Number: DOP-C02
Passing Score: 800
Time Limit: 120
File Version: 32.0

Exam Code: DOC-C02

Exam Name: AWS DevOps Engineer - Professional

Exam A

QUESTION 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. ymi 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. Down load Bundle

Correct Answer: C
Section:
Explanation:
This hook runs before the new application version is installed on the replacement instances. This is the best place to run the script because it ensures that the license file is downloaded and installed before the replacement instances start to handle request traffic. If you use any other hook, you may encounter errors or inconsistencies in your application.

QUESTION 2
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 togs. The DevOps engineer manages the Kinesis consumer application, which also runs on Amazon EC2.
Sudden increases of data cause the Kinesis consumer application to (all 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.

Correct Answer: B
Section:
Explanation:
https://docs.aws.amazon.com/streams/latest/dev/monitoring-with-cloudwatch.html
GetRecords.IteratorAgeMilliseconds - The age of the last record in all GetRecords calls made against a Kinesis stream, measured over the specified time period. Age is the difference between the current time and when the last record of the GetRecords call was written to the stream. The Minimum and Maximum statistics can be used to track the progress of Kinesis consumer applications. A value of zero indicates that the records being read are completely caught up.

QUESTION 3
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 protect to run within a VPC and use dedicated instances to increase throughput.

Correct Answer: C
Section:
Explanation:
https://docs.aws.amazon.com/codepipeline/latest/userguide/reference-pipeline-structure.html
AWS doc: 'To specify parallel actions, use the same integer for each action you want to run in parallel. For example, if you want three actions to run in sequence in a stage, you would give the first action the runOrder value of 1, the second action the runOrder value of 2, and the third the runOrder value of 3. However, if you want the second and third actions to run in parallel, you would give the first action the runOrder value of 1 and both the second and third actions the runOrder value of 2.'

QUESTION 4
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.

Correct Answer: B
Section:
Explanation:
AWS CloudWatch Events can be used to monitor events across different AWS resources, and a CloudWatch Event Rule can be created to trigger an AWS Lambda function when a deployment issue is detected in the pipeline. The Lambda function can then evaluate the issue and send a notification to the appropriate stakeholders through an Amazon SNS topic. This approach allows for real-time notifications and faster resolution times.

QUESTION 5
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 restriction for the specific repositories in the policy repositories in the policy statement with a condition that references the main branch. 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 mam 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.

Correct Answer: A
Section:
Explanation:

By default, the AWSCodeCommitPowerUser managed policy allows users to push changes to any branch in any repository in the AWS account. To restrict the developers' ability to push changes to the main branch directly, an additional policy is needed that explicitly denies these actions for the main branch.

The Deny rule should be included in a policy statement that targets the specific repositories and includes a condition that references the main branch. The policy statement should look something like this:

{

'Effect': 'Deny',

'Action': [

'codecommit:GitPush',

'codecommit:PutFile'

],

'Resource': 'arn:aws:codecommit:<region>::<repository-name>',

'Condition': {

'StringEqualsIfExists': {

'codecommit:References': [

'refs/heads/main'

]

}

}

QUESTION 6
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 client 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 CodePipelme 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 Create. Deployment 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.

Correct Answer: A
Section:
Explanation:
This solution would allow the company to automate the process of updating the SDK and making it available to web clients. By adding a CodePipeline action immediately after the deployment stage of the API, the Lambda

function will be invoked automatically each time the API is updated. The Lambda function should be able to download the new SDK from API Gateway, upload it to the S3 bucket and also create a CloudFront invalidation for the SDK path so that the latest version of the SDK is available for the web clients. This is the most straight forward solution and it will meet the requirements.

QUESTION 7
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.

Correct Answer: C
Section:
Explanation:
https://docs.aws.amazon.com/codedeploy/latest/userguide/reference-appspec-file-structure-hooks.html

QUESTION 8
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 tor Internet Security (CIS) Benchmarks tor 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 m Security Hub after the accounts are created.
Which combination of steps will meet these requirements in the MOST automated way? (Select 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 organ ration'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.

Correct Answer: A, C, E
Section:
Explanation:
https://docs.aws.amazon.com/securityhub/latest/userguide/accounts-orgs-auto-enable.html

QUESTION 9
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.

Correct Answer: B, C, F
Section:
Explanation:
The requirement is to create the infrastructure as code (IaC) and the CI/CD pipeline for the Lambda application that uses canary deployment and automated rollback. To do this, the DevOps team needs to use the following steps:
Create an AWS Serverless Application Model (AWS SAM) template for the application. AWS SAM is a framework that simplifies the development and deployment of serverless applications on AWS. AWS SAM allows customers to define Lambda functions and other resources in a template by using a simplified syntax. For each Lambda function, the DevOps team can include configurations for the AutoPublishAlias property and the DeploymentPreference property. The AutoPublishAlias property specifies the name of the alias that points to the latest version of the function. The DeploymentPreference property specifies how CodeDeploy deploys new versions of the function. By configuring the deployment configuration type to LambdaCanary10Percent10Minutes, the DevOps team can enable canary deployment with 10% of traffic shifted to the new version every 10 minutes.
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 SAM template. CodeCommit is a fully managed source control service that hosts Git repositories. CodePipeline is a fully managed continuous delivery service that automates the release process of software applications. CodeBuild is a fully managed continuous integration service that compiles source code and runs tests. By using these services, the DevOps team can create a CI/CD pipeline for the Lambda application. The pipeline should use the CodeCommit repository as the source stage, where the DevOps team can upload the SAM template and source code. The pipeline should also use a CodeBuild project as the build stage, where the SAM template can be built and deployed.
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. CloudWatch is a service that monitors and collects metrics from AWS resources and applications. CloudWatch alarms are actions that are triggered when a metric crosses a specified threshold. By creating CloudWatch alarms for each Lambda function, the DevOps team can monitor the health and performance of each function version during deployment. By configuring the alarms to enter the ALARM state if any errors are detected, the DevOps team can enable automated rollback if any failures are reported.

QUESTION 10
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 BeforeAllowTraffic 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.

Correct Answer: A
Section:
Explanation:
https://docs.aws.amazon.com/codedeploy/latest/userguide/reference-appspec-file-structure-hooks.html#appspec-hooks-lambda

QUESTION 11
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 Conflg 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.

Correct Answer: B
Section:
Explanation:
https://aws.amazon.com/blogs/mt/aws-config-auto-remediation-s3-compliance/ https://aws.amazon.com/blogs/aws/aws-config-rules-dynamic-compliance-checking-for-cloud-resources/

QUESTION 12
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.

Correct Answer: C
Section:
Explanation:
https://docs.aws.amazon.com/codedeploy/latest/userguide/reference-appspec-file-structure-hooks.html

QUESTION 13
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.

Correct Answer: A
Section:
Explanation:
https://docs.aws.amazon.com/storagegateway/latest/APIReference/API_RefreshCache.html ' It only updates the cached inventory to reflect changes in the inventory of the objects in the S3 bucket. This operation is only supported in the S3 File Gateway types.'

QUESTION 14
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.

Correct Answer: A, D, E
Section:
Explanation:
S3 cross-Region replication (CRR) automatically replicates data between buckets across different AWS Regions. To enable CRR, you need to add a replication configuration to your source bucket that specifies the destination bucket, the IAM role, and the encryption type (optional). You also need to grant permissions to the IAM role to perform replication actions on both the source and destination buckets. Additionally, you can choose the destination storage class and enable additional replication options such as S3 Replication Time Control (S3 RTC) or S3 Batch Replication. https://medium.com/cloud-techies/s3-same-region-replication-srr-and-cross-region-replication-crr-34d446806bab https://aws.amazon.com/getting-started/hands-on/replicate-data-using-amazon-s3-replication/ https://docs.aws.amazon.com/AmazonS3/latest/userguide/replication.html

QUESTION 15
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.

Correct Answer: B, C, E
Section:
Explanation:
https://aws.amazon.com/premiumsupport/knowledge-center/lambda-function-assume-iam-role/ https://kreuzwerker.de/post/aws-multi-account-setups-reloaded

QUESTION 16
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.

Correct Answer: C
Section:
Explanation:
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.

QUESTION 17
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.

Correct Answer: D
Section:
Explanation:
service catalog uses stacksets and can enforce tag and restrict resources AWS Customer case with tag enforcement https://aws.amazon.com/ko/blogs/apn/enforce-centralized-tag-compliance-using-aws-service-catalog-amazon-dynamodb-aws-lambda-and-amazon-cloudwatch-events/ And Youtube video showing how to restrict resources per user with portfolio https://www.youtube.com/watch?v=LzvhTcqqyog

QUESTION 18
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.

Correct Answer: B, D
Section:
Explanation:
https://docs.aws.amazon.com/vpc/latest/userguide/vpc-nat-gateway.html

QUESTION 19
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.

Correct Answer: B
Section:

QUESTION 20
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.

Correct Answer: D
Section:
Explanation:
because of 'as few maintenance tasks as possible on the underlying infrastructure'. Fargate does that better than 'one centralized application server'

QUESTION 21
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.

Correct Answer: C, D, F
Section:
Explanation:
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. This will allow users in the new Region to access the application with lower latency by reducing the network hops between the user and the application servers.
D) Create Amazon Route 53 records, health checks, and latency-based routing policies to route to the ALB. This will enable Route 53 to route user traffic to the nearest healthy ALB, based on the latency between the user and the ALBs.
F) Convert the DynamoDB table to a global table. This will enable reads and writes to the table in both Regions with low latency, improving the overall response time of the application

QUESTION 22
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.

Correct Answer: D
Section:

QUESTION 23
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.

Correct Answer: A
Section:
Explanation:
'Note that the metric filter is different from a log insights query, where the experience is interactive and provides immediate search results for the user to investigate. No automatic action can be invoked from an insights query. Metric filters, on the other hand, will generate metric data in the form of a time series. This lets you create alarms that integrate into your ITSM processes, execute AWS Lambda functions, or even create anomaly detection models.' https://aws.amazon.com/blogs/mt/quantify-custom-application-metrics-with-amazon-cloudwatch-logs-and-metric-filters/

QUESTION 24
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.

Correct Answer: C
Section:
Explanation:
The following are the steps that the DevOps engineer should take to reduce the latency of the Lambda function at all times of the day:
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.
The provisioned concurrency setting ensures that there is always a minimum number of Lambda function instances available to handle requests. The Application Auto Scaling setting will automatically scale the number of Lambda function instances up or down based on the demand for the application.
This solution will ensure that the Lambda function is able to handle the increased load during the middle of the day, while also keeping the cold-start latency low.
The following are the reasons why the other options are not correct:
Option A is incorrect because it will not reduce the cold-start latency of the Lambda function.
Option B is incorrect because it will not scale the number of Lambda function instances up or down based on demand.
Option D is incorrect because it will only configure reserved concurrency on the API Gateway API, which will not affect the Lambda function.

QUESTION 25
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.

Correct Answer: B
Section:
Explanation:
The following are the steps that the company can take to change the log level dynamically when the deployment occurs:
Create a script that uses the CodeDeploy environment variableDEPLOYMENT_GROUP_NAMEto identify which deployment group the instanceis part of.
Use this information to configure the log level settings.
Reference this script as part of theBeforeInstalllifecycle hook in theappspec.ymlfile.
The DEPLOYMENT_GROUP_NAME environment variable is automatically set by CodeDeploy when the deployment is triggered. This means that the script does not need to call the metadata service or the EC2 API to identify the deployment group.
This solution is the least complex and requires the least management overhead. It also does not require different script versions for each deployment group.
The following are the reasons why the other options are not correct:
Option A is incorrect because it would require tagging the Amazon EC2 instances, which would be a manual and time-consuming process.
Option C is incorrect because it would require creating a custom environment variable for each environment. This would be a complex and error-prone process.
Option D is incorrect because it would use theDEPLOYMENT_GROUP_IDenvironment variable. However, this variable is not automatically set by CodeDeploy, so the script would need to call the metadata service or the EC2 API to get the deployment group ID. This would add complexity and overhead to the solution.

QUESTION 26
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.

Correct Answer: B
Section:
Explanation:
The following are the steps that the DevOps engineer should take 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:
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.
The managed rule AWS::Config::EBSVolumesWithoutBackupTag will return a compliance failure for any EBS volume that does not have the Backup_Frequency tag applied. The remediation action will then use the Systems Manager Automation runbook to apply the Backup_Frequency tag with a value of weekly to the EBS volume.

QUESTION 27
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.

Correct Answer: C
Section:
Explanation:
To meet the requirements, the DevOps engineer should do the following:
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.
Turning on the Multi-AZ option will create a replica of the database in a different Availability Zone. This will ensure that the database remains available even if one of the Availability Zones is unavailable.
Updating the application to use the Aurora cluster endpoint for write operations will ensure that all writes are sent to both the primary and replica databases. This will ensure that the data is always consistent.
Updating the Aurora cluster's reader endpoint for reads will allow the application to read data from the replica database. This will improve the performance of the application during the maintenance window.

QUESTION 28
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.

Correct Answer: A, D, F
Section:
Explanation:
The Auto Scaling group service-linked role must have a specific grant in the source account in order to decrypt the encrypted AMI. This is because the service-linked role does not have permissions to assume the default IAM role in the source account.
The following steps are required to meet the requirements:
In the source account, copy the unencrypted AMI to an encrypted AMI. Specify the KMS key in the copy action.
In the source account, create a KMS grant that delegates permissions to the Auto Scaling group service-linked role in the target account.
In the source account, share the encrypted AMI with the target account.
In the target account, attach the KMS grant to the Auto Scaling group service-linked role.
The first three steps are the same as the steps that I described earlier. The fourth step is required to grant the Auto Scaling group service-linked role permissions to decrypt the AMI in the target account.

QUESTION 29
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.

Correct Answer: A, D
Section:
Explanation:
https://docs.aws.amazon.com/codedeploy/latest/userguide/integrations-aws-auto-scaling.html

QUESTION 30
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.

Correct Answer: A, C
Section:
Explanation:
If instead you want to automatically apply the policy to existing in-scope resources, choose Auto remediate any noncompliant resources. This option creates a web ACL in each applicable account within the AWS organization and associates the web ACL with the resources in the accounts. When you choose Auto remediate any noncompliant resources, you can also choose to remove existing web ACL associations from in-scope resources, for the web ACLs that aren't managed by another active Firewall Manager policy. If you choose this option, Firewall Manager first associates the policy's web ACL with the resources, and then removes the prior associations. If a resource has an association with another web ACL that's managed by a different active Firewall Manager policy, this choice doesn't affect that association.

QUESTION 31
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.

Correct Answer: C
Section:

Explanation:
https://aws.amazon.com/blogs/security/how-to-use-aws-config-to-determine-compliance-of-aws-kms-key-policies-to-your-specifications/

QUESTION 32
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.

Correct Answer: C
Section:
Explanation:
A bucket policy is a resource-based policy that defines who can access a specific S3 bucket and what actions they can perform on it. By removing unauthenticated access from the bucket policy, you can prevent anyone without valid credentials from accessing the bucket. A service role is an IAM role that allows an AWS service, such as CodeBuild, to perform actions on your behalf. By modifying the service role for the CodeBuild project to include Amazon S3 access, you can grant the project permission to read and write objects in the S3 bucket. The AWS CLI is a command-line tool that allows you to interact with AWS services, such as S3, using commands in your terminal. By using the AWS CLI to download the database population script, you can leverage the service role credentials and encryption to secure the data transfer.
For more information, you can refer to these web pages:
[Using bucket policies and user policies - Amazon Simple Storage Service]
[Create a service role for CodeBuild - AWS CodeBuild]
[AWS Command Line Interface]

QUESTION 33
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.

