance and specify a frequency to match the RP
- J. Create an AWS Step Functions task to call Amazon RDS to perform a cross-region snapshot copy into the failover region, and configure the state machine to execute the task when the RDS snapshot create state is complet
- K. Create an SNS topic subscribed to RDS availability events, and push these messages to an Amazon SQS queue located in the failover regio
- L. Configure an Auto Scaling group of worker nodes to poll the queue for new messages and make a call to Amazon RDS to restore a database from a snapshot after a checksum on the cross-region copied snapshot returns valid.
- M. Use Amazon RDS scheduled instance lifecycle events to create a snapshot and specify a frequency to match the RP
- N. Use Amazon RDS scheduled instance lifecycle event configuration to perform a cross-region snapshot copy into the failover region upon SnapshotCreateComplete event
- O. Configure Amazon CloudWatch to alert when the CloudWatch RDS namespace CPUUtilization metric for the database instance falls to 0% and make a call to Amazon RDS to restore the database snapshot in the failover region.

**Answer: B**

**Explanation:**

<https://aws.amazon.com/blogs/database/cross-region-automatic-disaster-recovery-on-amazon-rds-for-oracle-data>

**NEW QUESTION 2**

A company has microservices running in AWS Lambda that read data from Amazon DynamoDB. The Lambda code is manually deployed by Developers after successful testing. The company now needs the tests and deployments be automated and run in the cloud. Additionally, traffic to the new versions of each microservice should be incrementally shifted over time after deployment.

What solution meets all the requirements, ensuring the MOST developer velocity?

- A. Create an AWS CodePipeline configuration and set up a post-commit hook to trigger the pipeline after tests have passe
- B. Use AWS CodeDeploy and create a Canary deployment configuration that specifies the percentage of traffic and interval.
- C. Create an AWS CodeBuild configuration that triggers when the test code is pushe
- D. Use AWS CloudFormation to trigger an AWS CodePipeline configuration that deploys the new Lambda versions and specifies the traffic shift percentage and interval.
- E. Create an AWS CodePipeline configuration and set up the source code step to trigger when code is pushe
- F. Set up the build step to use AWS CodeBuild to run the test
- G. Set up an AWS CodeDeploy configuration to deploy, then select the CodeDeployDefault.LambdaLinear10PercentEvery3Minutes option.
- H. Use the AWS CLI to set up a post-commit hook that uploads the code to an Amazon S3 bucket after tests have passe
- I. Set up an S3 event trigger that runs a Lambda function that deploys the new versio
- J. Use an interval in the Lambda function to deploy the code over time at the required percentage.

**Answer: C**

**Explanation:**

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

**NEW QUESTION 3**

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 4**

A web application has been deployed using an AWS Elastic Beanstalk application The Application Developers are concerned that they are seeing high latency in two different areas of the application: HTTP client requests to a third-party API MySQL client library queries to an Amazon RDS database A DevOps Engineer

must gather trace data to diagnose the issues. Which steps will gather the trace information with the LEAST amount of changes and performance impacts to the application?

- A. Add additional logging to the application code
- B. Use the Amazon CloudWatch agent to stream the application logs into Amazon Elasticsearch Service
- C. Query the log data in Amazon ES.
- D. Instrument the application to use the AWS X-Ray SDK
- E. Post trace data to an Amazon Elasticsearch Service cluster
- F. Query the trace data for calls to the HTTP client and the MySQL client.
- G. On the AWS Elastic Beanstalk management page for the application, enable the AWS X-Ray daemon. View the trace data in the X-Ray console.
- H. Instrument the application using the AWS X-Ray SDK
- I. On the AWS Elastic Beanstalk management page for the application, enable the X-Ray daemon
- J. View the trace data in the X-Ray console.

**Answer: C**

#### NEW QUESTION 5

Management has reported an increase in the monthly bill from Amazon Web Services, and they are extremely concerned with this increased cost. Management has asked you to determine the exact cause of this increase. After reviewing the billing report, you notice an increase in the data transfer cost. How can you provide management with a better insight into data transfer use?

- A. Update your Amazon CloudWatch metrics to use five-second granularity, which will give better detailed metrics that can be combined with your billing data to pinpoint anomalies.
- B. Use Amazon CloudWatch Logs to run a map-reduce on your logs to determine high usage and data transfer.
- C. Deliver custom metrics to Amazon CloudWatch per application that breaks down application data transfer into multiple, more specific data points.
- D. Using Amazon CloudWatch metrics, pull your Elastic Load Balancing outbound data transfer metrics monthly, and include them with your billing report to show which application is causing higher bandwidth usage

**Answer: C**

#### Explanation:

You can publish your own metrics to CloudWatch using the AWS CLI or an API. You can view statistical graphs of your published metrics with the AWS Management Console.

CloudWatch stores data about a metric as a series of data points. Each data point has an associated time stamp. You can even publish an aggregated set of data points called a statistic set.

If you have custom metrics specific to your application, you can give a breakdown to the management on the exact issue.

Option A won't be sufficient to provide better insights.

Option B is an overhead when you can make the application publish custom metrics. Option D is invalid because just the ELB metrics will not give the entire picture.

For more information on custom metrics, please refer to the below document link: from AWS

➤ <http://docs.aws.amazon.com/AmazonCloudWatch/latest/monitoring/publishingMetrics.html>

#### NEW QUESTION 6

A company recently migrated its legacy application from on-premises to AWS. The application is hosted on Amazon EC2 instances behind an Application Load Balancer, which is behind Amazon API Gateway. The company wants to ensure users experience minimal disruptions during any deployment of a new version of the application. The company also wants to ensure it can quickly roll back updates if there is an issue.

Which solution will meet these requirements with MINIMAL changes to the application?

- A. Introduce changes as a separate environment parallel to the existing one
- B. Configure API Gateway to use a canary release deployment to send a small subset of user traffic to the new environment.
- C. Introduce changes as a separate environment parallel to the existing one
- D. Update the application's DNS alias records to point to the new environment.
- E. Introduce changes as a separate target group behind the existing Application Load Balancer
- F. Configure API Gateway to route user traffic to the new target group in steps.
- G. Introduce changes as a separate target group behind the existing Application Load Balancer
- H. Configure API Gateway to route all traffic to the Application Load Balancer, which then sends the traffic to the new target group.

