QUESTION 31
A company has its AWS accounts in an organization in AWS Organizations. AWS Config is manually configured in each AWS account. The company needs to implement a solution to centrally configure AWS Config for all accounts in the organization The solution also must record resource changes to a central account.
Which combination of actions should a DevOps engineer perform to meet these requirements? (Choose two.)

A. Configure a delegated administrator account for AWS Config. Enable trusted access for AWS Config in the organization.
B. Configure a delegated administrator account for AWS Config. Create a service-linked role for AWS Config in the organization's management account.
C. Create an AWS CloudFormation template to create an AWS Config aggregator. Configure a CloudFormation stack set to deploy the template to all accounts in the organization.
D. Create an AWS Config organization aggregator in the organization's management account. Configure data collection from all AWS accounts in the organization and from all AWS Regions.
E. Create an AWS Config organization aggregator in the delegated administrator account. Configure data collection from all AWS accounts in the organization and from all AWS Regions.

Answer: AE
Explanation:
https://docs.aws.amazon.com/organizations/latest/userguide/services-that-can-integrate-config.html

QUESTION 32
A company wants to migrate its content sharing web application hosted on Amazon EC2 to a serverless architecture. The company currently deploys changes to its application by creating a new Auto Scaling group of EC2 instances and a new Elastic Load Balancer, and then shifting the traffic away using an Amazon Route 53 weighted routing policy.
For its new serverless application, the company is planning to use Amazon API Gateway and AWS Lambda. The company will need to update its deployment processes to work with the new application. It will also need to retain the ability to test new features on a small number of users before rolling the features out to the entire user base.
Which deployment strategy will meet these requirements?

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

custom layer. When code needs to be changed, use OpsWorks to perform a blue/green deployment and shift traffic gradually.

Answer: B

QUESTION 33
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.

Answer: B

QUESTION 34
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.

Answer: B

QUESTION 35

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.

Answer: ACE

QUESTION 36
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.

Answer: B

QUESTION 37
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.

Answer: A

QUESTION 38
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.

Answer: D

QUESTION 39
A company has developed a serverless web application that is hosted on AWS. The application consists of Amazon S3. Amazon API Gateway, several AWS Lambda functions, and an Amazon RDS for MySQL database. The company is using AWS CodeCommit to store the source code. The source code is a combination of AWS Serverless Application Model (AWS SAM) templates and Python code.
A security audit and penetration test reveal that user names and passwords for authentication to the database are hardcoded within CodeCommit repositories. A DevOps engineer must implement a solution to automatically detect and prevent hardcoded secrets.
What is the MOST secure solution that meets these requirements?

A. Enable Amazon CodeGuru Profiler. Decorate the handler function with @with_lambda_profiler(). Manually review the recommendation report. Write the secret to AWS Systems Manager Parameter Store as a secure string. Update the SAM templates and the Python code to pull the secret from Parameter Store.
B. Associate the CodeCommit repository with Amazon CodeGuru Reviewer. Manually check the code review for any recommendations. Choose the option to protect the secret. Update the SAM templates and the Python code to pull the secret from AWS Secrets Manager.
C. Enable Amazon CodeGuru Profiler. Decorate the handler function with @with_lambda_profiler(). Manually review the recommendation report. Choose the option to protect the secret. Update the SAM templates and the Python code to pull the secret from AWS Secrets Manager.
D. Associate the CodeCommit repository with Amazon CodeGuru Reviewer. Manually check the code review for any recommendations. Write the secret to AWS Systems Manager Parameter Store as a string. Update the SAM templates and the Python code to pull the secret from Parameter Store.

Answer: B

QUESTION 40
You have been asked to de-risk deployments at your company. Specifically, the CEO is concerned about outages that occur because of accidental inconsistencies between Staging and Production, which sometimes cause unexpected behaviors in Production even when Staging tests pass. You already use Docker to get high consistency between Staging and Production for the application environment on your EC2 instances. How do you further de-risk the rest of the execution environment, since in AWS, there are many service components you may use beyond EC2 virtual machines?

A. Develop models of your entire cloud system in CloudFormation. Use this model in Staging and Production to achieve greater parity.
B. Use AWS Config to force the Staging and Production stacks to have configuration parity. Any differences will be detected for you so you are aware of risks.
C. Use AMIs to ensure the whole machine, including the kernel of the virual machines, is consistent, since Docker uses Linux Container (LXC) technology, and we need to make sure the container environment is consistent.
D. Use AWS ECS and Docker clustering. This will make sure that the AMIs and machine sizes are the same across both environments.

