Custom View Settings

Topic 1 - Exam A

Question #1 *Topic 1*

A company has a mobile application that makes HTTP API calls to an Application Load Balancer (ALB). The ALB routes requests to an AWS Lambda function. Many different versions of the application are in use at any given time, including versions that are in testing by a subset of users. The version of the application is defined in the user-agent header that is sent with all requests to the API.

After a series of recent changes to the API, the company has observed issues with the application. The company needs to gather a metric for each API operation by response code for each version of the application that is in use. A DevOps engineer has modified the Lambda function to extract the API operation name, version information from the user-agent header and response code.

Which additional set of actions should the DevOps engineer take to gather the required metrics?

A. Modify the Lambda function to write the API operation name, response code, and version number as a log line to an Amazon CloudWatch Logs log group. Configure a CloudWatch Logs metric filter that increments a metric for each API operation name. Specify response code and application version as dimensions for the metric.

B. Modify the Lambda function to write the API operation name, response code, and version number as a log line to an Amazon CloudWatch Logs log group. Configure a CloudWatch Logs Insights query to populate CloudWatch metrics from the log lines. Specify response code and application version as dimensions for the metric.

C. Configure the ALB access logs to write to an Amazon CloudWatch Logs log group. Modify the Lambda function to respond to the ALB with the API operation name, response code, and version number as response metadata. Configure a CloudWatch Logs metric filter that increments a metric for each API operation name. Specify response code and application version as dimensions for the metric.

D. Configure AWS X-Ray integration on the Lambda function. Modify the Lambda function to create an X-Ray subsegment with the API operation name, response code, and version number. Configure X-Ray insights to extract an aggregated metric for each API operation name and to publish the metric to Amazon CloudWatch. Specify response code and application version as dimensions for the metric.

Question #2 *Topic 1*

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

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

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

Which solution will meet these requirements?

A. Configure provisioned concurrency on the Lambda function with a concurrency value of 1. Delete the DAX cluster for the DynamoDB table.

B. Configure reserved concurrency on the Lambda function with a concurrency value of 0.

C. Configure provisioned concurrency on the Lambda function. Configure AWS Application Auto Scaling on the Lambda function with provisioned concurrency values set to a minimum of 1 and a maximum of 100.

D. Configure reserved concurrency on the Lambda function. Configure AWS Application Auto Scaling on the API Gateway API with a reserved concurrency maximum value of 100.

Question #3 *Topic 1*

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

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

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

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

Question #4 *Topic 1*

A company requires its developers to tag all Amazon Elastic Block Store (Amazon EBS) volumes in an account to indicate a desired backup frequency. This requirement Includes EBS volumes that do not require backups. The company uses custom tags named Backup_Frequency that have values of none, dally, or weekly that correspond to the desired backup frequency. An audit finds that developers are occasionally not tagging the EBS volumes.
A DevOps engineer needs to ensure that all EBS volumes always have the Backup_Frequency tag so that the company can perform backups at least weekly unless a different value is specified.
Which solution will meet these requirements?

A. Set up AWS Config in the account. Create a custom rule that returns a compliance failure for all Amazon EC2 resources that do not have a Backup Frequency tag applied. Configure a remediation action that uses a custom AWS Systems Manager Automation runbook to apply the Backup_Frequency tag with a value of weekly.

B. Set up AWS Config in the account. Use a managed rule that returns a compliance failure for EC2::Volume resources that do not have a Backup Frequency tag applied. Configure a remediation action that uses a custom AWS Systems Manager Automation runbook to apply the Backup_Frequency tag with a value of weekly.

C. Turn on AWS CloudTrail in the account. Create an Amazon EventBridge rule that reacts to EBS CreateVolume events. Configure a custom AWS Systems Manager Automation runbook to apply the Backup_Frequency tag with a value of weekly. Specify the runbook as the target of the rule.

D. Turn on AWS CloudTrail in the account. Create an Amazon EventBridge rule that reacts to EBS CreateVolume events or EBS ModifyVolume events. Configure a custom AWS Systems Manager Automation runbook to apply the Backup_Frequency tag with a value of weekly. Specify the runbook as the target of the rule.

Question #5 *Topic 1*

A company is using an Amazon Aurora cluster as the data store for its application. The Aurora cluster is configured with a single DB instance. The application performs read and write operations on the database by using the cluster's instance endpoint.
The company has scheduled an update to be applied to the cluster during an upcoming maintenance window. The cluster must remain available with the least possible interruption during the maintenance window.
What should a DevOps engineer do to meet these requirements?

A. Add a reader instance to the Aurora cluster. Update the application to use the Aurora cluster endpoint for write operations. Update the Aurora cluster's reader endpoint for reads.

B. Add a reader instance to the Aurora cluster. Create a custom ANY endpoint for the cluster. Update the application to use the Aurora cluster's custom ANY endpoint for read and write operations.

C. Turn on the Multi-AZ option on the Aurora cluster. Update the application to use the Aurora cluster endpoint for write operations. Update the Aurora cluster’s reader endpoint for reads.

D. Turn on the Multi-AZ option on the Aurora cluster. Create a custom ANY endpoint for the cluster. Update the application to use the Aurora cluster's custom ANY endpoint for read and write operations

Question #6 *Topic 1*

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

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

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

A. In the source account, copy the unencrypted AMI to an encrypted AMI. Specify the KMS key in the copy action.

B. In the source account, copy the unencrypted AMI to an encrypted AMI. Specify the default Amazon Elastic Block Store (Amazon EBS) encryption key in the copy action.

C. In the source account, create a KMS grant that delegates permissions to the Auto Scaling group service-linked role in the target account.

D. In the source account, modify the key policy to give the target account permissions to create a grant. In the target account, create a KMS grant that delegates permissions to the Auto Scaling group service-linked role.

E. In the source account, share the unencrypted AMI with the target account.

F. In the source account, share the encrypted AMI with the target account.

Question #7 *Topic 1*

A company uses AWS CodePipeline pipelines to automate releases of its application A typical pipeline consists of three stages build, test, and deployment. The company has been using a separate AWS CodeBuild project to run scripts for each stage. However, the company now wants to use AWS CodeDeploy to handle the deployment stage of the pipelines.

The company has packaged the application as an RPM package and must deploy the application to a fleet of Amazon EC2 instances. The EC2 instances are in an EC2 Auto Scaling group and are launched from a common AMI.

Which combination of steps should a DevOps engineer perform to meet these requirements? (Choose two.)

A. Create a new version of the common AMI with the CodeDeploy agent installed. Update the IAM role of the EC2 instances to allow access to CodeDeploy.

B. Create a new version of the common AMI with the CodeDeploy agent installed. Create an AppSpec file that contains application deployment scripts and grants access to CodeDeploy.

C. Create an application in CodeDeploy. Configure an in-place deployment type. Specify the Auto Scaling group as the deployment target. Add a step to the CodePipeline pipeline to use EC2 Image Builder to create a new AMI. Configure CodeDeploy to deploy the newly created AMI.

D. Create an application in CodeDeploy. Configure an in-place deployment type. Specify the Auto Scaling group as the deployment target. Update the CodePipeline pipeline to use the CodeDeploy action to deploy the application.

E. Create an application in CodeDeploy. Configure an in-place deployment type. Specify the EC2 instances that are launched from the common AMI as the deployment target. Update the CodePipeline pipeline to use the CodeDeploy action to deploy the application.

Question #8 *Topic 1*

A company's security team requires that all external Application Load Balancers (ALBs) and Amazon API Gateway APIs are associated with AWS WAF web ACLs. The company has hundreds of AWS accounts, all of which are included in a single organization in AWS Organizations. The company has configured AWS Config for the organization. During an audit, the company finds some externally facing ALBs that are not associated with AWS WAF web ACLs.

Which combination of steps should a DevOps engineer take to prevent future violations? (Choose two.)

A. Delegate AWS Firewall Manager to a security account.

B. Delegate Amazon GuardDuty to a security account.

C. Create an AWS Firewall Manager policy to attach AWS WAF web ACLs to any newly created ALBs and API Gateway APIs.

D. Create an Amazon GuardDuty policy to attach AWS WAF web ACLs to any newly created ALBs and API Gateway APIs.

E. Configure an AWS Config managed rule to attach AWS WAF web ACLs to any newly created ALBs and API Gateway APIs.

Question #9 *Topic 1*

A company uses AWS Key Management Service (AWS KMS) keys and manual key rotation to meet regulatory compliance requirements. The security team wants to be notified when any keys have not been rotated after 90 days.

Which solution will accomplish this?

A. Configure AWS KMS to publish to an Amazon Simple Notification Service (Amazon SNS) topic when keys are more than 90 days old.

B. Configure an Amazon EventBridge event to launch an AWS Lambda function to call the AWS Trusted Advisor API and publish to an Amazon Simple Notification Service (Amazon SNS) topic.

C. Develop an AWS Config custom rule that publishes to an Amazon Simple Notification Service (Amazon SNS) topic when keys are more than 90 days old.

D. Configure AWS Security Hub to publish to an Amazon Simple Notification Service (Amazon SNS) topic when keys are more than 90 days old.

Question #10 *Topic 1*

A security review has identified that an AWS CodeBuild project is downloading a database population script from an Amazon S3 bucket using an unauthenticated request. The security team does not allow unauthenticated requests to S3 buckets for this project.

How can this issue be corrected in the MOST secure manner?

A. Add the bucket name to the AllowedBuckets section of the CodeBuild project settings. Update the build spec to use the AWS CLI to download the database population script.

B. Modify the S3 bucket settings to enable HTTPS basic authentication and specify a token. Update the build spec to use cURL to pass the token and download the database population script.

C. Remove unauthenticated access from the S3 bucket with a bucket policy. Modify the service role for the CodeBuild project to include Amazon S3 access. Use the AWS CLI to download the database population script.

D. Remove unauthenticated access from the S3 bucket with a bucket policy. Use the AWS CLI to download the database population script using an IAM access key and a secret access key.

Question #11 *Topic 1*

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

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

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

A. Create IAM policies that include the required permissions. Include the aws:PrincipalTag condition key.

B. Create permission sets. Attach an inline policy that includes the required permissions and uses the aws:PrincipalTag condition key to scope the permissions.

C. Create a group in the IdP. Place users in the group. Assign the group to accounts and the permission sets in IAM Identity Center.

D. Create a group in the IdP. Place users in the group. Assign the group to OUs and IAM policies.

E. Enable attributes for access control in IAM Identity Center. Apply tags to users. Map the tags as key-value pairs.

F. Enable attributes for access control in IAM Identity Center. Map attributes from the IdP as key-value pairs.

Question #12 *Topic 1*

An ecommerce company is receiving reports that its order history page is experiencing delays in reflecting the processing status of orders. The order processing system consists of an AWS Lambda function that uses reserved concurrency. The Lambda function processes order messages from an Amazon Simple Queue Service (Amazon SQS) queue and inserts processed orders into an Amazon DynamoDB table. The DynamoDB table has auto scaling enabled for read and write capacity.

Which actions should a DevOps engineer take to resolve this delay? (Choose two.)

A. Check the ApproximateAgeOfOldestMessage metric for the SQS queue. Increase the Lambda function concurrency limit.

B. Check the ApproximateAgeOfOldestMessage metnc for the SQS queue Configure a redrive policy on the SQS queue.

C. Check the NumberOfMessagesSent metric for the SQS queue. Increase the SQS queue visibility timeout.

D. Check the WriteThrottleEvents metric for the DynamoDB table. Increase the maximum write capacity units (WCUs) for the table's scaling policy.

E. Check the Throttles metric for the Lambda function. Increase the Lambda function timeout.

Question #13 *Topic 1*

A company has a single AWS account that runs hundreds of Amazon EC2 instances in a single AWS Region. New EC2 instances are launched and terminated each hour in the account. The account also includes existing EC2 instances that have been running for longer than a week.

The company's security policy requires all running EC2 instances to use an EC2 instance profile. If an EC2 instance does not have an instance profile attached, the EC2 instance must use a default instance profile that has no IAM permissions assigned.

A DevOps engineer reviews the account and discovers EC2 instances that are running without an instance profile. During the review, the DevOps engineer also observes that new EC2 instances are being launched without an instance profile.

Which solution will ensure that an instance profile is attached to all existing and future EC2 instances in the Region?

A. Configure an Amazon EventBridge rule that reacts to EC2 RunInstances API calls. Configure the rule to invoke an AWS Lambda function to attach the default instance profile to the EC2 instances.

B. Configure the ec2-instance-profile-attached AWS Config managed rule with a trigger type of configuration changes. Configure an automatic remediation action that invokes an AWS Systems Manager Automation runbook to attach the default instance profile to the EC2 instances.

C. Configure an Amazon EventBridge rule that reacts to EC2 StartInstances API calls. Configure the rule to invoke an AWS Systems Manager Automation runbook to attach the default instance profile to the EC2 instances

D. Configure the iam-role-managed-policy-check AWS Config managed rule with a trigger type of configuration changes. Configure an automatic remediation action that invokes an AWS Lambda function to attach the default instance profile to the EC2 instances.

Question #14 *Topic 1*

A DevOps engineer is building a continuous deployment pipeline for a serverless application that uses AWS Lambda functions. The company wants to reduce the customer impact of an unsuccessful deployment. The company also wants to monitor for issues.

Which deploy stage configuration will meet these requirements?

A. Use an AWS Serverless Application Model (AWS SAM) template to define the serverless application. Use AWS CodeDeploy to deploy the Lambda functions with the Canary10Percent15Minutes Deployment Preference Type. Use Amazon CloudWatch alarms to monitor the health of the functions.

B. Use AWS CloudFormation to publish a new stack update, and include Amazon CloudWatch alarms on all resources. Set up an AWS CodePipeline approval action for a developer to verify and approve the AWS CloudFormation change set.

C. Use AWS CloudFormation to publish a new version on every stack update, and include Amazon CloudWatch alarms on all resources. Use the RoutingConfig property of the AWS::Lambda::Alias resource to update the traffic routing during the stack update.

D. Use AWS CodeBuild to add sample event payloads for testing to the Lambda functions. Publish a new version of the functions, and include Amazon CloudWatch alarms. Update the production alias to point to the new version. Configure rollbacks to occur when an alarm is in the ALARM state.

Question #15 *Topic 1*

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

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

A. Launch the instances in a public subnet with Elastic IP addresses attached. Once the application is installed and running, run a script to disassociate the Elastic IP addresses afterwards.

B. Set up a NAT gateway. Deploy the EC2 instances to a private subnet. Update the private subnet's route table to use the NAT gateway as the default route.

C. Publish the application artifacts to an Amazon S3 bucket and create a VPC endpoint for S3. Assign an IAM instance profile to the EC2 instances so they can read the application artifacts from the S3 bucket.

D. Create a security group for the application instances and allow only outbound traffic to the artifact repository. Remove the security group rule once the install is complete.

Question #16 *Topic 1*

A development team is using AWS CodeCommit to version control application code and AWS CodePipeline to orchestrate software deployments. The team has decided to use a remote main branch as the trigger for the pipeline to integrate code changes. A developer has pushed code changes to the CodeCommit repository, but noticed that the pipeline had no reaction, even after 10 minutes.

Which of the following actions should be taken to troubleshoot this issue?

A. Check that an Amazon EventBridge rule has been created for the main branch to trigger the pipeline.

B. Check that the CodePipeline service role has permission to access the CodeCommit repository.

C. Check that the developer’s IAM role has permission to push to the CodeCommit repository.

D. Check to see if the pipeline failed to start because of CodeCommit errors in Amazon CloudWatch Logs.

Question #17 *Topic 1*

A company's developers use Amazon EC2 instances as remote workstations. The company is concerned that users can create or modify EC2 security groups to allow unrestricted inbound access.

A DevOps engineer needs to develop a solution to detect when users create unrestricted security group rules. The solution must detect changes to security group rules in near real time, remove unrestricted rules, and send email notifications to the security team. The DevOps engineer has created an AWS Lambda function that checks for security group ID from input, removes rules that grant unrestricted access, and sends notifications through Amazon Simple Notification Service (Amazon SNS).

What should the DevOps engineer do next to meet the requirements?

A. Configure the Lambda function to be invoked by the SNS topic. Create an AWS CloudTrail subscription for the SNS topic. Configure a subscription filter for security group modification events.

B. Create an Amazon EventBridge scheduled rule to invoke the Lambda function. Define a schedule pattern that runs the Lambda function every hour.

C. Create an Amazon EventBridge event rule that has the default event bus as the source. Define the rule’s event pattern to match EC2 security group creation and modification events. Configure the rule to invoke the Lambda function.

D. Create an Amazon EventBridge custom event bus that subscribes to events from all AWS services. Configure the Lambda function to be invoked by the custom event bus.

Question #18 *Topic 1*

A DevOps engineer is creating an AWS CloudFormation template to deploy a web service. The web service will run on Amazon EC2 instances in a private subnet behind an Application Load Balancer (ALB). The DevOps engineer must ensure that the service can accept requests from clients that have IPv6 addresses.

What should the DevOps engineer do with the CloudFormation template so that IPv6 clients can access the web service?

A. Add an IPv6 CIDR block to the VPC and the private subnet for the EC2 instances. Create route table entries for the IPv6 network, use EC2 instance types that support IPv6, and assign IPv6 addresses to each EC2 instance.

B. Assign each EC2 instance an IPv6 Elastic IP address. Create a target group, and add the EC2 instances as targets. Create a listener on port 443 of the ALB, and associate the target group with the ALB.

C. Replace the ALB with a Network Load Balancer (NLB). Add an IPv6 CIDR block to the VPC and subnets for the NLB, and assign the NLB an IPv6 Elastic IP address.

D. Add an IPv6 CIDR block to the VPC and subnets for the ALB. Create a listener on port 443. and specify the dualstack IP address type on the ALB. Create a target group, and add the EC2 instances as targets. Associate the target group with the ALB.

Question #19 *Topic 1*

A company uses AWS Organizations and AWS Control Tower to manage all the company's AWS accounts. The company uses the Enterprise Support plan.

A DevOps engineer is using Account Factory for Terraform (AFT) to provision new accounts. When new accounts are provisioned, the DevOps engineer notices that the support plan for the new accounts is set to the Basic Support plan. The DevOps engineer needs to implement a solution to provision the new accounts with the Enterprise Support plan.

Which solution will meet these requirements?

A. Use an AWS Config conformance pack to deploy the account-part-of-organizations AWS Config rule and to automatically remediate any noncompliant accounts.

B. Create an AWS Lambda function to create a ticket for AWS Support to add the account to the Enterprise Support plan. Grant the Lambda function the support:ResolveCase permission.

C. Add an additional value to the control_tower_parameters input to set the AWSEnterpriseSupport parameter as the organization's management account number.

D. Set the aft_feature_enterprise_support feature flag to True in the AFT deployment input configuration. Redeploy AFT and apply the changes.

Question #20 *Topic 1*

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

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

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

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

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

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

Question #21 *Topic 1*

A company has containerized all of its in-house quality control applications. The company is running Jenkins on Amazon EC2 instances, which require patching and upgrading. The compliance officer has requested a DevOps engineer begin encrypting build artifacts since they contain company intellectual property.
What should the DevOps engineer do to accomplish this in the MOST maintainable manner?

A. Automate patching and upgrading using AWS Systems Manager on EC2 instances and encrypt Amazon EBS volumes by default.

B. Deploy Jenkins to an Amazon ECS cluster and copy build artifacts to an Amazon S3 bucket with default encryption enabled.

C. Leverage AWS CodePipeline with a build action and encrypt the artifacts using AWS Secrets Manager.

D. Use AWS CodeBuild with artifact encryption to replace the Jenkins instance running on EC2 instances.

Question #22 *Topic 1*

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 a DelelionPolicy 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 with an AWS Lambda function with the DependsOn attribute specifying the S3 bucket, and an IAM role. Write the Lambda function to delete all objects from the bucket when RequestType is Delete.

C. Identify the resource that was not deleted. Manually empty the S3 bucket and then delete it.

D. Replace the EC2 and S3 bucket resources with a single AWS OpsWorks Stacks resource. Define a custom recipe for the stack to create and delete the EC2 instance and the S3 bucket.

Question #23 *Topic 1*

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

The CodeBuild project uses the aws cloudformation package AWS CLI command to build an artifact that contains the Lambda function code's .zip file and the CloudFormation template. The CloudFormation deploy action references the CloudFormation template from the output artifact of the CodeBuild project's build action.

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

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

A. Modify the CloudFormation template to include a parameter for the Lambda function code's zip file location. Create a new CloudFormation deploy action for us-east-1 in the pipeline. Configure the new deploy action to pass in the us-east-1 artifact location as a parameter override.

B. Create a new CloudFormation deploy action for us-east-1 in the pipeline. Configure the new deploy action to use the CloudFormation template from the us-east-1 output artifact.

C. Create an S3 bucket in us-east-1. Configure the S3 bucket policy to allow CodePipeline to have read and write access.

D. Create an S3 bucket in us-east-1. Configure S3 Cross-Region Replication (CRR) from the S3 bucket in eu-west-1 to the S3 bucket in us-east-1.

E. Modify the pipeline to include the S3 bucket for us-east-1 as an artifact store. Create a new CloudFormation deploy action for us-east-1 in the pipeline. Configure the new deploy action to use the CloudFormation template from the us-east-1 output artifact.

Question #24 *Topic 1*

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

Which solution will meet these requirements?

A. Create an Amazon CloudWatch alarm for the StatusCheckFailed metric. Use the recover action to stop and start the instance. Use an S3 event notification to push the metadata to the instance when the instance is back up and running.

B. Configure AWS OpsWorks, and use the auto healing feature to stop and start the instance. Use a lifecycle event in OpsWorks to pull the metadata from Amazon S3 and update it on the instance.

C. Use EC2 Auto Recovery to automatically stop and start the instance in case of a failure. Use an S3 event notification to push the metadata to the instance when the instance is back up and running.

D. Use AWS CloudFormation to create an EC2 instance that includes the UserData property for the EC2 resource. Add a command in UserData to retrieve the application metadata from Amazon S3.

Question #25 *Topic 1*

A company has multiple AWS accounts. The company uses AWS IAM Identity Center (AWS Single Sign-On) that is integrated with AWS Toolkit for Microsoft Azure DevOps. The attributes for access control feature is enabled in IAM Identity Center.

The attribute mapping list contains two entries. The department key is mapped to ${path:enterprise.department}. The costCenter key is mapped to ${path:enterprise.costCenter}.

All existing Amazon EC2 instances have a department tag that corresponds to three company departments (d1, d2, d3). A DevOps engineer must create policies based on the matching attributes. The policies must minimize administrative effort and must grant each Azure AD user access to only the EC2 instances that are tagged with the user's respective department name.

Which condition key should the DevOps engineer include in the custom permissions policies to meet these requirements?

A.

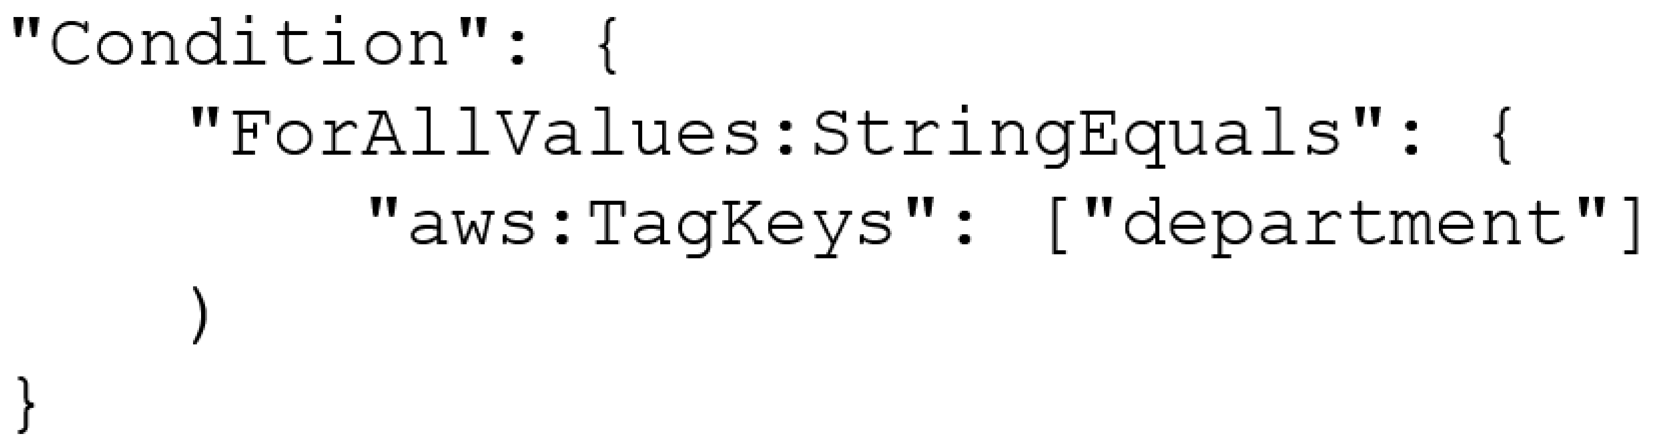

```
"Condition": {
    "ForAllValues:StringEquals": {
        "aws:TagKeys": ["department"]
    )
}
```

B.

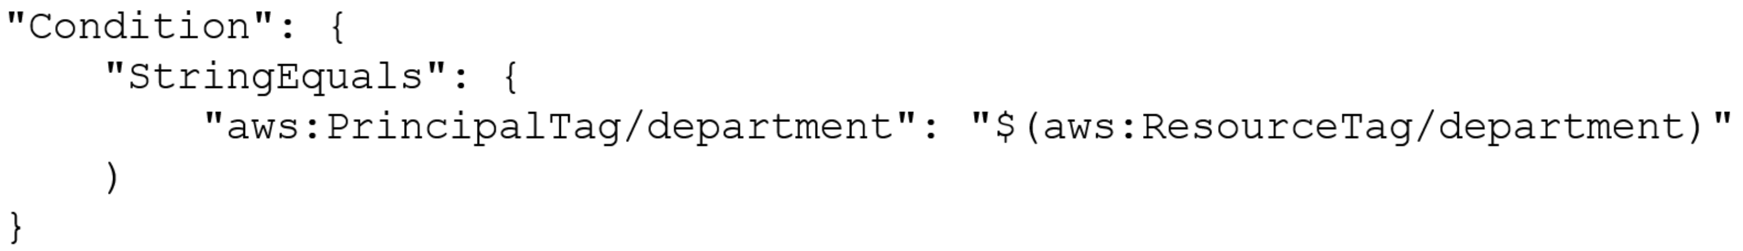

```
"Condition": {
    "StringEquals": {
        "aws:PrincipalTag/department": "$(aws:ResourceTag/department)"
    )
}
```

C.

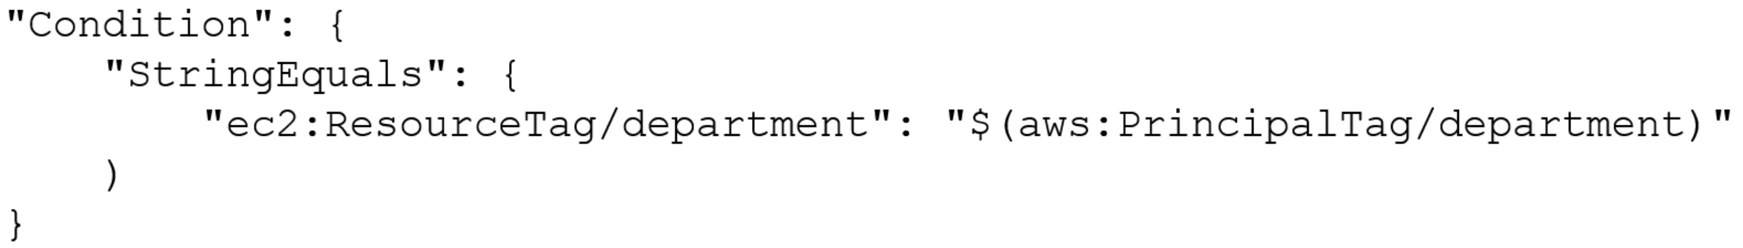

```
"Condition": {
    "StringEquals": {
        "ec2:ResourceTag/department": "$(aws:PrincipalTag/department)"
    )
}
```

D.

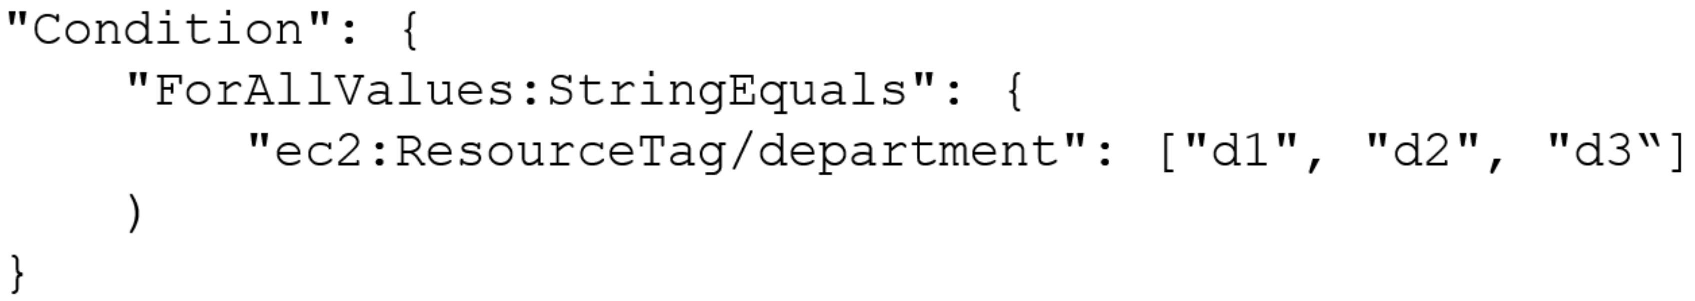

```
"Condition": {
    "ForAllValues:StringEquals": {
        "ec2:ResourceTag/department": ["d1", "d2", "d3“]
    )
}
```

Question #26 *Topic 1*

A company hosts a security auditing application in an AWS account. The auditing application uses an IAM role to access other AWS accounts. All the accounts are in the same organization in AWS Organizations.

A recent security audit revealed that users in the audited AWS accounts could modify or delete the auditing application's IAM role. The company needs to prevent any modification to the auditing application's IAM role by any entity other than a trusted administrator IAM role.

Which solution will meet these requirements?

A. Create an SCP that includes a Deny statement for changes to the auditing application's IAM role. Include a condition that allows the trusted administrator IAM role to make changes. Attach the SCP to the root of the organization.

B. Create an SCP that includes an Allow statement for changes to the auditing application's IAM role by the trusted administrator IAM role. Include a Deny statement for changes by all other IAM principals. Attach the SCP to the IAM service in each AWS account where the auditing application has an IAM role.

C. Create an IAM permissions boundary that includes a Deny statement for changes to the auditing application's IAM role. Include a condition that allows the trusted administrator IAM role to make changes. Attach the permissions boundary to the audited AWS accounts.

D. Create an IAM permissions boundary that includes a Deny statement for changes to the auditing application’s IAM role. Include a condition that allows the trusted administrator IAM role to make changes. Attach the permissions boundary to the auditing application's IAM role in the AWS accounts.

Question #27 *Topic 1*

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

Which solution will meet these requirements?

A. Deploy the application on an Amazon EC2 instance, and create an AMI of the instance. Use the AMI to create an automatic scaling launch configuration that is used in an Auto Scaling group. Use Elastic Load Balancing to distribute traffic. When changes are made to the application, a new AMI will be created, which will initiate an EC2 instance refresh.

B. Use Amazon Lightsail to deploy the application. Store the application in a zipped format in an Amazon S3 bucket. Use this zipped version to deploy new versions of the application to Lightsail. Use Lightsail deployment options to manage the deployment.

C. Use AWS CodeArtifact to store the application code. Use AWS CodeDeploy to deploy the application to a fleet of Amazon EC2 instances. Use Elastic Load Balancing to distribute the traffic to the EC2 instances. When making changes to the application, upload a new version to CodeArtifact and create a new CodeDeploy deployment.

D. Use AWS Elastic Beanstalk to host the application. Store a zipped version of the application in Amazon S3. Use that location to deploy new versions of the application. Use Elastic Beanstalk to manage the deployment options.