**Answer: A**

#### NEW QUESTION 7

A Developer is designing a continuous deployment workflow for a new Development team to facilitate the process for source code promotion in AWS. Developers would like to store and promote code for deployment from development to production while maintaining the ability to roll back that deployment if it fails.

Which design will incur the LEAST amount of downtime?

- A. Create one repository in AWS CodeCommit
- B. Create a development branch to hold merged changes
- C. Use AWS CodeBuild to build and test the code stored in the development branch triggered on a new commit
- D. Merge to the master and deploy to production by using AWS CodeDeploy for a blue/green deployment.
- E. Create one repository for each Developer in AWS CodeCommit and another repository to hold the production code
- F. Use AWS CodeBuild to merge development and production repositories, and deploy to production by using AWS CodeDeploy for a blue/green deployment.
- G. Create one repository for development code in AWS CodeCommit and another repository to hold the production code
- H. Use AWS CodeBuild to merge development and production repositories, and deploy to production by using AWS CodeDeploy for a blue/green deployment.
- I. Create a shared Amazon S3 bucket for the Development team to store their code
- J. Set up an Amazon CloudWatch Events rule to trigger an AWS Lambda function that deploys the code to production by using AWS CodeDeploy for a blue/green deployment.

**Answer: A**

#### NEW QUESTION 8

A company is beginning to move to the AWS Cloud. Internal customers are classified into two groups according to their AWS skills: beginners and experts. The DevOps Engineer needs to build a solution to allow beginners to deploy a restricted set of AWS architecture blueprints expressed as AWS CloudFormation templates. Deployment should only be possible on predetermined Virtual Private Clouds (VPCs). However, expert users should be able to deploy blueprints without constraints. Experts should also be able to access other AWS services, as needed.

How can the Engineer implement a solution to meet these requirements with the LEAST amount of overhead?

- A. Apply constraints to the parameters in the templates, limiting the VPCs available for deployment
- B. Store the templates on Amazon S3. Create an IAM group for beginners and give them access to the templates and CloudFormation
- C. Create a separate group for experts, giving them access to the templates, CloudFormation, and other AWS services.
- D. Store the templates on Amazon S3. Use AWS Service Catalog to create a portfolio of products based on those templates
- E. Apply template constraints to the products with rules limiting VPCs available for deployment
- F. Create an IAM group for beginners giving them access to the portfolio
- G. Create a separate group for experts giving them access to the templates, CloudFormation, and other AWS services.
- H. Store the templates on Amazon S3. Use AWS Service Catalog to create a portfolio of products based on those templates
- I. Create an IAM role restricting VPCs available for creation of AWS resources
- J. Apply a launch constraint to the products using this role
- K. Create an IAM group for beginners giving them access to the portfolio
- L. Create a separate group for experts giving them access to the portfolio and other AWS services.
- M. Create two templates for each architecture blueprint where only one of them limits the VPC available for deployment
- N. Store the templates in Amazon DynamoDB
- O. Create an IAM group for beginners giving them access to the constrained templates and CloudFormation
- P. Create a separate group for experts giving them access to the unconstrained templates, CloudFormation, and other AWS services.

**Answer: B**

#### NEW QUESTION 9

An ecommerce company is running an application on AWS. The company wants to create a standby disaster recovery solution in an additional Region that keeps the current application code. The application 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 database layer is hosted on an Amazon RDS MySQL Multi-AZ DB instance. Amazon Route 53 DNS records point to the ALB.

Which combination of actions will meet these requirements with the LOWEST cost? (Select THREE.)

- A. Configure a failover routing policy for the application DNS entry.
- B. Configure a geolocation routing policy for the application DNS entry.
- C. Create a cross-Region RDS read replica in the new standby Region.
- D. Migrate the database layer to Amazon DynamoDB and enable global replication to the new standby Region.
- E. Provision the ALB and Auto Scaling group in the new standby Region and set the desired capacity to match the active Region.
- F. Provision the ALB and Auto Scaling group in the new standby Region and set the desired capacity to 1.

**Answer: AEF**

#### NEW QUESTION 10

A company is adopting AWS CodeDeploy to automate its application deployments for a Java-Apache Tomcat application with an Apache webserver. The Development team started with a proof of concept, created a deployment group for a developer environment, and performed functional tests within the application. After completion, the team will create additional deployment groups for staging and production.

The current log level is configured within the Apache settings, but the team wants to change this configuration dynamically when the deployment occurs, so that they can set different log level configurations depending on the deployment group without having a different application revision for each group.

How can these requirements be met with the LEAST management overhead and without requiring different script versions for each deployment group?

- A. Tag the Amazon EC2 instances depending on the deployment group
- B. Then place a script into the application revision that calls the metadata service and the EC2 API to identify which deployment group the instance is part of
- C. Use this information to configure the log level setting
- D. Reference the script as part of the AfterInstall lifecycle hook in the appspec.yml file.
- E. Create a script that uses the CodeDeploy environment variable DEPLOYMENT\_GROUP\_NAME to identify which deployment group the instances are part of
- F. Use this information to configure the log level setting
- G. Reference this script as part of the BeforeInstall lifecycle hook in the appspec.yml file
- H. Create a CodeDeploy custom environment variable for each environment
- I. Then place a script into the application revision that checks this environment variable to identify which deployment group the instance is part of
- J. Use this information to configure the log level setting
- K. Reference this script as part of the ValidateService lifecycle hook in the appspec.yml file.
- L. Create a script that uses the CodeDeploy environment variable DEPLOYMENT\_GROUP\_ID to identify which deployment group the instance is part of to configure the log level setting
- M. Reference this script as part of the Install lifecycle hook in the appspec.yml file.

**Answer: B**

#### Explanation:

<https://docs.aws.amazon.com/codedeploy/latest/userguide/reference-appspec-file-structure-hooks.html>

#### NEW QUESTION 10

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

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

Which deployment strategy will meet these requirements?

