Number: DOP-C02
Passing Score: 800
Time Limit: 120
File Version: 11.0

Exam Code: DOC-C02

Exam Name: AWS DevOps Engineer - Professional

Exam A

QUESTION 1
A company has a mobile application that makes HTTP API calls to an Application Load Balancer (ALB). The ALB routes requests to an AWS Lambda function. Many different versions of the application are in use at any given time, including versions that are in testing by a subset of users. The version of the application is defined in the user-agent header that is sent with all requests to the API.
After a series of recent changes to the API, the company has observed issues with the application. The company needs to gather a metric for each API operation by response code for each version of the application that is in use. A DevOps engineer has modified the Lambda function to extract the API operation name, version information from the user-agent header and response code.
Which additional set of actions should the DevOps engineer take to gather the required metrics?

A. Modify the Lambda function to write the API operation name, response code, and version number as a log line to an Amazon CloudWatch Logs log group. Configure a CloudWatch Logs metric filter that increments a metric for each API operation name. Specify response code and application version as dimensions for the metric.
B. Modify the Lambda function to write the API operation name, response code, and version number as a log line to an Amazon CloudWatch Logs log group. Configure a CloudWatch Logs Insights query to populate CloudWatch metrics from the log lines. Specify response code and application version as dimensions for the metric.
C. Configure the ALB access logs to write to an Amazon CloudWatch Logs log group. Modify the Lambda function to respond to the ALB with the API operation name, response code, and version number as response metadata. Configure a CloudWatch Logs metric filter that increments a metric for each API operation name. Specify response code and application version as dimensions for the metric.
D. Configure AWS X-Ray integration on the Lambda function. Modify the Lambda function to create an X-Ray subsegment with the API operation name, response code, and version number. Configure X-Ray insights to extract an aggregated metric for each API operation name and to publish the metric to Amazon CloudWatch. Specify response code and application version as dimensions for the metric.

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 2
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 3
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 4
A company has a legacy application A DevOps engineer needs to automate the process of building the deployable artifact for the legacy application. The solution must store the deployable artifact in an existing Amazon S3 bucket for future deployments to reference
Which solution will meet these requirements in the MOST operationally efficient way?

A. Create a custom Docker image that contains all the dependencies tor the legacy application Store the custom Docker image in a new Amazon Elastic Container Registry (Amazon ECR) repository Configure a new AWS CodeBuild project to use the custom Docker image to build the deployable artifact and to save the artifact to the S3 bucket.
B. Launch a new Amazon EC2 instance Install all the dependencies (or the legacy application on the EC2 instance Use the EC2 instance to build the deployable artifact and to save the artifact to the S3 bucket.
C. Create a custom EC2 Image Builder image Install all the dependencies for the legacy application on the image Launch a new Amazon EC2 instance from the image Use the new EC2 instance to build the deployable artifact and to save the artifact to the S3 bucket.
D. Create an Amazon Elastic Kubernetes Service (Amazon EKS) cluster with an AWS Fargate profile that runs in multiple Availability Zones Create a custom Docker image that contains all the dependencies for the legacy

application Store the custom Docker image in a new Amazon Elastic Container Registry (Amazon ECR) repository Use the custom Docker image inside the EKS cluster to build the deployable artifact and to save the artifact to the S3 bucket.

Correct Answer: A
Section:
Explanation:
This approach is the most operationally efficient because it leverages the benefits of containerization, such as isolation and reproducibility, as well as AWS managed services. AWS CodeBuild is a fully managed build service that can compile your source code, run tests, and produce deployable software packages. By using a custom Docker image that includes all dependencies, you can ensure that the environment in which your code is built is consistent. Using Amazon ECR to store Docker images lets you easily deploy the images to any environment. Also, you can directly upload the build artifacts to Amazon S3 from AWS CodeBuild, which is beneficial for version control and archival purposes.