Question #28 *Topic 1*

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 state. Create an Amazon CloudWatch alarm for EC2 Instance Terminate Successful and trigger an AWS Lambda function that invokes an SSM Run Command script to collect logs, push them to Amazon S3, and complete the lifecycle action once logs are collected.

B. Use Auto Scaling lifecycle hooks to put instances in a Terminating:Wait state. Create an AWS Config rule for EC2 Instance-terminate Lifecycle Action and trigger a step function that invokes a script to collect logs, push them to Amazon S3, and complete the lifecycle action once logs are collected.

C. Use Auto Scaling lifecycle hooks to put instances in a Terminating:Wait state. Create an Amazon CloudWatch subscription filter for EC2 Instance Terminate Successful and trigger a CloudWatch agent that invokes a script to collect logs, push them to Amazon S3, and complete the lifecycle action once logs are collected.

D. Use Auto Scaling lifecycle hooks to put instances in a Terminating:Wait state. Create an Amazon EventBridge rule for EC2 Instance-terminate Lifecycle Action and trigger an AWS Lambda function that invokes an SSM Run Command script to collect logs, push them to Amazon S3, and complete the lifecycle action once logs are collected.

Question #29 *Topic 1*

A company has an organization in AWS Organizations. The organization includes workload accounts that contain enterprise applications. The company centrally manages users from an operations account. No users can be created in the workload accounts. The company recently added an operations team and must provide the operations team members with administrator access to each workload account.

Which combination of actions will provide this access? (Choose three.)

A. Create a SysAdmin role in the operations account. Attach the AdministratorAccess policy to the role. Modify the trust relationship to allow the sts:AssumeRole action from the workload accounts.

B. Create a SysAdmin role in each workload account. Attach the AdministratorAccess policy to the role. Modify the trust relationship to allow the sts:AssumeRole action from the operations account.

C. Create an Amazon Cognito identity pool in the operations account. Attach the SysAdmin role as an authenticated role.

D. In the operations account, create an IAM user for each operations team member.

E. In the operations account, create an IAM user group that is named SysAdmins. Add an IAM policy that allows the sts:AssumeRole action for the SysAdmin role in each workload account. Add all operations team members to the group.

F. Create an Amazon Cognito user pool in the operations account. Create an Amazon Cognito user for each operations team member.

Question #30

Topic 1

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

A. Designate an account to be the delegated Amazon GuardDuty administrator account. Turn on GuardDuty for all accounts across the organization. In the GuardDuty administrator account, create an SNS topic. Subscribe the SecOps team's email address to the SNS topic. In the same account, create an Amazon EventBridge rule that uses an event pattern for GuardDuty findings and a target of the SNS topic.

B. Create an AWS CloudFormation template that creates an SNS topic and subscribes the SecOps team’s email address to the SNS topic. In the template, include an Amazon EventBridge rule that uses an event pattern of CloudTrail activity for s3:PutBucketPublicAccessBlock and a target of the SNS topic. Deploy the stack to every account in the organization by using CloudFormation StackSets.

C. Turn on AWS Config across the organization. In the delegated administrator account, create an SNS topic. Subscribe the SecOps team's email address to the SNS topic. Deploy a conformance pack that uses the s3-bucket-level-public-access-prohibited AWS Config managed rule in each account and uses an AWS Systems Manager document to publish an event to the SNS topic to notify the SecOps team.

D. Turn on Amazon Inspector across the organization. In the Amazon Inspector delegated administrator account, create an SNS topic. Subscribe the SecOps team’s email address to the SNS topic. In the same account, create an Amazon EventBridge rule that uses an event pattern for public network exposure of the S3 bucket and publishes an event to the SNS topic to notify the SecOps team.

Question #31

Topic 1

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

A. Enable Amazon CloudWatch Logs to log the EKS components. Create a CloudWatch subscription filter for each component with Lambda as the subscription feed destination.

B. Enable Amazon CloudWatch Logs to log the EKS components. Create CloudWatch Logs Insights queries linked to Amazon EventBridge events that invoke Lambda.

C. Enable Amazon S3 logging for the EKS components. Configure an Amazon CloudWatch subscription filter for each component with Lambda as the subscription feed destination.

D. Enable Amazon S3 logging for the EKS components. Configure S3 PUT Object event notifications with AWS Lambda as the destination.

Question #32

Topic 1

A company is implementing an Amazon Elastic Container Service (Amazon ECS) cluster to run its workload. The company architecture will run multiple ECS services on the cluster. The architecture includes an Application Load Balancer on the front end and uses multiple target groups to route traffic.

A DevOps engineer must collect application and access logs. The DevOps engineer then needs to send the logs to an Amazon S3 bucket for near-real-time analysis.

Which combination of steps must the DevOps engineer take to meet these requirements? (Choose three.)

A. Download the Amazon CloudWatch Logs container instance from AWS. Configure this instance as a task. Update the application service definitions to include the logging task.

B. Install the Amazon CloudWatch Logs agent on the ECS instances. Change the logging driver in the ECS task definition to awslogs.

C. Use Amazon EventBridge to schedule an AWS Lambda function that will run every 60 seconds and will run the Amazon CloudWatch Logs create-export-task command. Then point the output to the logging S3 bucket.

D. Activate access logging on the ALB. Then point the ALB directly to the logging S3 bucket.

E. Activate access logging on the target groups that the ECS services use. Then send the logs directly to the logging S3 bucket.

F. Create an Amazon Kinesis Data Firehose delivery stream that has a destination of the logging S3 bucket. Then create an Amazon CloudWatch Logs subscription filter for Kinesis Data Firehose.

Question #33

Topic 1

A company that uses electronic health records is running a fleet of Amazon EC2 instances with an Amazon Linux operating system. As part of patient privacy requirements, the company must ensure continuous compliance for patches for operating system and applications running on the EC2 instances.

How can the deployments of the operating system and application patches be automated using a default and custom repository?

A. Use AWS Systems Manager to create a new patch baseline including the custom repository. Run the AWS-RunPatchBaseline document using the run command to verify and install patches.

B. Use AWS Direct Connect to integrate the corporate repository and deploy the patches using Amazon CloudWatch scheduled events, then use the CloudWatch dashboard to create reports.

C. Use yum-config-manager to add the custom repository under /etc/yum.repos.d and run yum-config-manager-enable to activate the repository.

D. Use AWS Systems Manager to create a new patch baseline including the corporate repository. Run the AWS-AmazonLinuxDefaultPatchBaseline document using the run command to verify and install patches.

Question #34 *Topic 1*

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

Which strategy will meet these requirements?

A. Add a stage to the CodePipeline pipeline between the source and deploy stages. Use AWS CodeBuild to create a runtime environment and build commands in the buildspec file to invoke test scripts. If errors are found, use the aws deploy stop-deployment command to stop the deployment.

B. Add a stage to the CodePipeline pipeline between the source and deploy stages. Use this stage to invoke an AWS Lambda function that will run the test scripts. If errors are found, use the aws deploy stop-deployment command to stop the deployment.

C. Add a hooks section to the CodeDeploy AppSpec file. Use the AfterAllowTestTraffic lifecycle event to invoke an AWS Lambda function to run the test scripts. If errors are found, exit the Lambda function with an error to initiate rollback.

D. Add a hooks section to the CodeDeploy AppSpec file. Use the AfterAllowTraffic lifecycle event to invoke the test scripts. If errors are found, use the aws deploy stop-deployment CLI command to stop the deployment.

Question #35 *Topic 1*

A company uses AWS Storage Gateway in file gateway mode in front of an Amazon S3 bucket that is used by multiple resources. In the morning when business begins, users do not see the objects processed by a third party the previous evening. When a DevOps engineer looks directly at the S3 bucket, the data is there, but it is missing in Storage Gateway.

Which solution ensures that all the updated third-party files are available in the morning?

A. Configure a nightly Amazon EventBridge event to invoke an AWS Lambda function to run the RefreshCache command for Storage Gateway.

B. Instruct the third party to put data into the S3 bucket using AWS Transfer for SFTP.

C. Modify Storage Gateway to run in volume gateway mode.

D. Use S3 Same-Region Replication to replicate any changes made directly in the S3 bucket to Storage Gateway.

Question #36 *Topic 1*

A DevOps engineer needs to back up sensitive Amazon S3 objects that are stored within an S3 bucket with a private bucket policy using S3 cross-Region replication functionality. The objects need to be copied to a target bucket in a different AWS Region and account.

Which combination of actions should be performed to enable this replication? (Choose three.)

A. Create a replication IAM role in the source account

B. Create a replication I AM role in the target account.

C. Add statements to the source bucket policy allowing the replication IAM role to replicate objects.

D. Add statements to the target bucket policy allowing the replication IAM role to replicate objects.

E. Create a replication rule in the source bucket to enable the replication.

F. Create a replication rule in the target bucket to enable the replication.

Question #37 *Topic 1*

A company has multiple member accounts that are part of an organization in AWS Organizations. The security team needs to review every Amazon EC2 security group and their inbound and outbound rules. The security team wants to programmatically retrieve this information from the member accounts using an AWS Lambda function in the management account of the organization.
Which combination of access changes will meet these requirements? (Choose three.)

A. Create a trust relationship that allows users in the member accounts to assume the management account IAM role.

B. Create a trust relationship that allows users in the management account to assume the IAM roles of the member accounts.

C. Create an IAM role in each member account that has access to the AmazonEC2ReadOnlyAccess managed policy.

D. Create an I AM role in each member account to allow the sts:AssumeRole action against the management account IAM role's ARN.

E. Create an I AM role in the management account that allows the sts:AssumeRole action against the member account IAM role's ARN.

F. Create an IAM role in the management account that has access to the AmazonEC2ReadOnlyAccess managed policy.

Question #38 *Topic 1*

A space exploration company receives telemetry data from multiple satellites. Small packets of data are received through Amazon API Gateway and are placed directly into an Amazon Simple Queue Service (Amazon SQS) standard queue. A custom application is subscribed to the queue and transforms the data into a standard format.
Because of inconsistencies in the data that the satellites produce, the application is occasionally unable to transform the data. In these cases, the messages remain in the SQS queue. A DevOps engineer must develop a solution that retains the failed messages and makes them available to scientists for review and future processing.
Which solution will meet these requirements?

A. Configure AWS Lambda to poll the SQS queue and invoke a Lambda function to check whether the queue messages are valid. If validation fails, send a copy of the data that is not valid to an Amazon S3 bucket so that the scientists can review and correct the data. When the data is corrected, amend the message in the SQS queue by using a replay Lambda function with the corrected data.

B. Convert the SQS standard queue to an SQS FIFO queue. Configure AWS Lambda to poll the SQS queue every 10 minutes by using an Amazon EventBridge schedule. Invoke the Lambda function to identify any messages with a SentTimestamp value that is older than 5 minutes, push the data to the same location as the application's output location, and remove the messages from the queue.

C. Create an SQS dead-letter queue. Modify the existing queue by including a redrive policy that sets the Maximum Receives setting to 1 and sets the dead-letter queue ARN to the ARN of the newly created queue. Instruct the scientists to use the dead-letter queue to review the data that is not valid. Reprocess this data at a later time.

D. Configure API Gateway to send messages to different SQS virtual queues that are named for each of the satellites. Update the application to use a new virtual queue for any data that it cannot transform, and send the message to the new virtual queue. Instruct the scientists to use the virtual queue to review the data that is not valid. Reprocess this data at a later time.

Question #39 *Topic 1*

A company wants to use AWS CloudFormation for infrastructure deployment. The company has strict tagging and resource requirements and wants to limit the deployment to two Regions. Developers will need to deploy multiple versions of the same application.
Which solution ensures resources are deployed in accordance with company policy?

A. Create AWS Trusted Advisor checks to find and remediate unapproved CloudFormation StackSets.

B. Create a Cloud Formation drift detection operation to find and remediate unapproved CloudFormation StackSets.

C. Create CloudFormation StackSets with approved CloudFormation templates.

D. Create AWS Service Catalog products with approved CloudFormation templates.

Question #40 *Topic 1*

A company requires that its internally facing web application be highly available. The architecture is made up of one Amazon EC2 web server instance and one NAT instance that provides outbound internet access for updates and accessing public data.
Which combination of architecture adjustments should the company implement to achieve high availability? (Choose two.)

A. Add the NAT instance to an EC2 Auto Scaling group that spans multiple Availability Zones. Update the route tables.

B. Create additional EC2 instances spanning multiple Availability Zones. Add an Application Load Balancer to split the load between them.

C. Configure an Application Load Balancer in front of the EC2 instance. Configure Amazon CloudWatch alarms to recover the EC2 instance upon host failure.

D. Replace the NAT instance with a NAT gateway in each Availability Zone. Update the route tables.

E. Replace the NAT instance with a NAT gateway that spans multiple Availability Zones. Update the route tables.

Question #41 *Topic 1*

A DevOps engineer is building a multistage pipeline with AWS CodePipeline to build, verify, stage, test, and deploy an application. A manual approval stage is required between the test stage and the deploy stage. The development team uses a custom chat tool with webhook support that requires near-real-time notifications.
How should the DevOps engineer configure status updates for pipeline activity and approval requests to post to the chat tool?

A. Create an Amazon CloudWatch Logs subscription that filters on CodePipeline Pipeline Execution State Change. Publish subscription events to an Amazon Simple Notification Service (Amazon SNS) topic. Subscribe the chat webhook URL to the SNS topic, and complete the subscription validation.

B. Create an AWS Lambda function that is invoked by AWS CloudTrail events. When a CodePipeline Pipeline Execution State Change event is detected, send the event details to the chat webhook URL.

C. Create an Amazon EventBridge rule that filters on CodePipeline Pipeline Execution State Change. Publish the events to an Amazon Simple Notification Service (Amazon SNS) topic. Create an AWS Lambda function that sends event details to the chat webhook URL. Subscribe the function to the SNS topic.

D. Modify the pipeline code to send the event details to the chat webhook URL at the end of each stage. Parameterize the URL so that each pipeline can send to a different URL based on the pipeline environment.

Question #42 *Topic 1*

A company's application development team uses Linux-based Amazon EC2 instances as bastion hosts. Inbound SSH access to the bastion hosts is restricted to specific IP addresses, as defined in the associated security groups. The company's security team wants to receive a notification if the security group rules are modified to allow SSH access from any IP address.
What should a DevOps engineer do to meet this requirement?

A. Create an Amazon EventBridge rule with a source of aws.cloudtrail and the event name AuthorizeSecurityGroupIngress. Define an Amazon Simple Notification Service (Amazon SNS) topic as the target.

B. Enable Amazon GuardDuty and check the findings for security groups in AWS Security Hub. Configure an Amazon EventBridge rule with a custom pattern that matches GuardDuty events with an output of NON_COMPLIANT. Define an Amazon Simple Notification Service (Amazon SNS) topic as the target.

C. Create an AWS Config rule by using the restricted-ssh managed rule to check whether security groups disallow unrestricted incoming SSH traffic. Configure automatic remediation to publish a message to an Amazon Simple Notification Service (Amazon SNS) topic.

D. Enable Amazon Inspector. Include the Common Vulnerabilities and Exposures-1.1 rules package to check the security groups that are associated with the bastion hosts. Configure Amazon Inspector to publish a message to an Amazon Simple Notification Service (Amazon SNS) topic.

Question #43 *Topic 1*

A DevOps team manages an API running on-premises that serves as a backend for an Amazon API Gateway endpoint. Customers have been complaining about high response latencies, which the development team has verified using the API Gateway latency metrics in Amazon CloudWatch. To identify the cause, the team needs to collect relevant data without introducing additional latency.
Which actions should be taken to accomplish this? (Choose two.)

A. Install the CloudWatch agent server side and configure the agent to upload relevant logs to CloudWatch.

B. Enable AWS X-Ray tracing in API Gateway, modify the application to capture request segments, and upload those segments to X-Ray during each request.

C. Enable AWS X-Ray tracing in API Gateway, modify the application to capture request segments, and use the X-Ray daemon to upload segments to X-Ray.

D. Modify the on-premises application to send log information back to API Gateway with each request.

E. Modify the on-premises application to calculate and upload statistical data relevant to the API service requests to CloudWatch metrics.

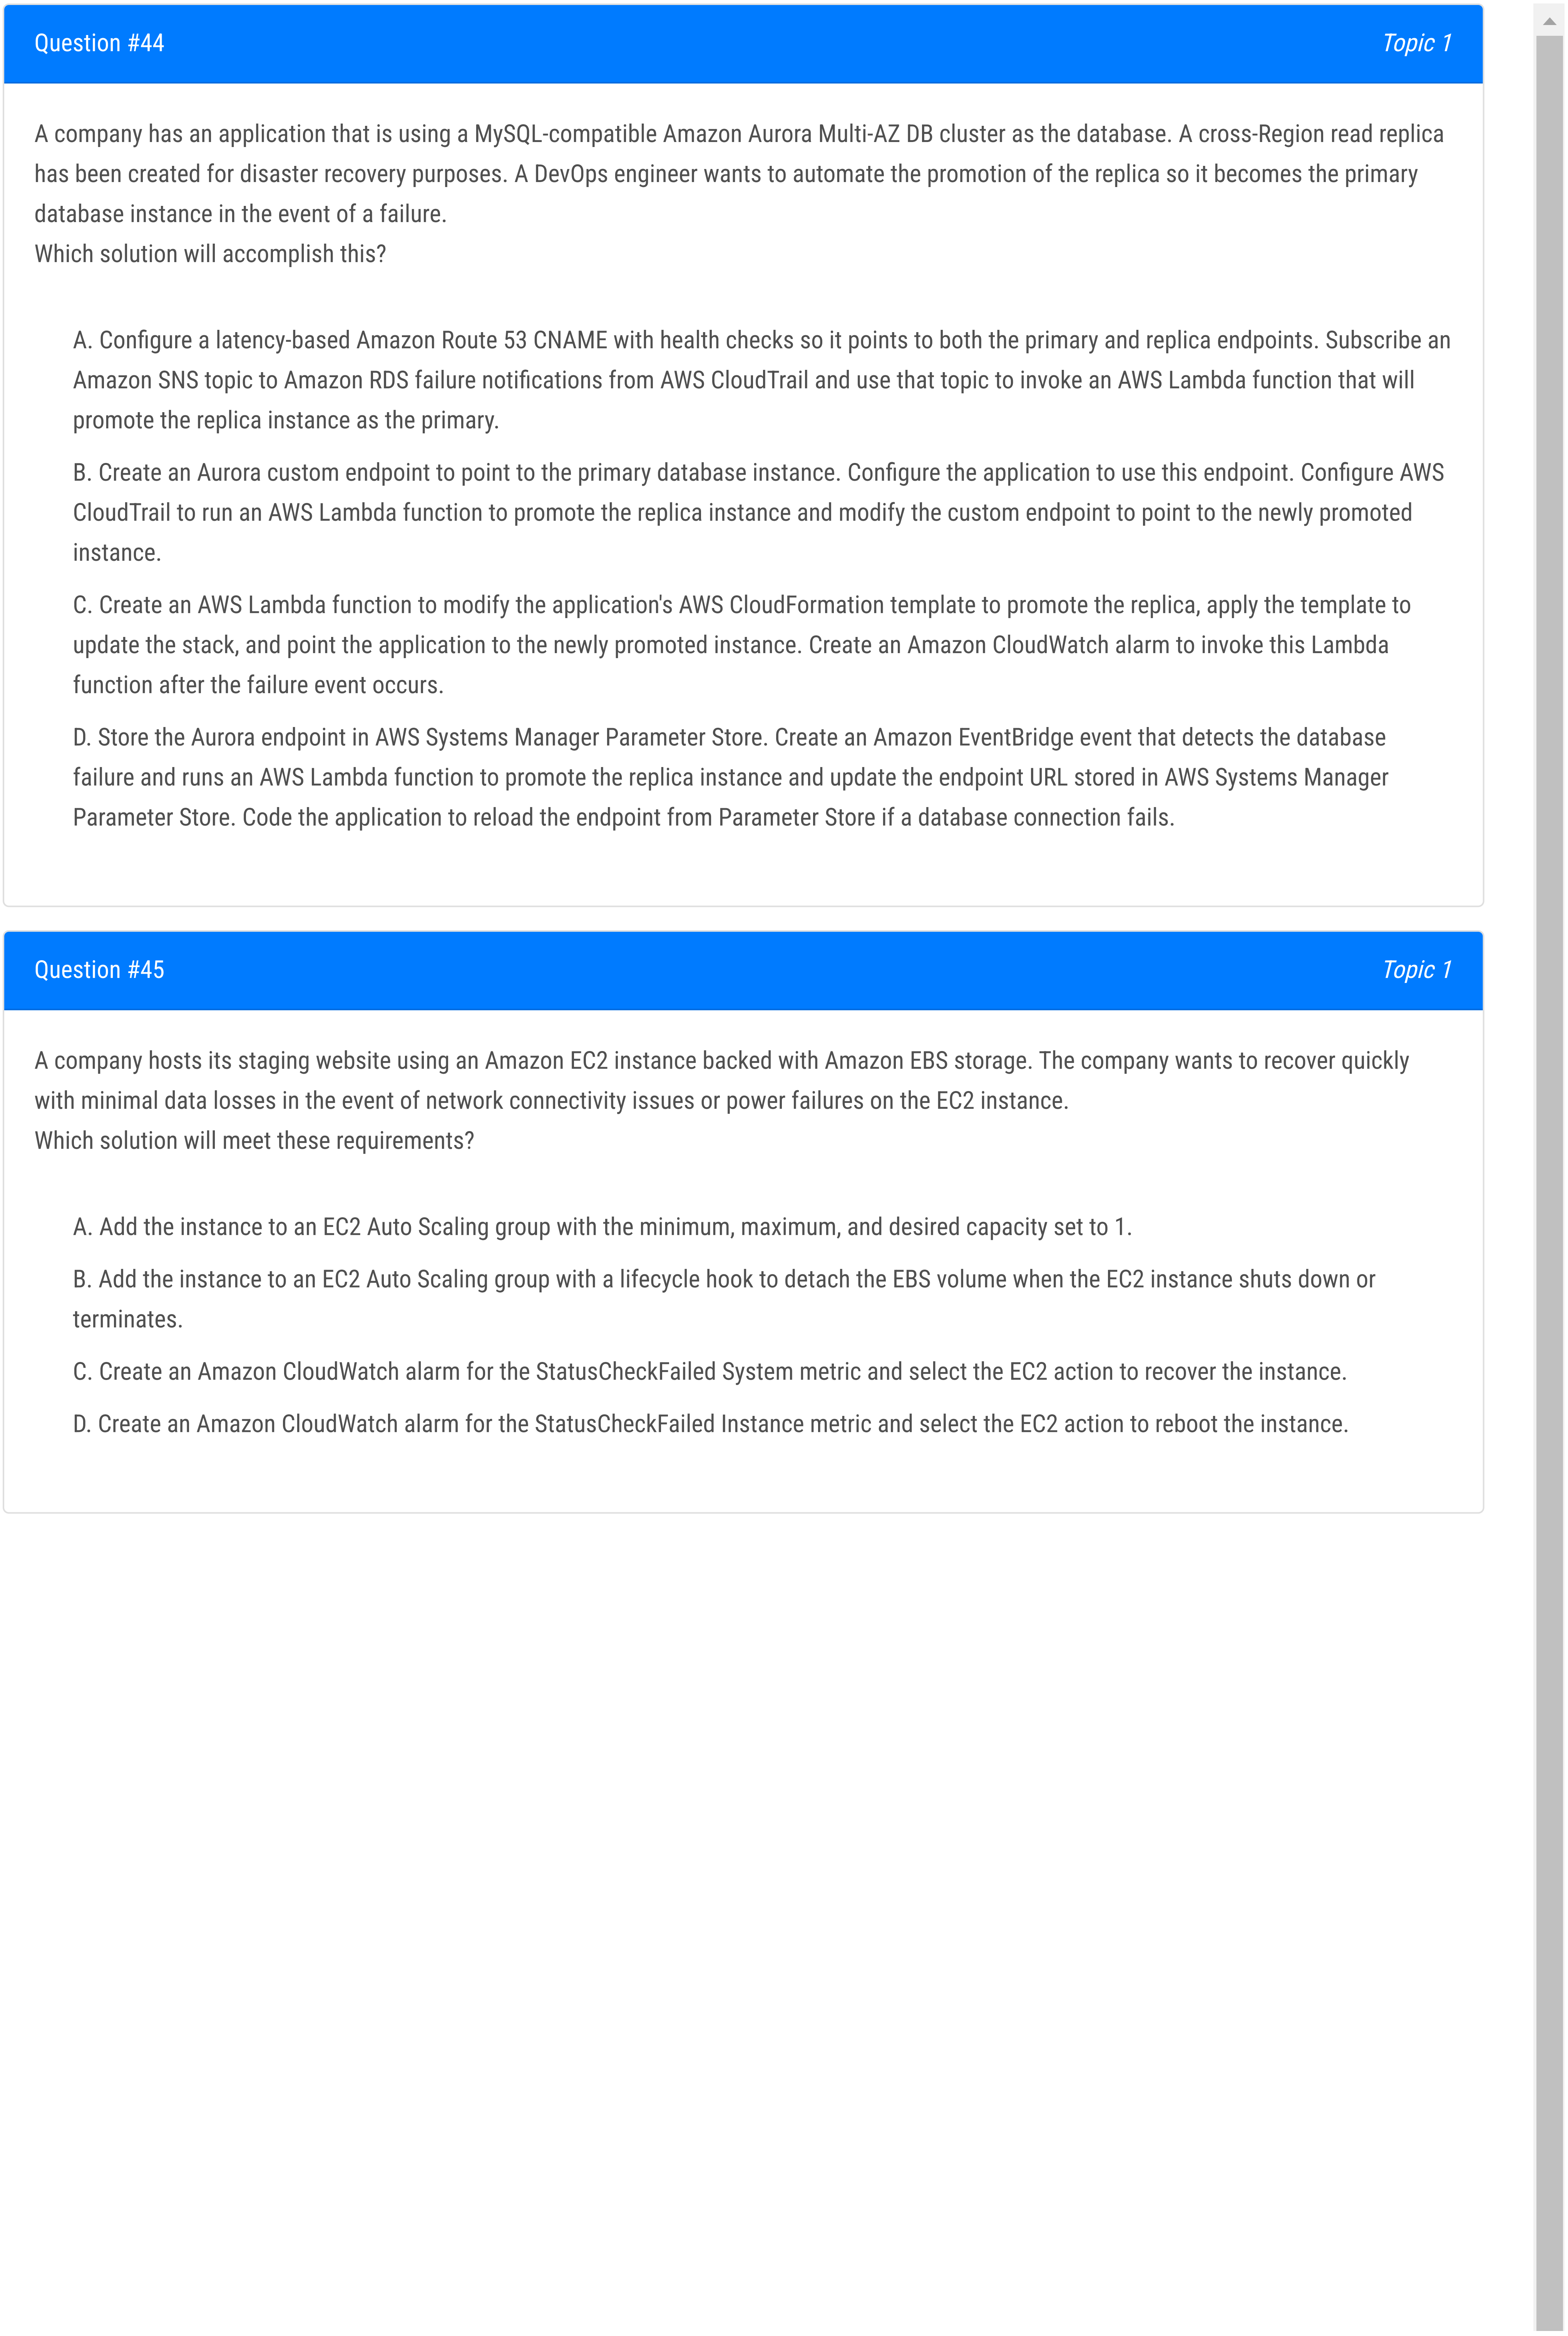

Question #44 *Topic 1*

A company has an application that is using a MySQL-compatible Amazon Aurora Multi-AZ DB cluster as the database. A cross-Region read replica has been created for disaster recovery purposes. A DevOps engineer wants to automate the promotion of the replica so it becomes the primary database instance in the event of a failure.
Which solution will accomplish this?

A. Configure a latency-based Amazon Route 53 CNAME with health checks so it points to both the primary and replica endpoints. Subscribe an Amazon SNS topic to Amazon RDS failure notifications from AWS CloudTrail and use that topic to invoke an AWS Lambda function that will promote the replica instance as the primary.

B. Create an Aurora custom endpoint to point to the primary database instance. Configure the application to use this endpoint. Configure AWS CloudTrail to run an AWS Lambda function to promote the replica instance and modify the custom endpoint to point to the newly promoted instance.

C. Create an AWS Lambda function to modify the application's AWS CloudFormation template to promote the replica, apply the template to update the stack, and point the application to the newly promoted instance. Create an Amazon CloudWatch alarm to invoke this Lambda function after the failure event occurs.

D. Store the Aurora endpoint in AWS Systems Manager Parameter Store. Create an Amazon EventBridge event that detects the database failure and runs an AWS Lambda function to promote the replica instance and update the endpoint URL stored in AWS Systems Manager Parameter Store. Code the application to reload the endpoint from Parameter Store if a database connection fails.

Question #45 *Topic 1*

A company hosts its staging website using an Amazon EC2 instance backed with Amazon EBS storage. The company wants to recover quickly with minimal data losses in the event of network connectivity issues or power failures on the EC2 instance.
Which solution will meet these requirements?

A. Add the instance to an EC2 Auto Scaling group with the minimum, maximum, and desired capacity set to 1.

B. Add the instance to an EC2 Auto Scaling group with a lifecycle hook to detach the EBS volume when the EC2 instance shuts down or terminates.

C. Create an Amazon CloudWatch alarm for the StatusCheckFailed System metric and select the EC2 action to recover the instance.

D. Create an Amazon CloudWatch alarm for the StatusCheckFailed Instance metric and select the EC2 action to reboot the instance.

Question #46 *Topic 1*

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

Which solution will meet these requirements?

A. Use AWS CodeBuild to test the application. Use bash scripts invoked by AWS CodeDeploy's appspec.yml file to restart services, and deregister and register instances with the ALB. Use the appspec.yml file to update file permissions without a custom script.

B. Use AWS CodePipeline to move the application from the AWS CodeCommit repository to AWS CodeDeploy. Use CodeDeploy's deployment group to test the application, unregister and re-register instances with the ALand restart services. Use the appspec.yml file to update file permissions without a custom script.

C. Use AWS CodePipeline to move the application source code from the AWS CodeCommit repository to AWS CodeDeploy. Use CodeDeploy to test the application. Use CodeDeploy's appspec.yml file to restart services and update permissions without a custom script. Use AWS CodeBuild to unregister and re-register instances with the ALB.

D. Use AWS CodePipeline to trigger AWS CodeBuild to test the application. Use bash scripts invoked by AWS CodeDeploy's appspec.yml file to restart services. Unregister and re-register the instances in the AWS CodeDeploy deployment group with the ALB. Update the appspec.yml file to update file permissions without a custom script.

Question #47 *Topic 1*

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

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

A. Add the physical machines into AWS Systems Manager using Systems Manager Hybrid Activations.

B. Attach an IAM role to the EC2 instances, allowing them to be managed by AWS Systems Manager.

C. Create IAM access keys for the on-premises machines to interact with AWS Systems Manager.

D. Run an AWS Systems Manager Automation document to patch the systems every hour

E. Use Amazon EventBridge scheduled events to schedule a patch window.

F. Use AWS Systems Manager Maintenance Windows to schedule a patch window.

Question #48 *Topic 1*

A company has chosen AWS to host a new application. The company needs to implement a multi-account strategy. A DevOps engineer creates a new AWS account and an organization in AWS Organizations. The DevOps engineer also creates the OU structure for the organization and sets up a landing zone by using AWS Control Tower.

The DevOps engineer must implement a solution that automatically deploys resources for new accounts that users create through AWS Control Tower Account Factory. When a user creates a new account, the solution must apply AWS CloudFormation templates and SCPs that are customized for the OU or the account to automatically deploy all the resources that are attached to the account. All the OUs are enrolled in AWS Control Tower.

Which solution will meet these requirements in the MOST automated way?

A. Use AWS Service Catalog with AWS Control Tower. Create portfolios and products in AWS Service Catalog. Grant granular permissions to provision these resources. Deploy SCPs by using the AWS CLI and JSON documents.

B. Deploy CloudFormation stack sets by using the required templates. Enable automatic deployment. Deploy stack instances to the required accounts. Deploy a CloudFormation stack set to the organization's management account to deploy SCPs.

C. Create an Amazon EventBridge rule to detect the CreateManagedAccount event. Configure AWS Service Catalog as the target to deploy resources to any new accounts. Deploy SCPs by using the AWS CLI and JSON documents.