- A. Use AWS CDK to deploy API Gateway and Lambda function
- B. When code needs to be changed, update the AWS CloudFormation stack and deploy the new version of the APIs and Lambda function

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

**Answer: B**

#### NEW QUESTION 11

An application running on a set of Amazon EC2 instances in an Auto Scaling group requires a configuration file to operate. The instances are created and maintained with AWS CloudFormation. A DevOps engineer wants the instances to have the latest configuration file when launched, and wants changes to the configuration file to be reflected on all the instances with a minimal delay when the CloudFormation template is updated. Company policy requires that application configuration files be maintained along with AWS infrastructure configuration files in source control.

Which solution will accomplish this?

- A. In the CloudFormation template, add an AWS Config rule.
- B. Place the configuration file content in the rule's InputParameters property, and set the Scope property to the EC2 Auto Scaling group.
- C. Add an AWS Systems Manager Resource Data Sync resource to the template to poll for updates to the configuration.
- D. In the CloudFormation template, add an EC2 launch template resource.
- E. Place the configuration file content in the launch template.
- F. Configure the cfn-init script to run when the instance is launched, and configure the cfn-hup script to poll for updates to the configuration.
- G. In the CloudFormation template, add an EC2 launch template resource.
- H. Place the configuration file content in the launch template.
- I. Add an AWS Systems Manager Resource Data Sync resource to the template to poll for updates to the configuration.
- J. In the CloudFormation template, add CloudFormation init metadata.
- K. Place the configuration file content in the metadata.
- L. Configure the cfn-init script to run when the instance is launched, and configure the cfn-hup script to poll for updates to the configuration.

**Answer: B**

#### NEW QUESTION 12

After presenting a working proof of concept for a new application that uses AWS API Gateway, a Developer must set up a team development environment for the project. Due to a tight timeline, the Developer wants to minimize time spent on infrastructure setup, and would like to reuse the code repository created for the proof of concept. Currently, all source code is stored in AWS CodeCommit.

Company policy mandates having alpha, beta, and production stages with separate Jenkins servers to build code and run tests for every stage. The Development Manager must have the ability to block code propagation between admins at any time. The Security team wants to make sure that users will not be able to modify the environment without permission.

How can this be accomplished?

- A. Create API Gateway alpha, beta, and production stage.
- B. Create a CodeCommit trigger to deploy code to the different stages using an AWS Lambda function.
- C. Create API Gateway alpha, beta, and production stage.
- D. Create an AWS CodePipeline that pulls code from the CodeCommit repository.
- E. Create CodePipeline actions to deploy code to the API Gateway stages.
- F. Create Jenkins servers for the alpha, beta, and production stages on Amazon EC2 instance.
- G. Create multiple CodeCommit triggers to deploy code to different stages using an AWS Lambda function.
- H. Create an AWS CodePipeline pipeline that pulls code from the CodeCommit repository.
- I. Create alpha, beta, and production stages with Jenkins servers on CodePipeline.

**Answer: D**

#### NEW QUESTION 15

A software company wants to automate the build process for a project where the code is stored in GitHub. When the repository is updated, source code should be compiled, tested, and pushed to Amazon S3.

Which combination of steps would address these requirements? (Select THREE.)

- A. Add a buildspec.yml file to the source code with build instructions.
- B. Configure a GitHub webhook to trigger a build every time a code change is pushed to the repository.
- C. Create an AWS CodeBuild project with GitHub as the source repository.
- D. Create an AWS CodeDeploy application with the Amazon EC2/On-Premises compute platform.
- E. Create an AWS OpsWorks deployment with the install dependencies command.
- F. Provision an Amazon EC2 instance to perform the build.

**Answer: ACD**

#### NEW QUESTION 18

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 19**

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

How can the requirement be met?

- A. At instance launch time, use EC2 user data to deploy the AWS CodeDeploy agent
- B. Use CodeDeploy to install the Puppet agent
- C. When the Auto Scaling group scales out, run a script to register the newly deployed instances to the Puppet master
- D. When the Auto Scaling group scales in, use the EC2 Auto Scaling lifecycle hook to trigger de-registration from the Puppet master
- E. EC2\_INSTANCE\_TERMINATING
- F. Bake the AWS CodeDeploy agent into the base AM
- G. When the Auto Scaling group scales out, use CodeDeploy to install the Puppet agent, and execute a script to register the newly deployed instances to the Puppet master
- H. When the Auto Scaling group scales in, use the CodeDeploy ApplicationStop lifecycle hook to run a script to de-register the instance from the Puppet master.
- I. At instance launch time, use EC2 user data to deploy the AWS CodeDeploy agent
- J. When the Auto Scaling group scales out, use CodeDeploy to install the Puppet agent, and run a script to register the newly deployed instances to the Puppet master
- K. When the Auto Scaling group scales in, use the EC2 user data instance stop script to run a script to de-register the instance from the Puppet master.
- L. Bake the AWS Systems Manager agent into the base AM
- M. When the Auto Scaling group scales out, use the AWS Systems Manager to install the Puppet agent, and run a script to register the newly deployed instances to the Puppet master
- N. When the Auto Scaling group scales in, use the Systems Manager instance stop lifecycle hook to run a script to de-register the instance from the Puppet master.

**Answer:** C

**NEW QUESTION 21**

A DevOps engineer must ensure all IAM entity configurations across multiple AWS accounts in AWS Organizations are compliant with corporate IAM policies. Which combination of steps will accomplish this? (Select TWO.)

- A. Enable AWS Trusted Advisor in Organizations for all accounts to report on noncompliant IAM entities.
- B. Configure an AWS Config aggregator in the Organizations master account for all accounts
- C. Deploy AWS Config rules to the master account in Organizations that match corporate IAM policies.
- D. Apply an SCP in Organizations to ensure compliance of IAM entities.
- E. Deploy AWS Config rules to all accounts in Organizations that match the corporate IAM policies.

**Answer:** BE

**NEW QUESTION 25**

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 28**

A company has thousands of Amazon EC2 instances as well as hundreds of virtual machines on-premises. Developers routinely sign in to the console for on-premises systems to perform troubleshooting. The developers want to sign in to AWS instances to run performance tools, but are unable to do so due to the lack of a central console logging system. A DevOps engineer wants to ensure that console access is logged on all systems.

Which combination of steps will meet these requirements? (Select TWO.)