Correct Answer: B, C, F
Section:
Explanation:
Using the principalTag in the Permission Set inline policy a logged in user belonging to a specific AD group in the IDP can be permitted access to perform operations on certain resources if their group matches the group used in the PrincipleTag. Basically you are narrowing the scope of privileges assigned via Permission policies conditionally based on whether the logged in user belongs to a specific AD Group in IDP. The mapping of the AD group to the request attributes can be done using SSO attributes where we can pass other attributes like the SAML token as well.
https://docs.aws.amazon.com/singlesignon/latest/userguide/abac.html

QUESTION 34
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.

Correct Answer: A, D
Section:
Explanation:
A: If the ApproximateAgeOfOldestMessages indicate that orders are remaining in the SQS queue for longer than expected, the reserved concurrency limit may be set too small to keep up with the number of orders entering the queue and is being throttled. D: The DynamoDB table is using Auto Scaling. With Auto Scaling, you create a scaling policy that specifies whether you want to scale read capacity or write capacity (or both), and the minimum and maximum provisioned capacity unit settings for the table. The ThottledWriteRequests metric will indicate if there is a throttling issue on the DynamoDB table, which can be resolved by increasing the maximum write capacity units for the table's Auto Scaling policy. https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/AutoScaling.html

QUESTION 35
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.

Correct Answer: B
Section:
Explanation:
https://docs.aws.amazon.com/config/latest/developerguide/ec2-instance-profile-attached.html

QUESTION 36
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.

Correct Answer: D
Section:
Explanation:
Use routing configuration on an alias to send a portion of traffic to a second function version. For example, you can reduce the risk of deploying a new version by configuring the alias to send most of the traffic to the existing version, and only a small percentage of traffic to the new version. https://docs.aws.amazon.com/lambda/latest/dg/configuration-aliases.html
The following are the steps involved in the deploy stage configuration that will meet the requirements:
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.
This configuration will help to reduce the customer impact of an unsuccessful deployment by deploying the new version of the functions to a staging environment first. This will allow the DevOps engineer to test the new version of the functions before deploying it to production.
The configuration will also help to monitor for issues by including Amazon CloudWatch alarms. These alarms will alert the DevOps engineer if there are any problems with the new version of the functions.

QUESTION 37
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.

Correct Answer: C
Section:
Explanation:
EC2 instances running in private subnets of a VPC can now have controlled access to S3 buckets, objects, and API functions that are in the same region as the VPC. You can use an S3 bucket policy to indicate which VPCs and which VPC Endpoints have access to your S3 buckets 1- https://aws.amazon.com/pt/blogs/aws/new-vpc-endpoint-for-amazon-s3/

QUESTION 38
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.

Correct Answer: A
Section:
Explanation:

When you create a pipeline from CodePipeline during the step-by-step it creates a CloudWatch Event rule for a given branch and repo like this:

```
{
'source': [
'aws.codecommit'
],
'detail-type': [
'CodeCommit Repository State Change'
],
'resources': [
'arn:aws:codecommit:us-east-1:xxxxx:repo-name'
],
'detail': {
'event': [
'referenceCreated',
'referenceUpdated'
],
'referenceType': [
'branch'
],
'referenceName': [
'master'
]
}
```

https://docs.aws.amazon.com/codepipeline/latest/userguide/pipelines-trigger-source-repo-changes-console.html

QUESTION 39

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.

Correct Answer: C
Section:
Explanation:

To meet the requirements, the DevOps engineer should create an Amazon EventBridge event rule that has the default event bus as the source. The rule's event pattern should match EC2 security group creation and modification events, and it should be configured to invoke the Lambda function. This solution will allow for near real-time detection of security group rule changes and will trigger the Lambda function to remove any unrestricted rules and send email notifications to the security team.

https://repost.aws/knowledge-center/monitor-security-group-changes-ec2

QUESTION 40

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.

Correct Answer: D
Section:
Explanation:
it involves adding an IPv6 CIDR block to the VPC and subnets for the ALB and specifying the dualstack IP address type on the ALB listener. This allows the ALB to listen on both IPv4 and IPv6 addresses, and forward requests to the EC2 instances that are added as targets to the target group associated with the ALB.

QUESTION 41
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.

Correct Answer: D
Section:
Explanation:
AWS Organizations is a service that helps to manage multiple AWS accounts. AWS Control Tower is a service that makes it easy to set up and govern secure, compliant multi-account AWS environments. Account Factory for Terraform (AFT) is an AWS Control Tower feature that provisions new accounts using Terraform templates. To provision new accounts with the Enterprise Support plan, the DevOps engineer can set the aft_feature_enterprise_support feature flag to True in the AFT deployment input configuration. This flag enables the Enterprise Support plan for newly provisioned accounts.
https://docs.aws.amazon.com/controltower/latest/userguide/aft-feature-options.html

QUESTION 42
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.

Correct Answer: C

Section:
Explanation:
AWS Health provides real-time events and information related to your AWS infrastructure. It can be integrated with Amazon EventBridge to act upon the health events automatically. If the maintenance notification from AWS Health indicates that an EC2 instance requires a restart, you can set up an EventBridge rule to respond to such events. In this case, the target of this rule would be a Lambda function that would trigger a Systems Manager automation to restart the EC2 instance during a maintenance window. Remember, AWS Health is the source of the events (not EC2 or Systems Manager), and AWS Lambda can be used to execute complex remediation tasks, such as scheduling maintenance tasks via Systems Manager.
The following are the steps involved in configuring the EventBridge rule to meet these requirements:
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.
The AWS Lambda function will be triggered by the event from AWS Health. The function will then register an automation task to restart the EC2 instance during the next maintenance window.

QUESTION 43
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.

Correct Answer: D
Section:
Explanation:
The following are the steps involved in accomplishing this in the most maintainable manner:
Use AWS CodeBuild with artifact encryption to replace the Jenkins instance running on EC2 instances.
Configure CodeBuild to encrypt the build artifacts using AWS Secrets Manager.
Deploy the containerized quality control applications to CodeBuild.
This approach is the most maintainable because it eliminates the need to manage Jenkins on EC2 instances. CodeBuild is a managed service, so the DevOps engineer does not need to worry about patching or upgrading the service.
https://docs.aws.amazon.com/codebuild/latest/userguide/security-encryption.html Build artifact encryption - CodeBuild requires access to an AWS KMS CMK in order to encrypt its build output artifacts. By default, CodeBuild uses an AWS Key Management Service CMK for Amazon S3 in your AWS account. If you do not want to use this CMK, you must create and configure a customer-managed CMK. For more information Creating keys.

QUESTION 44
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.

Correct Answer: B
Section:
Explanation:
https://aws.amazon.com/premiumsupport/knowledge-center/cloudformation-s3-custom-resources/

QUESTION 45
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.

Correct Answer: A, B
Section:
Explanation:
A) The CloudFormation template should be modified to include a parameter that indicates the location of the .zip file containing the Lambda function's code. This allows the CloudFormation deploy action to use the correct artifact depending on the region. This is critical because Lambda functions need to reference their code artifacts from the same region they are being deployed in. B. You would also need to create a new CloudFormation deploy action for the us-east-1 Region within the pipeline. This action should be configured to use the CloudFormation template from the artifact that was specifically created for us-east-1.

QUESTION 46
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.

Correct Answer: B
Section:
Explanation:
https://aws.amazon.com/blogs/mt/how-to-set-up-aws-opsworks-stacks-auto-healing-notifications-in-amazon-cloudwatch-events/

QUESTION 47
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. Option D

Correct Answer: C
Section:
Explanation:
https://docs.aws.amazon.com/singlesignon/latest/userguide/configure-abac.html

QUESTION 48
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.

Correct Answer: A
Section:
Explanation:
https://docs.aws.amazon.com/organizations/latest/userguide/orgs_manage_policies_scps.html?icmpid=docs_orgs_console

QUESTION 49
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.

Correct Answer: D
Section:
Explanation:
https://aws.amazon.com/quickstart/architecture/blue-green-deployment/

QUESTION 50
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.

Correct Answer: D
Section:
Explanation:
https://blog.fourninecloud.com/auto-scaling-lifecycle-hooks-to-export-server-logs-when-instance-terminating-58e06d7c0d6a

QUESTION 51
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.

Correct Answer: B, D, E
Section:
Explanation:
https://docs.aws.amazon.com/IAM/latest/UserGuide/tutorial_cross-account-with-roles.html

QUESTION 52
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.

Correct Answer: C
Section:
Explanation:
Amazon GuardDuty is primarily on threat detection and response, not configuration monitoring A conformance pack is a collection of AWS Config rules and remediation actions that can be easily deployed as a single entity in an account and a Region or across an organization in AWS Organizations. https://docs.aws.amazon.com/config/latest/developerguide/conformance-packs.html
https://docs.aws.amazon.com/config/latest/developerguide/s3-account-level-public-access-blocks.html

QUESTION 53
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.

Correct Answer: A
Section:

Explanation:
https://docs.aws.amazon.com/AmazonCloudWatch/latest/logs/SubscriptionFilters.html#LambdaFunctionExample
https://docs.aws.amazon.com/AmazonCloudWatch/latest/logs/SubscriptionFilters.html

QUESTION 54
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.

Correct Answer: B, D, F
Section:
Explanation:
https://docs.aws.amazon.com/AmazonECS/latest/developerguide/ecs-logging-monitoring.html

QUESTION 55
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.

Correct Answer: A
Section:
Explanation:
https://docs.aws.amazon.com/systems-manager/latest/userguide/patch-manager-how-it-works-alt-source-repository.html

QUESTION 56
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.

Correct Answer: C
Section:
Explanation:
https://aws.amazon.com/premiumsupport/knowledge-center/sns-lambda-webhooks-chime-slack-teams/

QUESTION 57
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.

Correct Answer: A
Section:
Explanation:
https://aws.amazon.com/premiumsupport/knowledge-center/monitor-security-group-changes-ec2/

QUESTION 58
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.

Correct Answer: A, C
Section:
Explanation:
https://docs.aws.amazon.com/AmazonCloudWatch/latest/monitoring/install-CloudWatch-Agent-on-premise.html https://docs.aws.amazon.com/xray/latest/devguide/xray-api-sendingdata.html

QUESTION 59
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.

Correct Answer: D
Section:
Explanation:
EventBridge is needed to detect the database failure. Lambda is needed to promote the replica as it's in another Region (manual promotion, otherwise). Storing and updating the endpoint in Parameter store is important in updating the application. Look at High Availability section of Aurora FAQ: https://aws.amazon.com/rds/aurora/faqs/

QUESTION 60
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.

Correct Answer: A
Section:
Explanation:
This solution will meet the requirements because it will use a CloudFormation custom resource to execute custom logic during the stack set operation. A custom resource is a resource that you define in your template and that is associated with an AWS Lambda function. The Lambda function runs whenever the custom resource is created, updated, or deleted, and can perform any actions that are supported by the AWS SDK. In this case, the Lambda function can use the GuardDuty API to check whether GuardDuty is already enabled in each target account, and if not, enable it. This way, the DevOps engineer can avoid deploying the stack set to accounts that already have GuardDuty enabled, and prevent failure during the deployment.

QUESTION 61
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.

Correct Answer: B
Section:
Explanation:
CodeCommit generates events in CloudWatch, CloudWatch triggers the CodeBuild https://aws.amazon.com/es/blogs/devops/complete-ci-cd-with-aws-codecommit-aws-codebuild-aws-codedeploy-and-aws-codepipeline/

QUESTION 62
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.

Correct Answer: D
Section:
Explanation:
This solution meets the requirements in the most operationally efficient manner by automating the application restart process on the instances without restarting them. When the CloudWatch alarm enters the ALARM state, the EventBridge rule is triggered, which in turn invokes the Systems Manager Automation runbook that contains the script to restart the application on the instances.

QUESTION 63
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.

Correct Answer: A, C, E
Section:

QUESTION 64
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.

Correct Answer: B
Section:
Explanation:
https://docs.aws.amazon.com/codedeploy/latest/userguide/deployment-steps.html

QUESTION 65
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.

Correct Answer: A
Section:
Explanation:
Option A is the most cost-effective solution. By configuring a warm pool of EC2 instances in the Stopped state, the company can reduce the time it takes for new instances to be ready to serve requests. When the Auto Scaling group launches a new instance, it can attach the stopped EC2 instance from the warm pool. The instance can then be started up immediately, rather than having to wait for the data to be downloaded and processed. This reduces the overall startup time for the application.

QUESTION 66
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.

Correct Answer: D
Section:
Explanation:
When setting the flag authenticated-read in the command line, the owner gets FULL_CONTROL. The AuthenticatedUsers group (Anyone with an AWS account) gets READ access.
Reference: https://docs.aws.amazon.com/AmazonS3/latest/userguide/acl-overview.html

QUESTION 67
A company has 20 service learns 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.

Correct Answer: B
Section:
Explanation:
https://aws.amazon.com/blogs/networking-and-content-delivery/connecting-networks-with-overlapping-ip-ranges/ Private link is the best option because Transit Gateway doesn't support overlapping CIDR ranges.

QUESTION 68
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 Cloud Format ion 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 Cloud Formation actions fail with an access denied error.
Which combination of actions must the DevOps engineer perform to resolve this error? (Select 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 CodePipelme account modify the artifacts S3 bucket policy to allow the roles access Configure the CodePipeline CloudFormation action to use the roles.

Correct Answer: B, E
Section:

QUESTION 69
A company has many AWS accounts. During AWS account creation the company uses automation to create an Amazon CloudWatch Logs log group in every AWS Region that the company operates in. The automaton configures new resources in the accounts to publish logs to the provisioned log groups in their Region.
The company has created a logging account to centralize the logging from all the other accounts. A DevOps engineer needs to aggregate the log groups from all the accounts to an existing Amazon S3 bucket in the logging account.
Which solution will meet these requirements in the MOST operationally efficient manner?

A. In the logging account create a CloudWatch Logs destination with a destination policy. For each new account subscribe the CloudWatch Logs log groups to the. Destination Configure a single Amazon Kinesis data stream and a single Amazon Kinesis Data Firehose delivery stream to deliver the logs from the CloudWatch Logs destination to the S3 bucket.

B. In the logging account create a CloudWatch Logs destination with a destination policy for each Region. For each new account subscribe the CloudWatch Logs log groups to the destination. Configure a single Amazon Kinesis data stream and a single Amazon Kinesis Data Firehose delivery stream to deliver the logs from all the CloudWatch Logs destinations to the S3 bucket.

C. In the logging account create a CloudWatch Logs destination with a destination policy for each Region. For each new account subscribe the CloudWatch Logs log groups to the destination Configure an Amazon Kinesis data stream and an Amazon Kinesis Data Firehose delivery stream for each Region to deliver the logs from the CloudWatch Logs destinations to the S3 bucket.

D. In the logging account create a CloudWatch Logs destination with a destination policy. For each new account subscribe the CloudWatch Logs log groups to the destination. Configure a single Amazon Kinesis data stream to deliver the logs from the CloudWatch Logs destination to the S3 bucket.

Correct Answer: C
Section:
Explanation:
This solution will meet the requirements in the most operationally efficient manner because it will use CloudWatch Logs destination to aggregate the log groups from all the accounts to a single S3 bucket in the logging account. However, unlike option A, this solution will create a CloudWatch Logs destination for each region, instead of a single destination for all regions. This will improve the performance and reliability of the log delivery, as it will avoid cross-region data transfer and latency issues. Moreover, this solution will use an Amazon Kinesis data stream and an Amazon Kinesis Data Firehose delivery stream for each region, instead of a single stream for all regions. This will also improve the scalability and throughput of the log delivery, as it will avoid bottlenecks and throttling issues that may occur with a single stream.