D. Deploy the Customizations for AWS Control Tower (CfCT) solution. Use an AWS CodeCommit repository as the source. In the repository, create a custom package that includes the CloudFormation templates and the SCP JSON documents.

Question #49 *Topic 1*

An online retail company based in the United States plans to expand its operations to Europe and Asia in the next six months. Its product currently runs on Amazon EC2 instances behind an Application Load Balancer. The instances run in an Amazon EC2 Auto Scaling group across multiple Availability Zones. All data is stored in an Amazon Aurora database instance.

When the product is deployed in multiple regions, the company wants a single product catalog across all regions, but for compliance purposes, its customer information and purchases must be kept in each region.

How should the company meet these requirements with the LEAST amount of application changes?

A. Use Amazon Redshift for the product catalog and Amazon DynamoDB tables for the customer information and purchases.

B. Use Amazon DynamoDB global tables for the product catalog and regional tables for the customer information and purchases.

C. Use Aurora with read replicas for the product catalog and additional local Aurora instances in each region for the customer information and purchases.

D. Use Aurora for the product catalog and Amazon DynamoDB global tables for the customer information and purchases.

Question #50 *Topic 1*

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

Amazon API Gateway -

AWS Lambda -

Amazon DynamoDB -
Which solution will meet the requirements?

A. Configure Amazon Route 53 to point to API Gateway APIs in North America and Europe using health checks. Configure the APIs to forward requests to a Lambda function in that Region. Configure the Lambda functions to retrieve and update the data in a DynamoDB table in the same Region as the Lambda function.

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

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

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

Question #51 *Topic 1*

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 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 values. Use CloudFormation StackSets for the development environments, using the Count input parameter to indicate the number of environments needed. Use the UpdateStackSet command to update existing development environments.

B. Use nested stacks to define common infrastructure components. To access the exported values, use TemplateURL to reference the networking team's template. To retrieve Virtual Private Cloud (VPC) and subnet values, use Fn::ImportValue intrinsic functions in the Parameters section of the root template. Use the CreateChangeSet and ExecuteChangeSet commands to update existing development environments.

C. Use nested stacks to define common infrastructure components. Use Fn::ImportValue intrinsic functions with the resources of the nested stack to retrieve Virtual Private Cloud (VPC) and subnet values. Use the CreateChangeSet and ExecuteChangeSet commands to update existing development environments.

D. Use Fn::ImportValue intrinsic functions in the Parameters section of the root template to retrieve Virtual Private Cloud (VPC) and subnet values. Define the development resources in the order they need to be created in the CloudFormation nested stacks. Use the CreateChangeSet. and ExecuteChangeSet commands to update existing development environments.

Question #52 *Topic 1*

A company uses AWS Organizations to manage multiple accounts. Information security policies require that all unencrypted Amazon EBS volumes be marked as non-compliant. A DevOps engineer needs to automatically deploy the solution and ensure that this compliance check is always present.
Which solution will accomplish this?

A. Create an AWS CloudFormation template that defines an AWS Inspector rule to check whether EBS encryption is enabled. Save the template to an Amazon S3 bucket that has been shared with all accounts within the company. Update the account creation script pointing to the CloudFormation template in Amazon S3.

B. Create an AWS Config organizational rule to check whether EBS encryption is enabled and deploy the rule using the AWS CLI. Create and apply an SCP to prohibit stopping and deleting AWS Config across the organization.

C. Create an SCP in Organizations. Set the policy to prevent the launch of Amazon EC2 instances without encryption on the EBS volumes using a conditional expression. Apply the SCP to all AWS accounts. Use Amazon Athena to analyze the AWS CloudTrail output, looking for events that deny an ec2:RunInstances action.

D. Deploy an IAM role to all accounts from a single trusted account. Build a pipeline with AWS CodePipeline with a stage in AWS Lambda to assume the IAM role, and list all EBS volumes in the account. Publish a report to Amazon S3.

Question #53 *Topic 1*

A company is performing vulnerability scanning for all Amazon EC2 instances across many accounts. The accounts are in an organization in AWS Organizations. Each account's VPCs are attached to a shared transit gateway. The VPCs send traffic to the internet through a central egress VPC. The company has enabled Amazon Inspector in a delegated administrator account and has enabled scanning for all member accounts.
A DevOps engineer discovers that some EC2 instances are listed in the "not scanning" tab in Amazon Inspector.
Which combination of actions should the DevOps engineer take to resolve this issue? (Choose three.)

A. Verify that AWS Systems Manager Agent is installed and is running on the EC2 instances that Amazon Inspector is not scanning.

B. Associate the target EC2 instances with security groups that allow outbound communication on port 443 to the AWS Systems Manager service endpoint.

C. Grant inspector:StartAssessmentRun permissions to the IAM role that the DevOps engineer is using.

D. Configure EC2 Instance Connect for the EC2 instances that Amazon Inspector is not scanning.

E. Associate the target EC2 instances with instance profiles that grant permissions to communicate with AWS Systems Manager.

F. Create a managed-instance activation. Use the Activation Code and the Activation ID to register the EC2 instances.

Question #54 *Topic 1*

A development team uses AWS CodeCommit for version control for applications. The development team uses AWS CodePipeline, AWS CodeBuild. and AWS CodeDeploy for CI/CD infrastructure. In CodeCommit, the development team recently merged pull requests that did not pass long-running tests in the code base. The development team needed to perform rollbacks to branches in the codebase, resulting in lost time and wasted effort.

A DevOps engineer must automate testing of pull requests in CodeCommit to ensure that reviewers more easily see the results of automated tests as part of the pull request review.

What should the DevOps engineer do to meet this requirement?

A. Create an Amazon EventBridge rule that reacts to the pullRequestStatusChanged event. Create an AWS Lambda function that invokes a CodePipeline pipeline with a CodeBuild action that runs the tests for the application. Program the Lambda function to post the CodeBuild badge as a comment on the pull request so that developers will see the badge in their code review.

B. Create an Amazon EventBridge rule that reacts to the pullRequestCreated event. Create an AWS Lambda function that invokes a CodePipeline pipeline with a CodeBuild action that runs the tests for the application. Program the Lambda function to post the CodeBuild test results as a comment on the pull request when the test results are complete.

C. Create an Amazon EventBridge rule that reacts to pullRequestCreated and pullRequestSourceBranchUpdated events. Create an AWS Lambda function that invokes a CodePipeline pipeline with a CodeBuild action that runs the tests for the application. Program the Lambda function to post the CodeBuild badge as a comment on the pull request so that developers will see the badge in their code review.

D. Create an Amazon EventBridge rule that reacts to the pullRequestStatusChanged event. Create an AWS Lambda function that invokes a CodePipeline pipeline with a CodeBuild action that runs the tests for the application. Program the Lambda function to post the CodeBuild test results as a comment on the pull request when the test results are complete.

Question #55 *Topic 1*

A company has deployed an application in a production VPC in a single AWS account. The application is popular and is experiencing heavy usage. The company’s security team wants to add additional security, such as AWS WAF, to the application deployment. However, the application's product manager is concerned about cost and does not want to approve the change unless the security team can prove that additional security is necessary.

The security team believes that some of the application's demand might come from users that have IP addresses that are on a deny list. The security team provides the deny list to a DevOps engineer. If any of the IP addresses on the deny list access the application, the security team wants to receive automated notification in near real time so that the security team can document that the application needs additional security.

The DevOps engineer creates a VPC flow log for the production VPC.

Which set of additional steps should the DevOps engineer take to meet these requirements MOST cost-effectively?

A. Create a log group in Amazon CloudWatch Logs. Configure the VPC flow log to capture accepted traffic and to send the data to the log group. Create an Amazon CloudWatch metric filter for IP addresses on the deny list. Create a CloudWatch alarm with the metric filter as input. Set the period to 5 minutes and the datapoints to alarm to 1. Use an Amazon Simple Notification Service (Amazon SNS) topic to send alarm notices to the security team.

B. Create an Amazon S3 bucket for log files. Configure the VPC flow log to capture all traffic and to send the data to the S3 bucket. Configure Amazon Athena to return all log files in the S3 bucket for IP addresses on the deny list. Configure Amazon QuickSight to accept data from Athena and to publish the data as a dashboard that the security team can access. Create a threshold alert of 1 for successful access. Configure the alert to automatically notify the security team as frequently as possible when the alert threshold is met.

C. Create an Amazon S3 bucket for log files. Configure the VPC flow log to capture accepted traffic and to send the data to the S3 bucket. Configure an Amazon OpenSearch Service cluster and domain for the log files. Create an AWS Lambda function to retrieve the logs from the S3 bucket, format the logs, and load the logs into the OpenSearch Service cluster. Schedule the Lambda function to run every 5 minutes. Configure an alert and condition in OpenSearch Service to send alerts to the security team through an Amazon Simple Notification Service (Amazon SNS) topic when access from the IP addresses on the deny list is detected.

D. Create a log group in Amazon CloudWatch Logs. Create an Amazon S3 bucket to hold query results. Configure the VPC flow log to capture all traffic and to send the data to the log group. Deploy an Amazon Athena CloudWatch connector in AWS Lambda. Connect the connector to the log group. Configure Athena to periodically query for all accepted traffic from the IP addresses on the deny list and to store the results in the S3 bucket. Configure an S3 event notification to automatically notify the security team through an Amazon Simple Notification Service (Amazon SNS) topic when new objects are added to the S3 bucket.

Question #56 *Topic 1*

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

1. An AWS CodeBuild project compiles the deployment artifact and runs unit tests.
2. An AWS CodeDeploy deployment group deploys the web service to Amazon EC2 instances in the staging environment.
3. A CodeDeploy deployment group deploys the web service to EC2 instances in the production environment.

The quality assurance (QA) team requests permission to inspect the build artifact before the deployment to the production environment occurs. The QA team wants to run an internal penetration testing tool to conduct manual tests. The tool will be invoked by a REST API call.

Which combination of actions should the DevOps engineer take to fulfill this request? (Choose two.)

A. Insert a manual approval action between the test actions and deployment actions of the pipeline.

B. Modify the buildspec.yml file for the compilation stage to require manual approval before completion.

C. Update the CodeDeploy deployment groups so that they require manual approval to proceed.

D. Update the pipeline to directly call the REST API for the penetration testing tool.

E. Update the pipeline to invoke an AWS Lambda function that calls the REST API for the penetration testing tool.

Question #57 *Topic 1*

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

A. In both regions, deploy the application on AWS Elastic Beanstalk and use Amazon DynamoDB global tables for session data. Use an Amazon Route 53 weighted routing policy with health checks to distribute the traffic across the regions.

B. In both regions, launch the application in Auto Scaling groups and use DynamoDB for session data. Use a Route 53 failover routing policy with health checks to distribute the traffic across the regions.

C. In both regions, deploy the application in AWS Lambda, exposed by Amazon API Gateway, and use Amazon RDS for PostgreSQL with cross-region replication for session data. Deploy the web application with client-side logic to call the API Gateway directly.

D. In both regions, launch the application in Auto Scaling groups and use DynamoDB global tables for session data. Enable an Amazon CloudFront weighted distribution across regions. Point the Amazon Route 53 DNS record at the CloudFront distribution.

Question #58 *Topic 1*

A company runs an application on Amazon EC2 instances. The company uses a series of AWS CloudFormation stacks to define the application resources. A developer performs updates by building and testing the application on a laptop and then uploading the build output and CloudFormation stack templates to Amazon S3. The developer's peers review the changes before the developer performs the CloudFormation stack update and installs a new version of the application onto the EC2 instances.
The deployment process is prone to errors and is time-consuming when the developer updates each EC2 instance with the new application. The company wants to automate as much of the application deployment process as possible while retaining a final manual approval step before the modification of the application or resources.
The company already has moved the source code for the application and the CloudFormation templates to AWS CodeCommit. The company also has created an AWS CodeBuild project to build and test the application.
Which combination of steps will meet the company's requirements? (Choose two.)

A. Create an application group and a deployment group in AWS CodeDeploy. Install the CodeDeploy agent on the EC2 instances.

B. Create an application revision and a deployment group in AWS CodeDeploy. Create an environment in CodeDeploy. Register the EC2 instances to the CodeDeploy environment.

C. Use AWS CodePipeline to invoke the CodeBuild job, run the CloudFormation update, and pause for a manual approval step. After approval, start the AWS CodeDeploy deployment.

D. Use AWS CodePipeline to invoke the CodeBuild job, create CloudFormation change sets for each of the application stacks, and pause for a manual approval step. After approval, run the CloudFormation change sets and start the AWS CodeDeploy deployment.

E. Use AWS CodePipeline to invoke the CodeBuild job, create CloudFormation change sets for each of the application stacks, and pause for a manual approval step. After approval, start the AWS CodeDeploy deployment.

Question #59 *Topic 1*

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

A. Use an AWS CloudFormation template with a retention policy for the ALB set to 1 hour. Update the Amazon Route 53 record to reflect the new ALB.

B. Use two AWS Elastic Beanstalk environments to perform a blue/green deployment from the original environment to the new one. Create an application version lifecycle policy to terminate the original environment in 1 hour.

C. Use AWS CodeDeploy with a deployment group configured with a blue/green deployment configuration Select the option Terminate the original instances in the deployment group with a waiting period of 1 hour.

D. Use AWS Elastic Beanstalk with the configuration set to Immutable. Create an .ebextension using the Resources key that sets the deletion policy of the ALB to 1 hour, and deploy the application.

Question #60 *Topic 1*

A video-sharing company stores its videos in Amazon S3. The company has observed a sudden increase in video access requests, but the company does not know which videos are most popular. The company needs to identify the general access pattern for the video files. This pattern includes the number of users who access a certain file on a given day, as well as the number of pull requests for certain files.
How can the company meet these requirements with the LEAST amount of effort?

A. Activate S3 server access logging. Import the access logs into an Amazon Aurora database. Use an Aurora SQL query to analyze the access patterns.

B. Activate S3 server access logging. Use Amazon Athena to create an external table with the log files. Use Athena to create a SQL query to analyze the access patterns.

C. Invoke an AWS Lambda function for every S3 object access event. Configure the Lambda function to write the file access information, such as user. S3 bucket, and file key, to an Amazon Aurora database. Use an Aurora SQL query to analyze the access patterns.

D. Record an Amazon CloudWatch Logs log message for every S3 object access event. Configure a CloudWatch Logs log stream to write the file access information, such as user, S3 bucket, and file key, to an Amazon Kinesis Data Analytics for SQL application. Perform a sliding window analysis.

Question #61 *Topic 1*

A development team wants to use AWS CloudFormation stacks to deploy an application. However, the developer IAM role does not have the required permissions to provision the resources that are specified in the AWS CloudFormation template. A DevOps engineer needs to implement a solution that allows the developers to deploy the stacks. The solution must follow the principle of least privilege.
Which solution will meet these requirements?

A. Create an IAM policy that allows the developers to provision the required resources. Attach the policy to the developer IAM role.

B. Create an IAM policy that allows full access to AWS CloudFormation. Attach the policy to the developer IAM role.

C. Create an AWS CloudFormation service role that has the required permissions. Grant the developer IAM role a cloudformation:* action. Use the new service role during stack deployments.

D. Create an AWS CloudFormation service role that has the required permissions. Grant the developer IAM role the iam:PassRole permission. Use the new service role during stack deployments.

Question #62 *Topic 1*

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

A. Create a CloudWatch Logs subscription to an AWS Step Functions application. Configure an AWS Lambda function to add a tag to the EC2 instance that produced the login event and mark the instance to be decommissioned. Create an Amazon EventBridge rule to invoke a second Lambda function once a day that will terminate all instances with this tag.

B. Create an Amazon CloudWatch alarm that will be invoked by the login event. Send the notification to an Amazon Simple Notification Service (Amazon SNS) topic that the operations team is subscribed to, and have them terminate the EC2 instance within 24 hours.

C. Create an Amazon CloudWatch alarm that will be invoked by the login event. Configure the alarm to send to an Amazon Simple Queue Service (Amazon SQS) queue. Use a group of worker instances to process messages from the queue, which then schedules an Amazon EventBridge rule to be invoked.

D. Create a CloudWatch Logs subscription to an AWS Lambda function. Configure the function to add a tag to the EC2 instance that produced the login event and mark the instance to be decommissioned. Create an Amazon EventBridge rule to invoke a daily Lambda function that terminates all instances with this tag.

Question #63 *Topic 1*

A company has enabled all features for its organization in AWS Organizations. The organization contains 10 AWS accounts. The company has turned on AWS CloudTrail in all the accounts. The company expects the number of AWS accounts in the organization to increase to 500 during the next year. The company plans to use multiple OUs for these accounts.

The company has enabled AWS Config in each existing AWS account in the organization. A DevOps engineer must implement a solution that enables AWS Config automatically for all future AWS accounts that are created in the organization.

Which solution will meet this requirement?

A. In the organization's management account, create an Amazon EventBridge rule that reacts to a CreateAccount API call. Configure the rule to invoke an AWS Lambda function that enables trusted access to AWS Config for the organization.

B. In the organization's management account, create an AWS CloudFormation stack set to enable AWS Config. Configure the stack set to deploy automatically when an account is created through Organizations.

C. In the organization's management account, create an SCP that allows the appropriate AWS Config API calls to enable AWS Config. Apply the SCP to the root-level OU.

D. In the organization's management account, create an Amazon EventBridge rule that reacts to a CreateAccount API call. Configure the rule to invoke an AWS Systems Manager Automation runbook to enable AWS Config for the account.

Question #64 *Topic 1*

A company has many applications. Different teams in the company developed the applications by using multiple languages and frameworks. The applications run on premises and on different servers with different operating systems. Each team has its own release protocol and process. The company wants to reduce the complexity of the release and maintenance of these applications.

The company is migrating its technology stacks, including these applications, to AWS. The company wants centralized control of source code, a consistent and automatic delivery pipeline, and as few maintenance tasks as possible on the underlying infrastructure.

What should a DevOps engineer do to meet these requirements?

A. Create one AWS CodeCommit repository for all applications. Put each application's code in a different branch. Merge the branches, and use AWS CodeBuild to build the applications. Use AWS CodeDeploy to deploy the applications to one centralized application server.

B. Create one AWS CodeCommit repository for each of the applications. Use AWS CodeBuild to build the applications one at a time. Use AWS CodeDeploy to deploy the applications to one centralized application server.

C. Create one AWS CodeCommit repository for each of the applications. Use AWS CodeBuild to build the applications one at a time and to create one AMI for each server. Use AWS CloudFormation StackSets to automatically provision and decommission Amazon EC2 fleets by using these AMIs.

D. Create one AWS CodeCommit repository for each of the applications. Use AWS CodeBuild to build one Docker image for each application in Amazon Elastic Container Registry (Amazon ECR). Use AWS CodeDeploy to deploy the applications to Amazon Elastic Container Service (Amazon ECS) on infrastructure that AWS Fargate manages.

Question #65 *Topic 1*

A company's application is currently deployed to a single AWS Region. Recently, the company opened a new office on a different continent. The users in the new office are experiencing high latency. The company's application runs on Amazon EC2 instances behind an Application Load Balancer (ALB) and uses Amazon DynamoDB as the database layer. The instances run in an EC2 Auto Scaling group across multiple Availability Zones. A DevOps engineer is tasked with minimizing application response times and improving availability for users in both Regions.
Which combination of actions should be taken to address the latency issues? (Choose three.)

A. Create a new DynamoDB table in the new Region with cross-Region replication enabled.

B. Create new ALB and Auto Scaling group global resources and configure the new ALB to direct traffic to the new Auto Scaling group.

C. Create new ALB and Auto Scaling group resources in the new Region and configure the new ALB to direct traffic to the new Auto Scaling group.

D. Create Amazon Route 53 records, health checks, and latency-based routing policies to route to the ALB.

E. Create Amazon Route 53 aliases, health checks, and failover routing policies to route to the ALB.

F. Convert the DynamoDB table to a global table.

Question #66 *Topic 1*

A DevOps engineer needs to apply a core set of security controls to an existing set of AWS accounts. The accounts are in an organization in AWS Organizations. Individual teams will administer individual accounts by using the AdministratorAccess AWS managed policy. For all accounts. AWS CloudTrail and AWS Config must be turned on in all available AWS Regions. Individual account administrators must not be able to edit or delete any of the baseline resources. However, individual account administrators must be able to edit or delete their own CloudTrail trails and AWS Config rules.
Which solution will meet these requirements in the MOST operationally efficient way?

A. Create an AWS CloudFormation template that defines the standard account resources. Deploy the template to all accounts from the organization's management account by using CloudFormation StackSets. Set the stack policy to deny Update:Delete actions.

B. Enable AWS Control Tower. Enroll the existing accounts in AWS Control Tower. Grant the individual account administrators access to CloudTrail and AWS Config.

C. Designate an AWS Config management account. Create AWS Config recorders in all accounts by using AWS CloudFormation StackSets. Deploy AWS Config rules to the organization by using the AWS Config management account. Create a CloudTrail organization trail in the organization's management account. Deny modification or deletion of the AWS Config recorders by using an SCP.

D. Create an AWS CloudFormation template that defines the standard account resources. Deploy the template to all accounts from the organization's management account by using Cloud Formation StackSets Create an SCP that prevents updates or deletions to CloudTrail resources or AWS Config resources unless the principal is an administrator of the organization's management account.

Question #67 *Topic 1*

A company has its AWS accounts in an organization in AWS Organizations. AWS Config is manually configured in each AWS account. The company needs to implement a solution to centrally configure AWS Config for all accounts in the organization The solution also must record resource changes to a central account.
Which combination of actions should a DevOps engineer perform to meet these requirements? (Choose two.)

A. Configure a delegated administrator account for AWS Config. Enable trusted access for AWS Config in the organization.

B. Configure a delegated administrator account for AWS Config. Create a service-linked role for AWS Config in the organization’s management account.

C. Create an AWS CloudFormation template to create an AWS Config aggregator. Configure a CloudFormation stack set to deploy the template to all accounts in the organization.

D. Create an AWS Config organization aggregator in the organization's management account. Configure data collection from all AWS accounts in the organization and from all AWS Regions.

E. Create an AWS Config organization aggregator in the delegated administrator account. Configure data collection from all AWS accounts in the organization and from all AWS Regions.

Question #68 *Topic 1*

A company wants to migrate its content sharing web application hosted on Amazon EC2 to a serverless architecture. The company currently deploys changes to its application by creating a new Auto Scaling group of EC2 instances and a new Elastic Load Balancer, and then shifting the traffic away using an Amazon Route 53 weighted routing policy.
For its new serverless application, the company is planning to use Amazon API Gateway and AWS Lambda. The company will need to update its deployment processes to work with the new application. It will also need to retain the ability to test new features on a small number of users before rolling the features out to the entire user base.
Which deployment strategy will meet these requirements?

A. Use AWS CDK to deploy API Gateway and Lambda functions. When code needs to be changed, update the AWS CloudFormation stack and deploy the new version of the APIs and Lambda functions. Use a Route 53 failover routing policy for the canary release strategy.

B. Use AWS CloudFormation to deploy API Gateway and Lambda functions using Lambda function versions. When code needs to be changed, update the CloudFormation stack with the new Lambda code and update the API versions using a canary release strategy. Promote the new version when testing is complete.

C. Use AWS Elastic Beanstalk to deploy API Gateway and Lambda functions. When code needs to be changed, deploy a new version of the API and Lambda functions. Shift traffic gradually using an Elastic Beanstalk blue/green deployment.

D. Use AWS OpsWorks to deploy API Gateway in the service layer and Lambda functions in a custom layer. When code needs to be changed, use OpsWorks to perform a blue/green deployment and shift traffic gradually.

Question #69 *Topic 1*

A development team uses AWS CodeCommit, AWS CodePipeline, and AWS CodeBuild to develop and deploy an application. Changes to the code are submitted by pull requests. The development team reviews and merges the pull requests, and then the pipeline builds and tests the application.

Over time, the number of pull requests has increased. The pipeline is frequently blocked because of failing tests. To prevent this blockage, the development team wants to run the unit and integration tests on each pull request before it is merged.

Which solution will meet these requirements?

A. Create a CodeBuild project to run the unit and integration tests. Create a CodeCommit approval rule template. Configure the template to require the successful invocation of the CodeBuild project. Attach the approval rule to the project's CodeCommit repository.

B. Create an Amazon EventBridge rule to match pullRequestCreated events from CodeCommit Create a CodeBuild project to run the unit and integration tests. Configure the CodeBuild project as a target of the EventBridge rule that includes a custom event payload with the CodeCommit repository and branch information from the event.

C. Create an Amazon EventBridge rule to match pullRequestCreated events from CodeCommit. Modify the existing CodePipeline pipeline to not run the deploy steps if the build is started from a pull request. Configure the EventBridge rule to run the pipeline with a custom payload that contains the CodeCommit repository and branch information from the event.

D. Create a CodeBuild project to run the unit and integration tests. Create a CodeCommit notification rule that matches when a pull request is created or updated. Configure the notification rule to invoke the CodeBuild project.

Question #70 *Topic 1*

A company has an application that runs on a fleet of Amazon EC2 instances. The application requires frequent restarts. The application logs contain error messages when a restart is required. The application logs are published to a log group in Amazon CloudWatch Logs.

An Amazon CloudWatch alarm notifies an application engineer through an Amazon Simple Notification Service (Amazon SNS) topic when the logs contain a large number of restart-related error messages. The application engineer manually restarts the application on the instances after the application engineer receives a notification from the SNS topic.

A DevOps engineer needs to implement a solution to automate the application restart on the instances without restarting the instances.

Which solution will meet these requirements in the MOST operationally efficient manner?

A. Configure an AWS Systems Manager Automation runbook that runs a script to restart the application on the instances. Configure the SNS topic to invoke the runbook.

B. Create an AWS Lambda function that restarts the application on the instances. Configure the Lambda function as an event destination of the SNS topic.

C. Configure an AWS Systems Manager Automation runbook that runs a script to restart the application on the instances. Create an AWS Lambda function to invoke the runbook. Configure the Lambda function as an event destination of the SNS topic.

D. Configure an AWS Systems Manager Automation runbook that runs a script to restart the application on the instances. Configure an Amazon EventBridge rule that reacts when the CloudWatch alarm enters ALARM state. Specify the runbook as a target of the rule.

Question #71 *Topic 1*

A DevOps engineer at a company is supporting an AWS environment in which all users use AWS IAM Identity Center (AWS Single Sign-On). The company wants to immediately disable credentials of any new IAM user and wants the security team to receive a notification.
Which combination of steps should the DevOps engineer take to meet these requirements? (Choose three.)

A. Create an Amazon EventBridge rule that reacts to an IAM CreateUser API call in AWS CloudTrail.

B. Create an Amazon EventBridge rule that reacts to an IAM GetLoginProfile API call in AWS CloudTrail.

C. Create an AWS Lambda function that is a target of the EventBridge rule. Configure the Lambda function to disable any access keys and delete the login profiles that are associated with the IAM user.

D. Create an AWS Lambda function that is a target of the EventBridge rule. Configure the Lambda function to delete the login profiles that are associated with the IAM user.

E. Create an Amazon Simple Notification Service (Amazon SNS) topic that is a target of the EventBridge rule. Subscribe the security team's group email address to the topic.

F. Create an Amazon Simple Queue Service (Amazon SQS) queue that is a target of the Lambda function. Subscribe the security team's group email address to the queue.

Question #72 *Topic 1*

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

A. Use AWS CodePipeline with Amazon ECS. Amazon EC2, and Lambda as deploy providers.

B. Use AWS CodePipeline with AWS CodeDeploy as the deploy provider.

C. Use AWS CodePipeline with AWS Elastic Beanstalk as the deploy provider.

D. Use AWS CodeDeploy with GitHub integration to deploy the application.

Question #73 *Topic 1*

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

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

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

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

B. Increase the maximum instance count of the Auto Scaling group. Configure an autoscaling:EC2_INSTANCE_LAUNCHING lifecycle hook on the Auto Scaling group. Modify the application to complete the lifecycle hook when the application is ready to serve requests.

C. Configure a warm pool for the Auto Scaling group with warmed EC2 instances in the Running state. Configure an autoscaling:EC2_INSTANCE_LAUNCHING lifecycle hook on the Auto Scaling group. Modify the application to complete the lifecycle hook when the application is ready to serve requests.

D. Increase the maximum instance count of the Auto Scaling group. Configure an autoscaling:EC2_INSTANCE_LAUNCHING lifecycle hook on the Auto Scaling group. Modify the application to complete the lifecycle hook and to place the new instance in the Standby state when the application is ready to serve requests.

Question #74 *Topic 1*

A company is using an AWS CodeBuild project to build and package an application. The packages are copied to a shared Amazon S3 bucket before being deployed across multiple AWS accounts.

The buildspec.yml file contains the following:

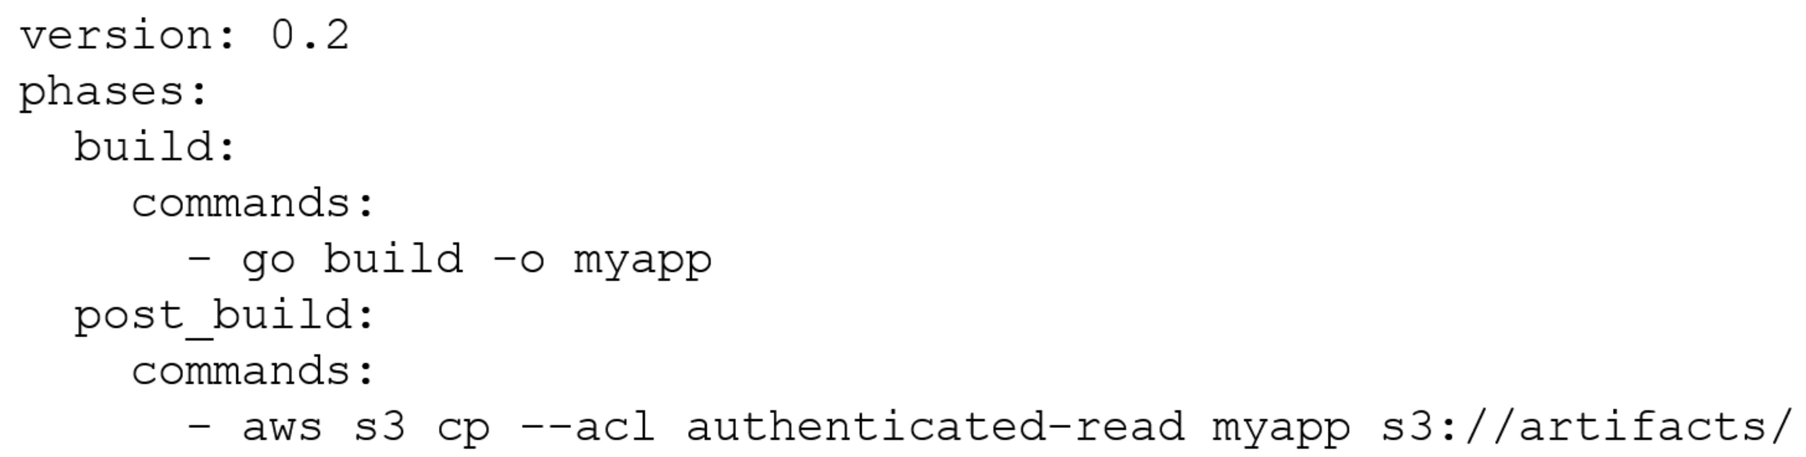

```
version: 0.2
phases:
  build:
    commands:
      - go build -o myapp
  post_build:
    commands:
      - aws s3 cp --acl authenticated-read myapp s3://artifacts/
```

The DevOps engineer has noticed that anybody with an AWS account is able to download the artifacts.

What steps should the DevOps engineer take to stop this?

A. Modify the post_build command to use --acl public-read and configure a bucket policy that grants read access to the relevant AWS accounts only.

B. Configure a default ACL for the S3 bucket that defines the set of authenticated users as the relevant AWS accounts only and grants read-only access.

C. Create an S3 bucket policy that grants read access to the relevant AWS accounts and denies read access to the principal “*”.

D. Modify the post_build command to remove --acl authenticated-read and configure a bucket policy that allows read access to the relevant AWS accounts only.