- A. Attach a role to all AWS instances that contains the appropriate permission
- B. Create an AWS Systems Manager managed-instance activation
- C. Install and configure Systems Manager Agent on on-premises machines.
- D. Enable AWS Systems Manager Session Manager logging to an Amazon S3 bucket
- E. Direct developers to connect to the systems with Session Manager only.
- F. Enable AWS Systems Manager Session Manager logging to AWS CloudTrail
- G. Direct developers to continue normal sign-in procedures for on-premise
- H. Use Session Manager for AWS instances.
- I. Install and configure an Amazon CloudWatch Logs agent on all systems
- J. Create an AWS Systems Manager managed-instance activation.
- K. Set up a Site-to-Site VPN connection between the on-premises and AWS network
- L. Set up a bastion instance to allow developers to sign in to the AWS instances.

**Answer:** AB

**NEW QUESTION 29**

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 30**

A rapidly growing company wants to scale for Developer demand for AWS development environments.

Development environments are created manually in the AWS Management Console. The Networking team uses AWS CloudFormation to manage the networking infrastructure, exporting stack output values for the

Amazon VPC and all subnets. The development environments have common standards, such as Application Load Balancers, Amazon EC2 Auto Scaling groups, security groups, and Amazon DynamoDB tables.

To keep up with the demand, the DevOps Engineer wants to automate the creation of development environments. Because the infrastructure required to support the application is expected to grow, there must be a way to easily update the deployed infrastructure. CloudFormation will be used to create a template for the development environments.

Which approach will meet these requirements and quickly provide consistent AWS environments for Developers?

- A. Use Fn::ImportValue intrinsic functions in the Resources section of the template to retrieve Virtual Private Cloud (VPC) and subnet value
- B. Use CloudFormation StackSets for the development environments, using the Count input parameter to indicate the number of environments needed
- C. use the command to update existing development environment
- D. UpdateStackSet
- E. Use nested stacks to define common infrastructure component
- F. To access the exported values, use TemplateURL to reference the Networking team's template
- G. To retrieve Virtual Private Cloud (VPC) and subnet values, use Fn::ImportValue intrinsic functions in the Parameters section of the master template
- H. Use the CreateChangeSet and ExecuteChangeSet commands to update existing development environments.
- I. Use nested stacks to define common infrastructure component
- J. Use Fn::ImportValue intrinsic functions with the resources of the nested stack to retrieve Virtual Private Cloud (VPC) and subnet value
- K. Use the CreateChangeSet and ExecuteChangeSet commands to update existing development environments.
- L. Use Fn::ImportValue intrinsic functions in the Parameters section of the master template to retrieve Virtual Private Cloud (VPC) and subnet value
- M. Define the development resources in the order they need to be created in the CloudFormation nested stack
- N. Use the CreateChangeSet and ExecuteChangeSet commands to update existing development environments.

**Answer: A**

**NEW QUESTION 33**

A company's application is running on Amazon EC2 instances in an Auto Scaling group. A DevOps engineer needs to ensure there are at least four application servers running at all times. Whenever an update has to be made to the application, the engineer creates a new AMI with the updated configuration and updates the AWS CloudFormation template with the new AMI ID. After the stack update finishes, the engineer manually terminates the old instances one by one, verifying that the new instance is operational before proceeding. The engineer needs to automate this process.

Which action will allow for the LEAST number of manual steps moving forward?

- A. Update the CloudFormation template to include the UpdatePolicy attribute with the AutoScalingRollingUpdate policy.
- B. Update the CloudFormation template to include the UpdatePolicy attribute with the AutoScalingReplacingUpdate policy.
- C. Use an Auto Scaling lifecycle hook to verify that the previous instance is operational before allowing the DevOps engineer's selected instance to terminate.
- D. Use an Auto Scaling lifecycle hook to confirm there are at least four running instances before allowing the DevOps engineer's selected instance to terminate.

**Answer: A**

**NEW QUESTION 38**

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

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

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

**Answer: C**

**Explanation:**

[https://docs.aws.amazon.com/zh\\_cn/AmazonCloudFront/latest/DeveloperGuide/lambda-examples.html](https://docs.aws.amazon.com/zh_cn/AmazonCloudFront/latest/DeveloperGuide/lambda-examples.html)

**NEW QUESTION 42**

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

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

**Answer:** DE

**NEW QUESTION 47**

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 51**

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 54**

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 59**

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 EC
- 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 60

A company is using tagging to allocate AWS costs. The company has Amazon EC2 instances that run in Auto Scaling groups. The Amazon Elastic Block Store (Amazon EBS) volumes that are attached to the EC2 instances are being created without the appropriate cost center tags. A DevOps engineer must ensure that the new EBS volumes are properly tagged.  
What is the MOST efficient solution that meets this requirement?

- A. Create a lifecycle hook on the autoscaling:EC2\_INSTANCE\_TERMINATING instance state that attaches the cost center tags to the EBS volumes.
- B. Update the Auto Scaling group launch template to include the cost center tags for EBS volumes.
- C. Update the Auto Scaling group to include the cost center tag
- D. Set the PropagateAtLaunch property to true.
- E. Use Tag Editor to search for EBS volumes that are missing the tags and to add the cost center tags to the volumes.

**Answer:** B

#### NEW QUESTION 61

A company runs an application on Amazon EC2 instances behind an Application Load Balancer. The instances run in an Amazon EC2 Auto Scaling group across multiple Availability Zones in us-east-1. The application stores data in an Amazon RDS MySQL Multi-AZ DB instance.  
A DevOps Engineer wants to modify the current solution and create a hot standby of the environment in another region to minimize downtime if a problem occurs in us-east-1.  
Which combination of steps should the DevOps Engineer take to meet these requirements? (Select THREE.)

- A. Add a health check to the Amazon Route 53 alias record to evaluate the health of the primary region. Use AWS Lambda, configured with an Amazon CloudWatch Events trigger, to elect the Amazon RDS master in the disaster recovery region.
- B. Create a new Application Load Balancer and Auto Scaling group in the disaster recovery region.
- C. Extend the current Auto Scaling group to the subnets in the disaster recovery region.
- D. Enable multi-region failover for the RDS configuration for the database instance.
- E. Deploy a read replica of the RDS instance in the disaster recovery region.
- F. Create an AWS Lambda function to evaluate the health of the primary region
- G. If it fails, modify the Amazon Route 53 record to point at the disaster recovery region and elect the RDS master.