QUESTION 5
A company has a data ingestion application that runs across multiple AWS accounts. The accounts are in an organization in AWS Organizations. The company needs to monitor the application and consolidate access to the application. Currently the company is running the application on Amazon EC2 instances from several Auto Scaling groups. The EC2 instances have no access to the internet because the data is sensitive Engineers have deployed the necessary VPC endpoints. The EC2 instances run a custom AMI that is built specifically tor the application.
To maintain and troubleshoot the application, system administrators need the ability to log in to the EC2 instances. This access must be automated and controlled centrally. The company's security team must receive a notification whenever the instances are accessed.
Which solution will meet these requirements?

A. Create an Amazon EventBridge rule to send notifications to the security team whenever a user logs in to an EC2 instance Use EC2 Instance Connect to log in to the instances. Deploy Auto Scaling groups by using AWS Cloud Formation Use the cfn-init helper script to deploy appropriate VPC routes for external access Rebuild the custom AMI so that the custom AMI includes AWS Systems Manager Agent.
B. Deploy a NAT gateway and a bastion host that has internet access Create a security group that allows incoming traffic on all the EC2 instances from the bastion host Install AWS Systems Manager Agent on all the EC2 instances Use Auto Scaling group lifecycle hooks for monitoring and auditing access Use Systems Manager Session Manager to log in to the instances Send logs to a log group m Amazon CloudWatch Logs. Export data to Amazon S3 for auditing Send notifications to the security team by using S3 event notifications.
C. Use EC2 Image Builder to rebuild the custom AMI Include the most recent version of AWS Systems Manager Agent in the Image Configure the Auto Scaling group to attach the AmazonSSMManagedinstanceCore role to all the EC2 instances Use Systems Manager Session Manager to log in to the instances Enable logging of session details to Amazon S3 Create an S3 event notification for new file uploads to send a message to the security team through an Amazon Simple Notification Service (Amazon SNS) topic.
D. Use AWS Systems Manager Automation to build Systems Manager Agent into the custom AMI Configure AWS Configure to attach an SCP to the root organization account to allow the EC2 instances to connect to Systems Manager Use Systems Manager Session Manager to log in to the instances Enable logging of session details to Amazon S3 Create an S3 event notification for new file uploads to send a message to the security team through an Amazon Simple Notification Service (Amazon SNS) topic.

Correct Answer: C
Section:
Explanation:
Even if AmazonSSMManagedInstanceCore is a managed policy and not an IAM role I will go with C because this policy is to be attached to an IAM role for EC2 to access System Manager.

QUESTION 6
A company recently migrated its legacy application from on-premises to AWS. The application is hosted on Amazon EC2 instances behind an Application Load Balancer which is behind Amazon API Gateway. The company wants to ensure users experience minimal disruptions during any deployment of a new version of the application. The company also wants to ensure it can quickly roll back updates if there is an issue.
Which solution will meet these requirements with MINIMAL changes to the application?

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

Correct Answer: A
Section:
Explanation:

API Gateway supports canary deployment on a deployment stage before you direct all traffic to that stage. A parallel environment means we will create a new ALB and a target group that will target a new set of EC2 instances on which the newer version of the app will be deployed. So the canary setting associated to the new version of the API will connect with the new ALB instance which in turn will direct the traffic to the new EC2 instances on which the newer version of the application is deployed.

QUESTION 7
A development team manually builds an artifact locally and then places it in an Amazon S3 bucket. The application has a local cache that must be cleared when a deployment occurs. The team runs a command to do this downloads the artifact from Amazon S3 and unzips the artifact to complete the deployment.
A DevOps team wants to migrate to a CI/CD process and build in checks to stop and roll back the deployment when a failure occurs. This requires the team to track the progression of the deployment.
Which combination of actions will accomplish this? (Select THREE)

A. Allow developers to check the code into a code repository Using Amazon EventBridge on every pull into the mam branch invoke an AWS Lambda function to build the artifact and store it in Amazon S3.
B. Create a custom script to clear the cache Specify the script in the BeforeInstall lifecycle hook in the AppSpec file.
C. Create user data for each Amazon EC2 instance that contains the clear cache script Once deployed test the application If it is not successful deploy it again.
D. Set up AWS CodePipeline to deploy the application Allow developers to check the code into a code repository as a source tor the pipeline.
E. Use AWS CodeBuild to build the artifact and place it in Amazon S3 Use AWS CodeDeploy to deploy the artifact to Amazon EC2 instances.
F. Use AWS Systems Manager to fetch the artifact from Amazon S3 and deploy it to all the instances.