Question #75 *Topic 1*

A company has developed a serverless web application that is hosted on AWS. The application consists of Amazon S3. Amazon API Gateway, several AWS Lambda functions, and an Amazon RDS for MySQL database. The company is using AWS CodeCommit to store the source code. The source code is a combination of AWS Serverless Application Model (AWS SAM) templates and Python code.

A security audit and penetration test reveal that user names and passwords for authentication to the database are hardcoded within CodeCommit repositories. A DevOps engineer must implement a solution to automatically detect and prevent hardcoded secrets.

What is the MOST secure solution that meets these requirements?

A. Enable Amazon CodeGuru Profiler. Decorate the handler function with @with_lambda_profiler(). Manually review the recommendation report. Write the secret to AWS Systems Manager Parameter Store as a secure string. Update the SAM templates and the Python code to pull the secret from Parameter Store.

B. Associate the CodeCommit repository with Amazon CodeGuru Reviewer. Manually check the code review for any recommendations. Choose the option to protect the secret. Update the SAM templates and the Python code to pull the secret from AWS Secrets Manager.

C. Enable Amazon CodeGuru Profiler. Decorate the handler function with @with_lambda_profiler(). Manually review the recommendation report. Choose the option to protect the secret. Update the SAM templates and the Python code to pull the secret from AWS Secrets Manager.

D. Associate the CodeCommit repository with Amazon CodeGuru Reviewer. Manually check the code review for any recommendations. Write the secret to AWS Systems Manager Parameter Store as a string. Update the SAM templates and the Python code to pull the secret from Parameter Store.

Question #76 *Topic 1*

A company is using Amazon S3 buckets to store important documents. The company discovers that some S3 buckets are not encrypted. Currently, the company's IAM users can create new S3 buckets without encryption. The company is implementing a new requirement that all S3 buckets must be encrypted.

A DevOps engineer must implement a solution to ensure that server-side encryption is enabled on all existing S3 buckets and all new S3 buckets. The encryption must be enabled on new S3 buckets as soon as the S3 buckets are created. The default encryption type must be 256-bit Advanced Encryption Standard (AES-256).

Which solution will meet these requirements?

A. Create an AWS Lambda function that is invoked periodically by an Amazon EventBridge scheduled rule. Program the Lambda function to scan all current S3 buckets for encryption status and to set AES-256 as the default encryption for any S3 bucket that does not have an encryption configuration.

B. Set up and activate the s3-bucket-server-side-encryption-enabled AWS Config managed rule. Configure the rule to use the AWS-EnableS3BucketEncryption AWS Systems Manager Automation runbook as the remediation action. Manually run the re-evaluation process to ensure that existing S3 buckets are compliant.

C. Create an AWS Lambda function that is invoked by an Amazon EventBridge event rule. Define the rule with an event pattern that matches the creation of new S3 buckets. Program the Lambda function to parse the EventBridge event, check the configuration of the S3 buckets from the event, and set AES-256 as the default encryption.

D. Configure an IAM policy that denies the s3:CreateBucket action if the s3:x-amz-server-side-encryption condition key has a value that is not AES-256. Create an IAM group for all the company's IAM users. Associate the IAM policy with the IAM group.

Question #77 *Topic 1*

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

Which architecture will meet these requirements with the LEAST amount of configuration?

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

B. Create a single AWS CodePipeline pipeline that deploys the application using a single AWS CodeDeploy application and single deployment group.

C. Create a single AWS CodePipeline pipeline that deploys the application in parallel using a single AWS CodeDeploy application and unique deployment group for each ALB-Auto Scaling group pair.

D. Create an AWS CodePipeline pipeline for each ALB-Auto Scaling group pair that deploys the application using an AWS CodeDeploy application and deployment group created for the same ALB-Auto Scaling group pair.

Question #78 *Topic 1*

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

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

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

A. Delete all objects, including previous versions, from the S3 bucket that contains the static website content.

B. Update the weighted DNS record entry that points to the S3 bucket. Apply a weight of 0. Specify the domain reset option to propagate changes immediately.

C. Configure webpage redirect requests on the S3 bucket with a hostname that redirects to the ALB.

D. Remove the weighted DNS record entry that points to the S3 bucket from the example.com hosted zone. Wait for DNS propagation to become complete.

Question #79 *Topic 1*

A company is implementing AWS CodePipeline to automate its testing process. The company wants to be notified when the execution state fails and used the following custom event pattern in Amazon EventBridge:

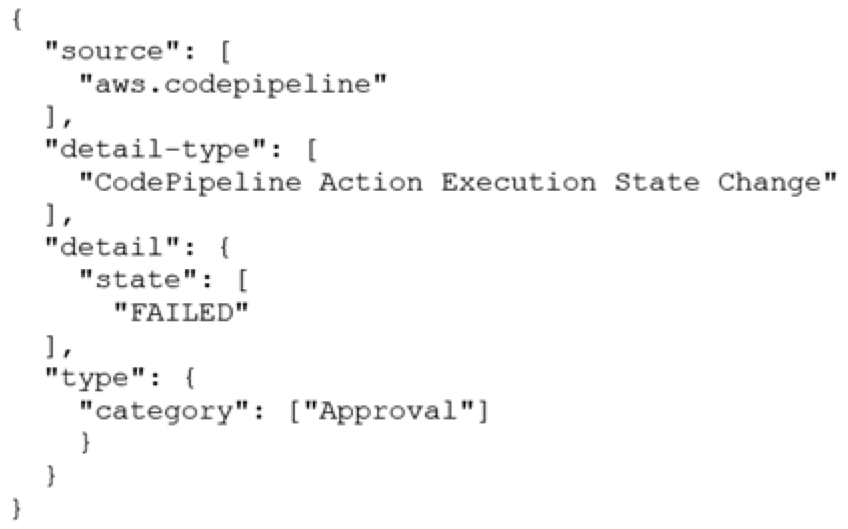

```
{
  "source": [
    "aws.codepipeline"
  ],
  "detail-type": [
    "CodePipeline Action Execution State Change"
  ],
  "detail": {
    "state": [
      "FAILED"
  ],
  "type": {
    "category": ["Approval"]
    }
  }
}
```

Which type of events will match this event pattern?

A. Failed deploy and build actions across all the pipelines

B. All rejected or failed approval actions across all the pipelines

C. All the events across all pipelines

D. Approval actions across all the pipelines

Question #80 *Topic 1*

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. Place the configuration file content in the rule's InputParameters property, and set the Scope property to the EC2 Auto Scaling group. Add an AWS Systems Manager Resource Data Sync resource to the template to poll for updates to the configuration.

B. In the CloudFormation template, add an EC2 launch template resource. Place the configuration file content in the launch template. 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.

C. In the CloudFormation template, add an EC2 launch template resource. Place the configuration file content in the launch template. 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 CloudFormation init metadata. Place the configuration file content in the metadata. 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.

Question #81 *Topic 1*

A company manages an application that stores logs in Amazon CloudWatch Logs. The company wants to archive the logs to an Amazon S3 bucket. Logs are rarely accessed after 90 days and must be retained for 10 years.

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

A. Configure a CloudWatch Logs subscription filter to use AWS Glue to transfer all logs to an S3 bucket.

B. Configure a CloudWatch Logs subscription filter to use Amazon Kinesis Data Firehose to stream all logs to an S3 bucket.

C. Configure a CloudWatch Logs subscription filter to stream all logs to an S3 bucket.

D. Configure the S3 bucket lifecycle policy to transition logs to S3 Glacier after 90 days and to expire logs after 3.650 days.

E. Configure the S3 bucket lifecycle policy to transition logs to Reduced Redundancy after 90 days and to expire logs after 3.650 days.

Question #82 *Topic 1*

A company is developing a new application. The application uses AWS Lambda functions for its compute tier. The company must use a canary deployment for any changes to the Lambda functions. Automated rollback must occur if any failures are reported.

The company's DevOps team needs to create the infrastructure as code (IaC) and the CI/CD pipeline for this solution.

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

A. Create an AWS CloudFormation template for the application. Define each Lambda function in the template by using the AWS::Lambda::Function resource type. In the template, include a version for the Lambda function by using the AWS::Lambda::Version resource type. Declare the CodeSha256 property. Configure an AWS::Lambda::Alias resource that references the latest version of the Lambda function.

B. Create an AWS Serverless Application Model (AWS SAM) template for the application. Define each Lambda function in the template by using the AWS::Serverless::Function resource type. For each function, include configurations for the AutoPublishAlias property and the DeploymentPreference property. Configure the deployment configuration type to LambdaCanary10Percent10Minutes.

C. Create an AWS CodeCommit repository. Create an AWS CodePipeline pipeline. Use the CodeCommit repository in a new source stage that starts the pipeline. Create an AWS CodeBuild project to deploy the AWS Serverless Application Model (AWS SAM) template. Upload the template and source code to the CodeCommit repository. In the CodeCommit repository, create a buildspec.yml file that includes the commands to build and deploy the SAM application.

D. Create an AWS CodeCommit repository. Create an AWS CodePipeline pipeline. Use the CodeCommit repository in a new source stage that starts the pipeline. Create an AWS CodeDeploy deployment group that is configured for canary deployments with a DeploymentPreference type of Canary10Percent10Minutes. Upload the AWS CloudFormation template and source code to the CodeCommit repository. In the CodeCommit repository, create an appspec.yml file that includes the commands to deploy the CloudFormation template.

E. Create an Amazon CloudWatch composite alarm for all the Lambda functions. Configure an evaluation period and dimensions for Lambda. Configure the alarm to enter the ALARM state if any errors are detected or if there is insufficient data.

F. Create an Amazon CloudWatch alarm for each Lambda function. Configure the alarms to enter the ALARM state if any errors are detected. Configure an evaluation period, dimensions for each Lambda function and version, and the namespace as AWS/Lambda on the Errors metric.

Question #83 *Topic 1*

A DevOps engineer is deploying a new version of a company's application in an AWS CodeDeploy deployment group associated with its Amazon EC2 instances. After some time, the deployment fails. The engineer realizes that all the events associated with the specific deployment ID are in a Skipped status, and code was not deployed in the instances associated with the deployment group.

What are valid reasons for this failure? (Choose two.)

A. The networking configuration does not allow the EC2 instances to reach the internet via a NAT gateway or internet gateway, and the CodeDeploy endpoint cannot be reached.

B. The IAM user who triggered the application deployment does not have permission to interact with the CodeDeploy endpoint.

C. The target EC2 instances were not properly registered with the CodeDeploy endpoint.

D. An instance profile with proper permissions was not attached to the target EC2 instances.

E. The appspec.yml file was not included in the application revision.

Question #84 *Topic 1*

A company has a guideline that every Amazon EC2 instance must be launched from an AMI that the company's security team produces. Every month, the security team sends an email message with the latest approved AMIs to all the development teams.

The development teams use AWS CloudFormation to deploy their applications. When developers launch a new service, they have to search their email for the latest AMIs that the security department sent. A DevOps engineer wants to automate the process that the security team uses to provide the AMI IDs to the development teams.

What is the MOST scalable solution that meets these requirements?

A. Direct the security team to use CloudFormation to create new versions of the AMIs and to list the AMI ARNs in an encrypted Amazon S3 object as part of the stack's Outputs section. Instruct the developers to use a cross-stack reference to load the encrypted S3 object and obtain the most recent AMI ARNs.

B. Direct the security team to use a CloudFormation stack to create an AWS CodePipeline pipeline that builds new AMIs and places the latest AMI ARNs in an encrypted Amazon S3 object as part of the pipeline output. Instruct the developers to use a cross-stack reference within their own CloudFormation template to obtain the S3 object location and the most recent AMI ARNs.

C. Direct the security team to use Amazon EC2 Image Builder to create new AMIs and to place the AMI ARNs as parameters in AWS Systems Manager Parameter Store. Instruct the developers to specify a parameter of type SSM in their CloudFormation stack to obtain the most recent AMI ARNs from Parameter Store.

D. Direct the security team to use Amazon EC2 Image Builder to create new AMIs and to create an Amazon Simple Notification Service (Amazon SNS) topic so that every development team can receive notifications. When the development teams receive a notification, instruct them to write an AWS Lambda function that will update their CloudFormation stack with the most recent AMI ARNs.

Question #85 *Topic 1*

An application runs on Amazon EC2 instances behind an Application Load Balancer (ALB). A DevOps engineer is using AWS CodeDeploy to release a new version. The deployment fails during the AllowTraffic lifecycle event, but a cause for the failure is not indicated in the deployment logs.

What would cause this?

A. The appspec.yml file contains an invalid script that runs in the AllowTraffic lifecycle hook.

B. The user who initiated the deployment does not have the necessary permissions to interact with the ALB.

C. The health checks specified for the ALB target group are misconfigured.

D. The CodeDeploy agent was not installed in the EC2 instances that are part of the ALB target group.

Question #86 *Topic 1*

A company has 20 service teams. Each service team is responsible for its own microservice. Each service team uses a separate AWS account for its microservice and a VPC with the 192.168.0.0/22 CIDR block. The company manages the AWS accounts with AWS Organizations.

Each service team hosts its microservice on multiple Amazon EC2 instances behind an Application Load Balancer. The microservices communicate with each other across the public internet. The company's security team has issued a new guideline that all communication between microservices must use HTTPS over private network connections and cannot traverse the public internet.

A DevOps engineer must implement a solution that fulfills these obligations and minimizes the number of changes for each service team.

Which solution will meet these requirements?

A. Create a new AWS account in AWS Organizations. Create a VPC in this account, and use AWS Resource Access Manager to share the private subnets of this VPC with the organization. Instruct the service teams to launch a new Network Load Balancer (NLB) and EC2 instances that use the shared private subnets. Use the NLB DNS names for communication between microservices.

B. Create a Network Load Balancer (NLB) in each of the microservice VPCs. Use AWS PrivateLink to create VPC endpoints in each AWS account for the NLBs. Create subscriptions to each VPC endpoint in each of the other AWS accounts. Use the VPC endpoint DNS names for communication between microservices.

C. Create a Network Load Balancer (NLB) in each of the microservice VPCs. Create VPC peering connections between each of the microservice VPCs. Update the route tables for each VPC to use the peering links. Use the NLB DNS names for communication between microservices.

D. Create a new AWS account in AWS Organizations. Create a transit gateway in this account, and use AWS Resource Access Manager to share the transit gateway with the organization. In each of the microservice VPCs, create a transit gateway attachment to the shared transit gateway. Update the route tables of each VPC to use the transit gateway. Create a Network Load Balancer (NLB) in each of the microservice VPCs. Use the NLB DNS names for communication between microservices.

Question #87 *Topic 1*

An Amazon EC2 instance is running in a VPC and needs to download an object from a restricted Amazon S3 bucket. When the DevOps engineer tries to download the object, an AccessDenied error is received.

What are the possible causes for this error? (Choose two.)

A. The S3 bucket default encryption is enabled.

B. There is an error in the S3 bucket policy.

C. The object has been moved to S3 Glacier.

D. There is an error in the IAM role configuration.

E. S3 Versioning is enabled.

Question #88 *Topic 1*

A company wants to use a grid system for a proprietary enterprise in-memory data store on top of AWS. This system can run in multiple server nodes in any Linux-based distribution. The system must be able to reconfigure the entire cluster every time a node is added or removed. When adding or removing nodes, an /etc/cluster/nodes.config file must be updated, listing the IP addresses of the current node members of that cluster.

The company wants to automate the task of adding new nodes to a cluster.

What can a DevOps engineer do to meet these requirements?

A. Use AWS OpsWorks Stacks to layer the server nodes of that cluster. Create a Chef recipe that populates the content of the /etc/cluster/nodes.config file and restarts the service by using the current members of the layer. Assign that recipe to the Configure lifecycle event.

B. Put the file nodes.config in version control. Create an AWS CodeDeploy deployment configuration and deployment group based on an Amazon EC2 tag value for the cluster nodes. When adding a new node to the cluster, update the file with all tagged instances, and make a commit in version control. Deploy the new file and restart the services.

C. Create an Amazon S3 bucket and upload a version of the /etc/cluster/nodes.config file. Create a crontab script that will poll for that S3 file and download it frequently. Use a process manager, such as Monit or systemd, to restart the cluster services when it detects that the new file was modified. When adding a node to the cluster, edit the file's most recent members. Upload the new file to the S3 bucket.

D. Create a user data script that lists all members of the current security group of the cluster and automatically updates the /etc/cluster/nodes.config file whenever a new instance is added to the cluster.

Question #89 *Topic 1*

A DevOps engineer is working on a data archival project that requires the migration of on-premises data to an Amazon S3 bucket. The DevOps engineer develops a script that incrementally archives on-premises data that is older than 1 month to Amazon S3. Data that is transferred to Amazon S3 is deleted from the on-premises location. The script uses the S3 PutObject operation.

During a code review, the DevOps engineer notices that the script does not verify whether the data was successfully copied to Amazon S3. The DevOps engineer must update the script to ensure that data is not corrupted during transmission. The script must use MD5 checksums to verify data integrity before the on-premises data is deleted.

Which solutions for the script will meet these requirements? (Choose two.)

A. Check the returned response for the VersionId. Compare the returned VersionId against the MD5 checksum.

B. Include the MD5 checksum within the Content-MD5 parameter. Check the operation call's return status to find out if an error was returned.

C. Include the checksum digest within the tagging parameter as a URL query parameter.

D. Check the returned response for the ETag. Compare the returned ETag against the MD5 checksum.

E. Include the checksum digest within the Metadata parameter as a name-value pair. After upload, use the S3 HeadObject operation to retrieve metadata from the object.

Question #90 *Topic 1*

A company deploys updates to its Amazon API Gateway API several times a week by using an AWS CodePipeline pipeline. As part of the update process, the company exports the JavaScript SDK for the API from the API Gateway console and uploads the SDK to an Amazon S3 bucket.

The company has configured an Amazon CloudFront distribution that uses the S3 bucket as an origin. Web clients then download the SDK by using the CloudFront distribution's endpoint. A DevOps engineer needs to implement a solution to make the new SDK available automatically during new API deployments.

Which solution will meet these requirements?

A. Create a CodePipeline action immediately after the deployment stage of the API. Configure the action to invoke an AWS Lambda function. Configure the Lambda function to download the SDK from API Gateway, upload the SDK to the S3 bucket, and create a CloudFront invalidation for the SDK path.

B. Create a CodePipeline action immediately after the deployment stage of the API. Configure the action to use the CodePipeline integration with API Gateway to export the SDK to Amazon S3. Create another action that uses the CodePipeline integration with Amazon S3 to invalidate the cache for the SDK path.

C. Create an Amazon EventBridge rule that reacts to UpdateStage events from aws.apigateway. Configure the rule to invoke an AWS Lambda function to download the SDK from API Gateway, upload the SDK to the S3 bucket, and call the CloudFront API to create an invalidation for the SDK path.

D. Create an Amazon EventBridge rule that reacts to CreateDeployment events from aws.apigateway. Configure the rule to invoke an AWS Lambda function to download the SDK from API Gateway, upload the SDK to the S3 bucket, and call the S3 API to invalidate the cache for the SDK path.

Question #91 *Topic 1*

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 has noticed 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 not having fully propagated before the Lambda function is invoked.

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.

Question #92 *Topic 1*

A company uses a single AWS account to test applications on Amazon EC2 instances. The company has turned on AWS Config in the AWS account and has activated the restricted-ssh AWS Config managed rule.

The company needs an automated monitoring solution that will provide a customized notification in real time if any security group in the account is not compliant with the restricted-ssh rule. The customized notification must contain the name and ID of the noncompliant security group.

A DevOps engineer creates an Amazon Simple Notification Service (Amazon SNS) topic in the account and subscribes the appropriate personnel to the topic.

What should the DevOps engineer do next to meet these requirements?

A. Create an Amazon EventBridge rule that matches an AWS Config evaluation result of NON_COMPLIANT for the restricted-ssh rule. Configure an input transformer for the EventBridge rule. Configure the EventBridge rule to publish a notification to the SNS topic.

B. Configure AWS Config to send all evaluation results for the restricted-ssh rule to the SNS topic. Configure a filter policy on the SNS topic to send only notifications that contain the text of NON_COMPLIANT in the notification to subscribers.

C. Create an Amazon EventBridge rule that matches an AWS Config evaluation result of NON_COMPLIANT for the restricted-ssh rule. Configure the EventBridge rule to invoke AWS Systems Manager Run Command on the SNS topic to customize a notification and to publish the notification to the SNS topic.

D. Create an Amazon EventBridge rule that matches all AWS Config evaluation results of NON_COMPLIANT. Configure an input transformer for the restricted-ssh rule. Configure the EventBridge rule to publish a notification to the SNS topic.

Question #93 *Topic 1*

A company requires an RPO of 2 hours and an RTO of 10 minutes for its data and application at all times. An application uses a MySQL database and Amazon EC2 web servers. The development team needs a strategy for failover and disaster recovery.

Which combination of deployment strategies will meet these requirements? (Choose two.)

A. Create an Amazon Aurora cluster in one Availability Zone across multiple Regions as the data store. Use Aurora's automatic recovery capabilities in the event of a disaster.

B. Create an Amazon Aurora global database in two Regions as the data store. In the event of a failure, promote the secondary Region as the primary for the application.

C. Create an Amazon Aurora multi-master cluster across multiple Regions as the data store. Use a Network Load Balancer to balance the database traffic in different Regions.

D. Set up the application in two Regions and use Amazon Route 53 failover-based routing that points to the Application Load Balancers in both Regions. Use health checks to determine the availability in a given Region. Use Auto Scaling groups in each Region to adjust capacity based on demand.

E. Set up the application in two Regions and use a multi-Region Auto Scaling group behind Application Load Balancers to manage the capacity based on demand. In the event of a disaster, adjust the Auto Scaling group's desired instance count to increase baseline capacity in the failover Region.

Question #94 *Topic 1*

A business has an application that consists of five independent AWS Lambda functions.

The DevOps engineer has built a CI/CD pipeline using AWS CodePipeline and AWS CodeBuild that builds, tests, packages, and deploys each Lambda function in sequence. The pipeline uses an Amazon EventBridge rule to ensure the pipeline starts as quickly as possible after a change is made to the application source code.

After working with the pipeline for a few months, the DevOps engineer has noticed the pipeline takes too long to complete.

What should the DevOps engineer implement to BEST improve the speed of the pipeline?

A. Modify the CodeBuild projects within the pipeline to use a compute type with more available network throughput.

B. Create a custom CodeBuild execution environment that includes a symmetric multiprocessing configuration to run the builds in parallel.

C. Modify the CodePipeline configuration to run actions for each Lambda function in parallel by specifying the same runOrder.

D. Modify each CodeBuild project to run within a VPC and use dedicated instances to increase throughput.

Question #95 *Topic 1*

A company uses AWS CloudFormation stacks to deploy updates to its application. The stacks consist of different resources. The resources include AWS Auto Scaling groups, Amazon EC2 instances, Application Load Balancers (ALBs), and other resources that are necessary to launch and maintain independent stacks. Changes to application resources outside of CloudFormation stack updates are not allowed.

The company recently attempted to update the application stack by using the AWS CLI. The stack failed to update and produced the following error message: “ERROR: both the deployment and the CloudFormation stack rollback failed. The deployment failed because the following resource(s) failed to update: [AutoScalingGroup].”

The stack remains in a status of UPDATE_ROLLBACK_FAILED.

Which solution will resolve this issue?

A. Update the subnet mappings that are configured for the ALBs. Run the aws cloudformation update-stack-set AWS CLI command.

B. Update the IAM role by providing the necessary permissions to update the stack. Run the aws cloudformation continue-update-rollback AWS CLI command.

C. Submit a request for a quota increase for the number of EC2 instances for the account. Run the aws cloudformation cancel-update-stack AWS CLI command.

D. Delete the Auto Scaling group resource. Run the aws cloudformation rollback-stack AWS CLI command.

Question #96 *Topic 1*

A company is deploying a new application that uses Amazon EC2 instances. The company needs a solution to query application logs and AWS account API activity.

Which solution will meet these requirements?

A. Use the Amazon CloudWatch agent to send logs from the EC2 instances to Amazon CloudWatch Logs. Configure AWS CloudTrail to deliver the API logs to Amazon S3. Use CloudWatch to query both sets of logs.

B. Use the Amazon CloudWatch agent to send logs from the EC2 instances to Amazon CloudWatch Logs. Configure AWS CloudTrail to deliver the API logs to CloudWatch Logs. Use CloudWatch Logs Insights to query both sets of logs.

C. Use the Amazon CloudWatch agent to send logs from the EC2 instances to Amazon Kinesis. Configure AWS CloudTrail to deliver the API logs to Kinesis. Use Kinesis to load the data into Amazon Redshift. Use Amazon Redshift to query both sets of logs.

D. Use the Amazon CloudWatch agent to send logs from the EC2 instances to Amazon S3. Use AWS CloudTrail to deliver the API logs to Amazon S3. Use Amazon Athena to query both sets of logs in Amazon S3.

Question #97 *Topic 1*

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

Which solution will meet these requirements?

A. Use AWS Systems Manager to detect vulnerabilities on the EC2 instances. Install the Amazon Kinesis Agent to capture system logs and deliver them to Amazon S3.

B. Use AWS Systems Manager to detect vulnerabilities on the EC2 instances. Install the Systems Manager Agent to capture system logs and view login activity in the CloudTrail console.

C. Configure Amazon CloudWatch to detect vulnerabilities on the EC2 instances. Install the AWS Config daemon to capture system logs and view them in the AWS Config console.

D. Configure Amazon Inspector to detect vulnerabilities on the EC2 instances. Install the Amazon CloudWatch Agent to capture system logs and record them via Amazon CloudWatch Logs.

Question #98 *Topic 1*

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.

Question #99 *Topic 1*

A company is using AWS Organizations to centrally manage its AWS accounts. The company has turned on AWS Config in each member account by using AWS CloudFormation StackSets. The company has configured trusted access in Organizations for AWS Config and has configured a member account as a delegated administrator account for AWS Config.

A DevOps engineer needs to implement a new security policy. The policy must require all current and future AWS member accounts to use a common baseline of AWS Config rules that contain remediation actions that are managed from a central account. Non-administrator users who can access member accounts must not be able to modify this common baseline of AWS Config rules that are deployed into each member account.

Which solution will meet these requirements?

A. Create a CloudFormation template that contains the AWS Config rules and remediation actions. Deploy the template from the Organizations management account by using CloudFormation StackSets.

B. Create an AWS Config conformance pack that contains the AWS Config rules and remediation actions. Deploy the pack from the Organizations management account by using CloudFormation StackSets.

C. Create a CloudFormation template that contains the AWS Config rules and remediation actions. Deploy the template from the delegated administrator account by using AWS Config.

D. Create an AWS Config conformance pack that contains the AWS Config rules and remediation actions. Deploy the pack from the delegated administrator account by using AWS Config.

Question #100 *Topic 1*

A DevOps engineer manages a large commercial website that runs on Amazon EC2. The website uses Amazon Kinesis Data Streams to collect and process web logs. The DevOps engineer manages the Kinesis consumer application, which also runs on Amazon EC2.

Sudden increases of data cause the Kinesis consumer application to fall behind, and the Kinesis data streams drop records before the records can be processed. The DevOps engineer must implement a solution to improve stream handling.

Which solution meets these requirements with the MOST operational efficiency?

A. Modify the Kinesis consumer application to store the logs durably in Amazon S3. Use Amazon EMR to process the data directly on Amazon S3 to derive customer insights. Store the results in Amazon S3.

B. Horizontally scale the Kinesis consumer application by adding more EC2 instances based on the Amazon CloudWatch GetRecords.IteratorAgeMilliseconds metric. Increase the retention period of the Kinesis data streams.

C. Convert the Kinesis consumer application to run as an AWS Lambda function. Configure the Kinesis data streams as the event source for the Lambda function to process the data streams.

D. Increase the number of shards in the Kinesis data streams to increase the overall throughput so that the consumer application processes the data faster.

Question #101 *Topic 1*

A company recently created a new AWS Control Tower landing zone in a new organization in AWS Organizations. The landing zone must be able to demonstrate compliance with the Center for Internet Security (CIS) Benchmarks for AWS Foundations.

The company's security team wants to use AWS Security Hub to view compliance across all accounts. Only the security team can be allowed to view aggregated Security Hub findings. In addition, specific users must be able to view findings from their own accounts within the organization. All accounts must be enrolled in Security Hub after the accounts are created.

Which combination of steps will meet these requirements in the MOST automated way? (Choose three.)

A. Turn on trusted access for Security Hub in the organization's management account. Create a new security account by using AWS Control Tower. Configure the new security account as the delegated administrator account for Security Hub. In the new security account, provide Security Hub with the CIS Benchmarks for AWS Foundations standards.

B. Turn on trusted access for Security Hub in the organization's management account. From the management account, provide Security Hub with the CIS Benchmarks for AWS Foundations standards.

C. Create an AWS IAM Identity Center (AWS Single Sign-On) permission set that includes the required permissions. Use the CreateAccountAssignment API operation to associate the security team users with the permission set and with the delegated security account.

D. Create an SCP that explicitly denies any user who is not on the security team from accessing Security Hub.

E. In Security Hub, turn on automatic enablement.

F. In the organization's management account, create an Amazon EventBridge rule that reacts to the CreateManagedAccount event. Create an AWS Lambda function that uses the Security Hub CreateMembers API operation to add new accounts to Security Hub. Configure the EventBridge rule to invoke the Lambda function.

Question #102 *Topic 1*

A company runs applications in AWS accounts that are in an organization in AWS Organizations. The applications use Amazon EC2 instances and Amazon S3.

The company wants to detect potentially compromised EC2 instances, suspicious network activity, and unusual API activity in its existing AWS accounts and in any AWS accounts that the company creates in the future. When the company detects one of these events, the company wants to use an existing Amazon Simple Notification Service (Amazon SNS) topic to send a notification to its operational support team for investigation and remediation.

Which solution will meet these requirements in accordance with AWS best practices?

A. In the organization's management account, configure an AWS account as the Amazon GuardDuty administrator account. In the GuardDuty administrator account, add the company's existing AWS accounts to GuardDuty as members. In the GuardDuty administrator account, create an Amazon EventBridge rule with an event pattern to match GuardDuty events and to forward matching events to the SNS topic.

B. In the organization's management account, configure Amazon GuardDuty to add newly created AWS accounts by invitation and to send invitations to the existing AWS accounts. Create an AWS CloudFormation stack set that accepts the GuardDuty invitation and creates an Amazon EventBridge rule. Configure the rule with an event pattern to match GuardDuty events and to forward matching events to the SNS topic. Configure the CloudFormation stack set to deploy into all AWS accounts in the organization.

C. In the organization's management account, create an AWS CloudTrail organization trail. Activate the organization trail in all AWS accounts in the organization. Create an SCP that enables VPC Flow Logs in each account in the organization. Configure AWS Security Hub for the organization. Create an Amazon EventBridge rule with an event pattern to match Security Hub events and to forward matching events to the SNS topic.

D. In the organization's management account, configure an AWS account as the AWS CloudTrail administrator account. In the CloudTrail administrator account, create a CloudTrail organization trail. Add the company's existing AWS accounts to the organization trail. Create an SCP that enables VPC Flow Logs in each account in the organization. Configure AWS Security Hub for the organization. Create an Amazon EventBridge rule with an event pattern to match Security Hub events and to forward matching events to the SNS topic.

Question #103 *Topic 1*

A company's DevOps engineer is working in a multi-account environment. The company uses AWS Transit Gateway to route all outbound traffic through a network operations account. In the network operations account, all account traffic passes through a firewall appliance for inspection before the traffic goes to an internet gateway.

The firewall appliance sends logs to Amazon CloudWatch Logs and includes event severities of CRITICAL, HIGH, MEDIUM, LOW, and INFO. The security team wants to receive an alert if any CRITICAL events occur.

What should the DevOps engineer do to meet these requirements?

A. Create an Amazon CloudWatch Synthetics canary to monitor the firewall state. If the firewall reaches a CRITICAL state or logs a CRITICAL event, use a CloudWatch alarm to publish a notification to an Amazon Simple Notification Service (Amazon SNS) topic. Subscribe the security team's email address to the topic.

B. Create an Amazon CloudWatch metric filter by using a search for CRITICAL events. Publish a custom metric for the finding. Use a CloudWatch alarm based on the custom metric to publish a notification to an Amazon Simple Notification Service (Amazon SNS) topic. Subscribe the security team's email address to the topic.