Answer: A
Explanation:
Only CloudFormation's JSON Templates allow declarative version control of repeatably deployable models of entire AWS clouds.
https://blogs.aws.amazon.com/application-management/blog/category/Best+practices

QUESTION 41
You are creating a new API for video game scores. Reads are 100 times more common than writes, and the top 1% of scores are read 100 times more frequently than the rest of the scores. What's the best design for this system, using DynamoDB?

A. DynamoDB table with 100x higher read than write throughput, with CloudFront caching.
B. DynamoDB table with roughly equal read and write throughput, with CloudFront caching.
C. DynamoDB table with 100x higher read than write throughput, with ElastiCache caching.
D. DynamoDB table with roughly equal read and write throughput, with ElastiCache caching.

Answer: D
Explanation:
Because the 100x read ratio is mostly driven by a small subset, with caching, only a roughly equal number of reads to writes will miss the cache, since the supermajority will hit the top 1% scores. Knowing we need to set the values roughly equal when using caching, we select AWS ElastiCache, because CloudFront cannot directly cache DynamoDB queries, and ElastiCache is an excellent in-memory cache for database queries, rather than a distributed proxy cache for content delivery. ... One solution would be to cache these reads at the application layer. Caching is a technique that is used in many high-throughput applications, offloading read activity on hot items to the cache rather than to the database. Your application can cache the most popular items in memory, or use a product such as ElastiCache to do the same.
http://docs.aws.amazon.com/amazondynamodb/latest/developerguide/GuidelinesForTables.html#GuidelinesForTables.CachePopularItem

QUESTION 42
Your system uses a multi-master, multi-region DynamoDB configuration spanning two regions to achieve high availablity. For the first time since launching your system, one of the AWS Regions in which you operate over went down for 3 hours, and the failover worked correctly. However, after recovery, your users are experiencing strange bugs, in which users on different sides of the globe see different data. What is a likely design issue that was not accounted for when launching?

A. The system does not have Lambda Functor Repair Automatons, to perform table scans and

chack for corrupted partition blocks inside the Table in the recovered Region.

B. The system did not implement DynamoDB Table Defragmentation for restoring partition performance in the Region that experienced an outage, so data is served stale.

C. The system did not include repair logic and request replay buffering logic for post-failure, to re-synchronize data to the Region that was unavailable for a number of hours.

D. The system did not use DynamoDB Consistent Read requests, so the requests in different areas are not utilizing consensus across Regions at runtime.

Answer: C
Explanation:
When using multi-region DynamoDB systems, it is of paramount importance to make sure that all requests made to one Region are replicated to the other. Under normal operation, the system in question would correctly perform write replays into the other Region. If a whole Region went down, the system would be unable to perform these writes for the period of downtime. Without buffering write requests somehow, there would be no way for the system to replay dropped cross-region writes, and the requests would be serviced differently depending on the Region from which they were served after recovery.
http://docs.aws.amazon.com/amazondynamodb/latest/developerguide/Streams.CrossRegionRepl.html

QUESTION 43
A DevOps Engineer needs to deploy a scalable three-tier Node.js application in AWS. The application must have zero downtime during deployments and be able to roll back to previous versions. Other applications will also connect to the same MySQL backend database. The CIO has provided the following guidance for logging:

```
- Centrally view all current web access server logs.
- Search and filter web and application logs in near-real time.
- Retain log data for three months.
```

How should these requirements be met?

A. Deploy the application using AWS Elastic Beanstalk. Configure the environment type for Elastic Load Balancing and Auto Scaling. Create an Amazon RDS MySQL instance inside the Elastic Beanstalk stack. Configure the Elastic Beanstalk log options to stream logs to Amazon CloudWatch Logs. Set retention to 90 days.

B. Deploy the application on Amazon EC2. Configure Elastic Load Balancing and Auto Scaling. Use an Amazon RDS MySQL instance for the database tier. Configure the application to store log files in Amazon S3. Use Amazon EMR to search and filter the data. Set an Amazon S3 lifecycle rule to expire objects after 90 days.

C. Deploy the application using AWS Elastic Beanstalk. Configure the environment type for Elastic Load Balancing and Auto Scaling. Create the Amazon RDS MySQL instance outside the Elastic Beanstalk stack. Configure the Elastic Beanstalk log options to stream logs to Amazon CloudWatch Logs. Set retention to 90 days.

D. Deploy the application on Amazon EC2. Configure Elastic Load Balancing and Auto Scaling. Use an Amazon RDS MySQL instance for the database tier. Configure the application to load streaming log data using Amazon Kinesis Data Firehouse into Amazon ES. Delete and create a new Amazon ES domain every 90 days.