**Answer:** ABE

#### Explanation:

<https://aws.amazon.com/blogs/database/implementing-a-disaster-recovery-strategy-with-amazon-rds/>

#### NEW QUESTION 65

A company is using AWS CodePipeline to automate its release pipeline. AWS CodeDeploy is being used in the pipeline to deploy an application to Amazon ECS using the blue/green deployment model. The company wants to implement scripts to shifting traffic. These scripts will complete in 5 minutes or less. If errors are discovered during these tests, the application must be rolled back.  
Which strategy will meet these requirements?

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

**Answer:** A

#### NEW QUESTION 66

A Development team is currently using AWS CodeDeploy to deploy an application revision to an Auto Scaling group. If the deployment process fails, it must be rolled back automatically and a notification must be sent.  
What is the MOST effective configuration that can satisfy all of the requirements?

- A. Create Amazon CloudWatch Events rules for CodeDeploy operation
- B. Configure a CloudWatch Events rule to send out an Amazon SNS message when the deployment fail
- C. Configure CodeDeploy to automatically roll back when the deployment fails.
- D. Use available Amazon CloudWatch metrics for CodeDeploy to create CloudWatch alarm
- E. Configure CloudWatch alarms to send out an Amazon SNS message when the deployment fail
- F. Use AWS CLI to redeploy a previously deployed revision.
- G. Configure a CodeDeploy agent to create a trigger that will send notification to Amazon SNS topics when the deployment fail
- H. Configure CodeDeploy to automatically roll back when the deployment fails.
- I. Use AWS CloudTrail to monitor API calls made by or on behalf of CodeDeploy in the AWS account. Send an Amazon SNS message when deployment fail
- J. Use AWS CLI to redeploy a previously deployed revision.

**Answer: C**

**Explanation:**

<https://docs.aws.amazon.com/codedeploy/latest/userguide/monitoring-sns-event-notifications-create-trigger.htm>

**NEW QUESTION 70**

A company uses federated access for its AWS environment. The available roles are created and managed using AWS CloudFormation from a CI/CD pipeline. All changes should be made to the IAM roles through the pipeline. The security team found that changes are being made to the roles out-of-band and would like to detect when this occurs.

Which action will accomplish this?

- A. Use Amazon Inspector rules to detect and notify when a CloudFormation stack has a configuration change.
- B. Use an AWS Trusted Advisor CloudWatch Events rule to detect and notify when a CloudFormation stack has a configuration change.
- C. Use AWS CloudTrail to detect and notify when a CloudFormation stack has detected a configuration change.
- D. Use an AWS Config rule to detect and notify when a CloudFormation stack has detected a configuration change.

**Answer: D**

**NEW QUESTION 74**

A company must ensure consistent behavior of an application running on Amazon Linux in its corporate ecosystem before moving into AWS. The company has an existing automated server build system using VMware. The goal is to demonstrate the functionality of the application and its prerequisites on the new target operating system.

The DevOps Engineer needs to use the existing corporate server pipeline and virtualization software to create a server image. The server image will be tested on-premises to resemble the build on Amazon EC2 as closely as possible.

How can this be accomplished?

- A. Download and integrate the latest ISO of CentOS 7 and execute the application deployment on the resulting server.
- B. Launch an Amazon Linux AMI using an AWS OpsWorks deployment agent onto the on-premises infrastructure, then execute the application deployment.
- C. Build an EC2 instance with the latest Amazon Linux operating system, and use the AWS Import/Export service to export the EC2 image to a VMware ISO in Amazon S3. Then import the resulting ISO onto the on-premises system.
- D. Download and integrate the latest ISO of Amazon Linux 2 and execute the application deployment on the resulting server.
- E. Confirm that operating system testing results are consistent with EC2 operating system behavior.

**Answer: D**

**NEW QUESTION 76**

A company has developed an AWS Lambda function that handles orders received through an API. The company is using AWS CodeDeploy to deploy the Lambda function as the final stage of a CI/CD pipeline. A DevOps Engineer notices there are intermittent failures of the ordering API for a few seconds after deployment. After some investigation, the DevOps Engineer believes the failures are due to database changes. The CloudFormation stack for the application lambda function begins executing. How should the DevOps Engineer overcome this?

- A. Add a BeforeAllowTraffic hook to the AppSpec file that tests and waits for any necessary database changes before traffic can flow to the new version of the Lambda function
- B. Add an AfterAllowTraffic hook to the AppSpec file that forces traffic to wait for any pending database changes before allowing the new version of the Lambda function to respond
- C. Add a BeforeInstall hook to the AppSpec file that tests and waits for any necessary database changes before deploying the new version of the Lambda function
- D. Add a ValidateService hook to the AppSpec file that inspects incoming traffic and rejects the payload if dependent services such as the database are not yet ready

**Answer: B**

**NEW QUESTION 77**

An application is being deployed with two Amazon EC2 Auto Scaling groups, each configured with an Application Load Balancer. The application is deployed to one of the Auto Scaling groups and an Amazon Route 53 alias record is pointed to the Application Load Balancer of the last deployed Auto Scaling group. Deployments alternate between the two Auto Scaling groups.

Home security devices are making requests into the application. The Development team notes that new requests are coming into the old stack days after the deployment. The issue is caused by devices that are not observing the Time to Live (TTL) setting on the Amazon Route 53 alias record.

What steps should the DevOps Engineer take to address the issue with requests coming to the old stacks, while creating minimal additional resources?

- A. Create a fleet of Amazon EC2 instances running HAProxy behind an Application Load Balance
- B. The HAProxy instances will proxy the requests to one of the existing Auto Scaling group
- C. After a deployment the HAProxy instances are updated to send requests to the newly deployed Auto Scaling group.
- D. Reduce the application to one Application Load Balance
- E. Create two target groups named Blue and Green
- F. Create a rule on the Application Load Balancer pointed to a single target group
- G. Add logic to the deployment to update the Application Load Balancer rule to the target group of the newly deployed Auto Scaling group.
- H. Move the application to an AWS Elastic Beanstalk application with two environments
- I. Perform new deployments on the non-live environment

- J. After a deployment, perform an Elastic Beanstalk CNAME swap to make the newly deployed environment the live environment.
- K. Create an Amazon CloudFront distributio
- L. Set the two existing Application Load Balancers as origins on the distributio
- M. After a deployment, update the CloudFront distribution behavior to send requests to the newly deployed Auto Scaling group.