QUESTION 70
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-xxxxxxxxxx"
                    ]
                }
            }
        }
    ]
}
```

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 Code Artifact? (Select 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.

Correct Answer: A, D
Section:
Explanation:
'AWS CodeArtifact operates in multiple Availability Zones and stores artifact data and metadata in Amazon S3 and Amazon DynamoDB. Your encrypted data is redundantly stored across multiple facilities and multiple devices in each facility, making it highly available and highly durable.' https://aws.amazon.com/codeartifact/features/ With no internet connectivity, a gateway endpoint becomes necessary to access S3.

QUESTION 71
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 togs to Amazon S3 Use Amazon Athena to query both sets of logs in Amazon S3.

Correct Answer: D
Section:
Explanation:
This solution will meet the requirements because it will use Amazon S3 as a common data lake for both the application logs and the API logs. Amazon S3 is a service that provides scalable, durable, and secure object storage for any type of data. You can use the Amazon CloudWatch agent to send logs from your EC2 instances to S3 buckets, and use AWS CloudTrail to deliver the API logs to S3 buckets as well. You can also use Amazon Athena to query both sets of logs in S3 using standard SQL, without loading or transforming them. Athena is a serverless interactive query service that allows you to analyze data in S3 using a variety of data formats, such as JSON, CSV, Parquet, and ORC.

QUESTION 72
A company wants to use a grid system for a proprietary enterprise m-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 thecluster 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 system, 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. Tile whenever a new instance is added to the cluster.

Correct Answer: A
Section:
Explanation:
You can run custom recipes manually, but the best approach is usually to have AWS OpsWorks Stacks run them automatically. Every layer has a set of built-in recipes assigned each of five lifecycle events---Setup, Configure, Deploy, Undeploy, and Shutdown. Each time an event occurs for an instance, AWS OpsWorks Stacks runs the associated recipes for each of the instance's layers, which handle the corresponding tasks. For example, when an instance finishes booting, AWS OpsWorks Stacks triggers a Setup event. This event runs the associated layer's Setup recipes, which typically handle tasks such as installing and configuring packages

QUESTION 73
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 m 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 tor an existing session.

B. Enable session sharing on the toad 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.

Correct Answer: D
Section:

Explanation:
https://aws.amazon.com/caching/session-management/

QUESTION 74
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? (Select 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.

Correct Answer: C, D
Section:
Explanation:
These solutions will meet the requirement because they will prevent the loss of notification messages in the future. An Amazon SQS queue is a service that provides a reliable, scalable, and secure message queue for asynchronous communication between distributed components. You can use an SQS queue to buffer messages from an SNS topic and ensure that they are delivered and processed by a Lambda function, even if the function or the database is temporarily unavailable.
Option C will configure an SQS dead-letter queue for the SNS topic. A dead-letter queue is a queue that receives messages that could not be delivered to any subscriber after a specified number of retries. You can use a dead-letter queue to store and analyze failed messages, or to reprocess them later. This way, you can avoid losing messages that could not be delivered to the Lambda function due to network errors, throttling, or other issues.
Option D will subscribe an SQS queue to the SNS topic and configure the Lambda function to process messages from the SQS queue. This will decouple the SNS topic from the Lambda function and provide more flexibility and control over the message delivery and processing. You can use an SQS queue to store messages from the SNS topic until they are ready to be processed by the Lambda function, and also to retry processing in case of failures. This way, you can avoid losing messages that could not be processed by the Lambda function due to database errors, timeouts, or other issues.

QUESTION 75
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.

Correct Answer: B
Section:
Explanation:
https://docs.aws.amazon.com/health/latest/ug/cloudwatch-events-health.html

QUESTION 76

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 seventies 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.

Correct Answer: B
Section:
Explanation:
'The firewall appliance sends logs to Amazon CloudWatch Logs and includes event severities of CRITICAL, HIGH, MEDIUM, LOW, and INFO'

QUESTION 77
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 m source control.
Which solution will accomplish this?

A. In the CloudFormaiion 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-mit 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 imt metadata. Place the configuration file content m 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.

Correct Answer: D
Section:
Explanation:
Use the AWS::CloudFormation::Init type to include metadata on an Amazon EC2 instance for the cfn-init helper script. If your template calls the cfn-init script, the script looks for resource metadata rooted in the AWS::CloudFormation::Init metadata key.
Reference: https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/aws-resource-init.html

QUESTION 78
A company updated the AWS Cloud Formation 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? (Select TWO.)

A. Attach the AWSC loud Formation FullAccess IAM policy to the AWS CtoudFormation 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.

Correct Answer: C, D
Section:
Explanation:
https://docs.aws.amazon.com/cli/latest/reference/cloudformation/continue-update-rollback.html For a specified stack that is in the UPDATE_ROLLBACK_FAILED state, continues rolling it back to the UPDATE_ROLLBACK_COMPLETE state. Depending on the cause of the failure, you can manually fix the error and continue the rollback. By continuing the rollback, you can return your stack to a working state (the UPDATE_ROLLBACK_COMPLETE state), and then try to update the stack again.

QUESTION 79
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 hearth 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 ma 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.

Correct Answer: D
Section:

QUESTION 80
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.

Correct Answer: B
Section:
Explanation:
https://docs.aws.amazon.com/autoscaling/ec2/userguide/ASGettingNotifications.html#auto-scaling-sns-notifications

QUESTION 81
A company requires an RPO of 2 hours and an RTO of 10 minutes for its data and application at all times. An application uses a MySQL database and Amazon EC2 web servers. The development team needs a strategy for failover and disaster recovery.
Which combination of deployment strategies will meet these requirements? (Select TWO.)

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

Correct Answer: B, D
Section:

QUESTION 82
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 (unction 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 to 1 hour. Create an Amazon EventBridge rule tor 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 StopLoggmg 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.
https://aws.amazon.com/blogs/mt/monitor-changes-and-auto-enable-logging-in-aws-cloudtrail/

Correct Answer: A
Section:

QUESTION 83
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 Pile System (Amazon EFS) as shared storage in Account A in the organization.
The company wants to continue to use Amazon EPS 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 EPS access point.
Which combination of steps should the DevOps engineer take to meet these requirements? (Select 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. Need to assume cross-account IAM role to describe the mounts so that a specific mount can be chosen.

Correct Answer: A, E, F

Section:
Explanation:

https://docs.aws.amazon.com/lambda/latest/dg/services-efs.html

https://aws.amazon.com/ru/blogs/storage/mount-amazon-efs-file-systems-cross-account-from-amazon-eks/

1. Need to update the file system policy on EFS to allow mounting the file system into Account B.

File System Policy

$ cat file-system-policy.json

{

'Statement': [

{

'Effect': 'Allow',

'Action': [

'elasticfilesystem:ClientMount',

'elasticfilesystem:ClientWrite'

],

'Principal': {

'AWS': 'arn:aws:iam::<aws-account-id-A>:root' # Replace with AWS account ID of EKS cluster

}

}

]

}

2. Need VPC peering between Account A and Account B as the pre-requisite

3. Need to assume cross-account IAM role to describe the mounts so that a specific mount can be chosen.

QUESTION 84

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.

Correct Answer: A, E
Section:

QUESTION 85

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.

Correct Answer: D
Section:
Explanation:

QUESTION 86
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.

Correct Answer: A, D
Section:
Explanation:
A- https://docs.aws.amazon.com/codedeploy/latest/userguide/codedeploy-agent.html D - This option correctly utilizes AWS CodePipeline to invoke the CodeBuild job and create CloudFormation change sets. It adds a manual approval step before executing the change sets and starting the AWS CodeDeploy deployment. This ensures that the deployment process is automated while retaining the final manual approval step.

QUESTION 87
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.

Correct Answer: C
Section:
Explanation:
https://docs.aws.amazon.com/codedeploy/latest/APIReference/API_BlueInstanceTerminationOption.html
The original revision termination settings are configured to wait 1 hour after traffic has been rerouted before terminating the blue task set. https://docs.aws.amazon.com/AmazonECS/latest/developerguide/deployment-type-bluegreen.html

QUESTION 88
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.

Correct Answer: A, E
Section:
Explanation:
https://aws.amazon.com/blogs/mt/org-aggregator-delegated-admin/ https://docs.aws.amazon.com/organizations/latest/userguide/services-that-can-integrate-config.html

QUESTION 89
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.

Correct Answer: B

Section:
Explanation:
https://docs.aws.amazon.com/serverless-application-model/latest/developerguide/automating-updates-to-serverless-apps.html

QUESTION 90
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 loT 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 m 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 loT devices and on-premises infrastructure? (Select THREE.)

A. Apply tags lo all the EC2 instances. AWS loT 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 loT Greengrass devices and on-premises servers.
C. Use Systems Manager Patch Manager to schedule patching loT the EC2 instances AWS loT 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 lo initiate a patch action for all EC2 instances AWS loT 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 loT 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 loT Greengrass IAM token exchange role Use the role to deploy SSM Agent on all the loT devices.

Correct Answer: C, E, F
Section:
Explanation:
https://aws.amazon.com/blogs/mt/how-to-centrally-manage-aws-iot-greengrass-devices-using-aws-systems-manager/?force_isolation=true

QUESTION 91
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 to 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 Cloud Formation 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 Cloud Formation 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 even 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.

Correct Answer: B
Section:
Explanation:
It allows the company 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 using Amazon GuardDuty. It also provides a solution for automatically adding future AWS accounts to GuardDuty by configuring GuardDuty to add newly created AWS accounts by invitation and to send invitations to the existing

AWS accounts.

QUESTION 92
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 pad of the ALB target group.

Correct Answer: C
Section:
Explanation:
This failure is typically due to incorrectly configured health checks in Elastic Load Balancing for the Classic Load Balancer, Application Load Balancer, or Network Load Balancer used to manage traffic for the deployment group. To resolve the issue, review and correct any errors in the health check configuration for the load balancer. https://docs.aws.amazon.com/codedeploy/latest/userguide/troubleshooting-deployments.html#troubleshooting-deployments-allowtraffic-no-logs

QUESTION 93
A company that runs many workloads on AWS has an Amazon EBS spend that has increased over time. The DevOps team notices there are many unattached
EBS volumes. Although there are workloads where volumes are detached, volumes over 14 days old are stale and no longer needed. A DevOps engineer has been tasked with creating automation that deletes unattached EBS volumes that have been unattached for 14 days.
Which solution will accomplish this?

A. Configure the AWS Config ec2-volume-inuse-check managed rule with a configuration changes trigger type and an Amazon EC2 volume resource target. Create a new Amazon CloudWatch Events rule scheduled to execute an AWS Lambda function in 14 days to delete the specified EBS volume.
B. Use Amazon EC2 and Amazon Data Lifecycle Manager to configure a volume lifecycle policy. Set the interval period for unattached EBS volumes to 14 days and set the retention rule to delete. Set the policy target volumes as *.
C. Create an Amazon CloudWatch Events rule to execute an AWS Lambda function daily. The Lambda function should find unattached EBS volumes and tag them with the current date, and delete unattached volumes that have tags with dates that are more than 14 days old.
D. Use AWS Trusted Advisor to detect EBS volumes that have been detached for more than 14 days. Execute an AWS Lambda function that creates a snapshot and then deletes the EBS volume.

Correct Answer: C
Section:
Explanation:
The requirement is to create automation that deletes unattached EBS volumes that have been unattached for 14 days. To do this, the DevOps engineer needs to use the following steps:
Create an Amazon CloudWatch Events rule to execute an AWS Lambda function daily. CloudWatch Events is a service that enables event-driven architectures by delivering events from various sources to targets. Lambda is a service that lets you run code without provisioning or managing servers. By creating a CloudWatch Events rule that executes a Lambda function daily, the DevOps engineer can schedule a recurring task to check and delete unattached EBS volumes.
The Lambda function should find unattached EBS volumes and tag them with the current date, and delete unattached volumes that have tags with dates that are more than 14 days old. The Lambda function can use the EC2 API to list and filter unattached EBS volumes based on their state and tags. The function can then tag each unattached volume with the current date using the create-tags command. The function can also compare the tag value with the current date and delete any unattached volume that has been tagged more than 14 days ago using the delete-volume command.

QUESTION 94
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 eng*neer 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 OrganzatlonAccountAccessR01e IAM role in the new member account.
B. In the management account, create a new SCR 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 AVVS 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 Organ zationAccountAccessRole IAM role. Grant the management account permission to assume the role.

Correct Answer: C
Section:
Explanation:
The problem is that the DevOps engineer cannot assume the OrganizationAccountAccessRole IAM role in the new member account that joined AnyCompany's management account through an Organizations invitation. The solution is to create a new IAM role with the same name and trust policy in the new member account.
Option A is incorrect, as it does not address the root cause of the error. The DevOps engineer's IAM user already has permission to assume the OrganizationAccountAccessRole IAM role in any member account, as this is the default role name that AWS Organizations creates when a new account joins an organization. The error occurs because the new member account does not have this role, as it was not created by AWS Organizations.
Option B is incorrect, as it does not address the root cause of the error. An SCP is a policy that defines the maximum permissions for account members of an organization or organizational unit (OU). An SCP does not grant permissions to IAM users or roles, but rather limits the permissions that identity-based policies or resource-based policies grant to them. An SCP also does not affect how IAM roles are assumed by other principals.
Option C is correct, as it addresses the root cause of the error. By creating a new IAM role with the same name and trust policy as the OrganizationAccountAccessRole IAM role in the new member account, the DevOps engineer can assume this role and access the account. The new role should have the AdministratorAccess AWS managed policy attached, which grants full access to all AWS resources in the account. The trust policy should allow the management account to assume the role, which can be done by specifying the management account ID as a principal in the policy statement.
Option D is incorrect, as it assumes that the new member account already has the OrganizationAccountAccessRole IAM role, which is not true. The new member account does not have this role, as it was not created by AWS Organizations. Editing the trust policy of a non-existent role will not solve the problem.

QUESTION 95
A company needs to implement failover for its application. The application includes an Amazon CloudFront distribution and a public Application Load Balancer (ALB) in an AWS Region. The company has configured the ALB as the default origin for the distribution.
After some recent application outages, the company wants a zero-second RTO. The company deploys the application to a secondary Region in a warm standby configuration. A DevOps engineer needs to automate the failover of the application to the secondary Region so that HTTP GET requests meet the desired RTO.
Which solution will meet these requirements?

A. Create a second CloudFront distribution that has the secondary ALB as the default origin. Create Amazon Route 53 alias records that have a failover policy and Evaluate Target Health set to Yes for both CloudFront distributions. Update the application to use the new record set.
B. Create a new origin on the distribution for the secondary ALB. Create a new origin group. Set the original ALB as the primary origin. Configure the origin group to fail over for HTTP 5xx status codes. Update the default behavior to use the origin group.
C. Create Amazon Route 53 alias records that have a failover policy and Evaluate Target Health set to Yes for both ALBs. Set the TTL of both records to 0. Update the distribution's origin to use the new record set.
D. Create a CloudFront function that detects HTTP 5xx status codes. Configure the function to return a 307 Temporary Redirect error response to the secondary ALB if the function detects 5xx status codes. Update the distribution's default behavior to send origin responses to the function.

Correct Answer: B
Section:
Explanation:
The best solution to implement failover for the application is to use CloudFront origin groups.Origin groups allow CloudFront to automatically switch to a secondary origin when the primary origin is unavailable or returns specific HTTP status codes that indicate a failure1. This way, CloudFront can serve the requests from the secondary ALB in the secondary Region without any delay or redirection. To set up origin groups, the DevOps engineer needs to create a new origin on the distribution for the secondary ALB, create a new origin group with the original ALB as the primary origin and the secondary ALB as the secondary origin, and configure the origin group to fail over for HTTP 5xx status codes.Then, the DevOps engineer needs to update the default behavior to use the origin group instead of the single origin2.
The other options are not as effective or efficient as the solution in option B. Option A is not suitable because creating a second CloudFront distribution will increase the complexity and cost of the application. Moreover, using

Route 53 alias records with a failover policy will introduce some delay in detecting and switching to the secondary CloudFront distribution, which may not meet the zero-second RTO requirement.Option C is not feasible because CloudFront does not support using Route 53 alias records as origins3. Option D is not advisable because using a CloudFront function to redirect the requests to the secondary ALB will add an extra round-trip and latency to the failover process, which may also not meet the zero-second RTO requirement.
1: Optimizing high availability with CloudFront origin failover - Amazon CloudFront
2: Creating an origin group - Amazon CloudFront
3: Values That You Specify When You Create or Update a Web Distribution - Amazon CloudFront

QUESTION 96
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? (Select 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.

Correct Answer: B, C
Section:
Explanation:
The correct answer is B and C. Option B is correct because inviting the acquired company's AWS accounts to join the organization and creating the OrganizationAccountAccessRole IAM role in the invited accounts allows the management account to assume the role and gain full administrative access to the member accounts. Option C is correct because using AWS Security Hub to collect and group findings across all accounts enables the DevOps team to monitor and improve the security posture of the organization. Security Hub can automatically detect new accounts as the accounts are added to the organization and enable Security Hub for them. Option A is incorrect because creating an SCP that has full administrative privileges and attaching it to the management account does not grant the management account access to the member accounts. SCPs are used to restrict the permissions of the member accounts, not to grant permissions to the management account. Option D is incorrect because using AWS Firewall Manager to collect and group findings across all accounts is not a valid use case for Firewall Manager. Firewall Manager is used to centrally configure and manage firewall rules across the organization, not to collect and group security findings. Option E is incorrect because using Amazon Inspector to collect and group findings across all accounts is not a valid use case for Amazon Inspector. Amazon Inspector is used to assess the security and compliance of applications running on Amazon EC2 instances, not to collect and group security findings across accounts.Reference:
Inviting an AWS account to join your organization
Enabling and disabling AWS Security Hub
Service control policies
AWS Firewall Manager
Amazon Inspector

QUESTION 97
A DevOps engineer is building an application that uses an AWS Lambda function to query an Amazon Aurora MySQL DB cluster. The Lambda function performs only read queries. Amazon EventBridge events invoke the Lambda function.
As more events invoke the Lambda function each second, the database's latency increases and the database's throughput decreases. The DevOps engineer needs to improve the performance of the application.
Which combination of steps will meet these requirements? (Select THREE.)

A. Use Amazon RDS Proxy to create a proxy. Connect the proxy to the Aurora cluster reader endpoint. Set a maximum connections percentage on the proxy.
B. Implement database connection pooling inside the Lambda code. Set a maximum number of connections on the database connection pool.
C. Implement the database connection opening outside the Lambda event handler code.
D. Implement the database connection opening and closing inside the Lambda event handler code.
E. Connect to the proxy endpoint from the Lambda function.
F. Connect to the Aurora cluster endpoint from the Lambda function.

Correct Answer: A, C, E
Section:
Explanation:
Short To improve the performance of the application, the DevOps engineer should use Amazon RDS Proxy, implement the database connection opening outside the Lambda event handler code, and connect to the proxy endpoint from the Lambda function.
Amazon RDS Proxy is a fully managed, highly available database proxy for Amazon Relational Database Service (RDS) that makes applications more scalable, more resilient to database failures, and more secure1. By using Amazon RDS Proxy, the DevOps engineer can reduce the overhead of opening and closing connections to the database, which can improve latency and throughput2.
The DevOps engineer should connect the proxy to the Aurora cluster reader endpoint, which allows read-only connections to one of the Aurora Replicas in the DB cluster3. This can help balance the load across multiple read replicas and improve performance for read-intensive workloads4.
The DevOps engineer should implement the database connection opening outside the Lambda event handler code, which means using a global variable to store the database connection object5. This can enable connection reuse across multiple invocations of the Lambda function, which can reduce latency and improve performance.
The DevOps engineer should connect to the proxy endpoint from the Lambda function, which is a unique URL that represents the proxy. This can allow the Lambda function to access the database through the proxy, which can provide benefits such as connection pooling, load balancing, failover handling, and enhanced security.
The other options are incorrect because:
Implementing database connection pooling inside the Lambda code is unnecessary and redundant when using Amazon RDS Proxy, which already provides connection pooling as a service.
Implementing the database connection opening and closing inside the Lambda event handler code is inefficient and costly, as it can increase latency and consume more resources for each invocation of the Lambda function.
Connecting to the Aurora cluster endpoint from the Lambda function is not optimal for read-only queries, as it can direct traffic to either the primary instance or one of the Aurora Replicas in the DB cluster. This can result in inconsistent performance and potential conflicts with write operations on the primary instance.

QUESTION 98
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? (Select 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": "Allow",
    "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"
        }
    }
}
```