Answer: C
Explanation:
The Amazon EC2 instances in your Elastic Beanstalk environment generate logs that you can view to troubleshoot issues with your application or configuration files. Logs created by the web server, application server, Elastic Beanstalk platform scripts, and AWS CloudFormation are stored locally on individual instances. You can easily retrieve them by using the environment management console or the EB CLI. You can also configure your environment to stream logs to Amazon CloudWatch Logs in real-time.
https://docs.aws.amazon.com/elasticbeanstalk/latest/dg/using-features.logging.html

QUESTION 44
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.

Answer: B
Explanation:
The DevOps engineer can 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. Then, a CloudWatch Logs metric filter can be configured to increment a metric for each API operation name, with response code and application version specified as dimensions for the metric.

QUESTION 45
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.

Answer: C
Explanation:
Lambda integrates with Application Auto Scaling, allowing you to manage provisioned concurrency on a schedule or

based on utilization.
https://docs.aws.amazon.com/lambda/latest/dg/provisioned-concurrency.html
https://docs.aws.amazon.com/autoscaling/application/userguide/services-that-can-integrate-lambda.html

QUESTION 46
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.

Answer: D

QUESTION 47
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.

Answer: B
Explanation:
https://aws.amazon.com/ru/blogs/mt/build-an-aws-config-custom-rule-to-optimize-amazon-ebs-volume-types/

QUESTION 48
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

Answer: B

QUESTION 49
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.

Answer: ADF
Explanation:
https://docs.aws.amazon.com/autoscaling/ec2/userguide/key-policy-requirements-EBS-encryption.html
Service linked role must get specific grant on the account it cannot use key policy.

QUESTION 50
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.

Answer: AD

QUESTION 51
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.

Answer: AC
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 52
A company runs a database on a single Amazon EC2 instance in a development environment. The data is stored on separate Amazon EBS volumes that are attached to the EC2 instance. An Amazon Route 53 A record has been created and configured to point to the EC2 instance. The company would like to automate the recovery of the database instance when an instance or Availability Zone (AZ) fails. The company also wants to keep its costs low. The RTO is 4 hours and RPO is 12 hours. Which solution should a DevOps Engineer implement to meet these requirements?

A. Run the database in an Auto Scaling group with a minimum and maximum instance count of 1 in multiple AZs. Add a lifecycle hook to the Auto Scaling group and define an Amazon CloudWatch Events rule that is triggered when a lifecycle event occurs. Have the CloudWatch Events rule invoke an AWS Lambda function to detach or attach the Amazon EBS data volumes from the

EC2 instance based on the event. Configure the EC2 instance UserData to mount the data volumes (retry on failure with a short delay), then start the database and update the Route 53 record.

B. Run the database on two separate EC2 instances in different AZs with one active and the other as a standby. Attach the data volumes to the active instance. Configure an Amazon CloudWatch Events rule to invoke an AWS Lambda function on EC2 instance termination. The Lambda function launches a replacement EC2 instance. If the terminated instance was the active node, then the function attaches the data volumes to the standby node. Start the database and update the Route 53 record.

C. Run the database in an Auto Scaling group with a minimum and maximum instance count of 1 in multiple AZs. Create an AWS Lambda function that is triggered by a scheduled Amazon CloudWatch Events rule every 4 hours to take a snapshot of the data volume and apply a tag. Have the instance UserData get the latest snapshot, create a new volume from it, and attach and mount the volume. Then start the database and update the Route 53 record.

D. Run the database on two separate EC2 instances in different AZs. Configure one of the instances as a master and the other as a standby. Set up replication between the master and standby instances. Point the Route 53 record to the master. Configure an Amazon CloudWatch Events rule to invoke an AWS Lambda function upon the EC2 instance termination. The Lambda function launches a replacement EC2 instance. If the terminated instance was the active node, the function promotes the standby to master and points the Route 53 record to it.

Answer: C
Explanation:
It restores EBS volumes from snapshot and snapshot is not AZ independent.
Two instances means more costs.
Also Auto Scaling group with min 1 max 1 IS the preferred method for HA.

QUESTION 53
A new zero-day vulnerability was found in OpenSSL requiring the immediate patching of a production web fleet running on Amazon Linux. Currently, OS updates are performed manually on a monthly basis and deployed using updates to the production Auto Scaling Group's launch configuration. Which method should a DevOps Engineer use to update packages in-place without downtime?

A. Use AWS CodePipline and AWS CodeBuild to generate new copies of these packages, and update the Auto Scaling group's launch configuration.