C. Enable Amazon GuardDuty in the network operations account. Configure GuardDuty to monitor flow logs. Create an Amazon EventBridge event rule that is invoked by GuardDuty events that are CRITICAL. Define an Amazon Simple Notification Service (Amazon SNS) topic as a target. Subscribe the security team's email address to the topic.

D. Use AWS Firewall Manager to apply consistent policies across all accounts. Create an Amazon EventBridge event rule that is invoked by Firewall Manager events that are CRITICAL. Define an Amazon Simple Notification Service (Amazon SNS) topic as a target. Subscribe the security team's email address to the topic.

Question #104 *Topic 1*

A company is divided into teams. Each team has an AWS account, and all the accounts are in an organization in AWS Organizations. Each team must retain full administrative rights to its AWS account. Each team also must be allowed to access only AWS services that the company approves for use. AWS services must gain approval through a request and approval process.

How should a DevOps engineer configure the accounts to meet these requirements?

A. Use AWS CloudFormation StackSets to provision IAM policies in each account to deny access to restricted AWS services. In each account, configure AWS Config rules that ensure that the policies are attached to IAM principals in the account.

B. Use AWS Control Tower to provision the accounts into OUs within the organization. Configure AWS Control Tower to enable AWS IAM Identity Center (AWS Single Sign-On). Configure IAM Identity Center to provide administrative access. Include deny policies on user roles for restricted AWS services.

C. Place all the accounts under a new top-level OU within the organization. Create an SCP that denies access to restricted AWS services. Attach the SCP to the OU.

D. Create an SCP that allows access to only approved AWS services. Attach the SCP to the root OU of the organization. Remove the FullAWSAccess SCP from the root OU of the organization.

Question #105 *Topic 1*

A DevOps engineer used an AWS CloudFormation custom resource to set up AD Connector. The AWS Lambda function ran 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.

Question #106 *Topic 1*

A company uses AWS CodeCommit for source code control. Developers apply their changes to various feature branches and create pull requests to move those changes to the main branch when the changes are ready for production.

The developers should not be able to push changes directly to the main branch. The company applied the AWSCodeCommitPowerUser managed policy to the developers' IAM role, and now these developers can push changes to the main branch directly on every repository in the AWS account.

What should the company do to restrict the developers' ability to push changes to the main branch directly?

A. Create an additional policy to include a Deny rule for the GitPush and PutFile actions. Include a restriction for the specific repositories in the policy statement with a condition that references the main branch.

B. Remove the IAM policy, and add an AWSCodeCommitReadOnly managed policy. Add an Allow rule for the GitPush and PutFile actions for the specific repositories in the policy statement with a condition that references the main branch.

C. Modify the IAM policy. Include a Deny rule for the GitPush and PutFile actions for the specific repositories in the policy statement with a condition that references the main branch.

D. Create an additional policy to include an Allow rule for the GitPush and PutFile actions. Include a restriction for the specific repositories in the policy statement with a condition that references the feature branches.

Question #107 *Topic 1*

A company manages a web application that runs on Amazon EC2 instances behind an Application Load Balancer (ALB). The EC2 instances run in an Auto Scaling group across multiple Availability Zones. The application uses an Amazon RDS for MySQL DB instance to store the data. The company has configured Amazon Route 53 with an alias record that points to the ALB.

A new company guideline requires a geographically isolated disaster recovery (DR) site with an RTO of 4 hours and an RPO of 15 minutes.

Which DR strategy will meet these requirements with the LEAST change to the application stack?

A. Launch a replica environment of everything except Amazon RDS in a different Availability Zone. Create an RDS read replica in the new Availability Zone, and configure the new stack to point to the local RDS DB instance. Add the new stack to the Route 53 record set by using a health check to configure a failover routing policy.

B. Launch a replica environment of everything except Amazon RDS in a different AWS Region. Create an RDS read replica in the new Region, and configure the new stack to point to the local RDS DB instance. Add the new stack to the Route 53 record set by using a health check to configure a latency routing policy.

C. Launch a replica environment of everything except Amazon RDS in a different AWS Region. In the event of an outage, copy and restore the latest RDS snapshot from the primary Region to the DR Region. Adjust the Route 53 record set to point to the ALB in the DR Region.

D. Launch a replica environment of everything except Amazon RDS in a different AWS Region. Create an RDS read replica in the new Region, and configure the new environment to point to the local RDS DB instance. Add the new stack to the Route 53 record set by using a health check to configure a failover routing policy. In the event of an outage, promote the read replica to primary.

Question #108 *Topic 1*

A large enterprise is deploying a web application on AWS. The application runs on Amazon EC2 instances behind an Application Load Balancer. The instances run in an Auto Scaling group across multiple Availability Zones. The application stores data in an Amazon RDS for Oracle DB instance and Amazon DynamoDB. There are separate environments for development, testing, and production.

What is the MOST secure and flexible way to obtain password credentials during deployment?

A. Retrieve an access key from an AWS Systems Manager SecureString parameter to access AWS services. Retrieve the database credentials from a Systems Manager SecureString parameter.

B. Launch the EC2 instances with an EC2 IAM role to access AWS services. Retrieve the database credentials from AWS Secrets Manager.

C. Retrieve an access key from an AWS Systems Manager plaintext parameter to access AWS services. Retrieve the database credentials from a Systems Manager SecureString parameter.

D. Launch the EC2 instances with an EC2 IAM role to access AWS services. Store the database passwords in an encrypted config file with the application artifacts.

Question #109 *Topic 1*

The security team depends on AWS CloudTrail to detect sensitive security issues in the company's AWS account. The DevOps engineer needs a solution to auto-remediate CloudTrail being turned off in an AWS account.

What solution ensures the LEAST amount of downtime for the CloudTrail log deliveries?

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

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

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

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

Question #110 *Topic 1*

A company uses AWS CodeArtifact to centrally store Python packages. The CodeArtifact repository is configured with the following repository policy:

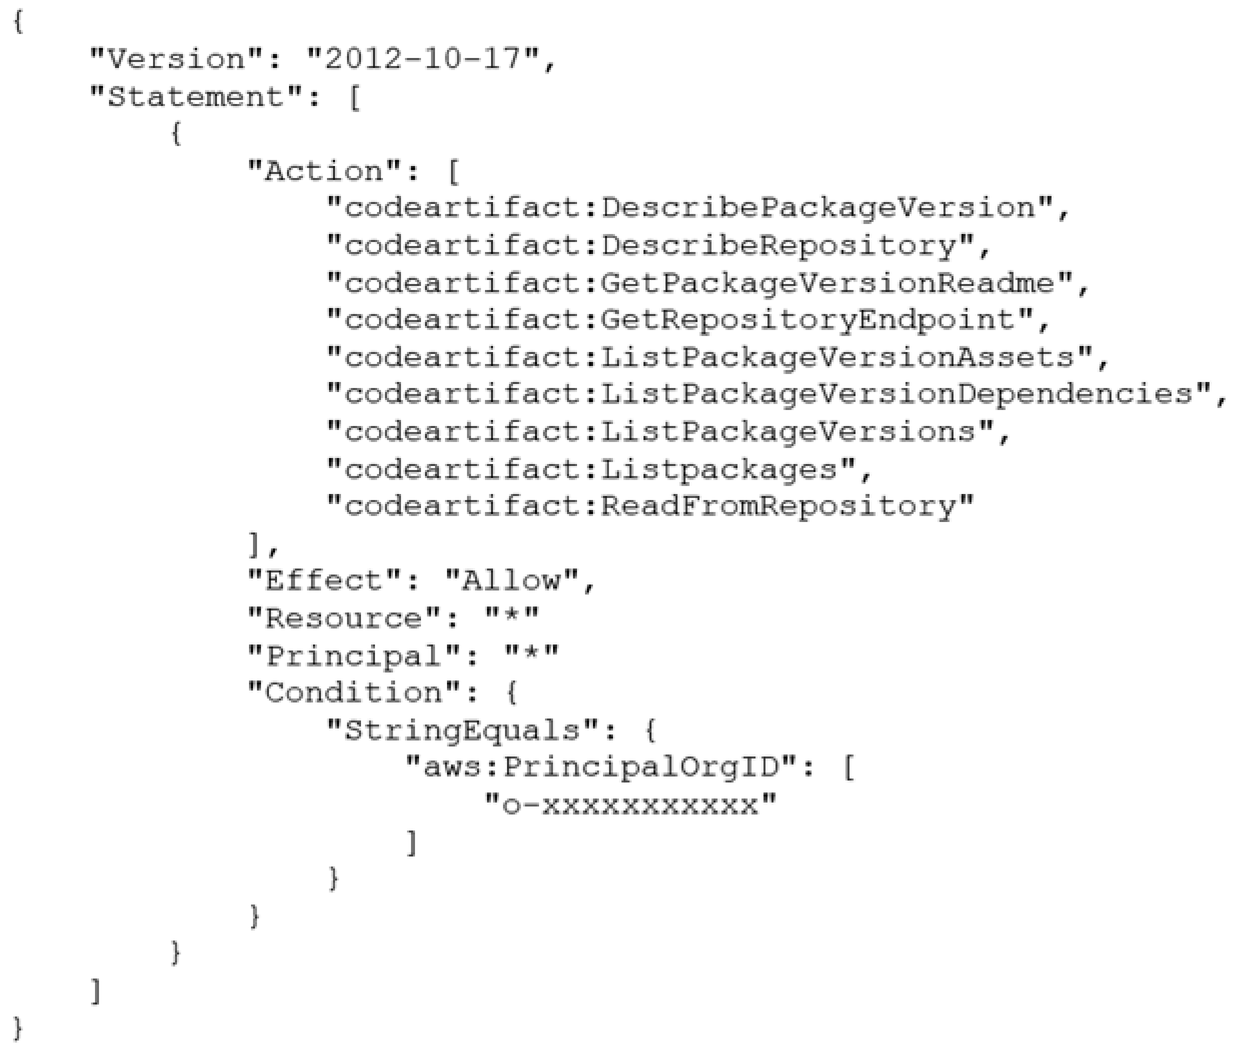

```
{
    "Version": "2012-10-17",
    "Statement": [
        {
            "Action": [
                "codeartifact:DescribePackageVersion",
                "codeartifact:DescribeRepository",
                "codeartifact:GetPackageVersionReadme",
                "codeartifact:GetRepositoryEndpoint",
                "codeartifact:ListPackageVersionAssets",
                "codeartifact:ListPackageVersionDependencies",
                "codeartifact:ListPackageVersions",
                "codeartifact:Listpackages",
                "codeartifact:ReadFromRepository"
            ],
            "Effect": "Allow",
            "Resource": "*"
            "Principal": "*"
            "Condition": {
                "StringEquals": {
                    "aws:PrincipalOrgID": [
                        "o-xxxxxxxxxxx"
                    ]
                }
            }
        }
    ]
}
```

A development team is building a new project in an account that is in an organization in AWS Organizations. The development team wants to use a Python library that has already been stored in the CodeArtifact repository in the organization. The development team uses AWS CodePipeline and AWS CodeBuild to build the new application. The CodeBuild job that the development team uses to build the application is configured to run in a VPC. Because of compliance requirements, the VPC has no internet connectivity.

The development team creates the VPC endpoints for CodeArtifact and updates the CodeBuild buildspec.yaml file. However, the development team cannot download the Python library from the repository.

Which combination of steps should a DevOps engineer take so that the development team can use CodeArtifact? (Choose two.)

A. Create an Amazon S3 gateway endpoint. Update the route tables for the subnets that are running the CodeBuild job.

B. Update the repository policy's Principal statement to include the ARN of the role that the CodeBuild project uses.

C. Share the CodeArtifact repository with the organization by using AWS Resource Access Manager (AWS RAM).

D. Update the role that the CodeBuild project uses so that the role has sufficient permissions to use the CodeArtifact repository.

E. Specify the account that hosts the repository as the delegated administrator for CodeArtifact in the organization.

Question #111 *Topic 1*

A company uses a series of individual Amazon CloudFormation templates to deploy its multi-Region applications. These templates must be deployed in a specific order. The company is making more changes to the templates than previously expected and wants to deploy new templates more efficiently. Additionally, the data engineering team must be notified of all changes to the templates.

What should the company do to accomplish these goals?

A. Create an AWS Lambda function to deploy the CloudFormation templates in the required order. Use stack policies to alert the data engineering team.

B. Host the CloudFormation templates in Amazon S3. Use Amazon S3 events to directly trigger CloudFormation updates and Amazon SNS notifications.

C. Implement CloudFormation StackSets and use drift detection to trigger update alerts to the data engineering team.

D. Leverage CloudFormation nested stacks and stack sets for deployments. Use Amazon SNS to notify the data engineering team.

Question #112 *Topic 1*

A DevOps engineer has implemented a CI/CD pipeline to deploy an AWS CloudFormation template that provisions a web application. The web application consists of an Application Load Balancer (ALB), a target group, a launch template that uses an Amazon Linux 2 AMI, an Auto Scaling group of Amazon EC2 instances, a security group, and an Amazon RDS for MySQL database. The launch template includes user data that specifies a script to install and start the application.

The initial deployment of the application was successful. The DevOps engineer made changes to update the version of the application with the user data. The CI/CD pipeline has deployed a new version of the template. However, the health checks on the ALB are now failing. The health checks have marked all targets as unhealthy.

During investigation, the DevOps engineer notices that the CloudFormation stack has a status of UPDATE_COMPLETE. However, when the DevOps engineer connects to one of the EC2 instances and checks /var/log/messages, the DevOps engineer notices that the Apache web server failed to start successfully because of a configuration error.

How can the DevOps engineer ensure that the CloudFormation deployment will fail if the user data fails to successfully finish running?

A. Use the cfn-signal helper script to signal success or failure to CloudFormation. Use the WaitOnResourceSignals update policy within the CloudFormation template. Set an appropriate timeout for the update policy.

B. Create an Amazon CloudWatch alarm for the UnhealthyHostCount metric. Include an appropriate alarm threshold for the target group. Create an Amazon Simple Notification Service (Amazon SNS) topic as the target to signal success or failure to CloudFormation.

C. Create a lifecycle hook on the Auto Scaling group by using the AWS::AutoScaling::LifecycleHook resource. Create an Amazon Simple Notification Service (Amazon SNS) topic as the target to signal success or failure to CloudFormation. Set an appropriate timeout on the lifecycle hook.

D. Use the Amazon CloudWatch agent to stream the cloud-init logs. Create a subscription filter that includes an AWS Lambda function with an appropriate invocation timeout. Configure the Lambda function to use the SignalResource API operation to signal success or failure to CloudFormation.

Question #113

Topic 1

A company has a data ingestion application that runs across multiple AWS accounts. The accounts are in an organization in AWS Organizations. The company needs to monitor the application and consolidate access to the application. Currently, the company is running the application on Amazon EC2 instances from several Auto Scaling groups. The EC2 instances have no access to the internet because the data is sensitive. Engineers have deployed the necessary VPC endpoints. The EC2 instances run a custom AMI that is built specifically for the application.

To maintain and troubleshoot the application, system administrators need the ability to log in to the EC2 instances. This access must be automated and controlled centrally. The company's security team must receive a notification whenever the instances are accessed.

Which solution will meet these requirements?

A. Create an Amazon EventBridge rule to send notifications to the security team whenever a user logs in to an EC2 instance. Use EC2 Instance Connect to log in to the instances. Deploy Auto Scaling groups by using AWS CloudFormation. Use the cfn-init helper script to deploy appropriate VPC routes for external access. Rebuild the custom AMI so that the custom AMI includes AWS Systems Manager Agent.

B. Deploy a NAT gateway and a bastion host that has internet access. Create a security group that allows incoming traffic on all the EC2 instances from the bastion host. Install AWS Systems Manager Agent on all the EC2 instances. Use Auto Scaling group lifecycle hooks for monitoring and auditing access. Use Systems Manager Session Manager to log in to the instances. Send logs to a log group in Amazon CloudWatch Logs. Export data to Amazon S3 for auditing. Send notifications to the security team by using S3 event notifications.

C. Use EC2 Image Builder to rebuild the custom AMI. Include the most recent version of AWS Systems Manager Agent in the image. Configure the Auto Scaling group to attach the AmazonSSMManagedInstanceCore role to all the EC2 instances. Use Systems Manager Session Manager to log in to the instances. Enable logging of session details to Amazon S3. Create an S3 event notification for new file uploads to send a message to the security team through an Amazon Simple Notification Service (Amazon SNS) topic.

D. Use AWS Systems Manager Automation to build Systems Manager Agent into the custom AMI. Configure AWS Config to attach an SCP to the root organization account to allow the EC2 instances to connect to Systems Manager. Use Systems Manager Session Manager to log in to the instances. Enable logging of session details to Amazon S3. Create an S3 event notification for new file uploads to send a message to the security team through an Amazon Simple Notification Service (Amazon SNS) topic.

Question #114

Topic 1

A company uses Amazon S3 to store proprietary information. The development team creates buckets for new projects on a daily basis. The security team wants to ensure that all existing and future buckets have encryption, logging, and versioning enabled. Additionally, no buckets should ever be publicly read or write accessible.

What should a DevOps engineer do to meet these requirements?

A. Enable AWS CloudTrail and configure automatic remediation using AWS Lambda.

B. Enable AWS Config rules and configure automatic remediation using AWS Systems Manager documents.

C. Enable AWS Trusted Advisor and configure automatic remediation using Amazon EventBridge.

D. Enable AWS Systems Manager and configure automatic remediation using Systems Manager documents.

Question #115 *Topic 1*

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 volume and can tolerate interruptions. Configuring the cluster software from a generic EC2 Linux image takes 30 minutes.

What is the MOST cost-effective solution?

A. Use Amazon EFS for checkpoint data. To complete the job, use an EC2 Auto Scaling group and an On-Demand pricing model to provision EC2 instances temporarily.

B. Use GlusterFS on EC2 instances for checkpoint data. To run the batch job, configure EC2 instances manually. When the job completes, shut down the instances manually.

C. Use Amazon EFS for checkpoint data. Use EC2 Fleet to launch EC2 Spot Instances, and utilize user data to configure the EC2 Linux instance on startup.

D. Use Amazon EFS for checkpoint data. Use EC2 Fleet to launch EC2 Spot Instances. Create a custom AMI for the cluster and use the latest AMI when creating instances.

Question #116 *Topic 1*

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. Configure API Gateway to use a canary release deployment to send a small subset of user traffic to the new environment.

B. Introduce changes as a separate environment parallel to the existing one. Update the application's DNS alias records to point to the new environment.

C. Introduce changes as a separate target group behind the existing Application Load Balancer. Configure API Gateway to route user traffic to the new target group in steps.

D. Introduce changes as a separate target group behind the existing Application Load Balancer. Configure API Gateway to route all traffic to the Application Load Balancer, which then sends the traffic to the new target group.

Question #117 *Topic 1*

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

Which combination of steps will meet these requirements with the LEAST amount of effort? (Choose three.)

A. Filter the data through AWS X-Ray to visualize the data.

B. Filter the data through Amazon QuickSight to visualize the data.

C. Query the data with Amazon Athena.

D. Query the data with Amazon Redshift.

E. Use the AWS Glue Data Catalog as the persistent metadata store.

F. Use Amazon DynamoDB as the persistent metadata store.

Question #118 *Topic 1*

A company deploys its corporate infrastructure on AWS across multiple AWS Regions and Availability Zones. The infrastructure is deployed on Amazon EC2 instances and connects with AWS IoT Greengrass devices. The company deploys additional resources on on-premises servers that are located in the corporate headquarters.

The company wants to reduce the overhead involved in maintaining and updating its resources. The company's DevOps team plans to use AWS Systems Manager to implement automated management and application of patches. The DevOps team confirms that Systems Manager is available in the Regions that the resources are deployed in. Systems Manager also is available in a Region near the corporate headquarters.

Which combination of steps must the DevOps team take to implement automated patch and configuration management across the company's EC2 instances, IoT devices, and on-premises infrastructure? (Choose three.)

A. Apply tags to all the EC2 instances, AWS IoT Greengrass devices, and on-premises servers. Use Systems Manager Session Manager to push patches to all the tagged devices.

B. Use Systems Manager Run Command to schedule patching for the EC2 instances, AWS IoT Greengrass devices, and on-premises servers.

C. Use Systems Manager Patch Manager to schedule patching for the EC2 instances, AWS IoT Greengrass devices, and on-premises servers as a Systems Manager maintenance window task.

D. Configure Amazon EventBridge to monitor Systems Manager Patch Manager for updates to patch baselines. Associate Systems Manager Run Command with the event to initiate a patch action for all EC2 instances, AWS IoT Greengrass devices, and on-premises servers.

E. Create an IAM instance profile for Systems Manager. Attach the instance profile to all the EC2 instances in the AWS account. For the AWS IoT Greengrass devices and on-premises servers, create an IAM service role for Systems Manager.

F. Generate a managed-instance activation. Use the Activation Code and Activation ID to install Systems Manager Agent (SSM Agent) on each server in the on-premises environment. Update the AWS IoT Greengrass IAM token exchange role. Use the role to deploy SSM Agent on all the IoT devices.

Question #119 *Topic 1*

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 operationally 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.

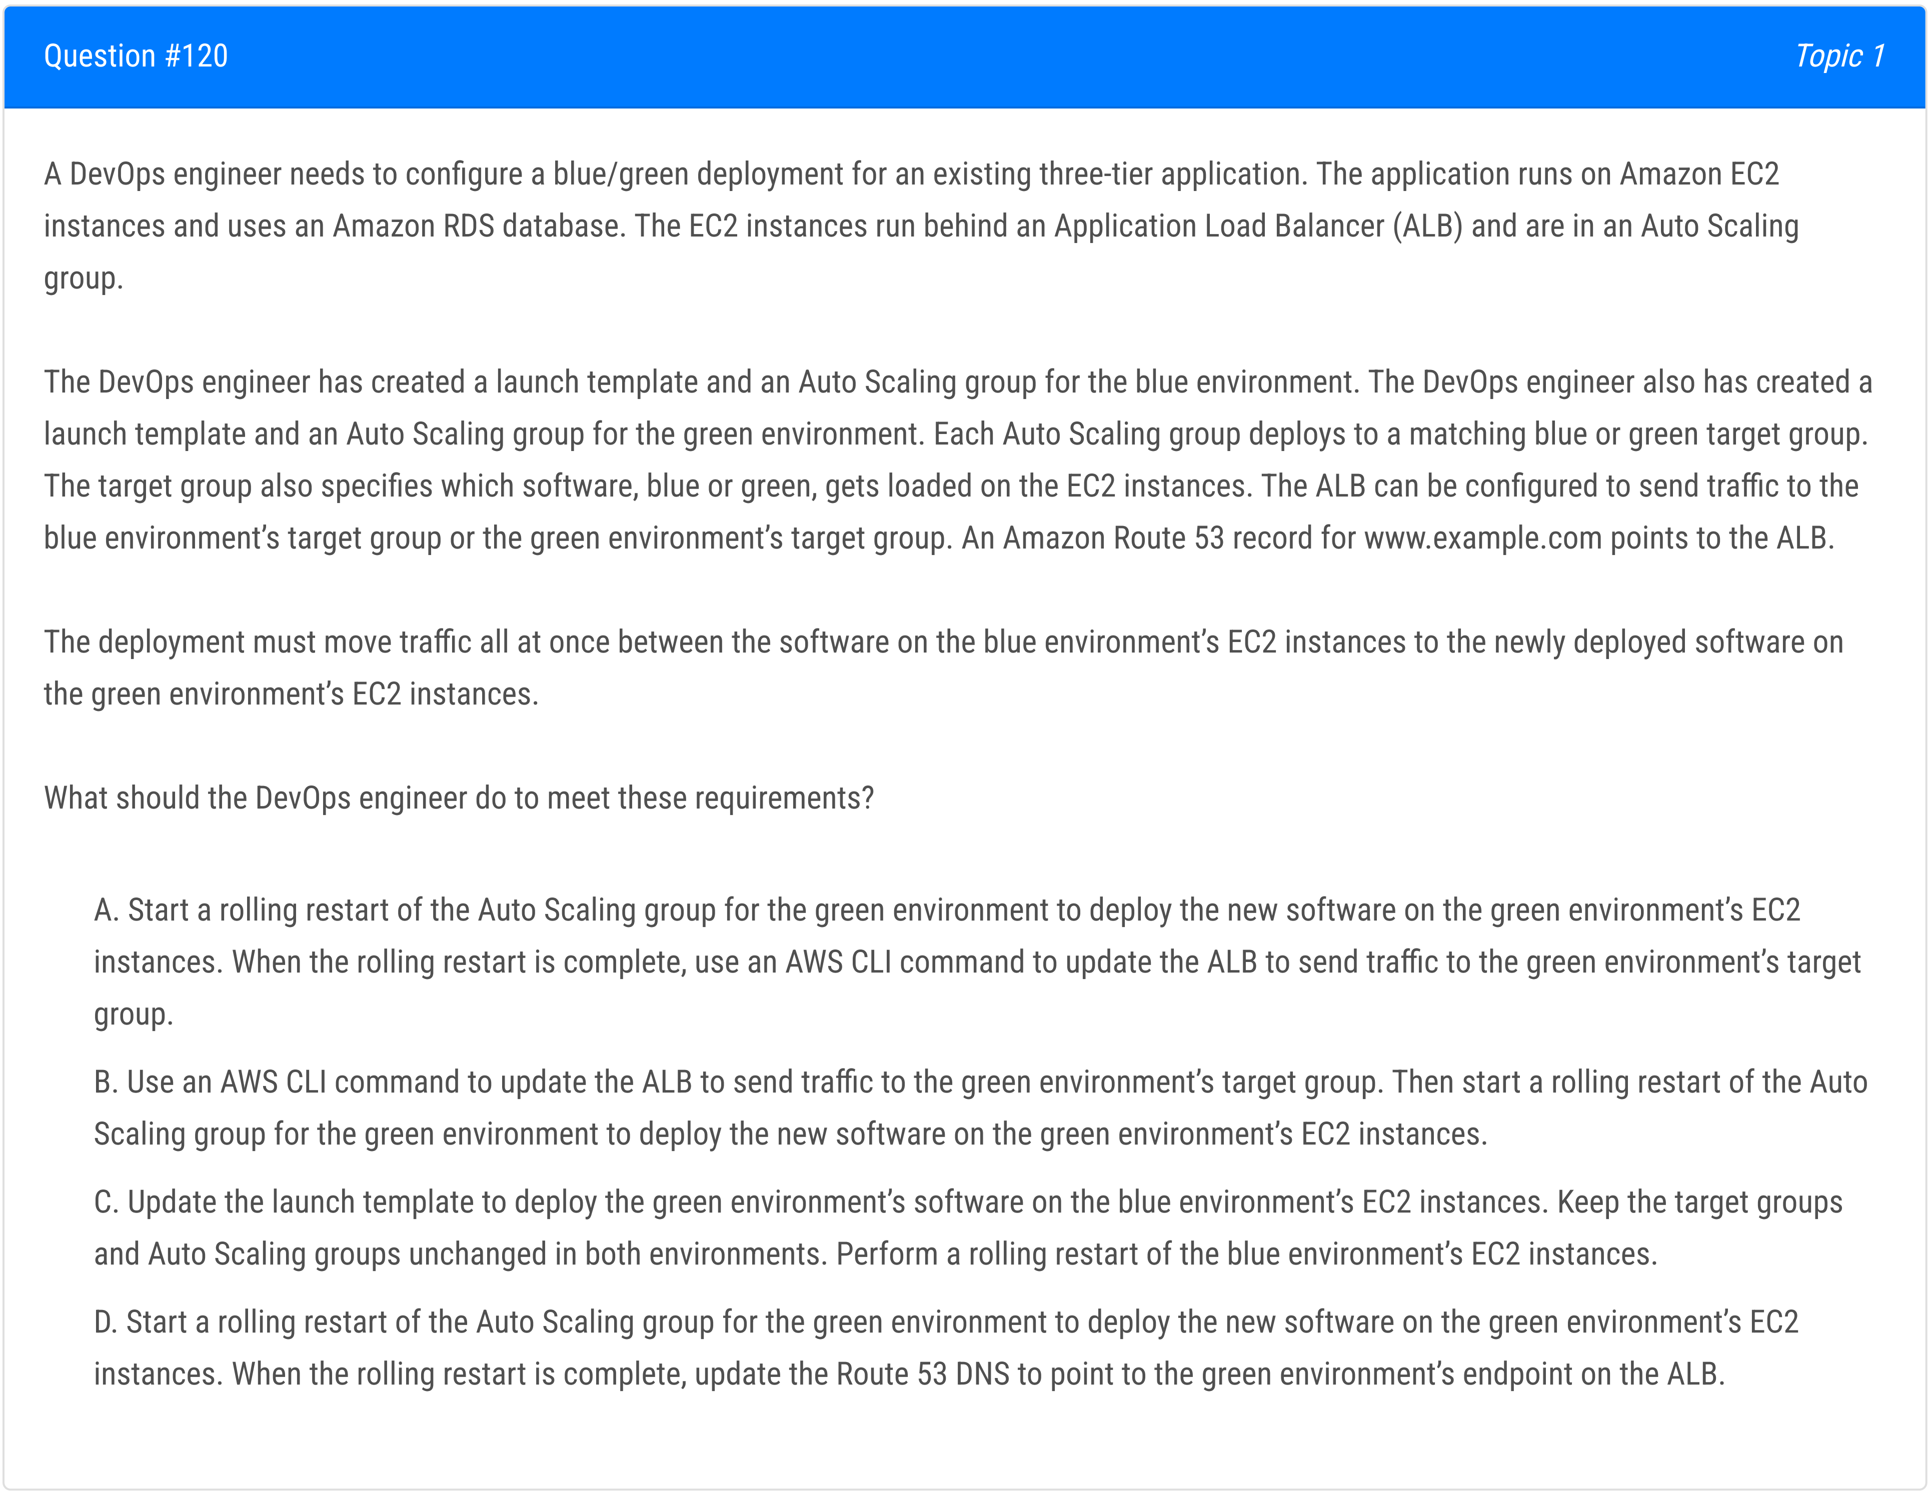

Question #120 *Topic 1*

A DevOps engineer needs to configure a blue/green deployment for an existing three-tier application. The application runs on Amazon EC2 instances and uses an Amazon RDS database. The EC2 instances run behind an Application Load Balancer (ALB) and are in an Auto Scaling group.

The DevOps engineer has created a launch template and an Auto Scaling group for the blue environment. The DevOps engineer also has created a launch template and an Auto Scaling group for the green environment. Each Auto Scaling group deploys to a matching blue or green target group. The target group also specifies which software, blue or green, gets loaded on the EC2 instances. The ALB can be configured to send traffic to the blue environment's target group or the green environment's target group. An Amazon Route 53 record for www.example.com points to the ALB.

The deployment must move traffic all at once between the software on the blue environment's EC2 instances to the newly deployed software on the green environment's EC2 instances.

What should the DevOps engineer do to meet these requirements?

A. Start a rolling restart of the Auto Scaling group for the green environment to deploy the new software on the green environment's EC2 instances. When the rolling restart is complete, use an AWS CLI command to update the ALB to send traffic to the green environment's target group.

B. Use an AWS CLI command to update the ALB to send traffic to the green environment's target group. Then start a rolling restart of the Auto Scaling group for the green environment to deploy the new software on the green environment's EC2 instances.

C. Update the launch template to deploy the green environment's software on the blue environment's EC2 instances. Keep the target groups and Auto Scaling groups unchanged in both environments. Perform a rolling restart of the blue environment's EC2 instances.

D. Start a rolling restart of the Auto Scaling group for the green environment to deploy the new software on the green environment's EC2 instances. When the rolling restart is complete, update the Route 53 DNS to point to the green environment's endpoint on the ALB.

Question #121 *Topic 1*

A company is building a new pipeline by using AWS CodePipeline and AWS CodeBuild in a build account. The pipeline consists of two stages. The first stage is a CodeBuild job to build and package an AWS Lambda function. The second stage consists of deployment actions that operate on two different AWS accounts: a development environment account and a production environment account. The deployment stages use the AWS CloudFormation action that CodePipeline invokes to deploy the infrastructure that the Lambda function requires.

A DevOps engineer creates the CodePipeline pipeline and configures the pipeline to encrypt build artifacts by using the AWS Key Management Service (AWS KMS) AWS managed key for Amazon S3 (the aws/s3 key). The artifacts are stored in an S3 bucket. When the pipeline runs, the CloudFormation actions fail with an access denied error.