**Answer: B**

#### NEW QUESTION 81

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 dat
- B. To complete the jo
- 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 dat
- E. To run the batch jo
- F. configure EC2 instances manuell
- G. When the job completes, shut down the instances manually.
- H. Use Amazon EFS for checkpoint dat
- 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 dat
- 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 83

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 regio
- B. In the second region, enable cross-region replication between the original S3 bucket and a new S3 bucke
- C. To fail over, promote the read replica as maste
- 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 regio
- 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 snapsho
- I. In the second region, enable cross-region replication between the original S3 bucket and a new S3 bucke
- 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 regio
- 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 84

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 89

An application has microservices spread across different AWS accounts and is integrated with an on-premises legacy system for some of its functionality. Because of the segmented architecture and missing logs, every time the application experiences issues, it is taking too long to gather the logs to identify the issues. A DevOps Engineer must fix the log aggregation process and provide a way to centrally analyze the logs.

Which is the MOST efficient and cost-effective solution?

- A. Collect system logs and application logs by using the Amazon CloudWatch Logs agen
- B. Use the Amazon S3 API to export on-premises logs, and store the logs in an S3 bucket in a central accoun
- C. Build an Amazon EMR cluster to reduce the logs and derive the root cause.
- D. Collect system logs and application logs by using the Amazon CloudWatch Logs agen

- E. Use the Amazon S3 API to import on-premises log
- F. Store all logs in S3 buckets in individual account
- G. Use Amazon Macie to write a query to search for the required specific event-related data point.
- H. Collect system logs and application logs using the Amazon CloudWatch Logs agent
- I. Install the CloudWatch Logs agent on the on-premises server
- J. Transfer all logs from AWS to the on-premises data center
- K. Use an Amazon Elasticsearch Logstash Kibana stack to analyze logs on premises.
- L. Collect system logs and application logs by using the Amazon CloudWatch Logs agent
- M. Install a CloudWatch Logs agent for on-premises resource
- N. Store all logs in an S3 bucket in a central account
- O. Set up an Amazon S3 trigger and an AWS Lambda function to analyze incoming logs and automatically identify anomalies
- P. Use Amazon Athena to run ad hoc queries on the logs in the central account.

**Answer:** D

**NEW QUESTION 92**

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"	"ttt"	"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 key
- 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 key
- 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 key
- 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 key
- 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 97**

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

- A. Generate credentials for an IAM user with a policy attached to allow the actions on AWS service
- B. Store credentials as encrypted environment variables for the build project
- C. As part of the build script, obtain the credentials to run the integration tests.
- D. Have CodeBuild run only the integration tests as a build job on a Jenkins server
- E. Create a role that has a policy attached to allow the actions on AWS service
- F. Generate credentials for an IAM user that is allowed to assume the role
- G. Configure the credentials as secrets in Jenkins, and allow the build job to use them to run the integration tests.
- H. Create a service role in IAM to be assumed by CodeBuild with a policy attached to allow the actions on AWS service
- I. Configure the build project to use the role created.
- J. Use AWS managed credential
- K. Encrypt the credentials with AWS KMS
- L. As part of the build script, decrypt with AWS KMS and use these credentials to run the integration tests.

**Answer: B**

#### NEW QUESTION 101

A company is testing a web application that runs on Amazon EC2 instances behind an Application Load Balancer. The instances run in an Auto Scaling group across multiple Availability Zones. The company uses a blue/green deployment process with immutable instances when deploying new software. During testing, users are being automatically logged out of the application at random times. Testers also report that, when a new version of the application is deployed, all users are logged out. The Development team needs a solution to ensure users remain logged in across scaling events and application deployments. What is the MOST efficient way to ensure users remain logged in?

- A. Enable smart sessions on the load balancer and modify the application to check for an existing session.
- B. Enable session sharing on the load balancer and modify the application to read from the session store.
- C. Store user session information in an Amazon S3 bucket and modify the application to read session information from the bucket.
- D. Modify the application to store user session information in an Amazon ElastiCache cluster.

**Answer: D**

#### NEW QUESTION 106

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 109

A mobile application running on eight Amazon EC2 instances is relying on a third-party API endpoint. The thirdparty service has a high failure rate because of limited capacity, which is expected to be resolved in a few weeks. In the meantime, the mobile application developers have added a retry mechanism and are logging failed API requests. A DevOps Engineer must automate the monitoring of application logs and count the specific error messages; if there are more than 10 errors within a 1-minute window, the system must issue an alert. How can the requirements be met with MINIMAL management overhead?

- A. Install the Amazon CloudWatch Logs agent on all instances to push the application logs to CloudWatch Log
- B. Use metric filters to count the error messages every minute, and trigger a CloudWatch alarm if the count exceeds 10 errors.
- C. Install the Amazon CloudWatch Logs agent on all instances to push the access logs to CloudWatch Log
- D. Create CloudWatch Events rule to count the error messages every minute, and trigger a CloudWatch alarm if the count exceeds 10 errors.
- E. Install the Amazon CloudWatch Logs agent on all instances to push the application logs to CloudWatchLog
- F. Use a metric filter to generate a custom CloudWatch metric that records the number of failures and triggers a CloudWatch alarm if the custom metric reaches 10 errors in a 1-minute period.
- G. Deploy a custom script on all instances to check application logs regularly in a cron job
- H. Count the number of error messages every minute, and push a data point to a custom CloudWatch metri
- I. CloudWatch metri
- J. Trigger a CloudWatch alarm if the custom metric reaches 10 errors in a 1-minute period.

**Answer: C**

#### NEW QUESTION 111

A DevOps Engineer must create a Linux AMI in an automated fashion. The newly created AMI identification must be stored in a location where other build pipelines can access the new identification programmatically. What is the MOST cost-effective way to do this?