Correct Answer: B, D, E
Section:

QUESTION 8
A DevOps engineer is working on a data archival project that requires the migration of on-premises data to an Amazon S3 bucket. The DevOps engineer develops a script that incrementally archives on-premises data that is older than 1 month to Amazon S3. Data that is transferred to Amazon S3 is deleted from the on-premises location The script uses the S3 PutObject operation.
During a code review the DevOps engineer notices that the script does not verity whether the data was successfully copied to Amazon S3. The DevOps engineer must update the script to ensure that data is not corrupted during transmission. The script must use MD5 checksums to verify data integrity before the on-premises data is deleted.
Which solutions for the script will meet these requirements'? (Select TWO.)

A. Check the returned response for the Versioned Compare the returned Versioned against the MD5 checksum.
B. Include the MD5 checksum within the Content-MD5 parameter. Check the operation call's return status to find out if an error was returned.
C. Include the checksum digest within the tagging parameter as a URL query parameter.
D. Check the returned response for the ETag. Compare the returned ETag against the MD5 checksum.
E. Include the checksum digest within the Metadata parameter as a name-value pair After upload use the S3 HeadObject operation to retrieve metadata from the object.

Correct Answer: B, D
Section:
Explanation:
https://docs.aws.amazon.com/AmazonS3/latest/userguide/checking-object-integrity.html

QUESTION 9
A DevOps engineer used an AWS Cloud Formation custom resource to set up AD Connector. The AWS Lambda function ran and created AD Connector, but Cloud Formation is not transitioning from CREATE_IN_PROGRESS to CREATE_COMPLETE.
Which action should the engineer take to resolve this issue?

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

Correct Answer: B
Section:

QUESTION 10
A DevOps engineer is working on a project that is hosted on Amazon Linux and has failed a security review. The DevOps manager has been asked to review the company buildspec. yaml die for an AWS CodeBuild project and provide recommendations. The buildspec. yaml file is configured as follows:

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

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

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

Correct Answer: B, C, E
Section:
Explanation:
B) Update the CodeBuild project role with the necessary permissions and then remove the AWS credentials from the environment variable. C. Store the DB_PASSWORD as a SecureString value in AWS Systems Manager Parameter Store and then remove the DB_PASSWORD from the environment variables. E. Use AWS Systems Manager run command versus scp and ssh commands directly to the instance.

QUESTION 11
A highly regulated company has a policy that DevOps engineers should not log in to their Amazon EC2 instances except in emergencies. It a DevOps engineer does log in the security team must be notified within 15 minutes of the occurrence.
Which solution will meet these requirements'?

A. Install the Amazon Inspector agent on each EC2 instance Subscribe to Amazon EventBridge notifications Invoke an AWS Lambda function to check if a message is about user logins If it is send a notification to the security team using Amazon SNS.
B. Install the Amazon CloudWatch agent on each EC2 instance Configure the agent to push all logs to Amazon CloudWatch Logs and set up a CloudWatch metric filter that searches for user logins. If a login is found send a notification to the security team using Amazon SNS.
C. Set up AWS CloudTrail with Amazon CloudWatch Logs. Subscribe CloudWatch Logs to Amazon Kinesis Attach AWS Lambda to Kinesis to parse and determine if a log contains a user login If it does, send a notification to the security team using Amazon SNS.
D. Set up a script on each Amazon EC2 instance to push all logs to Amazon S3 Set up an S3 event to invoke an AWS Lambda function which invokes an Amazon Athena query to run. The Athena query checks tor logins and sends the output to the security team using Amazon SNS.