Which combination of actions must the DevOps engineer perform to resolve this error? (Choose two.)

A. Create an S3 bucket in each AWS account for the artifacts. Allow the pipeline to write to the S3 buckets. Create a CodePipeline S3 action to copy the artifacts to the S3 bucket in each AWS account. Update the CloudFormation actions to reference the artifacts S3 bucket in the production account.

B. Create a customer managed KMS key. Configure the KMS key policy to allow the IAM roles used by the CloudFormation action to perform decrypt operations. Modify the pipeline to use the customer managed KMS key to encrypt artifacts.

C. Create an AWS managed KMS key. Configure the KMS key policy to allow the development account and the production account to perform decrypt operations. Modify the pipeline to use the KMS key to encrypt artifacts.

D. In the development account and in the production account, create an IAM role for CodePipeline. Configure the roles with permissions to perform CloudFormation operations and with permissions to retrieve and decrypt objects from the artifacts S3 bucket. In the CodePipeline account, configure the CodePipeline CloudFormation action to use the roles.

E. In the development account and in the production account, create an IAM role for CodePipeline. Configure the roles with permissions to perform CloudFormation operations and with permissions to retrieve and decrypt objects from the artifacts S3 bucket. In the CodePipeline account, modify the artifacts S3 bucket policy to allow the roles access. Configure the CodePipeline CloudFormation action to use the roles.

Question #122 *Topic 1*

A company is using an organization in AWS Organizations to manage multiple AWS accounts. The company's development team wants to use AWS Lambda functions to meet resiliency requirements and is rewriting all applications to work with Lambda functions that are deployed in a VPC. The development team is using Amazon Elastic File System (Amazon EFS) as shared storage in Account A in the organization.

The company wants to continue to use Amazon EFS with Lambda. Company policy requires all serverless projects to be deployed in Account B.

A DevOps engineer needs to reconfigure an existing EFS file system to allow Lambda functions to access the data through an existing EFS access point.

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

A. Update the EFS file system policy to provide Account B with access to mount and write to the EFS file system in Account A.

B. Create SCPs to set permission guardrails with fine-grained control for Amazon EFS.

C. Create a new EFS file system in Account B. Use AWS Database Migration Service (AWS DMS) to keep data from Account A and Account B synchronized.

D. Update the Lambda execution roles with permission to access the VPC and the EFS file system.

E. Create a VPC peering connection to connect Account A to Account B.

F. Configure the Lambda functions in Account B to assume an existing IAM role in Account A.

Question #123 *Topic 1*

A media company has several thousand Amazon EC2 instances in an AWS account. The company is using Slack and a shared email inbox for team communications and important updates. A DevOps engineer needs to send all AWS-scheduled EC2 maintenance notifications to the Slack channel and the shared inbox. The solution must include the instances' Name and Owner tags.

Which solution will meet these requirements?

A. Integrate AWS Trusted Advisor with AWS Config. Configure a custom AWS Config rule to invoke an AWS Lambda function to publish notifications to an Amazon Simple Notification Service (Amazon SNS) topic. Subscribe a Slack channel endpoint and the shared inbox to the topic.

B. Use Amazon EventBridge to monitor for AWS Health events. Configure the maintenance events to target an Amazon Simple Notification Service (Amazon SNS) topic. Subscribe an AWS Lambda function to the SNS topic to send notifications to the Slack channel and the shared inbox.

C. Create an AWS Lambda function that sends EC2 maintenance notifications to the Slack channel and the shared inbox. Monitor EC2 health events by using Amazon CloudWatch metrics. Configure a CloudWatch alarm that invokes the Lambda function when a maintenance notification is received.

D. Configure AWS Support integration with AWS CloudTrail. Create a CloudTrail lookup event to invoke an AWS Lambda function to pass EC2 maintenance notifications to Amazon Simple Notification Service (Amazon SNS). Configure Amazon SNS to target the Slack channel and the shared inbox.

Question #124 *Topic 1*

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

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

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

A. Implement Amazon CloudWatch Logs for CodePipeline and CodeDeploy, create an AWS Config rule to evaluate code deployment issues, and create an Amazon Simple Notification Service (Amazon SNS) topic to notify stakeholders of deployment issues.

B. Implement Amazon EventBridge for CodePipeline and CodeDeploy, create an AWS Lambda function to evaluate code deployment issues, and create an Amazon Simple Notification Service (Amazon SNS) topic to notify stakeholders of deployment issues.

C. Implement AWS CloudTrail to record CodePipeline and CodeDeploy API call information, create an AWS Lambda function to evaluate code deployment issues, and create an Amazon Simple Notification Service (Amazon SNS) topic to notify stakeholders of deployment issues.

D. Implement Amazon EventBridge for CodePipeline and CodeDeploy, create an Amazon Inspector assessment target to evaluate code deployment issues, and create an Amazon Simple Notification Service (Amazon SNS) topic to notify stakeholders of deployment issues.

Question #125 *Topic 1*

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

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

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

Which solution will resolve this error?

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

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

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

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

Question #126 *Topic 1*

A highly regulated company has a policy that DevOps engineers should not log in to their Amazon EC2 instances except in emergencies. If a DevOps engineer does log in, the security team must be notified within 15 minutes of the occurrence.

Which solution will meet these requirements?

A. Install the Amazon Inspector agent on each EC2 instance. Subscribe to Amazon EventBridge notifications. Invoke an AWS Lambda function to check if a message is about user logins. If it is, send a notification to the security team using Amazon SNS.

B. Install the Amazon CloudWatch agent on each EC2 instance. Configure the agent to push all logs to Amazon CloudWatch Logs and set up a CloudWatch metric filter that searches for user logins. If a login is found, send a notification to the security team using Amazon SNS.

C. Set up AWS CloudTrail with Amazon CloudWatch Logs. Subscribe CloudWatch Logs to Amazon Kinesis. Attach AWS Lambda to Kinesis to parse and determine if a log contains a user login. If it does, send a notification to the security team using Amazon SNS.

D. Set up a script on each Amazon EC2 instance to push all logs to Amazon S3. Set up an S3 event to invoke an AWS Lambda function, which invokes an Amazon Athena query to run. The Athena query checks for logins and sends the output to the security team using Amazon SNS.

Question #127 *Topic 1*

A company updated the AWS CloudFormation template for a critical business application. The stack update process failed due to an error in the updated template, and AWS CloudFormation automatically began the stack rollback process. Later, a DevOps engineer discovered that the application was still unavailable and that the stack was in the UPDATE_ROLLBACK_FAILED state.

Which combination of actions should the DevOps engineer perform so that the stack rollback can complete successfully? (Choose two.)

A. Attach the AWSCloudFormationFullAccess IAM policy to the AWS CloudFormation role.

B. Automatically recover the stack resources by using AWS CloudFormation drift detection.

C. Issue a ContinueUpdateRollback command from the AWS CloudFormation console or the AWS CLI.

D. Manually adjust the resources to match the expectations of the stack.

E. Update the existing AWS CloudFormation stack by using the original template.

Question #128 *Topic 1*

A development team manually builds an artifact locally and then places it in an Amazon S3 bucket. The application has a local cache that must be cleared when a deployment occurs. The team runs a command to do this, downloads the artifact from Amazon S3, and unzips the artifact to complete the deployment.

A DevOps team wants to migrate to a CI/CD process and build in checks to stop and roll back the deployment when a failure occurs. This requires the team to track the progression of the deployment.

Which combination of actions will accomplish this? (Choose three.)

A. Allow developers to check the code into a code repository. Using Amazon EventBridge, on every pull into the main branch, invoke an AWS Lambda function to build the artifact and store it in Amazon S3.

B. Create a custom script to clear the cache. Specify the script in the BeforeInstall lifecycle hook in the AppSpec file.

C. Create user data for each Amazon EC2 instance that contains the clear cache script. Once deployed, test the application. If it is not successful, deploy it again.

D. Set up AWS CodePipeline to deploy the application. Allow developers to check the code into a code repository as a source for the pipeline.

E. Use AWS CodeBuild to build the artifact and place it in Amazon S3. Use AWS CodeDeploy to deploy the artifact to Amazon EC2 instances.

F. Use AWS Systems Manager to fetch the artifact from Amazon S3 and deploy it to all the instances.

Question #129 *Topic 1*

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:

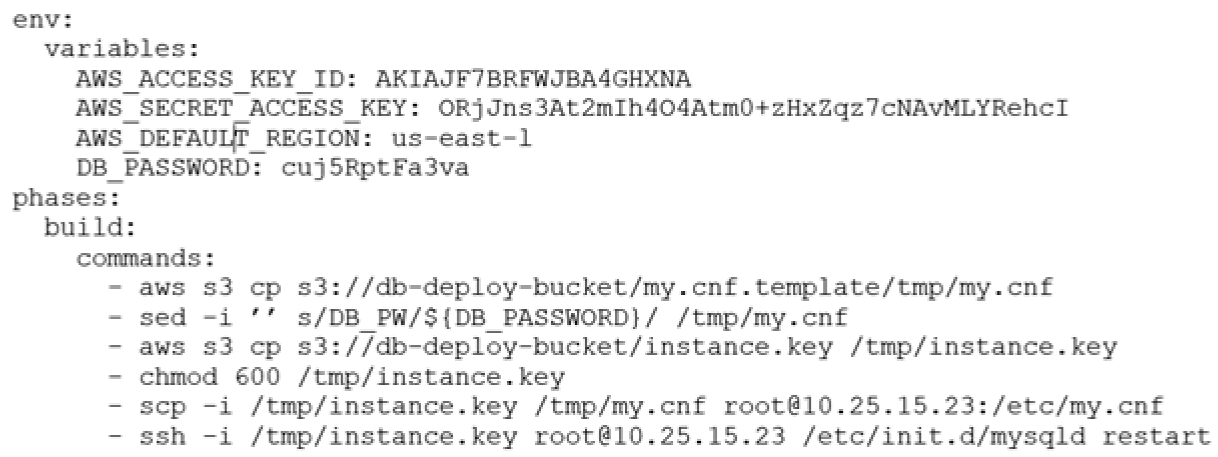

```
env:
  variables:
    AWS_ACCESS_KEY_ID: AKIAJF7BRFWJBA4GHXNA
    AWS_SECRET_ACCESS_KEY: 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
    AWS_DEFAULT_REGION: us-east-1
    DB_PASSWORD: cuj5RptFa3va
phases:
  build:
    commands:
      - aws s3 cp s3://db-deploy-bucket/my.cnf.template/tmp/my.cnf
      - sed -i '' s/DB_PW/${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? (Choose 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.

Question #130 *Topic 1*

A company has a legacy application. A DevOps engineer needs to automate the process of building the deployable artifact for the legacy application. The solution must store the deployable artifact in an existing Amazon S3 bucket for future deployments to reference.

Which solution will meet these requirements in the MOST operationally efficient way?

A. Create a custom Docker image that contains all the dependencies for the legacy application. Store the custom Docker image in a new Amazon Elastic Container Registry (Amazon ECR) repository. Configure a new AWS CodeBuild project to use the custom Docker image to build the deployable artifact and to save the artifact to the S3 bucket.

B. Launch a new Amazon EC2 instance. Install all the dependencies for the legacy application on the EC2 instance. Use the EC2 instance to build the deployable artifact and to save the artifact to the S3 bucket.

C. Create a custom EC2 Image Builder image. Install all the dependencies for the legacy application on the image. Launch a new Amazon EC2 instance from the image. Use the new EC2 instance to build the deployable artifact and to save the artifact to the S3 bucket.

D. Create an Amazon Elastic Kubernetes Service (Amazon EKS) cluster with an AWS Fargate profile that runs in multiple Availability Zones. Create a custom Docker image that contains all the dependencies for the legacy application. Store the custom Docker image in a new Amazon Elastic Container Registry (Amazon ECR) repository. Use the custom Docker image inside the EKS cluster to build the deployable artifact and to save the artifact to the S3 bucket.

Question #131 *Topic 1*

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

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

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

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

A. Update the buildspec.yml file to log in to the ECR repository by using the aws ecr get-login-password AWS CLI command to obtain an authentication token. Update the docker login command to use the authentication token to access the ECR repository.

B. Add an environment variable of type SECRETS_MANAGER to the CodeBuild project. In the environment variable, include the ARN of the CodeBuild project's IAM service role. Update the buildspec.yml file to use the new environment variable to log in with the docker login command to access the ECR repository.

C. Update the ECR repository to be a public image repository. Add an ECR repository policy that allows the IAM service role to have access.

D. Update the buildspec.yml file to use the AWS CLI to assume the IAM service role for ECR operations. Add an ECR repository policy that allows the IAM service role to have access.

Question #132 *Topic 1*

A company manually provisions IAM access for its employees. The company wants to replace the manual process with an automated process. The company has an existing Active Directory system configured with an external SAML 2.0 identity provider (IdP).

The company wants employees to use their existing corporate credentials to access AWS. The groups from the existing Active Directory system must be available for permission management in AWS Identity and Access Management (IAM). A DevOps engineer has completed the initial configuration of AWS IAM Identity Center (AWS Single Sign-On) in the company's AWS account.

What should the DevOps engineer do next to meet the requirements?

A. Configure an external IdP as an identity source. Configure automatic provisioning of users and groups by using the SCIM protocol.

B. Configure AWS Directory Service as an identity source. Configure automatic provisioning of users and groups by using the SAML protocol.

C. Configure an AD Connector as an identity source. Configure automatic provisioning of users and groups by using the SCIM protocol.

D. Configure an external IdP as an identity source Configure automatic provisioning of users and groups by using the SAML protocol.

Question #133 *Topic 1*

A company is using AWS to run digital workloads. Each application team in the company has its own AWS account for application hosting. The accounts are consolidated in an organization in AWS Organizations.

The company wants to enforce security standards across the entire organization. To avoid noncompliance because of security misconfiguration, the company has enforced the use of AWS CloudFormation. A production support team can modify resources in the production environment by using the AWS Management Console to troubleshoot and resolve application-related issues.

A DevOps engineer must implement a solution to identify in near real time any AWS service misconfiguration that results in noncompliance. The solution must automatically remediate the issue within 15 minutes of identification. The solution also must track noncompliant resources and events in a centralized dashboard with accurate timestamps.

Which solution will meet these requirements with the LEAST development overhead?

A. Use CloudFormation drift detection to identify noncompliant resources. Use drift detection events from CloudFormation to invoke an AWS Lambda function for remediation. Configure the Lambda function to publish logs to an Amazon CloudWatch Logs log group. Configure an Amazon CloudWatch dashboard to use the log group for tracking.

B. Turn on AWS CloudTrail in the AWS accounts. Analyze CloudTrail logs by using Amazon Athena to identify noncompliant resources. Use AWS Step Functions to track query results on Athena for drift detection and to invoke an AWS Lambda function for remediation. For tracking, set up an Amazon QuickSight dashboard that uses Athena as the data source.

C. Turn on the configuration recorder in AWS Config in all the AWS accounts to identify noncompliant resources. Enable AWS Security Hub with the --no-enable-default-standards option in all the AWS accounts. Set up AWS Config managed rules and custom rules. Set up automatic remediation by using AWS Config conformance packs. For tracking, set up a dashboard on Security Hub in a designated Security Hub administrator account.

D. Turn on AWS CloudTrail in the AWS accounts. Analyze CloudTrail logs by using Amazon CloudWatch Logs to identify noncompliant resources. Use CloudWatch Logs filters for drift detection. Use Amazon EventBridge to invoke the Lambda function for remediation. Stream filtered CloudWatch logs to Amazon OpenSearch Service. Set up a dashboard on OpenSearch Service for tracking.

Question #134 *Topic 1*

A company uses AWS Organizations to manage its AWS accounts. The organization root has an OU that is named Environments. The Environments OU has two child OUs that are named Development and Production, respectively.

The Environments OU and the child OUs have the default FullAWSAccess policy in place. A DevOps engineer plans to remove the FullAWSAccess policy from the Development OU and replace the policy with a policy that allows all actions on Amazon EC2 resources.

What will be the outcome of this policy replacement?

A. All users in the Development OU will be allowed all API actions on all resources.

B. All users in the Development OU will be allowed all API actions on EC2 resources. All other API actions will be denied.

C. All users in the Development OU will be denied all API actions on all resources.

D. All users in the Development OU will be denied all API actions on EC2 resources. All other API actions will be allowed.

Question #135 *Topic 1*

A company is examining its disaster recovery capability and wants the ability to switch over its daily operations to a secondary AWS Region. The company uses AWS CodeCommit as a source control tool in the primary Region.

A DevOps engineer must provide the capability for the company to develop code in the secondary Region. If the company needs to use the secondary Region, developers can add an additional remote URL to their local Git configuration.

Which solution will meet these requirements?

A. Create a CodeCommit repository in the secondary Region. Create an AWS CodeBuild project to perform a Git mirror operation of the primary Region's CodeCommit repository to the secondary Region's CodeCommit repository. Create an AWS Lambda function that invokes the CodeBuild project. Create an Amazon EventBridge rule that reacts to merge events in the primary Region's CodeCommit repository. Configure the EventBridge rule to invoke the Lambda function.

B. Create an Amazon S3 bucket in the secondary Region. Create an AWS Fargate task to perform a Git mirror operation of the primary Region's CodeCommit repository and copy the result to the S3 bucket. Create an AWS Lambda function that initiates the Fargate task. Create an Amazon EventBridge rule that reacts to merge events in the CodeCommit repository. Configure the EventBridge rule to invoke the Lambda function.

C. Create an AWS CodeArtifact repository in the secondary Region. Create an AWS CodePipeline pipeline that uses the primary Region's CodeCommit repository for the source action. Create a cross-Region stage in the pipeline that packages the CodeCommit repository contents and stores the contents in the CodeArtifact repository when a pull request is merged into the CodeCommit repository.

D. Create an AWS Cloud9 environment and a CodeCommit repository in the secondary Region. Configure the primary Region's CodeCommit repository as a remote repository in the AWS Cloud9 environment. Connect the secondary Region's CodeCommit repository to the AWS Cloud9 environment.

Question #136 *Topic 1*

A DevOps team is merging code revisions for an application that uses an Amazon RDS Multi-AZ DB cluster for its production database. The DevOps team uses continuous integration to periodically verify that the application works. The DevOps team needs to test the changes before the changes are deployed to the production database.

Which solution will meet these requirements?

A. Use a buildspec file in AWS CodeBuild to restore the DB cluster from a snapshot of the production database, run integration tests, and drop the restored database after verification.

B. Deploy the application to production. Configure an audit log of data control language (DCL) operations to capture database activities to perform if verification fails.

C. Create a snapshot of the DB cluster before deploying the application. Use the Update requires:Replacement property on the DB instance in AWS CloudFormation to deploy the application and apply the changes.

D. Ensure that the DB cluster is a Multi-AZ deployment. Deploy the application with the updates. Fail over to the standby instance if verification fails.

Question #137 *Topic 1*

A company manages a multi-tenant environment in its VPC and has configured Amazon GuardDuty for the corresponding AWS account. The company sends all GuardDuty findings to AWS Security Hub.

Traffic from suspicious sources is generating a large number of findings. A DevOps engineer needs to implement a solution to automatically deny traffic across the entire VPC when GuardDuty discovers a new suspicious source.

Which solution will meet these requirements?

A. Create a GuardDuty threat list. Configure GuardDuty to reference the list. Create an AWS Lambda function that will update the threat list. Configure the Lambda function to run in response to new Security Hub findings that come from GuardDuty.

B. Configure an AWS WAF web ACL that includes a custom rule group. Create an AWS Lambda function that will create a block rule in the custom rule group. Configure the Lambda function to run in response to new Security Hub findings that come from GuardDuty.

C. Configure a firewall in AWS Network Firewall. Create an AWS Lambda function that will create a Drop action rule in the firewall policy. Configure the Lambda function to run in response to new Security Hub findings that come from GuardDuty.

D. Create an AWS Lambda function that will create a GuardDuty suppression rule. Configure the Lambda function to run in response to new Security Hub findings that come from GuardDuty.

Question #138 *Topic 1*

A company uses AWS Secrets Manager to store a set of sensitive API keys that an AWS Lambda function uses. When the Lambda function is invoked the Lambda function retrieves the API keys and makes an API call to an external service. The Secrets Manager secret is encrypted with the default AWS Key Management Service (AWS KMS) key.

A DevOps engineer needs to update the infrastructure to ensure that only the Lambda function’s execution role can access the values in Secrets Manager. The solution must apply the principle of least privilege.

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

A. Update the default KMS key for Secrets Manager to allow only the Lambda function’s execution role to decrypt

B. Create a KMS customer managed key that trusts Secrets Manager and allows the Lambda function's execution role to decrypt. Update Secrets Manager to use the new customer managed key

C. Create a KMS customer managed key that trusts Secrets Manager and allows the account's root principal to decrypt. Update Secrets Manager to use the new customer managed key

D. Ensure that the Lambda function’s execution role has the KMS permissions scoped on the resource level. Configure the permissions so that the KMS key can encrypt the Secrets Manager secret

E. Remove all KMS permissions from the Lambda function’s execution role

Question #139 *Topic 1*

A company's DevOps engineer is creating an AWS Lambda function to process notifications from an Amazon Simple Notification Service (Amazon SNS) topic. The Lambda function will process the notification messages and will write the contents of the notification messages to an Amazon RDS Multi-AZ DB instance.

During testing, a database administrator accidentally shut down the DB instance. While the database was down the company lost several of the SNS notification messages that were delivered during that time.

The DevOps engineer needs to prevent the loss of notification messages in the future.

Which solutions will meet this requirement? (Choose two.)

A. Replace the RDS Multi-AZ DB instance with an Amazon DynamoDB table.

B. Configure an Amazon Simple Queue Service (Amazon SQS) queue as a destination of the Lambda function.

C. Configure an Amazon Simple Queue Service (Amazon SQS) dead-letter queue for the SNS topic.

D. Subscribe an Amazon Simple Queue Service (Amazon SQS) queue to the SNS topic. Configure the Lambda function to process messages from the SQS queue.

E. Replace the SNS topic with an Amazon EventBridge event bus. Configure an EventBridge rule on the new event bus to invoke the Lambda function for each event.

Question #140 *Topic 1*

A company has an application that runs on Amazon EC2 instances. The company uses an AWS CodePipeline pipeline to deploy the application into multiple AWS Regions. The pipeline is configured with a stage for each Region. Each stage contains an AWS CloudFormation action for each Region.

When the pipeline deploys the application to a Region, the company wants to confirm that the application is in a healthy state before the pipeline moves on to the next Region. Amazon Route 53 record sets are configured for the application in each Region. A DevOps engineer creates a Route 53 health check that is based on an Amazon CloudWatch alarm for each Region where the application is deployed.

What should the DevOps engineer do next to meet the requirements?

A. Create an AWS Step Functions workflow to check the state of the CloudWatch alarm. Configure the Step Functions workflow to exit with an error if the alarm is in the ALARM state. Create a new stage in the pipeline between each Region deployment stage. In each new stage, include an action to invoke the Step Functions workflow.

B. Configure an AWS CodeDeploy application to deploy a CloudFormation template with automatic rollback. Configure the CloudWatch alarm as the instance health check for the CodeDeploy application. Remove the CloudFormation actions from the pipeline. Create a CodeDeploy action in the pipeline stage for each Region.

C. Create a new pipeline stage for each Region where the application is deployed. Configure a CloudWatch alarm action for the new stage to check the state of the CloudWatch alarm and to exit with an error if the alarm is in the ALARM state

D. Configure the CloudWatch agent on the EC2 instances to report the application status to the Route 53 health check. Create a new pipeline stage for each Region where the application is deployed. Configure a CloudWatch alarm action to exit with an error if the CloudWatch alarm is in the ALARM state.

Question #141 *Topic 1*

A company plans to use Amazon CloudWatch to monitor its Amazon EC2 instances. The company needs to stop EC2 instances when the average of the NetworkPacketsIn metric is less than 5 for at least 3 hours in a 12-hour time window. The company must evaluate the metric every hour. The EC2 instances must continue to run if there is missing data for the NetworkPacketsIn metric during the evaluation period.

A DevOps engineer creates a CloudWatch alarm for the NetworkPacketsIn metric. The DevOps engineer configures a threshold value of 5 and an evaluation period of 1 hour.

Which set of additional actions should the DevOps engineer take to meet these requirements?

A. Configure the Datapoints to Alarm value to be 3 out of 12. Configure the alarm to treat missing data as breaching the threshold. Add an AWS Systems Manager action to stop the instance when the alarm enters the ALARM state.

B. Configure the Datapoints to Alarm value to be 3 out of 12. Configure the alarm to treat missing data as not breaching the threshold. Add an EC2 action to stop the instance when the alarm enters the ALARM state.

C. Configure the Datapoints to Alarm value to be 9 out of 12. Configure the alarm to treat missing data as breaching the threshold. Add an EC2 action to stop the instance when the alarm enters the ALARM state.

D. Configure the Datapoints to Alarm value to be 9 out of 12. Configure the alarm to treat missing data as not breaching the threshold. Add an AWS Systems Manager action to stop the instance when the alarm enters the ALARM state.

Question #142 *Topic 1*

A company manages 500 AWS accounts that are in an organization in AWS Organizations. The company discovers many unattached Amazon Elastic Block Store (Amazon EBS) volumes in all the accounts. The company wants to automatically tag the unattached EBS volumes for investigation.

A DevOps engineer needs to deploy an AWS Lambda function to all the AWS accounts. The Lambda function must run every 30 minutes to tag all the EBS volumes that have been unattached for a period of 7 days or more.

Which solution will meet these requirements in the MOST operationally efficient manner?

A. Configure a delegated administrator account for the organization. Create an AWS CloudFormation template that contains the Lambda function. Use CloudFormation StackSets to deploy the CloudFormation template from the delegated administrator account to all the member accounts in the organization. Create an Amazon EventBridge event bus in the delegated administrator account to invoke the Lambda function in each member account every 30 minutes.

B. Create a cross-account IAM role in the organization's member accounts. Attach the AWSLambda_FullAccess policy and the AWSCloudFormationFullAccess policy to the role. Create an AWS CloudFormation template that contains the Lambda function and an Amazon EventBridge scheduled rule to invoke the Lambda function every 30 minutes. Create a custom script in the organization's management account that assumes the role and deploys the CloudFormation template to the member accounts.

C. Configure a delegated administrator account for the organization. Create an AWS CloudFormation template that contains the Lambda function and an Amazon EventBridge scheduled rule to invoke the Lambda function every 30 minutes. Use CloudFormation StackSets to deploy the CloudFormation template from the delegated administrator account to all the member accounts in the organization

D. Create a cross-account IAM role in the organization's member accounts. Attach the AmazonS3FullAccess policy and the AWSCodeDeployDeployerAccess policy to the role. Use AWS CodeDeploy to assume the role to deploy the Lambda function from the organization's management account. Configure an Amazon EventBridge scheduled rule in the member accounts to invoke the Lambda function every 30 minutes.

Question #143 *Topic 1*

A company's production environment uses an AWS CodeDeploy blue/green deployment to deploy an application. The deployment incudes Amazon EC2 Auto Scaling groups that launch instances that run Amazon Linux 2.

A working appspec.yml file exists in the code repository and contains the following text:

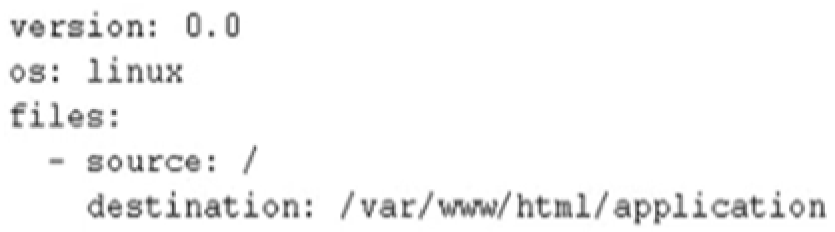

```
version: 0.0
os: linux
files:
  - source: /
    destination: /var/www/html/application
```

A DevOps engineer needs to ensure that a script downloads and installs a license file onto the instances before the replacement instances start to handle request traffic. The DevOps engineer adds a hooks section to the appspec.yml file.

Which hook should the DevOps engineer use to run the script that downloads and installs the license file?

A. AfterBlockTraffic

B. BeforeBlockTraffic

C. BeforeInstall

D. DownloadBundle

Question #144 *Topic 1*

A company has an application that includes AWS Lambda functions. The Lambda functions run Python code that is stored in an AWS CodeCommit repository. The company has recently experienced failures in the production environment because of an error in the Python code. An engineer has written unit tests for the Lambda functions to help avoid releasing any future defects into the production environment.

The company's DevOps team needs to implement a solution to integrate the unit tests into an existing AWS CodePipeline pipeline. The solution must produce reports about the unit tests for the company to view.

Which solution will meet these requirements?

A. Associate the CodeCommit repository with Amazon CodeGuru Reviewer. Create a new AWS CodeBuild project. In the CodePipeline pipeline, configure a test stage that uses the new CodeBuild project. Create a buildspec.yml file in the CodeCommit repository. In the buildspec yml file, define the actions to run a CodeGuru review.

B. Create a new AWS CodeBuild project. In the CodePipeline pipeline, configure a test stage that uses the new CodeBuild project. Create a CodeBuild report group. Create a buildspec.yml file in the CodeCommit repository. In the buildspec.yml file, define the actions to run the unit tests with an output of JUNITXML in the build phase section. Configure the test reports to be uploaded to the new CodeBuild report group.

C. Create a new AWS CodeArtifact repository. Create a new AWS CodeBuild project. In the CodePipeline pipeline, configure a test stage that uses the new CodeBuild project. Create an appspec.yml file in the original CodeCommit repository. In the appspec.yml file, define the actions to run the unit tests with an output of CUCUMBERJSON in the build phase section. Configure the tests reports to be sent to the new CodeArtifact repository.

D. Create a new AWS CodeBuild project. In the CodePipeline pipeline, configure a test stage that uses the new CodeBuild project. Create a new Amazon S3 bucket. Create a buildspec.yml file in the CodeCommit repository. In the buildspec yml file, define the actions to run the unit tests with an output of HTML in the phases section. In the reports section, upload the test reports to the S3 bucket.

Question #145 *Topic 1*

A company manages multiple AWS accounts in AWS Organizations. The company's security policy states that AWS account root user credentials for member accounts must not be used. The company monitors access to the root user credentials.

A recent alert shows that the root user in a member account launched an Amazon EC2 instance. A DevOps engineer must create an SCP at the organization's root level that will prevent the root user in member accounts from making any AWS service API calls.

Which SCP will meet these requirements?

A.

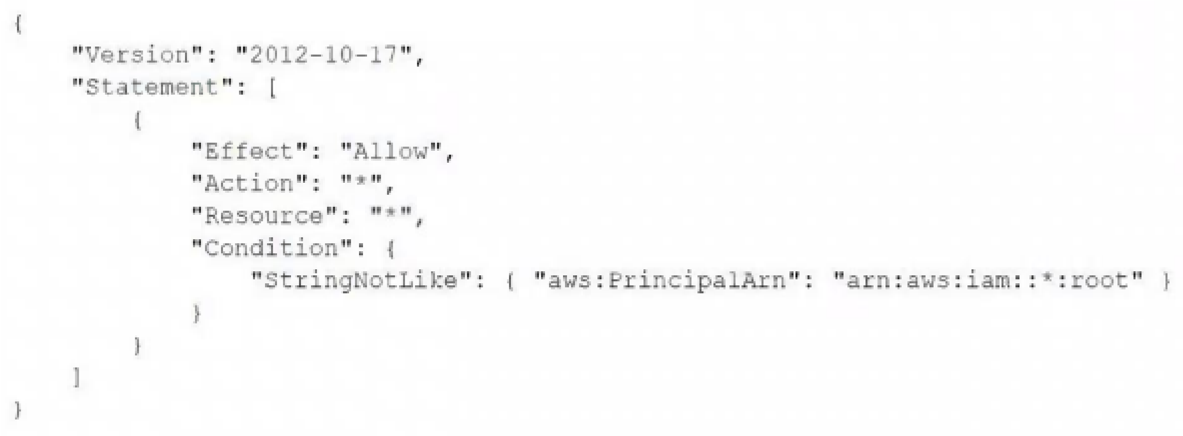