- A. Build a pipeline in AWS CodePipeline to download and save the latest operating system Open Virtualization Format (OVF) image to an Amazon S3 bucket, then customize the image using the guestfish utility
- B. Use the virtual machine (VM) import command to convert the OVF to an AMI, and store the AMI identification output as an AWS Systems Manager parameter.
- C. Create an AWS Systems Manager automation document with values instructing how the image should be created
- D. Then build a pipeline in AWS CodePipeline to execute the automation document to build the AMI when triggered
- E. Store the AMI identification output as a Systems Manager parameter.
- F. Build a pipeline in AWS CodePipeline to take a snapshot of an Amazon EC2 instance running the latest version of the application
- G. Then start a new EC2 instance from the snapshot and update the running instance using an AWS Lambda function
- H. Take a snapshot of the updated instance, then convert it to an AMI
- I. Store the AMI identification output in an Amazon DynamoDB table.
- J. Launch an Amazon EC2 instance and install Packer
- K. Then configure a Packer build with values defining how the image should be created
- L. Build a Jenkins pipeline to invoke the Packer build when triggered to build an AMI
- M. Store the AMI identification output in an Amazon DynamoDB table.

**Answer: D**

**NEW QUESTION 112**

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 databases
- 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 SNS
- 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 encryption
- J. Configure the application to force SSL for incoming connections and configure RDS to only grant access if the session is encrypted
- 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 databases
- M. Store RDS credentials in EC2 Parameter Store
- 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 115**

A company is building a web and mobile application that uses a serverless architecture powered by AWS Lambda and Amazon API Gateway. The company wants to fully automate the backend Lambda deployment based on code that is pushed to the appropriate environment branch in an AWS CodeCommit repository.

The deployment must have the following:

\*Separate environment pipelines for testing and production.

\*Automatic deployment that occurs for test environments only. Which steps should be taken to meet these requirements?

- A. Configure a new AWS CodePipeline service
- B. Create a CodeCommit repository for each environment. Set up CodePipeline to retrieve the source code from the appropriate repository
- C. Set up a deployment step to deploy the Lambda functions with AWS CloudFormation.
- D. Create two AWS CodePipeline configurations for test and production environment
- E. Configure the production pipeline to have a manual approval step
- F. Create a CodeCommit repository for each environment
- G. Set up each CodePipeline to retrieve the source code from the appropriate repository
- H. Set up the deployment step to deploy the Lambda functions with AWS CloudFormation.
- I. Create two AWS CodePipeline configurations for test and production environment
- J. Configure the production pipeline to have a manual approval step
- K. Create one CodeCommit repository with a branch for each environment
- L. Set up each CodePipeline to retrieve the source code from the appropriate branch in the repository
- M. Set up the deployment step to deploy the Lambda functions with AWS CloudFormation.
- N. Create an AWS CodeBuild configuration for test and production environment
- O. Configure the production pipeline to have a manual approval step
- P. Create one CodeCommit repository with a branch for each environment
- Q. Push the Lambda function code to an Amazon S3 bucket
- R. Set up the deployment step to deploy the Lambda functions from the S3 bucket.

**Answer: B**

**NEW QUESTION 118**

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

How can the Engineer improve the performance of the website?

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

**Answer: C**

**Explanation:**

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

<https://aws.amazon.com/dynamodb/dax/>

**NEW QUESTION 123**

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 125

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 127

A DevOps Engineer is working with an application deployed to 12 Amazon EC2 instances across 3 Availability Zones. New instances can be started from an AMI image. On a typical day, each EC2 instance has 30% utilization during business hours and 10% utilization after business hours. The CPU utilization has an immediate spike in the first few minutes of business hours. Other increases in CPU utilization rise gradually.

The Engineer has been asked to reduce costs while retaining the same or higher reliability. Which solution meets these requirements?

A. Create two Amazon CloudWatch Events rules with schedules before and after business hours begin and end

B. Create two AWS Lambda functions, one invoked by each rule

C. The first function should stop nine instances after business hours end, the second function should restart the nine instances before the business day begins.

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

E. Create two Amazon CloudWatch Events rules with schedules before and after business hours begin and end

F. Create an AWS CloudFormation stack, which creates an EC2 Auto Scaling group, with a parameter for the number of instance

G. Invoke the stack from each rule, passing a parameter value of three in the morning, and six in the evening.

H. Create an EC2 Auto Scaling group using the AMI image, with a scaling action based on the Auto Scaling group's CPU Utilization average with a target of 75%. Create a scheduled action to terminate nine instances each evening after the close of business.

**Answer: B**

#### NEW QUESTION 129

A company uses a complex system that consists of networking, IAM policies, and multiple three-tier applications. Requirements are still being defined for a new system, so the number of AWS components present in the final design is not known. The DevOps Engineer needs to begin defining AWS resources using AWS CloudFormation to automate and version-control the new infrastructure.

What is the best practice for using CloudFormation to create new environments?

A. Manually construct the networking layer using Amazon VPC and then define all other resources using CloudFormation.

B. Create a single template to encompass all resources that are required for the system so there is only one template to version-control.

- C. Create multiple separate templates for each logical part of the system, use cross-stack references in CloudFormation, and maintain several templates in version control.
- D. Create many separate templates for each logical part of the system, and provide the outputs from one to the next using an Amazon EC2 instance running SDK for granular control.

**Answer:** C

#### NEW QUESTION 130

A DevOps engineer used an AWS CloudFormation custom resource to set up AD Connector. The AWS Lambda function executed and created AD Connector, but CloudFormation is not transitioning from CREATE\_IN\_PROGRESS to CREATE\_COMPLETE. Which action should the engineer take to resolve this issue?

- A. Ensure the Lambda function code has exited successfully.
- B. Ensure the Lambda function code returns a response to the pre-signed URL.
- C. Ensure the Lambda function IAM role has cloudformation:UpdateStack permissions for the stack ARN.
- D. Ensure the Lambda function IAM role has ds:ConnectDirectory permissions for the AWS account.

**Answer:** A

#### NEW QUESTION 135

A company is running an application on Amazon EC2 instances behind an ELB Application Load Balancer. The instances run in an EC2 Auto Scaling group across multiple Availability Zones.

After a recent application update, users are getting HTTP 502 Bad Gateway errors from the application URL. The DevOps Engineer cannot analyze the problem because Auto Scaling is terminating all EC2 instances shortly after launch for being unhealthy.

What steps will allow the DevOps Engineer access to one of the unhealthy instances to troubleshoot the deployed application?