Correct Answer: B
Section:
Explanation:
https://aws.amazon.com/blogs/security/how-to-monitor-and-visualize-failed-ssh-access-attempts-to-amazon-ec2-linux-instances/

QUESTION 12
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 13
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 14
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 15
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 16
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 17
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 18
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 19
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 20
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 21
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 22
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 23
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 24
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 25
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 26
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 27
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 28
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 29
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 30
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 31
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 32
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 33

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 34

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 35
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 36
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 37
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 38
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 39
A company is hosting a static website from an Amazon S3 bucket. The website is available to customers at example.com. The company uses an Amazon Route 53 weighted routing policy with a TTL of 1 day. The company has decided to replace the existing static website with a dynamic web application. The dynamic web application uses an Application Load Balancer (ALB) in front of a fleet of Amazon EC2 instances.
On the day of production launch to customers, the company creates an additional Route 53 weighted DNS record entry that points to the ALB with a weight of 255 and a TTL of 1 hour. Two days later, a DevOps engineer notices that the previous static website is displayed sometimes when customers navigate to example.com.
How can the DevOps engineer ensure that the company serves only dynamic content for example.com?

A. Delete all objects, including previous versions, from the S3 bucket that contains the static website content.
B. Update the weighted DNS record entry that points to the S3 bucket. Apply a weight of 0. Specify the domain reset option to propagate changes immediately.

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

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

Correct Answer: D
Section:

QUESTION 40
A company manages multiple AWS accounts in AWS Organizations. The company's security policy states that AWS account root user credentials for member accounts must not be used. The company monitors access to the root user credentials.
A recent alert shows that the root user in a member account launched an Amazon EC2 instance. A DevOps engineer must create an SCP at the organization's root level that will prevent the root user in member accounts from making any AWS service API calls.
Which SCP will meet these requirements?
A)

```
{
    "Version": "2012-10-17",
    "Statement": [
        {
            "Effect": "Allow",
            "Action": "*",
            "Resource": "*",
            "Condition": {
                "StringNotLike": { "aws:PrincipalArn": "arn:aws:iam::*:root" }
            }
        }
    ]
}
```

B)

```
{
    "Version": "2012-10-17",
    "Statement": [
        {
            "Effect": "Deny",
            "Action": "*",
            "Resource": "*",
            "Principal": { "AWS": "arn:aws:iam::*:root" }
        }
    ]
}
```

C)

```
{
    "Version": "2012-10-17",
    "Statement": [
        {
            "Effect": "Deny",
            "Action": "*",
            "Resource": "*",
            "Condition": {
                "StringLike": { "aws:PrincipalArn": "arn:aws:iam::*:root" }
            }
        }
    ]
}
```

D)

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


A. Option A
B. Option B
C. Option C
D. Option D

Correct Answer: D
Section:

QUESTION 41
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 42
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 43
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 44
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 45
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 46
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 47
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 48
A global company manages multiple AWS accounts by using AWS Control Tower. The company hosts internal applications and public applications.
Each application team in the company has its own AWS account for application hosting. The accounts are consolidated in an organization in AWS Organizations. One of the AWS Control Tower member accounts serves as a centralized DevOps account with CI/CD pipelines that application teams use to deploy applications to their respective target AWS accounts. An 1AM role for deployment exists in the centralized DevOps account.
An application team is attempting to deploy its application to an Amazon Elastic Kubernetes Service (Amazon EKS) cluster in an application AWS account. An 1AM role for deployment exists in the application AWS account.
The deployment is through an AWS CodeBuild project that is set up in the centralized DevOps account. The CodeBuild project uses an 1AM service role for CodeBuild. The deployment is failing with an Unauthorized error during attempts to connect to the cross-account EKS cluster from CodeBuild.
Which solution will resolve this error?

A. Configure the application account's deployment 1AM role to have a trust relationship with the centralized DevOps account. Configure the trust relationship to allow the sts:AssumeRole action. Configure the application account's deployment 1AM role to have the required access to the EKS cluster. Configure the EKS cluster aws-auth ConfigMap to map the role to the appropriate system permissions.
B. Configure the centralized DevOps account's deployment I AM role to have a trust relationship with the application account. Configure the trust relationship to allow the sts:AssumeRole action. Configure the centralized DevOps account's deployment 1AM role to allow the required access to CodeBuild.
C. Configure the centralized DevOps account's deployment 1AM role to have a trust relationship with the application account. Configure the trust relationship to allow the sts:AssumeRoleWithSAML action. Configure the centralized DevOps account's deployment 1AM role to allow the required access to CodeBuild.
D. Configure the application account's deployment 1AM role to have a trust relationship with the AWS Control Tower management account. Configure the trust relationship to allow the sts:AssumeRole action. Configure the application account's deployment 1AM role to have the required access to the EKS cluster. Configure the EKS cluster aws-auth ConfigMap to map the role to the appropriate system permissions.