```
{
    "Version": "2012-10-17",
    "Statement": [
        {
            "Effect": "Allow",
            "Action": "*",
            "Resource": "*",
            "Condition": {
                "StringNotLike": { "aws:PrincipalArn": "arn:aws:iam::*:root" }
            }
        }
    ]
}
```

B.

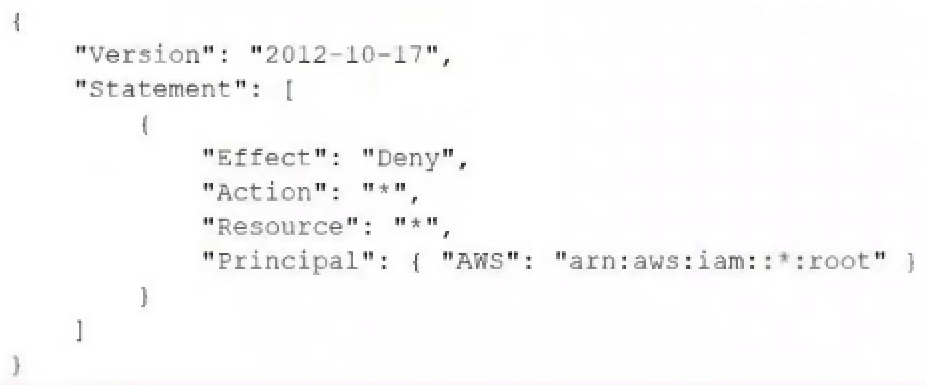

```
{
    "Version": "2012-10-17",
    "Statement": [
        {
            "Effect": "Deny",
            "Action": "*",
            "Resource": "*",
            "Principal": { "AWS": "arn:aws:iam::*:root" }
        }
    ]
}
```

C.

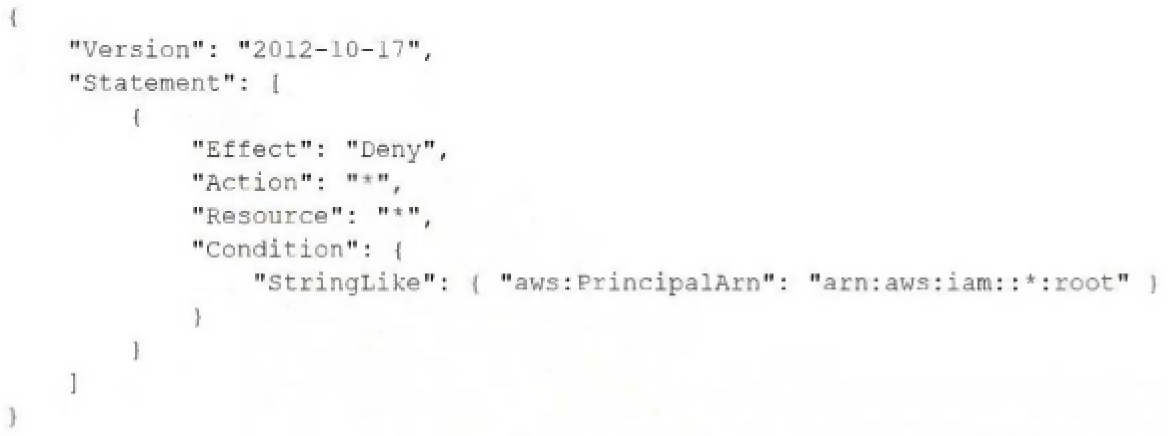

```
{
    "Version": "2012-10-17",
    "Statement": [
        {
            "Effect": "Deny",
            "Action": "*",
            "Resource": "*",
            "Condition": {
                "StringLike": { "aws:PrincipalArn": "arn:aws:iam::*:root" }
            }
        }
    ]
}
```

D.

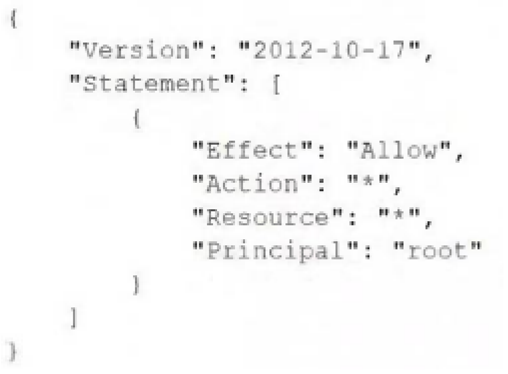