- A. Create an image from the terminated instance and create a new instance from that image.
- B. The Application team can then log into the new instance.
- C. As soon as a new instance is created by AutoScaling, put the instance into a Standby state as this will prevent the instance from being terminated.
- D. Add a lifecycle hook to your Auto Scaling group to move instances in the Terminating state to the Terminating:Wait state.
- E. Edit the Auto Scaling group to enable termination protection as this will protect unhealthy instances from being terminated.

**Answer:** B

#### Explanation:

<https://aws.amazon.com/blogs/aws/auto-scaling-update-lifecycle-standby-detach/>

#### NEW QUESTION 140

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 key.
- 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 type.
- 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 Store.
- 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 144

A company uses AWS CodePipeline to manage and deploy infrastructure as code. The infrastructure is defined in AWS CloudFormation templates and is primarily comprised of multiple Amazon EC2 instances and Amazon RDS databases. The Security team has observed many operators creating inbound security group rules with a source CIDR of 0.0.0.0/0 and would like to proactively stop the deployment of rules with open CIDRs.

The DevOps Engineer will implement a predeployment step that runs some security checks over the CloudFormation template before the pipeline processes it. This check should allow only inbound security group rules with a source CIDR of 0.0.0.0/0 if the rule has the description "Security Approval Ref XXXXX (where XXXXX is a preallocated reference)". The pipeline step should fail if this condition is not met and the deployment should be blocked. How should this be accomplished?

- A. Enable a SCP in AWS Organization.
- B. The policy should deny access to the API call CreateSecurityGroupRule if the rule specifies 0.0.0.0/0 without a description referencing a security approval.
- C. Add an initial stage to CodePipeline called Security Check.
- D. This stage should call an AWS Lambda function that scans the CloudFormation template and fails the pipeline if it finds 0.0.0.0/0 in a security group without a description referencing a security approval.
- E. Create an AWS Config rule that is triggered on creation or edit of resource type EC2 SecurityGroup. This rule should call an AWS Lambda function to send a failure notification if the security group has any rules with a source CIDR of 0.0.0.0/0 without a description referencing a security approval.
- F. Modify the IAM role used by CodePipeline.
- G. The IAM policy should deny access.

**Answer:** B

#### NEW QUESTION 147

A company needs to implement a robust CI/CD pipeline to automate the deployment of an application in AWS. The pipeline must support continuous integration,

continuous delivery, and automatic rollback upon deployment failure. The entire CI/CD pipeline must be capable of being re-provisioned in alternate AWS accounts or Regions within minutes. A DevOps engineer has already created an AWS CodeCommit repository to store the source code. Which combination of actions should be taken when building this pipeline to meet these requirements? (Select THREE.)

- A. Configure an AWS CodePipeline pipeline with a build stage using AWS CodeBuild.
- B. Copy the build artifact from CodeCommit to Amazon S3.
- C. Create an Auto Scaling group of Amazon EC2 instances behind an Application Load Balancer (ALB) and set the ALB as the deployment target in AWS CodePipeline.
- D. Create an AWS Elastic Beanstalk environment as the deployment target in AWS CodePipeline.
- E. Implement an Amazon SQS queue to decouple the pipeline components.
- F. Provision all resources using AWS CloudFormation.

**Answer:** ABD

#### NEW QUESTION 152

A company wants to use Amazon ECS to provide a Docker container runtime environment. For compliance reasons, all Amazon EBS volumes used in the ECS cluster must be encrypted. Rolling updates will be made to the cluster instances and the company wants the instances drained of all tasks before being terminated. How can these requirements be met? (Select TWO.)

- A. Modify the default ECS AMI user data to create a script that executes `docker rm -f {id}` for all running container instance
- B. Copy the script to the `/etc/init.d/rc.d` directory and execute `chconfig` enabling the script to run during operating system shutdown.
- C. Use AWS CodePipeline to build a pipeline that discovers the latest Amazon-provided ECS AMI, then copies the image to an encrypted AMI outputting the encrypted AMI ID
- D. Use the encrypted AMI ID when deploying the cluster.
- E. Copy the default AWS CloudFormation template that ECS uses to deploy cluster instance
- F. Modify the template resource EBS configuration setting to set `'Encrypted: True'` and include the AWS KMS alias: `'aws/ebs'` to encrypt the AMI.
- G. Create an Auto Scaling lifecycle hook backed by an AWS Lambda function that uses the AWS SDK to mark a terminating instance as DRAINING
- H. Prevent the lifecycle hook from completing until the running tasks on the instance are zero.
- I. Create an IAM role that allows the action `ECS::EncryptedImage`
- J. Configure the AWS CLI and a profile to use this role
- K. Start the cluster using the AWS CLI providing the `--use-encrypted-image` and `--kms-key` arguments to the `create-cluster` ECS command.

**Answer:** CD

#### NEW QUESTION 153

A company wants to use AWS Systems Manager documents to bootstrap physical laptops for developers. The bootstrap code is stored in GitHub. A DevOps engineer has already created a Systems Manager activation, installed the Systems Manager agent with the registration code, and installed an activation ID on all the laptops.

Which set of steps should be taken next?

- A. Configure the Systems Manager document to use the `AWS-RunShellScript` command to copy the files from GitHub to Amazon S3, then use the `aws-downloadContent` plugin with a source Type of S3.
- B. Configure the Systems Manager document to use the `aws-configurePackage` plugin with an install action and point to the Git repository.
- C. Configure the Systems Manager document to use the `aws-downloadContent` plugin with a sourceType of GitHub and sourceInfo with the repository details.
- D. Configure the Systems Manager document to use the `aws:softwareInventory` plugin and run the script from the Git repository.

**Answer:** D

#### NEW QUESTION 156

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 161

.....

## Thank You for Trying Our Product

\* 100% Pass or Money Back

All our products come with a 90-day Money Back Guarantee.

\* One year free update

You can enjoy free update one year. 24x7 online support.

\* Trusted by Millions

We currently serve more than 30,000,000 customers.

\* Shop Securely

All transactions are protected by VeriSign!

**100% Pass Your DOP-C01 Exam with Our Prep Materials Via below:**

<https://www.certleader.com/DOP-C01-dumps.html>