Correct Answer: A
Section:
Explanation:
In the source AWS account, the IAM role used by the CI/CD pipeline should have permissions to access the source code repository, build artifacts, and any other resources required for the build process. In the destination AWS accounts, the IAM role used for deployment should have permissions to access the AWS resources required for deploying the application, such as EC2 instances, RDS databases, S3 buckets, etc. The exact permissions required will depend on the specific resources being used by the application. the IAM role used for deployment in the destination accounts should also have permissions to assume the IAM role for deployment in the centralized DevOps account. This is typically done using an IAM role trust policy that allows the destination account to assume the DevOps account role.

QUESTION 49
A company is implementing AWS CodePipeline to automate its testing process The company wants to be notified when the execution state fails and used the following custom event pattern in Amazon EventBridge:

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

Which type of events will match this event pattern?

A. Failed deploy and build actions across all the pipelines
B. All rejected or failed approval actions across all the pipelines
C. All the events across all pipelines
D. Approval actions across all the pipelines

Correct Answer: B
Section:
Explanation:
Action-level states in events
Action state Description
STARTED The action is currently running.
SUCCEEDED The action was completed successfully.
FAILED For Approval actions, the FAILED state means the action was either rejected by the reviewer or failed due to an incorrect action configuration.
CANCELED The action was canceled because the pipeline structure was updated.

QUESTION 50
A DevOps engineer has implemented a Cl/CO pipeline to deploy an AWS Cloud Format ion template that provisions a web application. The web application consists of an Application Load Balancer (ALB) a target group, a launch template that uses an Amazon Linux 2 AMI an Auto Scaling group of Amazon EC2 instances, a security group and an Amazon RDS for MySQL database The launch template includes user data that specifies a script to install and start the application.
The initial deployment of the application was successful. The DevOps engineer made changes to update the version of the application with the user data. The CI/CD pipeline has deployed a new version of the template However, the health checks on the ALB are now failing The health checks have marked all targets as unhealthy.
During investigation the DevOps engineer notices that the Cloud Formation stack has a status of UPDATE_COMPLETE. However, when the DevOps engineer connects to one of the EC2 instances and checks /varar/log messages, the DevOps engineer notices that the Apache web server failed to start successfully because of a configuration error

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

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

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

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

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

Correct Answer: A
Section:
Explanation:
https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/aws-attribute-updatepolicy.html

QUESTION 51
A company wants to ensure that their EC2 instances are secure. They want to be notified if any new vulnerabilities are discovered on their instances and they also want an audit trail of all login activities on the instances.
Which solution will meet these requirements'?

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

Correct Answer: D
Section:
Explanation:
This solution will meet the requirements because it will use Amazon Inspector to scan the EC2 instances for any new vulnerabilities and generate findings that can be viewed in the Inspector console or sent as notifications via Amazon Simple Notification Service (SNS). It will also use the Amazon CloudWatch Agent to collect and send system logs from the EC2 instances to Amazon CloudWatch Logs, where they can be stored, searched, and analyzed. The system logs can provide an audit trail of all login activities on the instances, as well as other useful information such as performance metrics, errors, and events.
https://docs.aws.amazon.com/inspector/latest/user/what-is-inspector.html

QUESTION 52
A company uses a series of individual Amazon Cloud Formation templates to deploy its multi-Region Applications. These templates must be deployed in a specific order. The company is making more changes to the templates than previously expected and wants to deploy new templates more efficiently. Additionally, the data engineering team must be notified of all changes to the templates.
What should the company do to accomplish these goals?