```
{
    "Version": "2012-10-17",
    "Statement": [
        {
            "Effect": "Allow",
            "Action": "*",
            "Resource": "*",
            "Principal": "root"
        }
    ]
}
```

Question #146 *Topic 1*

A company uses AWS and has a VPC that contains critical compute infrastructure with predictable traffic patterns. The company has configured VPC flow logs that are published to a log group in Amazon CloudWatch Logs.

The company's DevOps team needs to configure a monitoring solution for the VPC flow logs to identify anomalies in network traffic to the VPC over time. If the monitoring solution detects an anomaly, the company needs the ability to initiate a response to the anomaly.

How should the DevOps team configure the monitoring solution to meet these requirements?

A. Create an Amazon Kinesis data stream. Subscribe the log group to the data stream. Configure Amazon Kinesis Data Analytics to detect log anomalies in the data stream. Create an AWS Lambda function to use as the output of the data stream. Configure the Lambda function to write to the default Amazon EventBridge event bus in the event of an anomaly finding.

B. Create an Amazon Kinesis Data Firehose delivery stream that delivers events to an Amazon S3 bucket. Subscribe the log group to the delivery stream. Configure Amazon Lookout for Metrics to monitor the data in the S3 bucket for anomalies. Create an AWS Lambda function to run in response to Lookout for Metrics anomaly findings. Configure the Lambda function to publish to the default Amazon EventBridge event bus.

C. Create an AWS Lambda function to detect anomalies. Configure the Lambda function to publish an event to the default Amazon EventBridge event bus if the Lambda function detects an anomaly. Subscribe the Lambda function to the log group.

D. Create an Amazon Kinesis data stream. Subscribe the log group to the data stream. Create an AWS Lambda function to detect log anomalies. Configure the Lambda function to write to the default Amazon EventBridge event bus if the Lambda function detects an anomaly. Set the Lambda function as the processor for the data stream.

Question #147 *Topic 1*

AnyCompany is using AWS Organizations to create and manage multiple AWS accounts. AnyCompany recently acquired a smaller company, Example Corp. During the acquisition process, Example Corp's single AWS account joined AnyCompany's management account through an Organizations invitation. AnyCompany moved the new member account under an OU that is dedicated to Example Corp.

AnyCompany's DevOps engineer has an IAM user that assumes a role that is named OrganizationAccountAccessRole to access member accounts. This role is configured with a full access policy. When the DevOps engineer tries to use the AWS Management Console to assume the role in Example Corp's new member account, the DevOps engineer receives the following error message: "Invalid information in one or more fields. Check your information or contact your administrator."

Which solution will give the DevOps engineer access to the new member account?

A. In the management account, grant the DevOps engineer's IAM user permission to assume the OrganizationAccountAccessRole IAM role in the new member account.

B. In the management account, create a new SCP. In the SCP, grant the DevOps engineer's IAM user full access to all resources in the new member account. Attach the SCP to the OU that contains the new member account.

C. In the new member account, create a new IAM role that is named OrganizationAccountAccessRole. Attach the AdministratorAccess AWS managed policy to the role. In the role's trust policy, grant the management account permission to assume the role.

D. In the new member account, edit the trust policy for the OrganizationAccountAccessRole IAM role. Grant the management account permission to assume the role.

Question #148 *Topic 1*

A DevOps engineer is designing an application that integrates with a legacy REST API. The application has an AWS Lambda function that reads records from an Amazon Kinesis data stream. The Lambda function sends the records to the legacy REST API.

Approximately 10% of the records that the Lambda function sends from the Kinesis data stream have data errors and must be processed manually. The Lambda function event source configuration has an Amazon Simple Queue Service (Amazon SQS) dead-letter queue as an on-failure destination. The DevOps engineer has configured the Lambda function to process records in batches and has implemented retries in case of failure.

During testing, the DevOps engineer notices that the dead-letter queue contains many records that have no data errors and that already have been processed by the legacy REST API. The DevOps engineer needs to configure the Lambda function's event source options to reduce the number of errorless records that are sent to the dead-letter queue.

Which solution will meet these requirements?

A. Increase the retry attempts.

B. Configure the setting to split the batch when an error occurs.

C. Increase the concurrent batches per shard.

D. Decrease the maximum age of record.

Question #149 *Topic 1*

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 passed. Use AWS CodeDeploy and create a Canary deployment configuration that specifies the percentage of traffic and interval.

B. Create an AWS CodeBuild configuration that triggers when the test code is pushed. Use AWS CloudFormation to trigger an AWS CodePipeline configuration that deploys the new Lambda versions and specifies the traffic shift percentage and interval.

C. Create an AWS CodePipeline configuration and set up the source code step to trigger when code is pushed. Set up the build step to use AWS CodeBuild to run the tests. Set up an AWS CodeDeploy configuration to deploy, then select the CodeDeployDefault.LambdaLinear10PercentEvery3Minutes option.

D. Use the AWS CLI to set up a post-commit hook that uploads the code to an Amazon S3 bucket after tests have passed Set up an S3 event trigger that runs a Lambda function that deploys the new version. Use an interval in the Lambda function to deploy the code over time at the required percentage.

Question #150 *Topic 1*

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

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

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

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

Question #151 *Topic 1*

A DevOps engineer wants to find a solution to migrate an application from on premises to AWS. The application is running on Linux and needs to run on specific versions of Apache Tomcat, HAProxy, and Varnish Cache to function properly. The application's operating system-level parameters require tuning. The solution must include a way to automate the deployment of new application versions. The infrastructure should be scalable and faulty servers should be replaced automatically.

Which solution should the DevOps engineer use?

A. Upload the application as a Docker image that contains all the necessary software to Amazon ECR. Create an Amazon ECS cluster using an AWS Fargate launch type and an Auto Scaling group. Create an AWS CodePipeline pipeline that uses Amazon ECR as a source and Amazon ECS as a deployment provider.

B. Upload the application code to an AWS CodeCommit repository with a saved configuration file to configure and install the software. Create an AWS Elastic Beanstalk web server tier and a load balanced-type environment that uses the Tomcat solution stack. Create an AWS CodePipeline pipeline that uses CodeCommit as a source and Elastic Beanstalk as a deployment provider.

C. Upload the application code to an AWS CodeCommit repository with a set of .ebextensions files to configure and install the software. Create an AWS Elastic Beanstalk worker tier environment that uses the Tomcat solution stack. Create an AWS CodePipeline pipeline that uses CodeCommit as a source and Elastic Beanstalk as a deployment provider.

D. Upload the application code to an AWS CodeCommit repository with an appspec.yml file to configure and install the necessary software. Create an AWS CodeDeploy deployment group associated with an Amazon EC2 Auto Scaling group. Create an AWS CodePipeline pipeline that uses CodeCommit as a source and CodeDeploy as a deployment provider.

Question #152 *Topic 1*

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

What is likely causing this issue?

A. The two affected instances failed to fetch the new deployment.

B. A failed AfterInstall lifecycle event hook caused the CodeDeploy agent to roll back to the previous version on the affected instances.

C. The CodeDeploy agent was not installed in two affected instances.

D. EC2 Auto Scaling launched two new instances while the new deployment had not yet finished, causing the previous version to be deployed on the affected instances.

Question #153 *Topic 1*

A security team is concerned that a developer can unintentionally attach an Elastic IP address to an Amazon EC2 instance in production. No developer should be allowed to attach an Elastic IP address to an instance. The security team must be notified if any production server has an Elastic IP address at any time.

How can this task be automated?

A. Use Amazon Athena to query AWS CloudTrail logs to check for any associate-address attempts. Create an AWS Lambda function to disassociate the Elastic IP address from the instance, and alert the security team.

B. Attach an IAM policy to the developers' IAM group to deny associate-address permissions. Create a custom AWS Config rule to check whether an Elastic IP address is associated with any instance tagged as production, and alert the security team.

C. Ensure that all IAM groups associated with developers do not have associate-address permissions. Create a scheduled AWS Lambda function to check whether an Elastic IP address is associated with any instance tagged as production, and alert the security team if an instance has an Elastic IP address associated with it.

D. Create an AWS Config rule to check that all production instances have EC2 IAM roles that include deny associate-address permissions. Verify whether there is an Elastic IP address associated with any instance, and alert the security team if an instance has an Elastic IP address associated with it.

Question #154 *Topic 1*

A company is using AWS Organizations to create separate AWS accounts for each of its departments. The company needs to automate the following tasks:

- Update the Linux AMIs with new patches periodically and generate a golden image
- Install a new version of Chef agents in the golden image, if available
- Provide the newly generated AMIs to the department's accounts

Which solution meets these requirements with the LEAST management overhead?

A. Write a script to launch an Amazon EC2 instance from the previous golden image. Apply the patch updates. Install the new version of the Chef agent, generate a new golden image, and then modify the AMI permissions to share only the new image with the department's accounts.

B. Use Amazon EC2 Image Builder to create an image pipeline that consists of the base Linux AMI and components to install the Chef agent. Use AWS Resource Access Manager to share EC2 Image Builder images with the department's accounts.

C. Use an AWS Systems Manager Automation runbook to update the Linux AMI by using the previous image. Provide the URL for the script that will update the Chef agent. Use AWS Organizations to replace the previous golden image in the department's accounts.

D. Use Amazon EC2 Image Builder to create an image pipeline that consists of the base Linux AMI and components to install the Chef agent. Create a parameter in AWS Systems Manager Parameter Store to store the new AMI ID that can be referenced by the department's accounts.

Question #155 *Topic 1*

A company has a mission-critical application on AWS that uses automatic scaling. The company wants the deployment lifecycle to meet the following parameters:

- The application must be deployed one instance at a time to ensure the remaining fleet continues to serve traffic.
- The application is CPU intensive and must be closely monitored.
- The deployment must automatically roll back if the CPU utilization of the deployment instance exceeds 85%.

Which solution will meet these requirements?

A. Use AWS CloudFormation to create an AWS Step Functions state machine and Auto Scaling lifecycle hooks to move to one instance at a time into a wait state. Use AWS Systems Manager automation to deploy the update to each instance and move it back into the Auto Scaling group using the heartbeat timeout.

B. Use AWS CodeDeploy with Amazon EC2 Auto Scaling Configure an alarm tied to the CPU utilization metric. Use the CodeDeployDefault OneAtAtime configuration as a deployment strategy. Configure automatic rollbacks within the deployment group to roll back the deployment if the alarm thresholds are breached.

C. Use AWS Elastic Beanstalk for load balancing and AWS Auto Scaling. Configure an alarm tied to the CPU utilization metric. Configure rolling deployments with a fixed batch size of one instance. Enable enhanced health to monitor the status of the deployment and roll back based on the alarm previously created.

D. Use AWS Systems Manager to perform a blue/green deployment with Amazon EC2 Auto Scaling. Configure an alarm tied to the CPU utilization metric. Deploy updates one at a time. Configure automatic rollbacks within the Auto Scaling group to roll back the deployment if the alarm thresholds are breached.

Question #156 *Topic 1*

A company has a single developer writing code for an automated deployment pipeline. The developer is storing source code in an Amazon S3 bucket for each project. The company wants to add more developers to the team but is concerned about code conflicts and lost work. The company also wants to build a test environment to deploy newer versions of code for testing and allow developers to automatically deploy to both environments when code is changed in the repository.

What is the MOST efficient way to meet these requirements?

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

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

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

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

Question #157 *Topic 1*

A DevOps engineer notices that all Amazon EC2 instances running behind an Application Load Balancer in an Auto Scaling group are failing to respond to user requests. The EC2 instances are also failing target group HTTP health checks.

Upon inspection, the engineer notices the application process was not running in any EC2 instances. There are a significant number of out of memory messages in the system logs. The engineer needs to improve the resilience of the application to cope with a potential application memory leak. Monitoring and notifications should be enabled to alert when there is an issue.

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

A. Change the Auto Scaling configuration to replace the instances when they fail the load balancer's health checks.

B. Change the target group health check HealthCheckIntervalSeconds parameter to reduce the interval between health checks.

C. Change the target group health checks from HTTP to TCP to check if the port where the application is listening is reachable.

D. Enable the available memory consumption metric within the Amazon CloudWatch dashboard for the entire Auto Scaling group. Create an alarm when the memory utilization is high. Associate an Amazon SNS topic to the alarm to receive notifications when the alarm goes off.

E. Use the Amazon CloudWatch agent to collect the memory utilization of the EC2 instances in the Auto Scaling group. Create an alarm when the memory utilization is high and associate an Amazon SNS topic to receive a notification.

Question #158 *Topic 1*

An ecommerce company uses a large number of Amazon Elastic Block Store (Amazon EBS) backed Amazon EC2 instances. To decrease manual work across all the instances, a DevOps engineer is tasked with automating restart actions when EC2 instance retirement events are scheduled.

How can this be accomplished?

A. Create a scheduled Amazon EventBridge rule to run an AWS Systems Manager Automation runbook that checks if any EC2 instances are scheduled for retirement once a week. If the instance is scheduled for retirement, the runbook will hibernate the instance.

B. Enable EC2 Auto Recovery on all of the instances. Create an AWS Config rule to limit the recovery to occur during a maintenance window only.

C. Reboot all EC2 instances during an approved maintenance window that is outside of standard business hours. Set up Amazon CloudWatch alarms to send a notification in case any instance is failing EC2 instance status checks.

D. Set up an AWS Health Amazon EventBridge rule to run AWS Systems Manager Automation runbooks that stop and start the EC2 instance when a retirement scheduled event occurs.

Question #159 *Topic 1*

A company manages AWS accounts for application teams in AWS Control Tower. Individual application teams are responsible for securing their respective AWS accounts.

A DevOps engineer needs to enable Amazon GuardDuty for all AWS accounts in which the application teams have not already enabled GuardDuty. The DevOps engineer is using AWS CloudFormation StackSets from the AWS Control Tower management account.

How should the DevOps engineer configure the CloudFormation template to prevent failure during the StackSets deployment?

A. Create a CloudFormation custom resource that invokes an AWS Lambda function. Configure the Lambda function to conditionally enable GuardDuty if GuardDuty is not already enabled in the accounts.

B. Use the Conditions section of the CloudFormation template to enable GuardDuty in accounts where GuardDuty is not already enabled.

C. Use the CloudFormation Fn::GetAtt intrinsic function to check whether GuardDuty is already enabled. If GuardDuty is not already enabled, use the Resources section of the CloudFormation template to enable GuardDuty.

D. Manually discover the list of AWS account IDs where GuardDuty is not enabled. Use the CloudFormation Fn::ImportValue intrinsic function to import the list of account IDs into the CloudFormation template to skip deployment for the listed AWS accounts.

Question #160 *Topic 1*

A company has an AWS Control Tower landing zone. The company's DevOps team creates a workload OU. A development OU and a production OU are nested under the workload OU. The company grants users full access to the company's AWS accounts to deploy applications.

The DevOps team needs to allow only a specific management IAM role to manage the IAM roles and policies of any AWS accounts in only the production OU.

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

A. Create an SCP that denies full access with a condition to exclude the management IAM role for the organization root.

B. Ensure that the FullAWSAccess SCP is applied at the organization root.

C. Create an SCP that allows IAM related actions. Attach the SCP to the development OU.

D. Create an SCP that denies IAM related actions with a condition to exclude the management IAM role. Attach the SCP to the workload OU.

E. Create an SCP that denies IAM related actions with a condition to exclude the management IAM role. Attach the SCP to the production OU.

Question #161 *Topic 1*

A company hired a penetration tester to simulate an internal security breach. The tester performed port scans on the company's Amazon EC2 instances. The company's security measures did not detect the port scans.

The company needs a solution that automatically provides notification when port scans are performed on EC2 instances. The company creates and subscribes to an Amazon Simple Notification Service (Amazon SNS) topic.

What should the company do next to meet the requirement?

A. Ensure that Amazon GuardDuty is enabled. Create an Amazon CloudWatch alarm for detected EC2 and port scan findings. Connect the alarm to the SNS topic.

B. Ensure that Amazon Inspector is enabled. Create an Amazon EventBridge event for detected network reachability findings that indicate port scans. Connect the event to the SNS topic.

C. Ensure that Amazon Inspector is enabled. Create an Amazon EventBridge event for detected CVEs that cause open port vulnerabilities. Connect the event to the SNS topic.

D. Ensure that AWS CloudTrail is enabled. Create an AWS Lambda function to analyze the CloudTrail logs for unusual amounts of traffic from an IP address range. Connect the Lambda function to the SNS topic.

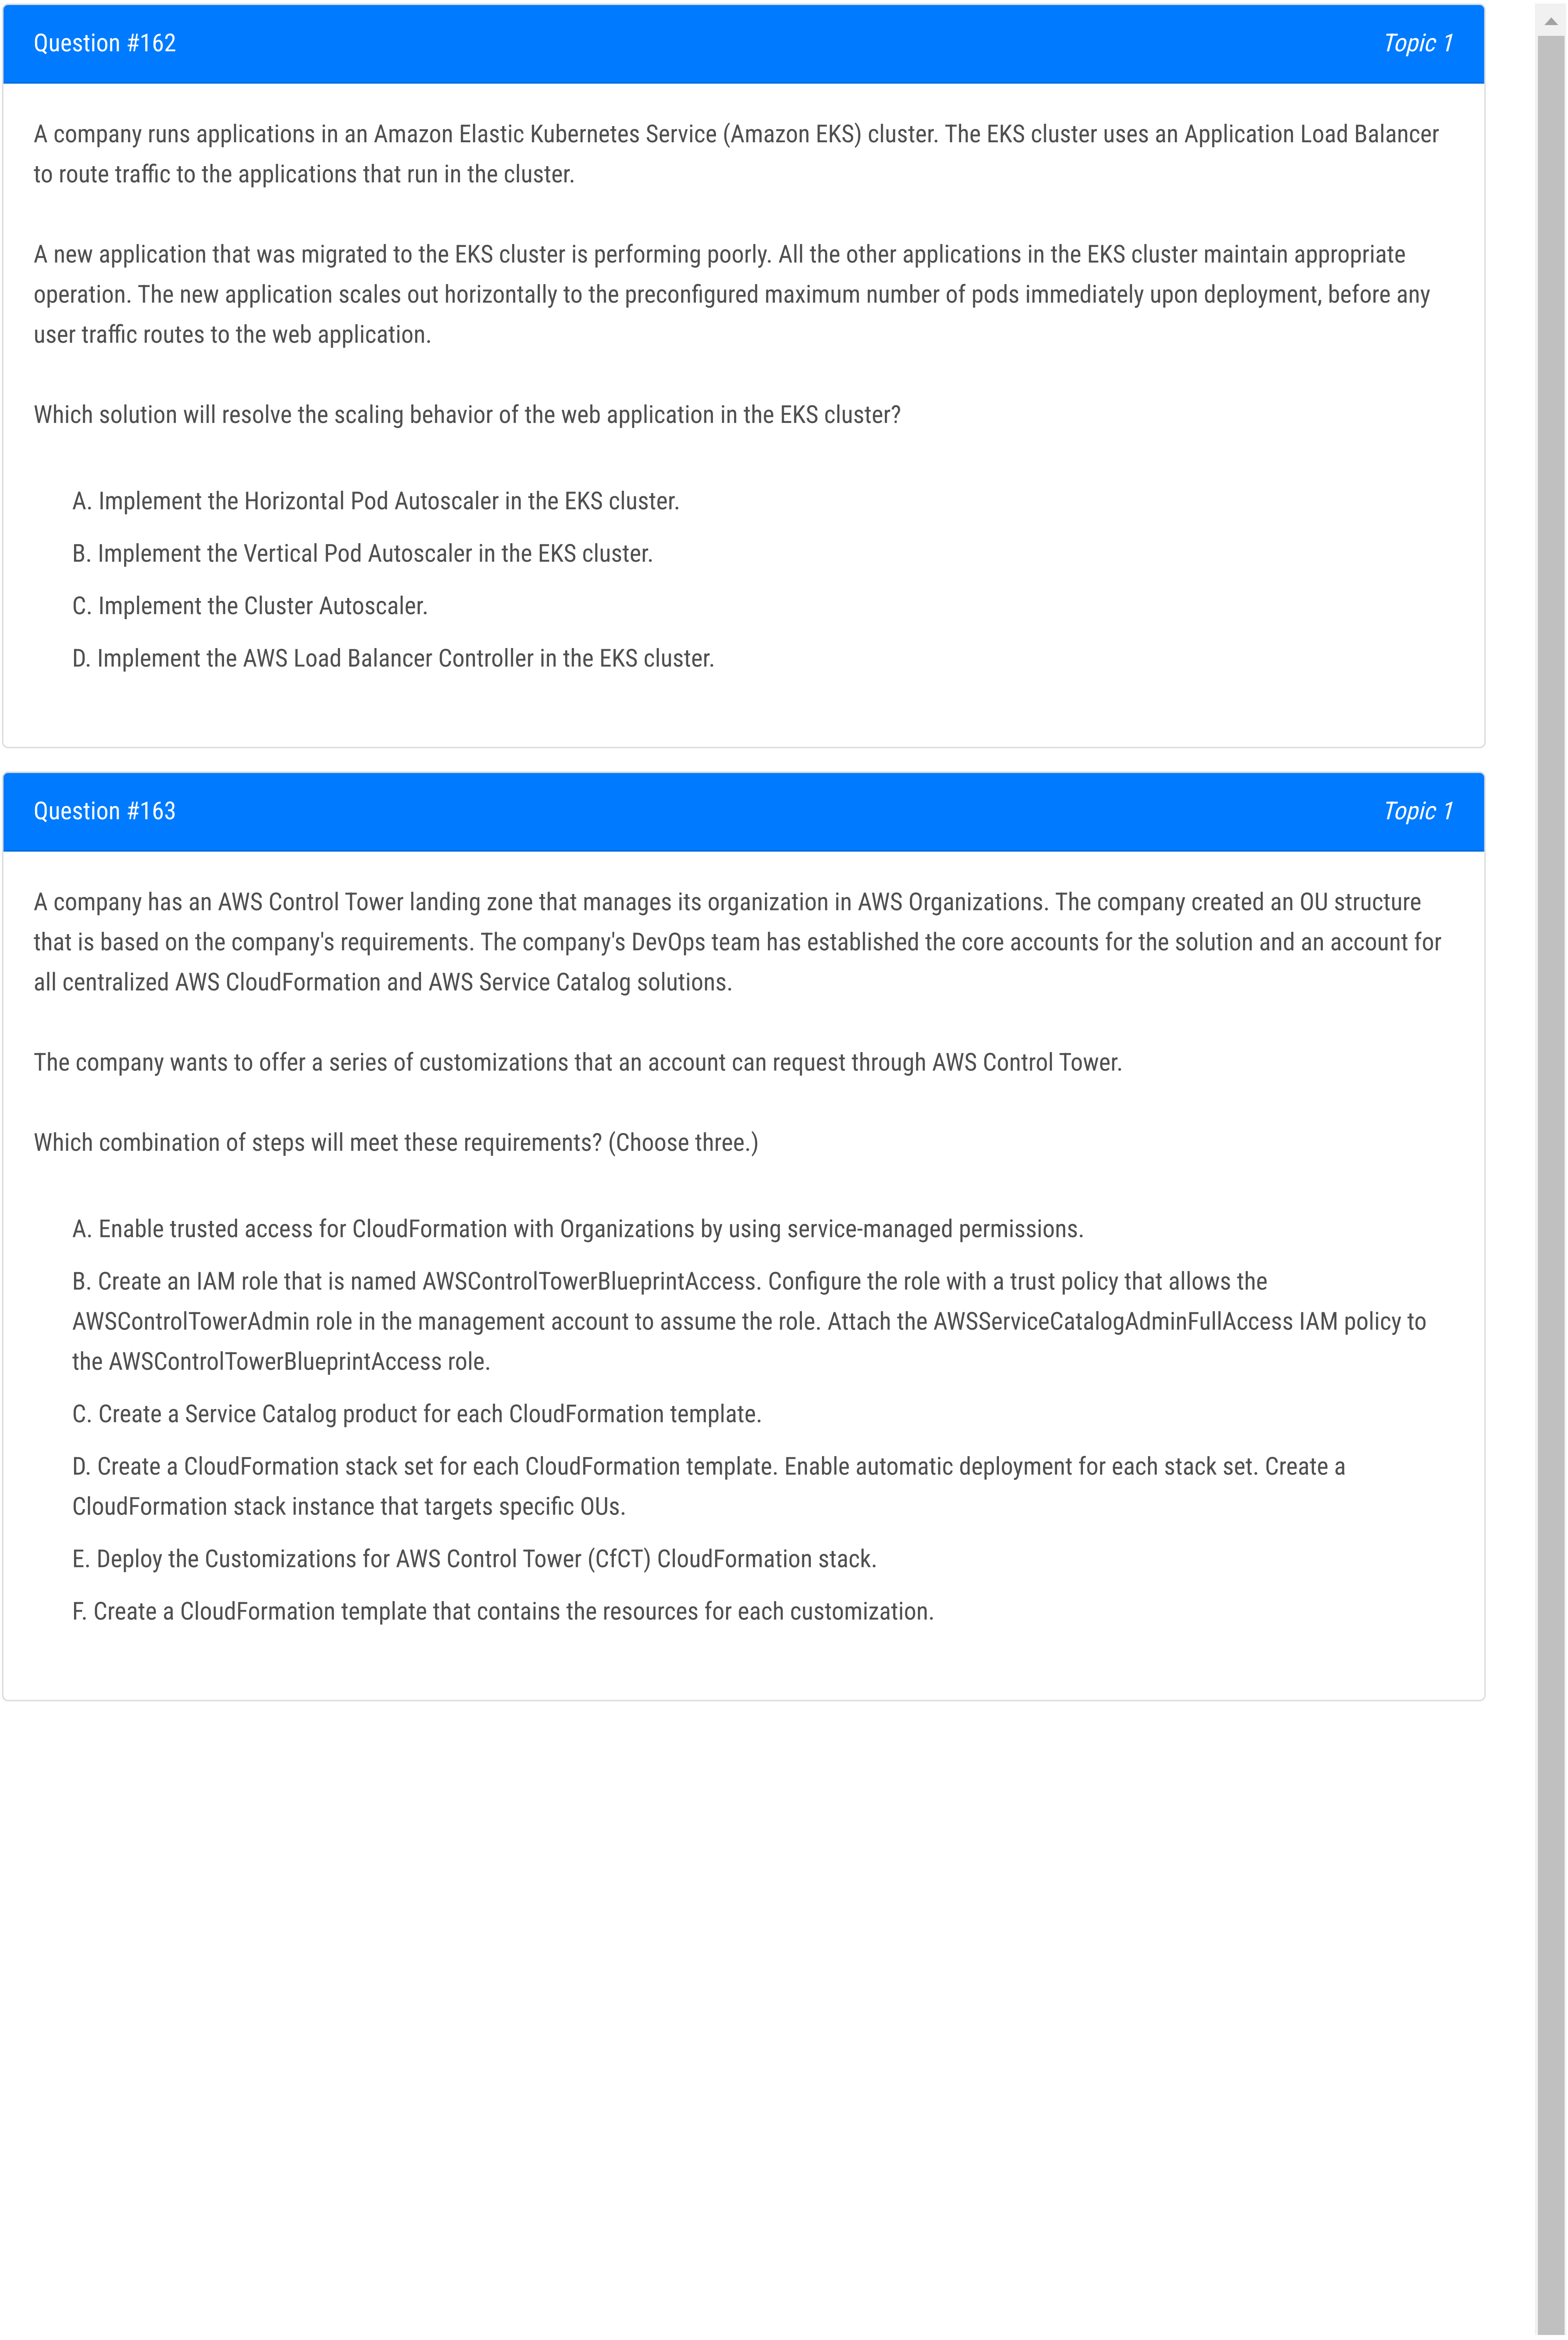

Question #162 *Topic 1*

A company runs applications in an Amazon Elastic Kubernetes Service (Amazon EKS) cluster. The EKS cluster uses an Application Load Balancer to route traffic to the applications that run in the cluster.

A new application that was migrated to the EKS cluster is performing poorly. All the other applications in the EKS cluster maintain appropriate operation. The new application scales out horizontally to the preconfigured maximum number of pods immediately upon deployment, before any user traffic routes to the web application.

Which solution will resolve the scaling behavior of the web application in the EKS cluster?

A. Implement the Horizontal Pod Autoscaler in the EKS cluster.

B. Implement the Vertical Pod Autoscaler in the EKS cluster.

C. Implement the Cluster Autoscaler.

D. Implement the AWS Load Balancer Controller in the EKS cluster.

Question #163 *Topic 1*

A company has an AWS Control Tower landing zone that manages its organization in AWS Organizations. The company created an OU structure that is based on the company's requirements. The company's DevOps team has established the core accounts for the solution and an account for all centralized AWS CloudFormation and AWS Service Catalog solutions.

The company wants to offer a series of customizations that an account can request through AWS Control Tower.

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

A. Enable trusted access for CloudFormation with Organizations by using service-managed permissions.

B. Create an IAM role that is named AWSControlTowerBlueprintAccess. Configure the role with a trust policy that allows the AWSControlTowerAdmin role in the management account to assume the role. Attach the AWSServiceCatalogAdminFullAccess IAM policy to the AWSControlTowerBlueprintAccess role.

C. Create a Service Catalog product for each CloudFormation template.

D. Create a CloudFormation stack set for each CloudFormation template. Enable automatic deployment for each stack set. Create a CloudFormation stack instance that targets specific OUs.

E. Deploy the Customizations for AWS Control Tower (CfCT) CloudFormation stack.

F. Create a CloudFormation template that contains the resources for each customization.

Question #164

Topic 1

A company runs a workload on Amazon EC2 instances. The company needs a control that requires the use of Instance Metadata Service Version 2 (IMDSv2) on all EC2 instances in the AWS account. If an EC2 instance does not prevent the use of Instance Metadata Service Version 1 (IMDSv1), the EC2 instance must be terminated.

Which solution will meet these requirements?

A. Set up AWS Config in the account. Use a managed rule to check EC2 instances. Configure the rule to remediate the findings by using AWS Systems Manager Automation to terminate the instance.

B. Create a permissions boundary that prevents the ec2:RunInstance action if the ec2:MetadataHttpTokens condition key is not set to a value of required. Attach the permissions boundary to the IAM role that was used to launch the instance.

C. Set up Amazon Inspector in the account. Configure Amazon Inspector to activate deep inspection for EC2 instances. Create an Amazon EventBridge rule for an Inspector2 finding. Set an AWS Lambda function as the target to terminate the instance.

D. Create an Amazon EventBridge rule for the EC2 instance launch successful event. Send the event to an AWS Lambda function to inspect the EC2 metadata and to terminate the instance.

Question #165

Topic 1

A company builds an application that uses an Application Load Balancer in front of Amazon EC2 instances that are in an Auto Scaling group. The application is stateless. The Auto Scaling group uses a custom AMI that is fully prebuilt. The EC2 instances do not have a custom bootstrapping process.

The AMI that the Auto Scaling group uses was recently deleted. The Auto Scaling group's scaling activities show failures because the AMI ID does not exist.

Which combination of steps should a DevOps engineer take to meet these requirements? (Choose three.)

A. Create a new launch template that uses the new AMI.

B. Update the Auto Scaling group to use the new launch template.

C. Reduce the Auto Scaling group's desired capacity to 0.

D. Increase the Auto Scaling group's desired capacity by 1.

E. Create a new AMI from a running EC2 instance in the Auto Scaling group.

F. Create a new AMI by copying the most recent public AMI of the operating system that the EC2 instances use.

Question #166 *Topic 1*

A company deploys a web application on Amazon EC2 instances that are behind an Application Load Balancer (ALB). The company stores the application code in an AWS CodeCommit repository. When code is merged to the main branch, an AWS Lambda function invokes an AWS CodeBuild project. The CodeBuild project packages the code, stores the packaged code in AWS CodeArtifact, and invokes AWS Systems Manager Run Command to deploy the packaged code to the EC2 instances.

Previous deployments have resulted in defects, EC2 instances that are not running the latest version of the packaged code, and inconsistencies between instances.

Which combination of actions should a DevOps engineer take to implement a more reliable deployment solution? (Choose two.)

A. Create a pipeline in AWS CodePipeline that uses the CodeCommit repository as a source provider. Configure pipeline stages that run the CodeBuild project in parallel to build and test the application. In the pipeline, pass the CodeBuild project output artifact to an AWS CodeDeploy action.

B. Create a pipeline in AWS CodePipeline that uses the CodeCommit repository as a source provider. Create separate pipeline stages that run a CodeBuild project to build and then test the application. In the pipeline, pass the CodeBuild project output artifact to an AWS CodeDeploy action.

C. Create an AWS CodeDeploy application and a deployment group to deploy the packaged code to the EC2 instances. Configure the ALB for the deployment group.

D. Create individual Lambda functions that use AWS CodeDeploy instead of Systems Manager to run build, test, and deploy actions.

E. Create an Amazon S3 bucket. Modify the CodeBuild project to store the packages in the S3 bucket instead of in CodeArtifact. Use deploy actions in CodeDeploy to deploy the artifact to the EC2 instances.

Question #167 *Topic 1*

A company uses an organization in AWS Organizations to manage its AWS accounts. The company's automation account contains a CI/CD pipeline that creates and configures new AWS accounts.

The company has a group of internal service teams that provide services to accounts in the organization. The service teams operate out of a set of services accounts. The service teams want to receive an AWS CloudTrail event in their services accounts when the CreateAccount API call creates a new account.

How should the company share this CloudTrail event with the service accounts?

A. Create an Amazon EventBridge rule in the automation account to send account creation events to the default event bus in the services accounts. Update the default event bus in the services accounts to allow events from the automation account.

B. Create a custom Amazon EventBridge event bus in the services accounts. Update the custom event bus to allow events from the automation account. Create an EventBridge rule in the services account that directly listens to CloudTrail events from the automation account.

C. Create a custom Amazon EventBridge event bus in the automation account and the services accounts. Create an EventBridge rule and policy that connects the custom event buses that are in the automation account and the services accounts.

D. Create a custom Amazon EventBridge event bus in the automation account. Create an EventBridge rule and policy that connects the custom event bus to the default event buses in the services accounts.

Question #168 *Topic 1*

A DevOps engineer is building a solution that uses Amazon Simple Queue Service (Amazon SQS) standard queues. The solution also includes an AWS Lambda function and an Amazon DynamoDB table. The Lambda function pulls content from an SQS queue event source and writes the content to the DynamoDB table.

The solution must maximize the scalability of Lambda and must prevent successfully processed SQS messages from being processed multiple times.

Which solution will meet these requirements?

A. Decrease the batch window to 1 second when configuring the Lambda function's event source mapping.

B. Decrease the batch size to 1 when configuring the Lambda function's event source mapping.

C. Include the ReportBatchItemFailures value in the FunctionResponseTypes list in the Lambda function's event source mapping.

D. Set the queue visibility timeout on the Lambda function's event source mapping to account for invocation throttling of the Lambda function.

Question #169 *Topic 1*

A company has a new AWS account that teams will use to deploy various applications. The teams will create many Amazon S3 buckets for application-specific purposes and to store AWS CloudTrail logs. The company has enabled Amazon Macie for the account.

A DevOps engineer needs to optimize the Macie costs for the account without compromising the account's functionality.

Which solutions will meet these requirements? (Choose two.)

A. Exclude S3 buckets that contain CloudTrail logs from automated discovery.

B. Exclude S3 buckets that have public read access from automated discovery.

C. Configure scheduled daily discovery jobs for all S3 buckets in the account.

D. Configure discovery jobs to include S3 objects based on the last modified criterion.

E. Configure discovery jobs to include S3 objects that are tagged as production only.

Question #170 *Topic 1*

A company uses an organization in AWS Organizations to manage its AWS accounts. The company recently acquired another company that has standalone AWS accounts. The acquiring company's DevOps team needs to consolidate the administration of the AWS accounts for both companies and retain full administrative control of the accounts. The DevOps team also needs to collect and group findings across all the accounts to implement and maintain a security posture.

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

A. Invite the acquired company's AWS accounts to join the organization. Create an SCP that has full administrative privileges. Attach the SCP to the management account.

B. Invite the acquired company's AWS accounts to join the organization. Create the OrganizationAccountAccessRole IAM role in the invited accounts. Grant permission to the management account to assume the role.

C. Use AWS Security Hub to collect and group findings across all accounts. Use Security Hub to automatically detect new accounts as the accounts are added to the organization.

D. Use AWS Firewall Manager to collect and group findings across all accounts. Enable all features for the organization. Designate an account in the organization as the delegated administrator account for Firewall Manager.

E. Use Amazon Inspector to collect and group findings across all accounts. Designate an account in the organization as the delegated administrator account for Amazon Inspector.

Question #171 *Topic 1*

A company has an application and a CI/CD pipeline. The CI/CD pipeline consists of an AWS CodePipeline pipeline and an AWS CodeBuild project. The CodeBuild project runs tests against the application as part of the build process and outputs a test report. The company must keep the test reports for 90 days.

Which solution will meet these requirements?

A. Add a new stage in the CodePipeline pipeline after the stage that contains the CodeBuild project. Create an Amazon S3 bucket to store the reports. Configure an S3 deploy action type in the new CodePipeline stage with the appropriate path and format for the reports.

B. Add a report group in the CodeBuild project buildspec file with the appropriate path and format for the reports. Create an Amazon S3 bucket to store the reports. Configure an Amazon EventBridge rule that invokes an AWS Lambda function to copy the reports to the S3 bucket when a build is completed. Create an S3 Lifecycle rule to expire the objects after 90 days.

C. Add a new stage in the CodePipeline pipeline. Configure a test action type with the appropriate path and format for the reports. Configure the report expiration time to be 90 days in the CodeBuild project buildspec file.

D. Add a report group in the CodeBuild project buildspec file with the appropriate path and format for the reports. Create an Amazon S3 bucket to store the reports. Configure the report group as an artifact in the CodeBuild project buildspec file. Configure the S3 bucket as the artifact destination. Set the object expiration to 90 days.

Question #172 *Topic 1*

A company uses an Amazon API Gateway regional REST API to host its application API. The REST API has a custom domain. The REST API's default endpoint is deactivated.

The company's internal teams consume the API. The company wants to use mutual TLS between the API and the internal teams as an additional layer of authentication.

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

A. Use AWS Certificate Manager (ACM) to create a private certificate authority (CA). Provision a client certificate that is signed by the private CA.

B. Provision a client certificate that is signed by a public certificate authority (CA). Import the certificate into AWS Certificate Manager (ACM).

C. Upload the provisioned client certificate to an Amazon S3 bucket. Configure the API Gateway mutual TLS to use the client certificate that is stored in the S3 bucket as the trust store.

D. Upload the provisioned client certificate private key to an Amazon S3 bucket. Configure the API Gateway mutual TLS to use the private key that is stored in the S3 bucket as the trust store.

E. Upload the root private certificate authority (CA) certificate to an Amazon S3 bucket. Configure the API Gateway mutual TLS to use the private CA certificate that is stored in the S3 bucket as the trust store.

Question #173 *Topic 1*

A company uses AWS Directory Service for Microsoft Active Directory as its identity provider (IdP). The company requires all infrastructure to be defined and deployed by AWS CloudFormation.

A DevOps engineer needs to create a fleet of Windows-based Amazon EC2 instances to host an application. The DevOps engineer has created a CloudFormation template that contains an EC2 launch template, IAM role, EC2 security group, and EC2 Auto Scaling group. The DevOps engineer must implement a solution that joins all EC2 instances to the domain of the AWS Managed Microsoft AD directory.

Which solution will meet these requirements with the MOST operational efficiency?

A. In the CloudFormation template, create an AWS::SSM::Document resource that joins the EC2 instance to the AWS Managed Microsoft AD domain by using the parameters for the existing directory. Update the launch template to include the SSMAssociation property to use the new SSM document. Attach the AmazonSSMManagedInstanceCore and AmazonSSMDirectoryServiceAccess AWS managed policies to the IAM role that the EC2 instances use.

B. In the CloudFormation template, update the launch template to include specific tags that propagate on launch. Create an AWS::SSM::Association resource to associate the AWS-JoinDirectoryServiceDomain Automation runbook with the EC2 instances that have the specified tags. Define the required parameters to join the AWS Managed Microsoft AD directory. Attach the AmazonSSMManagedInstanceCore and AmazonSSMDirectoryServiceAccess AWS managed policies to the IAM role that the EC2 instances use.

C. Store the existing AWS Managed Microsoft AD domain connection details in AWS Secrets Manager. In the CloudFormation template, create an AWS::SSM::Association resource to associate the AWS-CreateManagedWindowsInstanceWithApproval Automation runbook with the EC2 Auto Scaling group. Pass the ARNs for the parameters from Secrets Manager to join the domain. Attach the AmazonSSMDirectoryServiceAccess and SecretsManagerReadWrite AWS managed policies to the IAM role that the EC2 instances use.

D. Store the existing AWS Managed Microsoft AD domain administrator credentials in AWS Secrets Manager. In the CloudFormation template, update the EC2 launch template to include user data. Configure the user data to pull the administrator credentials from Secrets Manager and to join the AWS Managed Microsoft AD domain. Attach the AmazonSSMManagedInstanceCore and SecretsManagerReadWrite AWS managed policies to the IAM role that the EC2 instances use.

Question #174 *Topic 1*

A company uses AWS Organizations to manage its AWS accounts. The company has a root OU that has a child OU. The root OU has an SCP that allows all actions on all resources. The child OU has an SCP that allows all actions for Amazon DynamoDB and AWS Lambda, and denies all other actions.

The company has an AWS account that is named vendor-data in the child OU. A DevOps engineer has an IAM user that is attached to the Administrator Access IAM policy in the vendor-data account. The DevOps engineer attempts to launch an Amazon EC2 instance in the vendor-data account but receives an access denied error.

Which change should the DevOps engineer make to launch the EC2 instance in the vendor-data account?

A. Attach the AmazonEC2FullAccess IAM policy to the IAM user.

B. Create a new SCP that allows all actions for Amazon EC2. Attach the SCP to the vendor-data account.

C. Update the SCP in the child OU to allow all actions for Amazon EC2.

D. Create a new SCP that allows all actions for Amazon EC2. Attach the SCP to the root OU.

Question #175 *Topic 1*

A company's security policies require the use of security hardened AMIs in production environments. A DevOps engineer has used EC2 Image Builder to create a pipeline that builds the AMIs on a recurring schedule.

The DevOps engineer needs to update the launch templates of the company's Auto Scaling groups. The Auto Scaling groups must use the newest AMIs during the launch of Amazon EC2 instances.

Which solution will meet these requirements with the MOST operational efficiency?

A. Configure an Amazon EventBridge rule to receive new AMI events from Image Builder. Target an AWS Systems Manager Run Command document that updates the launch templates of the Auto Scaling groups with the newest AMI ID.

B. Configure an Amazon EventBridge rule to receive new AMI events from Image Builder. Target an AWS Lambda function that updates the launch templates of the Auto Scaling groups with the newest AMI ID.

C. Configure the launch template to use a value from AWS Systems Manager Parameter Store for the AMI ID. Configure the Image Builder pipeline to update the Parameter Store value with the newest AMI ID.

D. Configure the Image Builder distribution settings to update the launch templates with the newest AMI IConfigure the Auto Scaling groups to use the newest version of the launch template.

Question #176

Topic 1

A company has configured an Amazon S3 event source on an AWS Lambda function. The company needs the Lambda function to run when a new object is created or an existing object is modified in a particular S3 bucket. The Lambda function will use the S3 bucket name and the S3 object key of the incoming event to read the contents of the created or modified S3 object. The Lambda function will parse the contents and save the parsed contents to an Amazon DynamoDB table.

The Lambda function's execution role has permissions to read from the S3 bucket and to write to the DynamoDB table. During testing, a DevOps engineer discovers that the Lambda function does not run when objects are added to the S3 bucket or when existing objects are modified.

Which solution will resolve this problem?

A. Increase the memory of the Lambda function to give the function the ability to process large files from the S3 bucket.

B. Create a resource policy on the Lambda function to grant Amazon S3 the permission to invoke the Lambda function for the S3 bucket.

C. Configure an Amazon Simple Queue Service (Amazon SQS) queue as an OnFailure destination for the Lambda function.

D. Provision space in the /tmp folder of the Lambda function to give the function the ability to process large files from the S3 bucket.

Question #177

Topic 1

A company has deployed a critical application in two AWS Regions. The application uses an Application Load Balancer (ALB) in both Regions. The company has Amazon Route 53 alias DNS records for both ALBs.

The company uses Amazon Route 53 Application Recovery Controller to ensure that the application can fail over between the two Regions. The Route 53 ARC configuration includes a routing control for both Regions. The company uses Route 53 ARC to perform quarterly disaster recovery (DR) tests.

During the most recent DR test, a DevOps engineer accidentally turned off both routing controls. The company needs to ensure that at least one routing control is turned on at all times.

Which solution will meet these requirements?

A. In Route 53 ARC, create a new assertion safety rule. Apply the assertion safety rule to the two routing controls. Configure the rule with the ATLEAST type with a threshold of 1.

B. In Route 53 ARC, create a new gating safety rule. Apply the assertion safety rule to the two routing controls. Configure the rule with the OR type with a threshold of 1.

C. In Route 53 ARC, create a new resource set. Configure the resource set with an AWS::Route53::HealthCheck resource type. Specify the ARNs of the two routing controls as the target resource. Create a new readiness check for the resource set.

D. In Route 53 ARC, create a new resource set. Configure the resource set with an AWS::Route53RecoveryReadiness::DNSTargetResource resource type. Add the domain names of the two Route 53 alias DNS records as the target resource. Create a new readiness check for the resource set.

Question #178 *Topic 1*

A healthcare services company is concerned about the growing costs of software licensing for an application for monitoring patient wellness. The company wants to create an audit process to ensure that the application is running exclusively on Amazon EC2 Dedicated Hosts. A DevOps engineer must create a workflow to audit the application to ensure compliance.

What steps should the engineer take to meet this requirement with the LEAST administrative overhead?

A. Use AWS Systems Manager Configuration Compliance. Use calls to the put-compliance-items API action to scan and build a database of noncompliant EC2 instances based on their host placement configuration. Use an Amazon DynamoDB table to store these instance IDs for fast access. Generate a report through Systems Manager by calling the list-compliance-summaries API action.

B. Use custom Java code running on an EC2 instance. Set up EC2 Auto Scaling for the instance depending on the number of instances to be checked. Send the list of noncompliant EC2 instance IDs to an Amazon SQS queue. Set up another worker instance to process instance IDs from the SQS queue and write them to Amazon DynamoDUse an AWS Lambda function to terminate noncompliant instance IDs obtained from the queue, and send them to an Amazon SNS email topic for distribution.

C. Use AWS Config. Identify all EC2 instances to be audited by enabling Config Recording on all Amazon EC2 resources for the region. Create a custom AWS Config rule that triggers an AWS Lambda function by using the "config-rule-change -triggered" blueprint. Modify the Lambda evaluateCompliance() function to verify host placement to return a NON_COMPLIANT result if the instance is not running on an EC2 Dedicated Host. Use the AWS Config report to address noncompliant instances.

D. Use AWS CloudTrail. Identify all EC2 instances to be audited by analyzing all calls to the EC2 RunCommand API action. Invoke an AWS Lambda function that analyzes the host placement of the instance. Store the EC2 instance ID of noncompliant resources in an Amazon RDS for MySQL DB instance. Generate a report by querying the RDS instance and exporting the query results to a CSV text file.

Question #179 *Topic 1*

A DevOps engineer is planning to deploy a Ruby-based application to production. The application needs to interact with an Amazon RDS for MySQL database and should have automatic scaling and high availability. The stored data in the database is critical and should persist regardless of the state of the application stack.

The DevOps engineer needs to set up an automated deployment strategy for the application with automatic rollbacks. The solution also must alert the application team when a deployment fails.

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

A. Deploy the application on AWS Elastic Beanstalk. Deploy an Amazon RDS for MySQL DB instance as part of the Elastic Beanstalk configuration.

B. Deploy the application on AWS Elastic Beanstalk. Deploy a separate Amazon RDS for MySQL DB instance outside of Elastic Beanstalk.

C. Configure a notification email address that alerts the application team in the AWS Elastic Beanstalk configuration.

D. Configure an Amazon EventBridge rule to monitor AWS Health events. Use an Amazon Simple Notification Service (Amazon SNS) topic as a target to alert the application team.

E. Use the immutable deployment method to deploy new application versions.

F. Use the rolling deployment method to deploy new application versions.

Question #180 *Topic 1*

A company is using AWS CodePipeline to deploy an application. According to a new guideline, a member of the company's security team must sign off on any application changes before the changes are deployed into production. The approval must be recorded and retained.

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

A. Configure CodePipeline to write actions to Amazon CloudWatch Logs.

B. Configure CodePipeline to write actions to an Amazon S3 bucket at the end of each pipeline stage.

C. Create an AWS CloudTrail trail to deliver logs to Amazon S3.

D. Create a CodePipeline custom action to invoke an AWS Lambda function for approval. Create a policy that gives the security team access to manage CodePipeline custom actions.

E. Create a CodePipeline manual approval action before the deployment step. Create a policy that grants the security team access to approve manual approval stages.

Question #181 *Topic 1*

A company requires its internal business teams to launch resources through pre-approved AWS CloudFormation templates only. The security team requires automated monitoring when resources drift from their expected state.

Which strategy should be used to meet these requirements?

A. Allow users to deploy CloudFormation stacks using a CloudFormation service role only. Use CloudFormation drift detection to detect when resources have drifted from their expected state.

B. Allow users to deploy CloudFormation stacks using a CloudFormation service role only. Use AWS Config rules to detect when resources have drifted from their expected state.

C. Allow users to deploy CloudFormation stacks using AWS Service Catalog only. Enforce the use of a launch constraint. Use AWS Config rules to detect when resources have drifted from their expected state.

D. Allow users to deploy CloudFormation stacks using AWS Service Catalog only. Enforce the use of a template constraint. Use Amazon EventBridge notifications to detect when resources have drifted from their expected state.

Question #182 *Topic 1*

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

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

A. Create an Amazon EventBridge rule that runs periodically and targets an AWS Lambda function. Within the Lambda function, evaluate the current state of the AWS environment and compare deployed resource values to resource limits on the account. Notify the senior manager if the account is approaching a service limit.

B. Deploy an AWS Lambda function that refreshes AWS Trusted Advisor checks, and configure an Amazon EventBridge rule to run the Lambda function periodically. Create another EventBridge rule with an event pattern matching Trusted Advisor events and a target Lambda function. In the target Lambda function, notify the senior manager.

C. Deploy an AWS Lambda function that refreshes AWS Health Dashboard checks, and configure an Amazon EventBridge rule to run the Lambda function periodically. Create another EventBridge rule with an event pattern matching Health Dashboard events and a target Lambda function. In the target Lambda function, notify the senior manager.

D. Add an AWS Config custom rule that runs periodically, checks the AWS service limit status, and streams notifications to an Amazon Simple Notification Service (Amazon SNS) topic. Deploy an AWS Lambda function that notifies the senior manager, and subscribe the Lambda function to the SNS topic.

Question #183 *Topic 1*

A DevOps engineer is setting up a container-based architecture. The engineer has decided to use AWS CloudFormation to automatically provision an Amazon ECS cluster and an Amazon EC2 Auto Scaling group to launch the EC2 container instances. After successfully creating the CloudFormation stack, the engineer noticed that, even though the ECS cluster and the EC2 instances were created successfully and the stack finished the creation, the EC2 instances were associating with a different cluster.

How should the DevOps engineer update the CloudFormation template to resolve this issue?

A. Reference the EC2 instances in the AWS::ECS::Cluster resource and reference the ECS cluster in the AWS::ECS::Service resource.

B. Reference the ECS cluster in the AWS::AutoScaling::LaunchConfiguration resource of the UserData property.

C. Reference the ECS cluster in the AWS::EC2::Instance resource of the UserData property.

D. Reference the ECS cluster in the AWS::CloudFormation::CustomResource resource to trigger an AWS Lambda function that registers the EC2 instances with the appropriate ECS cluster.

Question #184 *Topic 1*

A DevOps engineer is implementing governance controls for a company that requires its infrastructure to be housed within the United States. The engineer must restrict which AWS Regions can be used, and ensure an alert is sent as soon as possible if any activity outside the governance policy takes place. The controls should be automatically enabled on any new Region outside the United States (US).

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

A. Create an AWS Organizations SCP that denies access to all non-global services in non-US Regions. Attach the policy to the root of the organization.

B. Configure AWS CloudTrail to send logs to Amazon CloudWatch Logs and enable it for all Regions. Use a CloudWatch Logs metric filter to send an alert on any service activity in non-US Regions.

C. Use an AWS Lambda function that checks for AWS service activity and deploy it to all Regions. Write an Amazon EventBridge rule that runs the Lambda function every hour, sending an alert if activity is found in a non-US Region.

D. Use an AWS Lambda function to query Amazon Inspector to look for service activity in non-US Regions and send alerts if any activity is found.

E. Write an SCP using the aws:RequestedRegion condition key limiting access to US Regions. Apply the policy to all users, groups, and roles.

Question #185 *Topic 1*

A company sells products through an ecommerce web application. The company wants a dashboard that shows a pie chart of product transaction details. The company wants to integrate the dashboard with the company's existing Amazon CloudWatch dashboards.

Which solution will meet these requirements with the MOST operational efficiency?

A. Update the ecommerce application to emit a JSON object to a CloudWatch log group for each processed transaction. Use CloudWatch Logs Insights to query the log group and to visualize the results in a pie chart format. Attach the results to the desired CloudWatch dashboard.

B. Update the ecommerce application to emit a JSON object to an Amazon S3 bucket for each processed transaction. Use Amazon Athena to query the S3 bucket and to visualize the results in a pie chart format. Export the results from Athena. Attach the results to the desired CloudWatch dashboard.

C. Update the ecommerce application to use AWS X-Ray for instrumentation. Create a new X-Ray subsegment. Add an annotation for each processed transaction. Use X-Ray traces to query the data and to visualize the results in a pie chart format. Attach the results to the desired CloudWatch dashboard.

D. Update the ecommerce application to emit a JSON object to a CloudWatch log group for each processed transaction. Create an AWS Lambda function to aggregate and write the results to Amazon DynamoDB. Create a Lambda subscription filter for the log file. Attach the results to the desired CloudWatch dashboard.

Question #186 *Topic 1*

A company is launching an application. The application must use only approved AWS services. The account that runs the application was created less than 1 year ago and is assigned to an AWS Organizations OU.

The company needs to create a new Organizations account structure. The account structure must have an appropriate SCP that supports the use of only services that are currently active in the AWS account. The company will use AWS Identity and Access Management (IAM) Access Analyzer in the solution.

Which solution will meet these requirements?

A. Create an SCP that allows the services that IAM Access Analyzer identifies. Create an OU for the account. Move the account into the new OU. Attach the new SCP to the new OU. Detach the default FullAWSAccess SCP from the new OU.

B. Create an SCP that denies the services that IAM Access Analyzer identifies. Create an OU for the account. Move the account into the new OU. Attach the new SCP to the new OU.

C. Create an SCP that allows the services that IAM Access Analyzer identifies. Attach the new SCP to the organization's root.

D. Create an SCP that allows the services that IAM Access Analyzer identifies. Create an OU for the account. Move the account into the new OU. Attach the new SCP to the management account. Detach the default FullAWSAccess SCP from the new OU.

Question #187 *Topic 1*

A company has multiple development teams in different business units that work in a shared single AWS account. All Amazon EC2 resources that are created in the account must include tags that specify who created the resources. The tagging must occur within the first hour of resource creation.

A DevOps engineer needs to add tags to the created resources that include the user ID that created the resource and the cost center ID. The DevOps engineer configures an AWS Lambda function with the cost center mappings to tag the resources. The DevOps engineer also sets up AWS CloudTrail in the AWS account. An Amazon S3 bucket stores the CloudTrail event logs.

Which solution will meet the tagging requirements?

A. Create an S3 event notification on the S3 bucket to invoke the Lambda function for s3:ObjectTagging:Put events. Enable bucket versioning on the S3 bucket.

B. Enable server access logging on the S3 bucket. Create an S3 event notification on the S3 bucket for s3:ObjectTagging:* events.

C. Create a recurring hourly Amazon EventBridge scheduled rule that invokes the Lambda function. Modify the Lambda function to read the logs from the S3 bucket.

D. Create an Amazon EventBridge rule that uses Amazon EC2 as the event source. Configure the rule to match events delivered by CloudTrail. Configure the rule to target the Lambda function.

Question #188 *Topic 1*

A company runs an application for multiple environments in a single AWS account. An AWS CodePipeline pipeline uses a development Amazon Elastic Container Service (Amazon ECS) cluster to test an image for the application from an Amazon Elastic Container Registry (Amazon ECR) repository. The pipeline promotes the image to a production ECS cluster.

The company needs to move the production cluster into a separate AWS account in the same AWS Region. The production cluster must be able to download the images over a private connection.

Which solution will meet these requirements?

A. Use Amazon ECR VPC endpoints and an Amazon S3 gateway endpoint. In the separate AWS account, create an ECR repository. Set the repository policy to allow the production ECS tasks to pull images from the main AWS account. Configure the production ECS task execution role to have permission to download the image from the ECR repository.

B. Set a repository policy on the production ECR repository in the main AWS account. Configure the repository policy to allow the production ECS tasks in the separate AWS account to pull images from the main account. Configure the production ECS task execution role to have permission to download the image from the ECR repository.

C. Configure ECR private image replication in the main AWS account. Activate cross-account replication. Define the destination account ID of the separate AWS account.

D. Use Amazon ECR VPC endpoints and an Amazon S3 gateway endpoint. Set a repository policy on the production ECR repository in the main AWS account. Configure the repository policy to allow the production ECS tasks in the separate AWS account to pull images from the main account. Configure the production ECS task execution role to have permission to download the image from the ECR repository.

Question #189 *Topic 1*

A company needs to ensure that flow logs remain configured for all existing and new VPCs in its AWS account. The company uses an AWS CloudFormation stack to manage its VPCs. The company needs a solution that will work for any VPCs that any IAM user creates.

Which solution will meet these requirements?

A. Add the AWS::EC2::FlowLog resource to the CloudFormation stack that creates the VPCs.

B. Create an organization in AWS Organizations. Add the company's AWS account to the organization. Create an SCP to prevent users from modifying VPC flow logs.

C. Turn on AWS Config. Create an AWS Config rule to check whether VPC flow logs are turned on. Configure automatic remediation to turn on VPC flow logs.

D. Create an IAM policy to deny the use of API calls for VPC flow logs. Attach the IAM policy to all IAM users.

Question #190 *Topic 1*

A company's application teams use AWS CodeCommit repositories for their applications. The application teams have repositories in multiple AWS accounts. All accounts are in an organization in AWS Organizations.

Each application team uses AWS IAM Identity Center (AWS Single Sign-On) configured with an external IdP to assume a developer IAM role. The developer role allows the application teams to use Git to work with the code in the repositories.

A security audit reveals that the application teams can modify the main branch in any repository. A DevOps engineer must implement a solution that allows the application teams to modify the main branch of only the repositories that they manage.

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

A. Update the SAML assertion to pass the user's team name. Update the IAM role's trust policy to add an access-team session tag that has the team name.

B. Create an approval rule template for each team in the Organizations management account. Associate the template with all the repositories. Add the developer role ARN as an approver.

C. Create an approval rule template for each account. Associate the template with all repositories. Add the "aws:ResourceTag/access-team": "$;{aws:PrincipalTag/access-team}" condition to the approval rule template.

D. For each CodeCommit repository, add an access-team tag that has the value set to the name of the associated team.

E. Attach an SCP to the accounts. Include the following statement:

```
{
    "Effect": "Deny",
    "Action": [
        "codecommit:GitPush",
        "codecommit:PutFile",
        "codecommit:Merge*"
    ],
    "Resource": "*",
    "Condition": {
        "StringEqualsIfExists": {
            "codecommit:References": ["refs/heads/main"]
        },
        "StringNotEquals": {
            "aws:ResourceTag/access-team": "$ ;{aws:PrincipalTag/access-team}"
        },
        "Null": {
            "codecommit:References": "false"

        }
    }
}
```