F. Create an IAM permissions boundary in each account. Include the following statement:

```
{
    "Effect": "Allow",
    "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"
        }
    }
}
```

Correct Answer: A, D, E
Section:
Explanation:
Short To meet the requirements, the DevOps engineer should 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, create an IAM permissions boundary in each account, and for each CodeCommit repository, add an access-team tag that has the value set to the name of the associated team.
Updating the SAML assertion to pass the user's team name allows the DevOps engineer to use IAM tags to identify which team a user belongs to. This can help enforce fine-grained access control based on the user's team membership1.
Updating the IAM role's trust policy to add an access-team session tag that has the team name allows the DevOps engineer to use IAM condition keys to restrict access based on the session tag value2. For example, the DevOps engineer can use the aws:PrincipalTag condition key to match the access-team tag of the user with the access-team tag of the repository3.
Creating an IAM permissions boundary in each account allows the DevOps engineer to set the maximum permissions that an identity-based policy can grant to an IAM entity. An entity's permissions boundary allows it to perform only the actions that are allowed by both its identity-based policies and its permissions boundaries4. For example, the DevOps engineer can use a permissions boundary policy to limit the actions that a user can perform on CodeCommit repositories based on their access-team tag5.
For each CodeCommit repository, adding an access-team tag that has the value set to the name of the associated team allows the DevOps engineer to use resource tags to identify which team manages a repository. This can help enforce fine-grained access control based on the resource tag value6.
The other options are incorrect because:
Creating an approval rule template for each team in the Organizations management account is not a valid option, as approval rule templates are not supported by AWS Organizations. Approval rule templates are specific to CodeCommit and can only be associated with one or more repositories in the same AWS Region where they are created7.
Creating an approval rule template for each account is not a valid option, as approval rule templates are not designed to restrict access to modify branches. Approval rule templates are designed to require approvals from specified users or groups before merging pull requests8.
Attaching an SCP to the accounts is not a valid option, as SCPs are not designed to restrict access based on tags. SCPs are designed to restrict access based on service actions and resources across all users and roles in an organization's account9.

QUESTION 99
A company needs a strategy for failover and disaster recovery of its data and application. The application uses a MySQL database and Amazon EC2 instances. The company requires a maximum RPO of 2 hours and a maximum RTO of 10 minutes for its data and application at all times.

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

A. Create an Amazon Aurora Single-AZ cluster in multiple AWS 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 AWS Regions as the data store. In the event of a failure, promote the secondary Region to the primary for the application. Update the application to use the Aurora cluster endpoint in the secondary Region.
C. Create an Amazon Aurora cluster in multiple AWS 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 AWS Regions. Use Amazon Route 53 failover routing that points to Application Load Balancers in both Regions. Use health checks and Auto Scaling groups in each Region.
E. Set up the application in two AWS Regions. Configure AWS Global Accelerator to point to Application Load Balancers (ALBs) in both Regions. Add both ALBs to a single endpoint group. Use health checks and Auto Scaling groups in each Region.

Correct Answer: B, E
Section:
Explanation:
Short To meet the requirements of failover and disaster recovery, the company should use the following deployment strategies:
Create an Amazon Aurora global database in two AWS Regions as the data store. In the event of a failure, promote the secondary Region to the primary for the application. Update the application to use the Aurora cluster endpoint in the secondary Region. This strategy can provide a low RPO and RTO for the data, as Aurora global database replicates data with minimal latency across Regions and allows fast and easy failover12. The company can use the Amazon Aurora cluster endpoint to connect to the current primary DB cluster without needing to change any application code1.
Set up the application in two AWS Regions. Configure AWS Global Accelerator to point to Application Load Balancers (ALBs) in both Regions. Add both ALBs to a single endpoint group. Use health checks and Auto Scaling groups in each Region. This strategy can provide high availability and performance for the application, as AWS Global Accelerator uses the AWS global network to route traffic to the closest healthy endpoint3. The company can also use static IP addresses that are assigned by Global Accelerator as a fixed entry point for their application1. By using health checks and Auto Scaling groups, the company can ensure that their application can scale up or down based on demand and handle any instance failures4.
The other options are incorrect because:
Creating an Amazon Aurora Single-AZ cluster in multiple AWS Regions as the data store would not provide a fast failover or disaster recovery solution, as the company would need to manually restore data from backups or snapshots in another Region in case of a failure.
Creating an Amazon Aurora cluster in multiple AWS Regions as the data store and using a Network Load Balancer to balance the database traffic in different Regions would not work, as Network Load Balancers do not support cross-Region routing. Moreover, this strategy would not provide a consistent view of the data across Regions, as Aurora clusters do not replicate data automatically between Regions unless they are part of a global database.
Setting up the application in two AWS Regions and using Amazon Route 53 failover routing that points to Application Load Balancers in both Regions would not provide a low RTO, as Route 53 failover routing relies on DNS resolution, which can take time to propagate changes across different DNS servers and clients. Moreover, this strategy would not provide deterministic routing, as Route 53 failover routing depends on DNS caching behavior, which can vary depending on different factors.

QUESTION 100
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.

Correct Answer: B
Section:
Explanation:
To meet the requirements, the DevOps engineer needs to create a solution that joins all EC2 instances to the domain of the AWS Managed Microsoft AD directory with the most operational efficiency. The DevOps engineer can use AWS Systems Manager Automation to automate the domain join process using an existing runbook called AWS-JoinDirectoryServiceDomain. This runbook can join Windows instances to an AWS Managed Microsoft AD or Simple AD directory by using PowerShell commands. The DevOps engineer can create an AWS::SSM::Association resource in the CloudFormation template to associate the runbook with the EC2 instances that have specific tags. The tags can be defined in the launch template and propagated on launch to the EC2 instances. The DevOps engineer can also define the required parameters for the runbook, such as the directory ID, directory name, and organizational unit. The DevOps engineer can attach the AmazonSSMManagedInstanceCore and AmazonSSMDirectoryServiceAccess AWS managed policies to the IAM role that the EC2 instances use. These policies grant the necessary permissions for Systems Manager and Directory Service operations.

QUESTION 101
A company has multiple development groups working in a single shared AWS account. The Senior Manager of the groups wants to be alerted via a third-party API call when the creation of resources approaches the service limits for the account.
Which solution will accomplish this with the LEAST amount of development effort?

A. Create an Amazon CloudWatch Event rule that runs periodically and targets an AWS Lambda function. Within the Lambda function, evaluate the current state of the AWS environment and compare deployed resource values to resource limits on the account. 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 CloudWatch Events rule to run the Lambda function periodically. Create another CloudWatch Events 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 Personal Health Dashboard checks, and configure an Amazon CloudWatch Events rule to run the Lambda function periodically. Create another CloudWatch Events rule with an event pattern matching Personal 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 SNS topic. Deploy an AWS Lambda function that notifies the Senior Manager, and subscribe the Lambda function to the SNS topic.

Correct Answer: B
Section:
Explanation:
To meet the requirements, the company needs to create a solution that alerts the Senior Manager when the creation of resources approaches the service limits for the account with the least amount of development effort. The company can use AWS Trusted Advisor, which is a service that provides best practice recommendations for cost optimization, performance, security, and service limits. The company can deploy an AWS Lambda function that refreshes Trusted Advisor checks, and configure an Amazon CloudWatch Events rule to run the Lambda function periodically. This will ensure that Trusted Advisor checks are up to date and reflect the current state of the account. The company can then create another CloudWatch Events rule with an event pattern matching Trusted Advisor events and a target Lambda function. The event pattern can filter for events related to service limit checks and their status. The target Lambda function can notify the Senior Manager via a third-party API call if the event indicates that the account is approaching or exceeding a service limit.

QUESTION 102
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? (Select 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.

Correct Answer: A, D
Section:
Explanation:
To optimize the Macie costs for the account without compromising the account's functionality, the DevOps engineer needs to exclude S3 buckets that do not contain sensitive data from automated discovery. S3 buckets that contain CloudTrail logs are unlikely to have sensitive data, and Macie charges for scanning and monitoring data in S3 buckets. Therefore, excluding S3 buckets that contain CloudTrail logs from automated discovery can reduce Macie costs. Similarly, configuring discovery jobs to include S3 objects based on the last modified criterion can also reduce Macie costs, as it will only scan and monitor new or updated objects, rather than all objects in the bucket.

QUESTION 103
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.

Correct Answer: B
Section:
Explanation:
To implement a control that requires the use of IMDSv2 on all EC2 instances in the account, the DevOps engineer can use a permissions boundary. A permissions boundary is a policy that defines the maximum permissions that an IAM entity can have. The DevOps engineer can create a permissions boundary that prevents the ec2:RunInstance action if the ec2:MetadataHttpTokens condition key is not set to a value of required. This condition key enforces the use of IMDSv2 on EC2 instances. The DevOps engineer can attach the permissions boundary to the IAM role that was used to launch the instance. This way, any attempt to launch an EC2 instance without using IMDSv2 will be denied by the permissions boundary.

QUESTION 104
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? (Select 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

Correct Answer: A, B
Section:
Explanation:
To implement governance controls that restrict AWS service usage to within the United States and ensure alerts for any activity outside the governance policy, the following actions will meet the requirements:
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.This action will effectively prevent users and roles in all accounts within the organization from accessing services in non-US Regions12.
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.This action will allow monitoring of all AWS Regions and will trigger alerts if any activity is detected in non-US Regions, ensuring that the governance team is notified as soon as possible3.

AWS Documentation on Service Control Policies (SCPs) and how they can be used to manage permissions and restrict access based on Regions12.
AWS Documentation on monitoring CloudTrail log files with Amazon CloudWatch Logs to set up alerts for specific activities3.

QUESTION 105
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.