A. Create an AWS Lambda function to deploy the Cloud Formation templates m the required order Use stack policies to alert the data engineering team.
B. Host the Cloud Formation templates in Amazon S3 Use Amazon S3 events to directly trigger CloudFormation updates and Amazon SNS notifications.
C. Implement CloudFormation StackSets and use drift detection to trigger update alerts to the data engineering team.
D. Leverage CloudFormation nested stacks and stack sets (or deployments Use Amazon SNS to notify the data engineering team.

Correct Answer: D
Section:

Explanation:
This solution will meet the requirements because it will use CloudFormation nested stacks and stack sets to deploy the templates more efficiently and consistently across multiple regions. Nested stacks allow the company to separate out common components and reuse templates, while stack sets allow the company to create stacks in multiple accounts and regions with a single template. The company can also use Amazon SNS to send notifications to the data engineering team whenever a change is made to the templates or the stacks. Amazon SNS is a service that allows you to publish messages to subscribers, such as email addresses, phone numbers, or other AWS services. By using Amazon SNS, the company can ensure that the data engineering team is aware of all changes to the templates and can take appropriate actions if needed. What is Amazon SNS? - Amazon Simple Notification Service

QUESTION 53
A DevOps engineer is deploying a new version of a company's application in an AWS CodeDeploy deployment group associated with its Amazon EC2 instances. After some time, the deployment fails. The engineer realizes that all the events associated with the specific deployment ID are in a Skipped status and code was not deployed in the instances associated with the deployment group.
What are valid reasons for this failure? (Select TWO.).

A. The networking configuration does not allow the EC2 instances to reach the internet via a NAT gateway or internet gateway and the CodeDeploy endpoint cannot be reached.
B. The IAM user who triggered the application deployment does not have permission to interact with the CodeDeploy endpoint.
C. The target EC2 instances were not properly registered with the CodeDeploy endpoint.
D. An instance profile with proper permissions was not attached to the target EC2 instances.
E. The appspec. yml file was not included in the application revision.

Correct Answer: A, D
Section:
Explanation:
https://docs.aws.amazon.com/codedeploy/latest/userguide/troubleshooting-deployments.html#troubleshooting-skipped-lifecycle-events

QUESTION 54
A company manages an application that stores logs in Amazon CloudWatch Logs. The company wants to archive the logs to an Amazon S3 bucket Logs are rarely accessed after 90 days and must be retained tor 10 years.
Which combination of steps should a DevOps engineer take to meet these requirements? (Select TWO.)

A. Configure a CloudWatch Logs subscription filter to use AWS Glue to transfer all logs to an S3 bucket.
B. Configure a CloudWatch Logs subscription filter to use Amazon Kinesis Data Firehose to stream all logs to an S3 bucket.
C. Configure a CloudWatch Logs subscription fitter to stream all logs to an S3 bucket.
D. Configure the S3 bucket lifecycle policy to transition logs to S3 Glacier after 90 days and to expire logs after 3.650 days.
E. Configure the S3 bucket lifecycle policy to transition logs to Reduced Redundancy after 90 days and to expire logs after 3.650 days.

Correct Answer: B, D
Section:
Explanation:
https://docs.aws.amazon.com/AmazonCloudWatch/latest/logs/SubscriptionFilters.html

QUESTION 55
A DevOps engineer needs to configure a blue green deployment for an existing three-tier application. The application runs on Amazon EC2 instances and uses an Amazon RDS database The EC2 instances run behind an Application Load Balancer (ALB) and are in an Auto Scaling group.
The DevOps engineer has created a launch template and an Auto Scaling group for the blue environment. The DevOps engineer also has created a launch template and an Auto Scaling group for the green environment. Each Auto Scaling group deploys to a matching blue or green target group. The target group also specifies which software blue or green gets loaded on the EC2 instances. The ALB can be configured to send traffic to the blue environments target group or the green environments target group. An Amazon Route 53 record for www example com points to the ALB.
The deployment must move traffic all at once between the software on the blue environment's EC2 instances to the newly deployed software on the green environments EC2 instances
What should the DevOps engineer do to meet these requirements?

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

the ALB to send traffic to the green environment's target group.

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

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

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

Correct Answer: A
Section:
Explanation:
This solution will meet the requirements because it will use a rolling restart to gradually replace the EC2 instances in the green environment with new instances that have the new software version installed. A rolling restart is a process that terminates and launches instances in batches, ensuring that there is always a minimum number of healthy instances in service. This way, the green environment can be updated without affecting the availability or performance of the application. When the rolling restart is complete, the DevOps engineer can use an AWS CLI command to modify the listener rules of the ALB and change the default action to forward traffic to the green environment's target group. This will switch the traffic from the blue environment to the green environment all at once, as required by the question.

QUESTION 56
A large enterprise is deploying a web application on AWS. The application runs on Amazon EC2 instances behind an Application Load Balancer. The instances run in an Auto Scaling group across multiple Availability Zones. The application stores data in an Amazon RDS for Oracle DB instance and Amazon DynamoDB. There are separate environments tor development testing and production.
What is the MOST secure and flexible way to obtain password credentials during deployment?

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

B. Launch the EC2 instances with an EC2 1AM role to access AWS services Retrieve the database credentials from AWS Secrets Manager.

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

D. Launch the EC2 instances with an EC2 1AM role to access AWS services Store the database passwords in an encrypted config file with the application artifacts.

Correct Answer: B
Section:
Explanation:
AWS Secrets Manager is a secrets management service that helps you protect access to your applications, services, and IT resources. This service enables you to easily rotate, manage, and retrieve database credentials, API keys, and other secrets throughout their lifecycle. Using Secrets Manager, you can secure and manage secrets used to access resources in the AWS Cloud, on third-party services, and on-premises. SSM parameter store and AWS Secret manager are both a secure option. However, Secrets manager is more flexible and has more options like password generation.
Reference: https://www.1strategy.com/blog/2019/02/28/aws-parameter-store-vs-aws-secrets-manager/

QUESTION 57
A DevOps engineer is designing an application that integrates with a legacy REST API. The application has an AWS Lambda function that reads records from an Amazon Kinesis data stream. The Lambda function sends the records to the legacy REST API.
Approximately 10% of the records that the Lambda function sends from the Kinesis data stream have data errors and must be processed manually. The Lambda function event source configuration has an Amazon Simple Queue Service (Amazon SQS) dead-letter queue as an on-failure destination. The DevOps engineer has configured the Lambda function to process records in batches and has implemented retries in case of failure.
During testing the DevOps engineer notices that the dead-letter queue contains many records that have no data errors and that already have been processed by the legacy REST API. The DevOps engineer needs to configure the Lambda function's event source options to reduce the number of errorless records that are sent to the dead-letter queue.
Which solution will meet these requirements?

A. Increase the retry attempts

B. Configure the setting to split the batch when an error occurs

C. Increase the concurrent batches per shard

D. Decrease the maximum age of record

Correct Answer: B
Section:
Explanation:
This solution will meet the requirements because it will reduce the number of errorless records that are sent to the dead-letter queue. When you configure the setting to split the batch when an error occurs, Lambda will retry only the records that caused the error, instead of retrying the entire batch. This way, the records that have no data errors and have already been processed by the legacy REST API will not be retried and sent to the dead-letter queue unnecessarily.
https://docs.aws.amazon.com/lambda/latest/dg/with-kinesis.html

QUESTION 58
A company is storing 100 GB of log data in csv format in an Amazon S3 bucket SQL developers want to query this data and generate graphs to visualize it. The SQL developers also need an efficient automated way to store metadata from the csv file.
Which combination of steps will meet these requirements with the LEAST amount of effort? (Select THREE.)

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

Correct Answer: B, C, E
Section:
Explanation:
https://docs.aws.amazon.com/glue/latest/dg/components-overview.html

QUESTION 59
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 60
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 61
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 62
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 63
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 64
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 65
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 66
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 67
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 68
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 69
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 70
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 71
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 72
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 73
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 74
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 1AM policy to the developers' 1AM 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 1AM 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 1AM 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 75
A company is reviewing its 1AM 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 76
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 77
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 78
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 1AM user that is attached to the AdministratorAccess 1AM 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 1AM policy to the 1AM 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 79
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 80
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 81
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 82
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.