Correct Answer: C
Section:
Explanation:
The best solution is to use AWS Config and AWS Security Hub to identify and remediate noncompliant resources across multiple AWS accounts. AWS Config enables continuous monitoring of the configuration of AWS resources and evaluates them against desired configurations. AWS Config can also automatically remediate noncompliant resources by using conformance packs, which are a collection of AWS Config rules and remediation actions that can be deployed as a single entity. AWS Security Hub provides a comprehensive view of the security posture of AWS accounts and resources. AWS Security Hub can aggregate and normalize the findings from AWS Config and other AWS services, as well as from partner solutions. AWS Security Hub can also be used to create a dashboard for tracking noncompliant resources and events in a centralized location.
The other options are not optimal because they either require more development overhead, do not provide near real time detection and remediation, or do not provide a centralized dashboard for tracking.
Option A is not optimal because CloudFormation drift detection is not a near real time solution. Drift detection has to be manually initiated on each stack or resource, or scheduled using a cron expression. Drift detection also does not provide remediation actions, so a custom Lambda function has to be developed and invoked. CloudWatch Logs and dashboard can be used for tracking, but they do not provide a comprehensive view of the security posture of the AWS accounts and resources.
Option B is not optimal because CloudTrail logs analysis using Athena is not a near real time solution. Athena queries have to be manually run or scheduled using a cron expression. Athena also does not provide remediation actions, so a custom Lambda function has to be developed and invoked. Step Functions can be used to orchestrate the query and remediation workflow, but it adds more complexity and cost. QuickSight dashboard can be used for tracking, but it does not provide a comprehensive view of the security posture of the AWS accounts and resources.
Option D is not optimal because CloudTrail logs analysis using CloudWatch Logs is not a near real time solution. CloudWatch Logs filters have to be manually created or updated for each resource type and configuration change. CloudWatch Logs also does not provide remediation actions, so a custom Lambda function has to be developed and invoked. EventBridge can be used to trigger the Lambda function, but it adds more complexity and cost. OpenSearch Service dashboard can be used for tracking, but it does not provide a comprehensive view of the security posture of the AWS accounts and resources.
AWS Config conformance packs
Introducing AWS Config conformance packs
Managing conformance packs across all accounts in your organization

QUESTION 106
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 CodePipelme 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.

Correct Answer: C
Section:
Explanation:
The correct approach to meet the requirements for separate environment pipelines and automatic deployment for test environments is to create two AWS CodePipeline configurations, one for each environment. The production pipeline should have a manual approval step to ensure that changes are reviewed before being deployed to production. A single AWS CodeCommit repository with separate branches for each environment allows for organized and efficient code management. Each CodePipeline retrieves the source code from the appropriate branch in the repository. The deployment step utilizes AWS CloudFormation to deploy the Lambda functions, ensuring that the infrastructure as code is maintained and version-controlled.
AWS Lambda with Amazon API Gateway:Using AWS Lambda with Amazon API Gateway
Tutorial on using Lambda with API Gateway:Tutorial: Using Lambda with API Gateway
AWS CodePipeline automatic deployment:Set Up a Continuous Deployment Pipeline Using AWS CodePipeline
Building a pipeline for test and production stacks:Walkthrough: Building a pipeline for test and production stacks

QUESTION 107
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 Cloud Formation 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.

Correct Answer: D
Section:
Explanation:
The correct answer is D. Creating an AWS Config conformance pack that contains the AWS Config rules and remediation actions and deploying it from the delegated administrator account by using AWS Config will meet the requirements.A conformance pack is a collection of AWS Config rules and remediation actions that can be easily deployed as a single entity in an account and a region or across an organization in AWS Organizations1. By using the delegated administrator account, the DevOps engineer can centrally manage the conformance pack and prevent non-administrator users from modifying it in the member accounts. Option A is incorrect because creating a CloudFormation template that contains the AWS Config rules and remediation actions and deploying it from the Organizations management account by using CloudFormation StackSets will not prevent non-administrator users from modifying the AWS Config rules in the member accounts. Option B is incorrect because deploying the conformance pack from the Organizations management account by using CloudFormation StackSets will not use the trusted access feature of AWS Config and will require additional permissions and resources. Option C is incorrect because creating a CloudFormation template that contains the AWS Config rules and remediation actions and deploying it from the delegated administrator account by using AWS Config will not leverage the benefits of conformance packs, such as simplified deployment and management.Reference:
Conformance Packs - AWS Config
Certified DevOps Engineer - Professional (DOP-C02) Study Guide(page 176)

QUESTION 108
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 to Chef agents in the golden image, is 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

Correct Answer: B
Section:
Explanation:
Amazon EC2 Image Builder is a service that automates the creation, management, and deployment of customized, secure, and up-to-date server images that are pre-installed with software and configuration settings tailored to meet specific IT standards. EC2 Image Builder simplifies the creation and maintenance of golden images, and makes it easy to generate images for multiple platforms, such as Amazon EC2 and on-premises. EC2 Image Builder also integrates with AWS Resource Access Manager, which allows you to share your images across accounts within your organization or with external AWS accounts. This solution meets the requirements of automating the tasks of updating the Linux AMIs, installing the Chef agent, and providing the images to the department's accounts with the least management overhead.Reference:
Amazon EC2 Image Builder
Sharing EC2 Image Builder images

QUESTION 109
A company has an AWS CodeDeploy application. The application has a deployment group that uses a single tag group to identify instances for the deployment of Application A. The single tag group configuration identifies instances that have Environment=Production and Name=ApplicattonA tags for the deployment of ApplicationA. The company launches an additional Amazon EC2 instance with Department=Marketing Environment^Production. and Name=ApplicationB tags. On the next CodeDeploy deployment of ApplicationA. the additional instance has ApplicationA installed on it. A DevOps engineer needs to configure the existing deployment group to prevent ApplicationA from being installed on the additional instance Which solution will meet these requirements?

A. Change the current single tag group to include only the Environment=Production tag Add another single tag group that includes only the Name=ApplicationA tag.
B. Change the current single tag group to include the Department=Marketmg Environment=Production and Name=ApplicationAtags
C. Add another single tag group that includes only the Department=Marketing tag. Keep the Environment=Production and Name=ApplicationA tags with the current single tag group
D. Change the current single tag group to include only the Environment=Production tag Add another single tag group that includes only the Department=Marketing tag

Correct Answer: A
Section:
Explanation:
To prevent ApplicationA from being installed on the additional instance, the deployment group configuration needs to be more specific. By changing the current single tag group to include only theEnvironment=Productiontag and adding another single tag group that includes only theName=ApplicationAtag, the deployment process will target only the instances that match both tag groups. This ensures that only instances intended for ApplicationA with the correct environment and name tags will receive the deployment, thus excluding the additional instance with theDepartment=MarketingandName=ApplicationBtags.
AWS CodeDeploy Documentation:Working with instances for CodeDeploy
AWS CodeDeploy Documentation:Stop a deployment with CodeDeploy
Stack Overflow Discussion:CodeDeploy Deployment failed to stop Application

QUESTION 110
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 secunty 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.

Correct Answer: B
Section:
Explanation:
To prevent developers from unintentionally attaching an Elastic IP address to an Amazon EC2 instance in production, the best approach is to use IAM policies and AWS Config rules. By attaching an IAM policy that denies theassociate-addresspermission to the developers' IAM group, you ensure that developers cannot perform this action. Additionally, creating a custom AWS Config rule to check for Elastic IP addresses associated with instances tagged as production provides ongoing monitoring. If the rule detects an Elastic IP address, it can trigger an alert to notify the security team.This method is proactive and enforces the necessary permissions while also providing a mechanism for detection and notification.Reference: from Amazon DevOps sources

QUESTION 111
A company is reviewing its IAM policies. One policy written by the DevOps engineer has been (lagged as too permissive. The policy is used by an AWS Lambda function that issues a stop command to Amazon EC2 instances tagged with Environment: NonProduccion over the weekend. The current policy is:

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": "ec2:*",
      "Resource": "*"
    }
  ]
}
```

What changes should the engineer make to achieve a policy ot least permission? (Select THREE.)

A.

A. Add the following conditional expression:

```
"Condition": {
  "StringEquals": {
    "aws:principaltype": "lambda.amazonaws.com"
  }
}
```

B.

Change `"Resource": "*"` to `"Resource": "arn:aws:ec2:*:*:instance/*"`

C.

Add the following conditional expression:

```
"Condition": {
  "StringNotEquals": {
    "ec2:ResourceTag/Environment": "Production"
  }
}
```

D.

Add the following conditional expression:

```
"Condition": {
  "StringEquals": {
    "ec2:ResourceTag/Environment": "NonProduction"
  }
}
```

E.

Change `"Action": "ec2:*"` to `"Action": "ec2:StopInstances"`

F.

Add the following conditional expression:

```
"Condition" :  {
  "DateGreaterThan" : {
    "aws:CurrentTime" : "$ ;{aws:DateTime:Friday}"
  },
  "DateLessThan": {
    "aws:CurrentTime" : "$ ;{aws:DateTime:Monday}"
  }
}
```


Correct Answer: A, B, D
Section:
Explanation:
The engineer should make the following changes to achieve a policy of least permission:
A:Add a condition to ensure that the principal making the request is an AWS Lambda function. This ensures that only Lambda functions can execute this policy.
B:Narrow down the resources by specifying the ARN of EC2 instances instead of allowing all resources. This ensures that the policy only affects EC2 instances.
D:Add a condition to ensure that this policy only applies to EC2 instances tagged with ''Environment: NonProduction''. This ensures that production environments are not affected by this policy.
AWS Identity and Access Management (IAM) - AWS Documentation
Certified DevOps Engineer - Professional (DOP-C02) Study Guide(page 179)

QUESTION 112
A company has a mission-critical application on AWS that uses automatic scaling The company wants the deployment lilecycle 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 CloudFormalion to create an AWS Step Functions state machine and Auto Scaling hfecycle 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

Correct Answer: B
Section:
Explanation:
https://aws.amazon.com/about-aws/whats-new/2016/09/aws-codedeploy-introduces-deployment-monitoring-with-amazon-cloudwatch-alarms-and-automatic-deployment-rollback/

QUESTION 113
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.

Correct Answer: B
Section:
Explanation:
The correct answer is B. Creating a new AWS CodeBuild project and configuring a test stage in the AWS CodePipeline pipeline that uses the new CodeBuild project is the best way to integrate the unit tests into the existing pipeline. Creating a CodeBuild report group and uploading the test reports to the new CodeBuild report group will produce reports about the unit tests for the company to view. Using JUNITXML as the output format for the unit tests is supported by CodeBuild and will generate a valid report.
Option A is incorrect because Amazon CodeGuru Reviewer is a service that provides automated code reviews and recommendations for improving code quality and performance. It is not a tool for running unit tests or producing test reports. Therefore, option A will not meet the requirements.
Option C is incorrect because AWS CodeArtifact is a service that provides secure, scalable, and cost-effective artifact management for software development. It is not a tool for running unit tests or producing test reports. Moreover, option C uses CUCUMBERJSON as the output format for the unit tests, which is not supported by CodeBuild and will not generate a valid report.
Option D is incorrect because uploading the test reports to an Amazon S3 bucket is not the best way to produce reports about the unit tests for the company to view. CodeBuild has a built-in feature to create and manage test reports, which is more convenient and efficient than using S3. Furthermore, option D uses HTML as the output format for the unit tests, which is not supported by CodeBuild and will not generate a valid report.

QUESTION 114
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 AdministratorAccess 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.

Correct Answer: C
Section:
Explanation:
The correct answer is C. Updating the SCP in the child OU to allow all actions for Amazon EC2 will enable the DevOps engineer to launch the EC2 instance in the vendor-data account. SCPs are applied to OUs and accounts in a hierarchical manner, meaning that the SCPs attached to the parent OU are inherited by the child OU and accounts. Therefore, the SCP in the child OU overrides the SCP in the root OU and denies all actions except for DynamoDB and Lambda. By adding EC2 to the allowed actions in the child OU's SCP, the DevOps engineer can access EC2 resources in the vendor-data account.
Option A is incorrect because attaching the AmazonEC2FullAccess IAM policy to the IAM user will not grant the user access to EC2 resources. IAM policies are evaluated after SCPs, so even if the IAM policy allows EC2 actions, the SCP will still deny them.
Option B is incorrect because creating a new SCP that allows all actions for EC2 and attaching it to the vendor-data account will not work. SCPs are not cumulative, meaning that only one SCP is applied to an account at a time. The SCP attached to the account will be the SCP attached to the OU that contains the account. Therefore, option B will not change the SCP that is applied to the vendor-data account.
Option D is incorrect because creating a new SCP that allows all actions for EC2 and attaching it to the root OU will not work. As explained earlier, the SCP in the child OU overrides the SCP in the root OU and denies all actions except for DynamoDB and Lambda. Therefore, option D will not affect the SCP that is applied to the vendor-data account.

QUESTION 115
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.

Correct Answer: A
Section:
Explanation:
The correct solution is to create a new assertion safety rule in Route 53 ARC and apply it to the two routing controls. An assertion safety rule is a type of safety rule that ensures that a minimum number of routing controls are always enabled. The ATLEAST type of assertion safety rule specifies the minimum number of routing controls that must be enabled for the rule to evaluate as healthy. By setting the threshold to 1, the rule ensures that at least one routing control is always turned on. This prevents the scenario where both routing controls are accidentally turned off and the application becomes unavailable in both Regions.
The other solutions are incorrect because they do not use safety rules to prevent both routing controls from being turned off. A gating safety rule is a type of safety rule that prevents routing control state changes that violate the rule logic. The OR type of gating safety rule specifies that one or more routing controls must be enabled for the rule to evaluate as healthy. However, this rule does not prevent a user from turning off both routing controls manually. A resource set is a collection of resources that are tested for readiness by Route 53 ARC. A readiness check is a test that verifies that all the resources in a resource set are operational. However, these concepts are not related to routing control states or safety rules. Therefore, creating a new resource set and a new readiness check will not ensure that at least one routing control is turned on at all times.Reference:
Routing control in Amazon Route 53 Application Recovery Controller
Viewing and updating routing control states in Route 53 ARC
Creating a control panel in Route 53 ARC
Creating safety rules in Route 53 ARC

QUESTION 116
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 DynamoDB. Use 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 a 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.

Correct Answer: C
Section:
Explanation:
The correct answer is C. Using AWS Config to identify and audit all EC2 instances based on their host placement configuration is the most efficient and scalable solution to ensure compliance with the software licensing requirement. AWS Config is a service that enables you to assess, audit, and evaluate the configurations of your AWS resources. By creating a custom AWS Config rule that triggers a Lambda function to verify host placement, the DevOps engineer can automate the process of checking whether the instances are running on EC2 Dedicated Hosts or not. The Lambda function can return a NON_COMPLIANT result if the instance is not running on an EC2 Dedicated Host, and the AWS Config report can provide a summary of the compliance status of the instances. This solution requires the least administrative overhead compared to the other options.
Option A is incorrect because using AWS Systems Manager Configuration Compliance to scan and build a database of noncompliant EC2 instances based on their host placement configuration is a more complex and costly solution than using AWS Config. AWS Systems Manager Configuration Compliance is a feature of AWS Systems Manager that enables you to scan your managed instances for patch compliance and configuration inconsistencies. To use this feature, the DevOps engineer would need to install the Systems Manager Agent on each EC2 instance, create a State Manager association to run the put-compliance-items API action periodically, and use a DynamoDB table to store the instance IDs of noncompliant resources. This solution would also require more API calls and storage costs than using AWS Config.
Option B is incorrect because using custom Java code running on an EC2 instance to check and terminate noncompliant EC2 instances is a more cumbersome and error-prone solution than using AWS Config. This solution would require the DevOps engineer to write and maintain the Java code, set up EC2 Auto Scaling for the instance, use an SQS queue and another worker instance to process the instance IDs, use a Lambda function and an SNS topic to terminate and notify the noncompliant instances, and handle any potential failures or exceptions in the workflow. This solution would also incur more compute, storage, and messaging costs than using AWS Config.
Option D is incorrect because using AWS CloudTrail to identify and audit EC2 instances by analyzing the EC2 RunCommand API action is a less reliable and accurate solution than using AWS Config. AWS CloudTrail is a service that enables you to monitor and log the API activity in your AWS account. The EC2 RunCommand API action is used to execute commands on one or more EC2 instances. However, this API action does not necessarily indicate the host placement of the instance, and it may not capture all the instances that are running on EC2 Dedicated Hosts or not. Therefore, option D would not provide a comprehensive and consistent audit of the EC2 instances.

QUESTION 117
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.

Correct Answer: A
Section:
Explanation:

The best solution to meet the disaster recovery capability and allow developers to switch over to a secondary AWS Region for code development is option A. This involves creating aCodeCommit repositoryin the secondary Region and setting up anAWS CodeBuild projectto perform a Git mirror operation of the primary Region's CodeCommit repository to the secondary Region's repository. AnAWS Lambda functionis then created to invoke the CodeBuild project.Additionally, anAmazon EventBridge ruleis configured to react to merge events in the primary Region's CodeCommit repository and invoke the Lambda function12.This setup ensures that the secondary Region's repository is always up-to-date with the primary repository, allowing for a seamless transition in case of a disaster recovery event1.
AWS CodeCommit User Guide on resilience and disaster recovery1.
AWS Documentation on monitoring CodeCommit events in Amazon EventBridge and Amazon CloudWatch Events2.

QUESTION 118
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? (Select 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 O.
D. Increase the Auto Scaling group's desired capacity by I.
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.

Correct Answer: A, B, F
Section:
Explanation:
To restore the functionality of the Auto Scaling group after the AMI was deleted, the DevOps engineer needs to create a new AMI and update the Auto Scaling group to use it. The DevOps engineer can create a new AMI by copying the most recent public AMI of the operating system that the EC2 instances use. This will ensure that the new AMI has the same operating system as the custom AMI that was deleted. The DevOps engineer can then create a new launch template that uses the new AMI and update the Auto Scaling group to use the new launch template. This will allow the Auto Scaling group to launch new instances with the new AMI.

QUESTION 119
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.

Correct Answer: B
Section:
Explanation:
To meet the requirements, the DevOps engineer needs to configure the CloudWatch alarm to stop the EC2 instances when the average of the NetworkPacketsIn metric is less than 5 for at least 3 hours in a 12-hour time window. This means that the alarm should trigger when 3 out of 12 datapoints are below the threshold of 5. The alarm should also treat missing data as not breaching the threshold, so that the EC2 instances continue to run if there is no data for the metric during the evaluation period. The DevOps engineer can add an EC2 action to stop the instance when the alarm enters the ALARM state, which is a built-in action type for CloudWatch alarms.

QUESTION 120

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.

Correct Answer: D
Section:
Explanation:
To meet the requirements, the DevOps team needs to configure a monitoring solution for the VPC flow logs that can detect anomalies in network traffic over time and initiate a response to the anomaly. The DevOps team can use Amazon Kinesis Data Streams to ingest and process streaming data from CloudWatch Logs. The DevOps team can subscribe the log group to a Kinesis data stream, which will deliver log events from CloudWatch Logs to Kinesis Data Streams in near real-time. The DevOps team can then create an AWS Lambda function to detect log anomalies using machine learning or statistical methods. The Lambda function can be set as a processor for the data stream, which means that it will process each record from the stream before sending it to downstream applications or destinations. The Lambda function can also write to the default Amazon EventBridge event bus if it detects an anomaly, which will allow other AWS services or custom applications to respond to the anomaly event.

QUESTION 121
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.

Correct Answer: C
Section:
Explanation:
The correct answer is C, Allowing users to deploy CloudFormation stacks using AWS Service Catalog only and enforcing the use of a launch constraint is the best way to ensure that the internal business teams launch resources through pre-approved CloudFormation templates only. AWS Service Catalog is a service that enables organizations to create and manage catalogs of IT services that are approved for use on AWS. A launch constraint is a rule that specifies the role that AWS Service Catalog assumes when launching a product. By using a launch constraint, the DevOps engineer can control the permissions that the users have when launching a product. Using AWS Config rules to detect when resources have drifted from their expected state is the best way to automate the monitoring of the resources. AWS Config is a service that enables you to assess, audit, and evaluate the configurations of your AWS resources. AWS Config rules are custom or managed rules that AWS Config uses to evaluate whether your AWS resources comply with your desired configurations. By using AWS Config rules, the DevOps engineer can track the changes in the resources and identify any non-compliant resources.
Option A is incorrect because allowing users to deploy CloudFormation stacks using a CloudFormation service role only is not the best way to ensure that the internal business teams launch resources through pre-approved CloudFormation templates only. A CloudFormation service role is an IAM role that CloudFormation assumes to create, update, or delete the stack resources. By using a CloudFormation service role, the DevOps engineer can control the permissions that CloudFormation has when acting on the resources, but not the permissions that the users have when launching a stack. Therefore, option A does not prevent the users from launching resources that are not approved by the company. Using CloudFormation drift detection to detect when resources have drifted from their expected state is a valid way to monitor the resources, but it is not as automated and scalable as using AWS Config rules. CloudFormation drift detection is a feature that enables you to detect whether a stack's actual configuration differs, or has drifted, from its expected configuration. To use this feature, the DevOps engineer would need to manually initiate a drift detection operation on the stack or the stack resources, and then view the drift status and details in the CloudFormation console or API.

Option B is incorrect because allowing users to deploy CloudFormation stacks using a CloudFormation service role only is not the best way to ensure that the internal business teams launch resources through pre-approved CloudFormation templates only, as explained in option A. Using AWS Config rules to detect when resources have drifted from their expected state is a valid way to monitor the resources, as explained in option C,
Option D is incorrect because enforcing the use of a template constraint is not the best way to ensure that the internal business teams launch resources through pre-approved CloudFormation templates only. A template constraint is a rule that defines the values or properties that users can specify when launching a product. By using a template constraint, the DevOps engineer can control the parameters that the users can provide when launching a product, but not the permissions that the users have when launching a product. Therefore, option D does not prevent the users from launching resources that are not approved by the company. Using Amazon EventBridge notifications to detect when resources have drifted from their expected state is a less reliable and consistent solution than using AWS Config rules. Amazon EventBridge is a service that enables you to connect your applications with data from a variety of sources. Amazon EventBridge can deliver a stream of real-time data from event sources, such as AWS services, and route that data to targets, such as AWS Lambda functions. However, to use this solution, the DevOps engineer would need to configure the event source, the event bus, the event rule, and the event target for each resource type that needs to be monitored, which is more complex and error-prone than using AWS Config rules.

QUESTION 122
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: lnstance 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.

Correct Answer: B
Section:
Explanation:
The UserData property of the AWS: AutoScaling: LaunchConfiguration resource can be used to specify a script that runs when the EC2 instances are launched. This script can include the ECS cluster name as an environment variable for the ECS agent running on the EC2 instances. This way, the EC2 instances will register with the correct ECS cluster. Option A is incorrect because the AWS: ECS: Cluster resource does not have a property to reference the EC2 instances. Option C is incorrect because the EC2 instances are launched by the Auto Scaling group, not by the AWS: EC2: Instance resource. Option D is incorrect because using a custom resource and a Lambda function is unnecessary and overly complex for this scenario.Reference:AWS::AutoScaling::LaunchConfiguration,Amazon ECS Container Agent Configuration

QUESTION 123
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? (Select 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.

Correct Answer: B, D, E
Section:
Explanation:
For deploying a Ruby-based application with requirements for interaction with an Amazon RDS for MySQL database, automatic scaling, high availability, and data persistence, the following steps will meet the requirements:
B) Deploy the application on AWS Elastic Beanstalk. Deploy a separate Amazon RDS for MySQL DB instance outside of Elastic Beanstalk.This approach ensures that the database persists independently of the Elastic Beanstalk environment, which can be torn down and recreated without affecting the database123.

E) Use the immutable deployment method to deploy new application versions.Immutable deployments provide a zero-downtime deployment method that ensures that if any part of the deployment process fails, the environment is rolled back to the original state automatically4.
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.This setup allows for automated monitoring and alerting of the application team in case of deployment failures or other health events56.
AWS Elastic Beanstalk documentation on deploying Ruby applications1.
AWS documentation on application auto-scaling7.
AWS documentation on automated deployment strategies with automatic rollbacks and alerts456.

QUESTION 124
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? (Select 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.

Correct Answer: C, E
Section:
Explanation:
To meet the new guideline for application deployment, the company can use a combination of AWS CodePipeline and AWS CloudTrail. A manual approval action in CodePipeline allows the security team to review and approve changes before they are deployed. This action can be configured to pause the pipeline until approval is granted, ensuring that no changes move to production without the necessary sign-off. Additionally, by creating an AWS CloudTrail trail, all actions taken within CodePipeline, including approvals, are recorded and delivered to an Amazon S3 bucket. This provides an audit trail that can be retained for compliance and review purposes.
AWS CodePipeline's manual approval action provides a way to ensure that a member of the security team can review and approve changes before they are deployed1.
AWS CloudTrail integration with CodePipeline allows for the recording and retention of all pipeline actions, including approvals, which can be stored in Amazon S3 for record-keeping2.

QUESTION 125
A company runs a web application that extends across multiple Availability Zones. The company uses an Application Load Balancer (ALB) for routing. AWS Fargate (or the application and Amazon Aurora for the application data The company uses AWS CloudFormation templates to deploy the application The company stores all Docker images in an Amazon Elastic Container Registry (Amazon ECR) repository in the same AWS account and AWS Region.
A DevOps engineer needs to establish a disaster recovery (DR) process in another Region. The solution must meet an RPO of 8 hours and an RTO of 2 hours The company sometimes needs more than 2 hours to build the Docker images from the Dockerfile
Which solution will meet the RTO and RPO requirements MOST cost-effectively?

A. Copy the CloudFormation templates and the Dockerfile to an Amazon S3 bucket in the DR Region Use AWS Backup to configure automated Aurora cross-Region hourly snapshots In case of DR, build the most recent Docker image and upload the Docker image to an ECR repository in the DR Region Use the CloudFormation template that has the most recent Aurora snapshot and the Docker image from the ECR repository to launch a new CloudFormation stack in the DR Region Update the application DNS records to point to the new ALB
B. Copy the CloudFormation templates to an Amazon S3 bucket in the DR Region Configure Aurora automated backup Cross-Region Replication Configure ECR Cross-Region Replication. In case of DR use the CloudFormation template with the most recent Aurora snapshot and the Docker image from the local ECR repository to launch a new CloudFormation stack in the DR Region Update the application DNS records to point to the new ALB
C. Copy the CloudFormation templates to an Amazon S3 bucket in the DR Region. Use Amazon EventBridge to schedule an AWS Lambda function to take an hourly snapshot of the Aurora database and of the most recent Docker image in the ECR repository. Copy the snapshot and the Docker image to the DR Region in case of DR, use the CloudFormation template with the most recent Aurora snapshot and the Docker image from the local ECR repository to launch a new CloudFormation stack in the DR Region
D. Copy the CloudFormation templates to an Amazon S3 bucket in the DR Region. Deploy a second application CloudFormation stack in the DR Region. Reconfigure Aurora to be a global database Update both CloudFormation stacks when a new application release in the current Region is needed. In case of DR. update, the application DNS records to point to the new ALB.

Correct Answer: B

Section:
Explanation:
The most cost-effective solution to meet the RTO and RPO requirements is option B. This option involves copying the CloudFormation templates to an Amazon S3 bucket in the DR Region, configuring Aurora automated backup Cross-Region Replication, and configuring ECR Cross-Region Replication. In the event of a disaster, the CloudFormation template with the most recent Aurora snapshot and the Docker image from the local ECR repository can be used to launch a new CloudFormation stack in the DR Region. This approach avoids the need to build Docker images from the Dockerfile, which can sometimes take more than 2 hours, thus meeting the RTO requirement. Additionally, the use of automated backups and replication ensures that the RPO of 8 hours is met.
AWS Documentation on Disaster Recovery:Plan for Disaster Recovery (DR) - Reliability Pillar
AWS Blog on Establishing RPO and RTO Targets:Establishing RPO and RTO Targets for Cloud Applications
AWS Documentation on ECR Cross-Region Replication: Amazon ECR Cross-Region Replication
AWS Documentation on Aurora Cross-Region Replication: Replicating Amazon Aurora DB Clusters Across AWS Regions

QUESTION 126
A company's application runs on Amazon EC2 instances. The application writes to a log file that records the username, date, time: and source IP address of the login. The log is published to a log group in Amazon CloudWatch Logs
The company is performing a root cause analysis for an event that occurred on the previous day The company needs to know the number of logins for a specific user from the past 7 days
Which solution will provide this information'?

A. Create a CloudWatch Logs metric filter on the log group Use a filter pattern that matches the username. Publish a CloudWatch metric that sums the number of logins over the past 7 days.
B. Create a CloudWatch Logs subscription on the log group Use a filter pattern that matches the username Publish a CloudWatch metric that sums the number of logins over the past 7 days
C. Create a CloudWatch Logs Insights query that uses an aggregation function to count the number of logins for the username over the past 7 days. Run the query against the log group
D. Create a CloudWatch dashboard. Add a number widget that has a filter pattern that counts the number of logins for the username over the past 7 days directly from the log group

Correct Answer: C
Section:
Explanation:
To analyze and find the number of logins for a specific user from the past 7 days, a CloudWatch Logs Insights query is the most suitable solution. CloudWatch Logs Insights enables you to interactively search and analyze your log data in Amazon CloudWatch Logs. You can use the query language to perform queries that contain multiple commands, including aggregation functions, which can count the occurrences of logins for a specific username over a specified time period. This approach is more direct and efficient than creating a metric filter or subscription, which would require additional steps to publish and sum a metric.Reference:AWS Certified DevOps Engineer - Professional,CloudWatch Logs Insights query syntax,Tutorial: Run a query with an aggregation function,Add or remove a number widget from a CloudWatch dashboard.

QUESTION 127
A company runs applications on Windows and Linux Amazon EC2 instances The instances run across multiple Availability Zones In an AWS Region. The company uses Auto Scaling groups for each application.
The company needs a durable storage solution for the instances. The solution must use SMB for Windows and must use NFS for Linux. The solution must also have sub-millisecond latencies. All instances will read and write the data.
Which combination of steps will meet these requirements? (Select THREE.)

A. Create an Amazon Elastic File System (Amazon EFS) file system that has targets in multiple Availability Zones
B. Create an Amazon FSx for NetApp ONTAP Multi-AZ file system.
C. Create a General Purpose SSD (gp3) Amazon Elastic Block Store (Amazon EBS) volume to use for shared storage.
D. Update the user data for each application's launch template to mount the file system
E. Perform an instance refresh on each Auto Scaling group.
F. Update the EC2 instances for each application to mount the file system when new instances are launched

Correct Answer: A, B, D
Section:
Explanation:
* Create an Amazon Elastic File System (Amazon EFS) File System with Targets in Multiple Availability Zones:
Amazon EFS provides a scalable and highly available network file system that supports the NFS protocol. EFS is ideal for Linux instances as it allows multiple instances to read and write data concurrently.

Setting up EFS with targets in multiple Availability Zones ensures high availability and durability.
* Create an Amazon FSx for NetApp ONTAP Multi-AZ File System:
Amazon FSx for NetApp ONTAP offers a fully managed file storage solution that supports both SMB for Windows and NFS for Linux.
The Multi-AZ deployment ensures high availability and durability, providing sub-millisecond latencies suitable for the application's performance requirements.
* Update the User Data for Each Application's Launch Template to Mount the File System:
Updating the user data in the launch template ensures that every new instance launched by the Auto Scaling group will automatically mount the appropriate file system.
This step is necessary to ensure that all instances can access the shared storage without manual intervention.
Example user data for mounting EFS (Linux)
#!/bin/bash
sudo yum install -y amazon-efs-utils
sudo mount -t efs fs-12345678:/ /mnt/efs
Example user data for mounting FSx (Windows):
By implementing these steps, the company can provide a durable storage solution with sub-millisecond latencies that supports both SMB and NFS protocols, meeting the requirements for both Windows and Linux instances.
Mounting EFS File Systems
Mounting Amazon FSx File Systems

QUESTION 128
A company has an application that stores data that includes personally Identifiable Information (Pll) In an Amazon S3 bucket All data Is encrypted with AWS Key Management Service (AWS KMS) customer managed keys. All AWS resources are deployed from an AWS Cloud Formation template.
A DevOps engineer needs to set up a development environment for the application in a different AWS account The data in the development environment's S3 bucket needs to be updated once a week from the production environment's S3 bucket.
The company must not move Pll from the production environment without anonymizmg the Pll first The data in each environment must be encrypted with different KMS customer managed keys.
Which combination of steps should the DevOps engineer take to meet these requirements? (Select TWO)

A. Activate Amazon Macie on the S3 bucket In the production account Create an AWS Step Functions state machine to initiate a discovery job and redact all Pll before copying files to the S3 bucket in the development account. Give the state machine tasks decrypt permissions on the KMS key in the production account. Give the state machine tasks encrypt permissions on the KMS key in the development account
B. Set up S3 replication between the production S3 bucket and the development S3 bucket Activate Amazon Macie on the development S3 bucket Create an AWS Step Functions state machine to initiate a discovery job and redact all Pll as the files are copied to the development S3 bucket. Give the state machine tasks encrypt and decrypt permissions on the KMS key in the development account.
C. Set up an S3 Batch Operations job to copy files from the production S3 bucket to the development S3 bucket. In the development account, configure an AWS Lambda function to redact all Pll. Configure S3 Object Lambda to use the Lambda function for S3 GET requests Give the Lambda function's IAM role encrypt and decrypt permissions on the KMS key in the development account.
D. Create a development environment from the CloudFormatlon template in the development account. Schedule an Amazon EventBridge rule to start the AWS Step Functions state machine once a week
E. Create a development environment from the CloudFormation template in the development account. Schedule a cron job on an Amazon EC2 instance to run once a week to start the S3 Batch Operations job.

Correct Answer: A, D
Section:
Explanation:
Activate Amazon Macie on the Production S3 Bucket:
Macie can identify and protect sensitive data such as PII.
Create a Step Functions state machine to automate data discovery and redaction before copying it to the development environment.
Example Step Functions state machine:
{
'Comment': 'Anonymize PII and copy data',
'StartAt': 'MacieDiscoveryJob',
'States': {
'MacieDiscoveryJob': {
'Type': 'Task',
'Resource': 'arn:aws:states:::macie:startClassificationJob',
'End': true
}

}
}
Create a Development Environment from CloudFormation Template:
Deploy the development environment in a new account using the existing CloudFormation template.
Schedule an EventBridge rule to start the Step Functions state machine on a weekly basis.
EventBridge rule example:
{
'ScheduleExpression': 'rate(7 days)',
'StateMachineArn': 'arn:aws:states:<region>::stateMachine:AnonymizeAndCopyData'
}
By using Macie for data anonymization and Step Functions for automation, you ensure PII is properly handled before data transfer between environments.
Amazon Macie
AWS Step Functions
AWS CloudFormation Templates

QUESTION 129
A DevOps engineer needs to implement integration tests into an existing AWS CodePipelme CI/CD workflow for an Amazon Elastic Container Service (Amazon ECS) service. The CI/CD workflow retrieves new application code from an AWS CodeCommit repository and builds a container image. The CI/CD workflow then uploads the container image to Amazon Elastic Container Registry (Amazon ECR) with a new image tag version.
The integration tests must ensure that new versions of the service endpoint are reachable and that vanous API methods return successful response data The DevOps engineer has already created an ECS cluster to test the service
Which combination of steps will meet these requirements with the LEAST management overhead? (Select THREE.)

A. Add a deploy stage to the pipeline Configure Amazon ECS as the action provider
B. Add a deploy stage to the pipeline Configure AWS CodeDeploy as the action provider
C. Add an appspec.yml file to the CodeCommit repository
D. Update the image build pipeline stage to output an imagedefinitions json file that references the new image tag.
E. Create an AWS Lambda function that runs connectivity checks and API calls against the service. Integrate the Lambda function with CodePipeline by using aLambda action stage
F. Write a script that runs integration tests against the service. Upload the script to an Amazon S3 bucket. Integrate the script in the S3 bucket with CodePipeline by using an S3 action stage.

Correct Answer: A, D, E
Section:
Explanation:
* Add a Deploy Stage to the Pipeline, Configure Amazon ECS as the Action Provider:
By adding a deploy stage to the pipeline and configuring Amazon ECS as the action provider, the pipeline can automatically deploy the new container image to the ECS cluster.
This ensures that the service is updated with the new image tag, making the new version of the service endpoint reachable.
* Update the Image Build Pipeline Stage to Output an imagedefinitions.json File that Reference the New Image Tag:
The imagedefinitions.json file provides the necessary information about the container images and their tags for the ECS task definitions.
Updating the pipeline to output this file ensures that the correct image version is deployed.
Example imagedefinitions.json
[
{
'name': 'container-name',
'imageUri': '123456789012.dkr.ecr.region.amazonaws.com/my-repo:my-tag'
}
]
*
Reference: CodePipeline ECS Deployment
* Create an AWS Lambda Function that Runs Connectivity Checks and API Calls against the Service. Integrate the Lambda Function with CodePipeline by Using a Lambda Action Stage:
The Lambda function can perform the necessary integration tests by making connectivity checks and API calls to the deployed service endpoint.

Integrating this Lambda function into CodePipeline ensures that these tests are run automatically after deployment, providing near-real-time feedback on the new deployment's health.
Example Lambda function integration:
actions:
- name: TestService
actionTypeId:
category: Test
owner: AWS
provider: Lambda
runOrder: 2
configuration:
FunctionName: testServiceFunction
These steps ensure that the CI/CD workflow deploys the new container image to ECS, updates the image references, and performs integration tests, meeting the requirements with minimal management overhead.

QUESTION 130
A company wants to use AWS Systems Manager documents to bootstrap physical laptops for developers The bootstrap code Is stored in GitHub A DevOps engineer has already created a Systems Manager activation, installed the Systems Manager agent with the registration code, and installed an activation ID on all the laptops.
Which set of steps should be taken next?

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

Correct Answer: C
Section:
Explanation:
Configure the Systems Manager Document to Use the aws-downloadContent Plugin with a sourceType of GitHub and sourceInfo with the Repository Details:
The aws-downloadContent plugin can download content from various sources, including GitHub, which is necessary for bootstrapping the laptops with the code stored in the GitHub repository.
schemaVersion: '2.2'
description: 'Download and run bootstrap script from GitHub'
mainSteps:
- action: aws:downloadContent
name: downloadBootstrapScript
inputs:
sourceType: GitHub
sourceInfo: '{'owner':'my-org','repository':'my-repo','path':'scripts/bootstrap.sh','getOptions':'branch:main'}'
destinationPath: /tmp/bootstrap.sh
- action: aws:runShellScript
name: runBootstrapScript
inputs:
runCommand:
- chmod +x /tmp/bootstrap.sh
- /tmp/bootstrap.sh
This setup ensures that the bootstrap code is downloaded from GitHub and executed on the laptops using Systems Manager.
AWS Systems Manager aws-downloadContent Plugin
Running Commands Using Systems Manager

QUESTION 131
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.

Correct Answer: A
Section:
Explanation:
* Ensure that Amazon GuardDuty is Enabled:
Amazon GuardDuty is a threat detection service that continuously monitors for malicious activity and unauthorized behavior.
It can detect port scans and generate findings for these events.
* Create an Amazon CloudWatch Alarm for Detected EC2 and Port Scan Findings:
Configure GuardDuty to monitor for port scans and other threats.
Create a CloudWatch alarm that triggers when GuardDuty detects port scan activities.
* Connect the Alarm to the SNS Topic:
The CloudWatch alarm should be configured to send notifications to the SNS topic subscribed by the security team.
This setup ensures that the security team receives near-real-time notifications when a port scan is detected on the EC2 instances.
Example configuration steps:
Enable GuardDuty and ensure it is monitoring the relevant AWS accounts.
Create a CloudWatch alarm:

```
{
'AlarmName': 'GuardDutyPortScanAlarm',
'MetricName': 'ThreatIntelIndicator',
'Namespace': 'AWS/GuardDuty',
'Statistic': 'Sum',
'Dimensions': [
{
'Name': 'FindingType',
'Value': 'Recon:EC2/Portscan'
}
],
'Period': 300,
'EvaluationPeriods': 1,
'Threshold': 1,
'ComparisonOperator': 'GreaterThanOrEqualToThreshold',
'AlarmActions': ['arn:aws:sns:region:account-id:SecurityAlerts']
}
```

Amazon GuardDuty
Creating CloudWatch Alarms for GuardDuty Findings

QUESTION 132
A company uses Amazon EC2 as its primary compute platform. A DevOps team wants to audit the company's EC2 instances to check whether any prohibited applications have been installed on the EC2 instances. Which solution will meet these requirements with the MOST operational efficiency?

A. Configure AWS Systems Manager on each instance Use AWS Systems Manager Inventory Use Systems Manager resource data sync to synchronize and store findings in an Amazon S3 bucket Create an AWS Lambda function that runs when new objects are added to the S3 bucket. Configure the Lambda function to identify prohibited applications.

B. Configure AWS Systems Manager on each instance Use Systems Manager Inventory Create AWS Config rules that monitor changes from Systems Manager Inventory to identify prohibited applications.
C. Configure AWS Systems Manager on each instance. Use Systems Manager Inventory. Filter a trail in AWS CloudTrail for Systems Manager Inventory events to identify prohibited applications.
D. Designate Amazon CloudWatch Logs as the log destination for all application instances Run an automated script across all instances to create an inventory of installed applications Configure the script to forward the results to CloudWatch Logs Create a CloudWatch alarm that uses filter patterns to search log data to identify prohibited applications.

Correct Answer: A
Section:
Explanation:
* Configure AWS Systems Manager on Each Instance:
AWS Systems Manager provides a unified interface for managing AWS resources. Install the Systems Manager agent on each EC2 instance to enable inventory management and other features.
* Use AWS Systems Manager Inventory:
Systems Manager Inventory collects metadata about your instances and the software installed on them. This data includes information about applications, network configurations, and more.
Enable Systems Manager Inventory on all EC2 instances to gather detailed information about installed applications.
* Use Systems Manager Resource Data Sync to Synchronize and Store Findings in an Amazon S3 Bucket:
Resource Data Sync aggregates inventory data from multiple accounts and regions into a single S3 bucket, making it easier to query and analyze the data.
Configure Resource Data Sync to automatically transfer inventory data to an S3 bucket for centralized storage.
* Create an AWS Lambda Function that Runs When New Objects are Added to the S3 Bucket:
Use an S3 event to trigger a Lambda function whenever new inventory data is added to the S3 bucket.
The Lambda function can parse the inventory data and check for the presence of prohibited applications.
* Configure the Lambda Function to Identify Prohibited Applications:
The Lambda function should be programmed to scan the inventory data for any known prohibited applications and generate alerts or take appropriate actions if such applications are found.
Example Lambda function in Python

```
import json
import boto3
def lambda_handler(event, context):
s3 = boto3.client('s3')
bucket = event['Records'][0]['s3']['bucket']['name']
key = event['Records'][0]['s3']['object']['key']
response = s3.get_object(Bucket=bucket, Key=key)
inventory_data = json.loads(response['Body'].read().decode('utf-8'))
prohibited_apps = ['app1', 'app2']
for instance in inventory_data['Instances']:
for app in instance['Applications']:
if app['Name'] in prohibited_apps:
# Send notification or take action
print(f'Prohibited application found: {app['Name']} on instance {instance['InstanceId']}')
return {'statusCode': 200, 'body': json.dumps('Check completed')}
```

By leveraging AWS Systems Manager Inventory, Resource Data Sync, and Lambda, this solution provides an efficient and automated way to audit EC2 instances for prohibited applications.
AWS Systems Manager Inventory
AWS Systems Manager Resource Data Sync
S3 Event Notifications
AWS Lambda

QUESTION 133
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? {Select 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 I AM 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

Correct Answer: B, E
Section:
Explanation:
You need to understand how SCP inheritance works in AWS. The way it works for Deny policies is different that allow policies.
Allow polices are passing down to children ONLY if they don't have an allow policy.
Deny policies always pass down to children.
That's why there is always an SCP set to the Root to allow everything by default. If you limit this policy, the whole organization will be limited, not matter what other policies are saying for the other OUs. So it's not A. It's not D because it restricts the wrong OU.

QUESTION 134
A company has an application that runs on Amazon EC2 instances behind an Application Load Balancer (ALB) The EC2 Instances are in multiple Availability Zones The application was misconfigured in a single Availability Zone, which caused a partial outage of the application.
A DevOps engineer made changes to ensure that the unhealthy EC2 instances in one Availability Zone do not affect the healthy EC2 instances in the other Availability Zones. The DevOps engineer needs to test the application's failover and shift where the ALB sends traffic During failover. the ALB must avoid sending traffic to the Availability Zone where the failure has occurred.
Which solution will meet these requirements?

A. Turn off cross-zone load balancing on the ALB Use Amazon Route 53 Application Recovery Controller to start a zonal shift away from the Availability Zone
B. Turn off cross-zone load balancing on the ALB's target group Use Amazon Route 53 Application Recovery Controller to start a zonal shift away from the Availability Zone
C. Create an Amazon Route 53 Application Recovery Controller resource set that uses the DNS hostname of the ALB Start a zonal shift for the resource set away from the Availability Zone
D. Create an Amazon Route 53 Application Recovery Controller resource set that uses the ARN of the ALB's target group Create a readiness check that uses the ElbV2TargetGroupsCanServeTraffic rule

Correct Answer: B
Section:
Explanation:
* Turn off cross-zone load balancing on the ALB's target group:
Cross-zone load balancing distributes traffic evenly across all registered targets in all enabled Availability Zones. Turning this off will ensure that each target group only handles requests from its respective Availability Zone.
To disable cross-zone load balancing:
Go to the Amazon EC2 console.
Navigate to Load Balancers and select the ALB.
Choose the Target Groups tab, select the target group, and then select the Group details tab.
Click on Edit and turn off Cross-zone load balancing.
* Use Amazon Route 53 Application Recovery Controller to start a zonal shift away from the Availability Zone:
Amazon Route 53 Application Recovery Controller provides the ability to control traffic flow to ensure high availability and disaster recovery.
By using Route 53 Application Recovery Controller, you can perform a zonal shift to redirect traffic away from the unhealthy Availability Zone.
To start a zonal shift:
Configure Route 53 Application Recovery Controller by creating a cluster and control panel.
Create routing controls to manage traffic shifts between Availability Zones.
Use the routing control to shift traffic away from the affected Availability Zone.
Disabling cross-zone load balancing
Route 53 Application Recovery Controller

QUESTION 135
A software team is using AWS CodePipeline to automate its Java application release pipeline The pipeline consists of a source stage, then a build stage, and then a deploy stage. Each stage contains a single action that has a runOrder value of 1.
The team wants to integrate unit tests into the existing release pipeline. The team needs a solution that deploys only the code changes that pass all unit tests.
Which solution will meet these requirements?

A. Modify the build stage. Add a test action that has a runOrder value of 1. Use AWS CodeDeploy as the action provider to run unit tests.
B. Modify the build stage Add a test action that has a runOrder value of 2 Use AWS CodeBuild as the action provider to run unit tests
C. Modify the deploy stage Add a test action that has a runOrder value of 1 Use AWS CodeDeploy as the action provider to run unit tests
D. Modify the deploy stage Add a test action that has a runOrder value of 2 Use AWS CodeBuild as the action provider to run unit tests

Correct Answer: B
Section:
Explanation:
* Modify the Build Stage to Add a Test Action with a RunOrder Value of 2:
The build stage in AWS CodePipeline can have multiple actions. By adding a test action with a runOrder value of 2, the test action will execute after the initial build action completes.
* Use AWS CodeBuild as the Action Provider to Run Unit Tests:
AWS CodeBuild is a fully managed build service that compiles source code, runs tests, and produces software packages.
Using CodeBuild to run unit tests ensures that the tests are executed in a controlled environment and that only the code changes that pass the unit tests proceed to the deploy stage.
Example configuration in CodePipeline:

```
{
'name': 'BuildStage',
'actions': [
{
'name': 'Build',
'actionTypeId': {
'category': 'Build',
'owner': 'AWS',
'provider': 'CodeBuild',
'version': '1'
},
'runOrder': 1
},
{
'name': 'Test',
'actionTypeId': {
'category': 'Test',
'owner': 'AWS',
'provider': 'CodeBuild',
'version': '1'
},
'runOrder': 2
}
]
}
```

By integrating the unit tests into the build stage and ensuring they run after the build process, the pipeline guarantees that only code changes passing all unit tests are deployed.
AWS CodePipeline
AWS CodeBuild
Using CodeBuild with CodePipeline

QUESTION 136
A company has set up AWS CodeArtifact repositories with public upstream repositories The company's development team consumes open source dependencies from the repositories in the company's internal network.
The company's security team recently discovered a critical vulnerability in the most recent version of a package that the development team consumes. The security team has produced a patched version to fix the vulnerability.
The company needs to prevent the vulnerable version from being downloaded. The company also needs to allow the security team to publish the patched version.
Which combination of steps will meet these requirements? {Select TWO.)

A. Update the status of the affected CodeArtifact package version to unlisted
B. Update the status of the affected CodeArtifact package version to deleted
C. Update the status of the affected CodeArtifact package version to archived.
D. Update the CodeArtifact package origin control settings to allow direct publishing and to block upstream operations
E. Update the CodeArtifact package origin control settings to block direct publishing and to allow upstream operations.

Correct Answer: B, D
Section:
Explanation:
Update the status of the affected CodeArtifact package version to deleted:
Deleting the vulnerable package version prevents it from being available for download by any users or systems, ensuring that the compromised version is not consumed.
Update the CodeArtifact package origin control settings to allow direct publishing and to block upstream operations:
By allowing direct publishing, the security team can publish the patched version of the package directly to the CodeArtifact repository.
Blocking upstream operations prevents the repository from automatically fetching and serving the vulnerable package version from upstream public repositories.
By deleting the vulnerable version and configuring the origin control settings to allow direct publishing and block upstream operations, the company ensures that only the patched version is available and the vulnerable version cannot be downloaded.
Managing Package Versions in CodeArtifact
Package Origin Controls in CodeArtifact

QUESTION 137
A company deploys an application to Amazon EC2 instances. The application runs Amazon Linux 2 and uses AWS CodeDeploy. The application has the following file structure for its code repository:

```
appspec.yml
config/config.txt
application/web
```

The appspec.yml file has the following contents in the files section:

```
files:
   - source: config/config.txt
     destination: /usr/local/src/config.txt
   - source: /
     destination: /var/www/html
```

What will the result be for the deployment of the config.txt file?

A. The config.txt file will be deployed to only /var/www/html/config/config txt
B. The config.txt file will be deployed to /usr/local/src/config.txt and to /var/www/html/config/config txt.
C. The config.txt file will be deployed to only /usr/local/src/config txt
D. The config txt file will be deployed to /usr/local/src/config.txt and to /var/www/html/application/web/config txt

Correct Answer: C
Section:
Explanation:
Deployment of config.txt file based on the appspec.yml:

The appspec.yml file specifies that config/config.txt should be copied to /usr/local/src/config.txt.
The source: / directive in the appspec.yml indicates that the entire directory structure starting from the root of the application source should be copied to the specified destination, which is /var/www/html.
Result of the Deployment:
The config.txt file will be specifically deployed to /usr/local/src/config.txt as per the explicit file mapping.
The entire directory structure including application/web will be copied to /var/www/html, but this does not include config/config.txt since it has a specific destination defined.
Thus, the config.txt file will be deployed only to /usr/local/src/config.txt.
Therefore, the correct answer is:
C . The config.txt file will be deployed to only /usr/local/src/config.txt.
AWS CodeDeploy AppSpec File Reference
AWS CodeDeploy Deployment Process

QUESTION 138
A company gives its employees limited rights to AWS DevOps engineers have the ability to assume an administrator role. For tracking purposes, the security team wants to receive a near-real-time notification when the administrator role is assumed.
How should this be accomplished?

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

Correct Answer: D
Section:
Explanation:
* Create an Amazon EventBridge Rule Using an AWS CloudTrail Event Pattern:
AWS CloudTrail logs API calls made in your account, including actions performed by roles.
Create an EventBridge rule that matches CloudTrail events where the AssumeRole API call is made to assume the administrator role.
* Invoke an AWS Lambda Function:
Configure the EventBridge rule to trigger a Lambda function whenever the rule's conditions are met.
The Lambda function will handle the logic to send a notification.
* Publish a Message to an Amazon SNS Topic:
The Lambda function will publish a message to an SNS topic to notify the security team.
Subscribe the security team's email address to this SNS topic to receive real-time notifications.
Example EventBridge rule pattern:
{
'source': ['aws.cloudtrail'],
'detail-type': ['AWS API Call via CloudTrail'],
'detail': {
'eventSource': ['sts.amazonaws.com'],
'eventName': ['AssumeRole'],
'requestParameters': {
'roleArn': ['arn:aws:iam:::role/AdministratorRole']
}
}
}
Example Lambda function (Node.js) to publish to SNS:
const AWS = require('aws-sdk');
const sns = new AWS.SNS();

```
exports.handler = async (event) => {
const params = {
Message: `Administrator role assumed: ${JSON.stringify(event.detail)}`,
TopicArn: 'arn:aws:sns:<region>::<sns-topic>'
};
await sns.publish(params).promise();
};
```

Creating EventBridge Rules
Using AWS Lambda with Amazon SNS

QUESTION 139
A DevOps learn has created a Custom Lambda rule in AWS Config. The rule monitors Amazon Elastic Container Repository (Amazon ECR) policy statements for ecr:' actions. When a noncompliant repository is detected, Amazon EventBridge uses Amazon Simple Notification Service (Amazon SNS) to route the notification to a security team.
When the custom AWS Config rule is evaluated, the AWS Lambda function fails to run.
Which solution will resolve the issue?

A. Modify the Lambda function's resource policy to grant AWS Config permission to invoke the function.
B. Modify the SNS topic policy to include configuration changes for EventBridge to publish to the SNS topic.
C. Modify the Lambda function's execution role to include configuration changes for custom AWS Config rules.
D. Modify all the ECR repository policies to grant AWS Config access to the necessary ECR API actions.

Correct Answer: A
Section:
Explanation:
Step 1: Understanding Lambda Permissions and AWS Config The custom AWS Config rule evaluates resources and invokes an AWS Lambda function when a compliance check is triggered. For AWS Config to invoke the Lambda function, it requires permission to do so. Issue: The Lambda function fails to execute because AWS Config doesn't have permission to invoke it. Action: Modify the resource-based policy of the Lambda function to grant AWS Config permission to invoke the Lambda function. Why: Without this permission, AWS Config cannot trigger the Lambda function, which is why the evaluation fails.
This corresponds to Option A: Modify the Lambda function's resource policy to grant AWS Config permission to invoke the function.

QUESTION 140
A company's organization in AWS Organizations has a single OU. The company runs Amazon EC2 instances in the OU accounts. The company needs to limit the use of each EC2 instance's credentials to the specific EC2 instance that the credential is assigned to. A DevOps engineer must configure security for the EC2 instances.
Which solution will meet these requirements?

A. Create an SCP that specifies the VPC CIDR block. Configure the SCP to check whether the value of the aws:VpcSourceIp condition key is in the specified block. In the same SCP check, check whether the values of the aws:EC2InstanceSourcePrivateIPv4 and aws:SourceVpc condition keys are the same. Deny access if either condition is false. Apply the SCP to the OU.
B. Create an SCP that checks whether the values of the aws:EC2InstanceSourceVPC and aws:SourceVpc condition keys are the same. Deny access if the values are not the same. In the same SCP check, check whether the values of the aws:EC2InstanceSourcePrivateIPv4 and awsVpcSourceIp condition keys are the same. Deny access if the values are not the same. Apply the SCP to the OU.
C. Create an SCP that includes a list of acceptable VPC values and checks whether the value of the aws:SourceVpc condition key is in the list. In the same SCP check, define a list of acceptable IP address values and check whether the value of the aws:VpcSourceIp condition key is in the list. Deny access if either condition is false. Apply the SCP to each account in the organization.
D. Create an SCP that checks whether the values of the aws:EC2InstanceSourceVPC and aws:VpcSourceIp condition keys are the same. Deny access if the values are not the same. In the same SCP check, check whether the values of the aws:EC2InstanceSourcePrivatoIPv4 and aws:SourceVpc condition keys are the same. Deny access if the values are not the same. Apply the SCP to each account in the organization.

Correct Answer: B
Section:
Explanation:
Step 1: Using Service Control Policies (SCPs) for EC2 Security To limit the use of EC2 instance credentials to the specific EC2 instance they are assigned to, you can create a Service Control Policy (SCP) that verifies specific conditions, such as whether the EC2 instance's source VPC and private IP match expected values. Action: Create an SCP that checks whether the values of the aws:EC2InstanceSourceVPC and aws:SourceVpc condition keys are the same. Deny access if they are not. Why: This ensures that credentials cannot be used outside the designated EC2 instance or VPC.

Step 2: Further Validation with Private IPs The SCP should also verify that the EC2 instance's private IP matches the IP range specified for the VPC. If the instance's private IP does not match, access should be denied. Action: In the same SCP, check whether the values of the aws:EC2InstanceSourcePrivateIP and aws:VpcSourceIP condition keys are the same. Deny access if they are not. Why: This ensures that the credentials are only used within the specific EC2 instance and its associated VPC.
This corresponds to Option B: Create an SCP that checks whether the values of the aws:EC2InstanceSourceVPC and aws:SourceVpc condition keys are the same. Deny access if the values are not the same. In the same SCP check, check whether the values of the aws:EC2InstanceSourcePrivateIP and aws:VpcSourceIP condition keys are the same. Deny access if the values are not the same. Apply the SCP to the OU.

QUESTION 141
A DevOps engineer uses AWS CodeBuild to frequently produce software packages. The CodeBuild project builds large Docker images that the DevOps engineer can use across multiple builds. The DevOps engineer wants to improve build performance and minimize costs. Which solution will meet these requirements?

A. Store the Docker images in an Amazon Elastic Container Registry (Amazon ECR) repository. Implement a local Docker layer cache for CodeBuild.
B. Cache the Docker images in an Amazon S3 bucket that is available across multiple build hosts. Expire the cache by using an S3 Lifecycle policy.
C. Store the Docker images in an Amazon Elastic Container Registry (Amazon ECR) repository. Modify the CodeBuild project runtime configuration to always use the most recent image version.
D. Create custom AMIs that contain the cached Docker images. In the CodeBuild build, launch Amazon EC2 instances from the custom AMIs.

Correct Answer: A
Section:
Explanation:
Step 1: Storing Docker Images in Amazon ECR Docker images can be large, and storing them in a centralized, scalable location can greatly reduce build times. Amazon Elastic Container Registry (ECR) is a fully managed container registry that stores, manages, and deploys Docker container images. Action: Store the Docker images in an ECR repository. Why: Storing Docker images in ECR ensures that Docker images can be reused across multiple builds, improving build performance by avoiding the need to rebuild the images from scratch.
Step 2: Implementing Docker Layer Caching in CodeBuild Docker layer caching is essential for improving performance in continuous integration pipelines. CodeBuild supports local caching of Docker layers, which speeds up builds that reuse Docker images across multiple runs.
Action: Implement Docker layer caching within the CodeBuild project.
Why: This improves performance by allowing frequently used Docker layers to be cached locally, avoiding the need to pull or build the layers every time.
This corresponds to Option A: Store the Docker images in an Amazon Elastic Container Registry (Amazon ECR) repository. Implement a local Docker layer cache for CodeBuild.

QUESTION 142
A company uses an Amazon Aurora PostgreSQL global database that has two secondary AWS Regions. A DevOps engineer has configured the database parameter group to guarantee an RPO of 60 seconds. Write operations on the primary cluster are occasionally blocked because of the RPO setting.
The DevOps engineer needs to reduce the frequency of blocked write operations.
Which solution will meet these requirements?

A. Add an additional secondary cluster to the global database.
B. Enable write forwarding for the global database.
C. Remove one of the secondary clusters from the global database.
D. Configure synchronous replication for the global database.

Correct Answer: C
Section:
Explanation:
Step 1: Reducing Replication Lag in Aurora Global Databases In Amazon Aurora global databases, write operations on the primary cluster can be delayed due to the time it takes to replicate to secondary clusters, especially when there are multiple secondary regions involved. Issue: The write operations are occasionally blocked due to the RPO setting, which guarantees replication within 60 seconds. Action: Remove one of the secondary clusters from the global database. Why: Fewer secondary clusters will reduce the overall replication lag, improving write performance and reducing the frequency of blocked writes.
This corresponds to Option C: Remove one of the secondary clusters from the global database.

QUESTION 143
A company uses AWS WAF to protect its cloud infrastructure. A DevOps engineer needs to give an operations team the ability to analyze log messages from AWS WAR. The operations team needs to be able to create alarms for specific patterns in the log output.

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

A. Create an Amazon CloudWatch Logs log group. Configure the appropriate AWS WAF web ACL to send log messages to the log group. Instruct the operations team to create CloudWatch metric filters.
B. Create an Amazon OpenSearch Service cluster and appropriate indexes. Configure an Amazon Kinesis Data Firehose delivery stream to stream log data to the indexes. Use OpenSearch Dashboards to create filters and widgets.
C. Create an Amazon S3 bucket for the log output. Configure AWS WAF to send log outputs to the S3 bucket. Instruct the operations team to create AWS Lambda functions that detect each desired log message pattern. Configure the Lambda functions to publish to an Amazon Simple Notification Service (Amazon SNS) topic.
D. Create an Amazon S3 bucket for the log output. Configure AWS WAF to send log outputs to the S3 bucket. Use Amazon Athena to create an external table definition that fits the log message pattern. Instruct the operations team to write SOL queries and to create Amazon CloudWatch metric filters for the Athena queries.

Correct Answer: A
Section:
Explanation:
Step 1: Sending AWS WAF Logs to CloudWatch Logs AWS WAF allows you to log requests that are evaluated against your web ACLs. These logs can be sent directly to CloudWatch Logs, which enables real-time monitoring and analysis. Action: Configure the AWS WAF web ACL to send log messages to a CloudWatch Logs log group. Why: This allows the operations team to view the logs in real time and analyze patterns using CloudWatch metric filters.
Step 2: Creating CloudWatch Metric Filters CloudWatch metric filters can be used to search for specific patterns in log data. The operations team can create filters for certain log patterns and set up alarms based on these filters.
Action: Instruct the operations team to create CloudWatch metric filters to detect patterns in the WAF log output.
Why: Metric filters allow the team to trigger alarms based on specific patterns without needing to manually search through logs.
This corresponds to Option A: Create an Amazon CloudWatch Logs log group. Configure the appropriate AWS WAF web ACL to send log messages to the log group. Instruct the operations team to create CloudWatch metric filters.