B. Use AWS Inspector to run "yum upgrade" on all running production instances, and manually update the AMI for the next maintenance window.

C. Use Amazon EC2 Run Command to issue a package update command to all running production instances, and update the AMI for future deployments.

D. Define a new AWS OpsWorks layer to match the running production instances, and use a recipe to issue a package update command to all running production instances.

Answer: C
Explanation:
https://aws.amazon.com/blogs/aws/ec2-run-command-is-now-a-cloudwatch-events-target/
EC2 Run Command is part of EC2 Systems Manager. It allows you to operate on collections of EC2 instances and on-premises servers reliably and at scale, in a controlled and selective fashion. You can run scripts, install software, collect metrics and log files, manage patches, and much more, on both Windows and Linux.

QUESTION 54
A defect was discovered in production and a new sprint item has been created for deploying a hotfix. However, any code change must go through the following steps before going into production:
- Scan the code for security breaches, such as password and access key leaks.
- Run the code through extensive, long running unit tests.
Which source control strategy should a DevOps Engineer use in combination with AWS CodePipeline to complete this process?

A. Create a hotfix tag on the last commit of the master branch. Trigger the development pipeline from the hotfix tag. Use AWS CodeDeploy with Amazon ECS to do a content scan and run unit tests. Add a manual approval stage that merges the hotfix tag into the master branch.
B. Create a hotfix branch from the master branch. Triger the development pipeline from the hotfix branch.
Use AWS CodeBuild to do a content scan and run unit tests. Add a manual approval stage that merges the hotfix branch into the master branch.
C. Create a hotfix branch from the master branch. Triger the development pipeline from the hotfix branch.
Use AWS Lambda to do a content scan and run unit tests. Add a manual approval stage that merges the hotfix branch into the master branch.
D. Create a hotfix branch from the master branch. Create a separate source stage for the hotfix branch in the production pipeline. Trigger the pipeline from the hotfix branch. Use AWS Lambda to do a content scan and use AWS CodeBuild to run unit tests. Add a manual approval stage that merges the hotfix branch into the master branch.

Answer: B
Explanation:
We need to create a feature branch to test the fix and codebuild can do both the scan and unit tests.
https://docs.aws.amazon.com/codebuild/latest/userguide/how-to-create-pipeline.html#how-to-create-pipeline-add

QUESTION 55
A Development team creates a build project in AWS CodeBuild. The build project invokes automated tests of modules that access AWS services.
Which of the following will enable the tests to run the MOST securely?

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

Answer: C
Explanation:
CodeBuild requires a service role to access AWS resources being tested/verified.
https://aws.amazon.com/blogs/devops/access-resources-in-a-vpc-from-aws-codebuild-builds/

QUESTION 56
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.

Answer: B
Explanation:
CFN will not delete non-empty bucket. It must be emptied first. Custom resource will do it.

QUESTION 57
A company has an AWS CodePipeline pipeline that is configured with an Amazon S3 bucket in the eu-west-1 Region. The pipeline deploys an AWS Lambda application to the same Region. The pipeline consists of an AWS CodeBuild project build action and an AWS CloudFormation deploy action.
The CodeBuild project uses the aws cloudformation package AWS CLI command to build an artifact that contains the Lambda function code's .zip file and the CloudFormation template. The CloudFormation deploy action references the CloudFormation template from the output artifact of the CodeBuild project's build action.
The company wants to also deploy the Lambda application to the us-east-1 Region by using the pipeline in eu-west-1. A DevOps engineer has already updated the CodeBuild project to use the aws cloudformation package command to produce an additional output artifact for us-east-1.
Which combination of additional steps should the DevOps engineer take to meet these requirements? (Choose two.)

A. Modify the CloudFormation template to include a parameter for the Lambda function code's zip file location. Create a new CloudFormation deploy action for us-east-1 in the pipeline. Configure the new deploy action to pass in the us-east-1 artifact location as a parameter override.
B. Create a new CloudFormation deploy action for us-east-1 in the pipeline. Configure the new deploy action to use the CloudFormation template from the us-east-1 output artifact.
C. Create an S3 bucket in us-east-1. Configure the S3 bucket policy to allow CodePipeline to have read and write access.
D. Create an S3 bucket in us-east-1. Configure S3 Cross-Region Replication (CRR) from the S3 bucket in eu-west-1 to the S3 bucket in us-east-1.
E. Modify the pipeline to include the S3 bucket for us-east-1 as an artifact store. Create a new CloudFormation deploy action for us-east-1 in the pipeline. Configure the new deploy action to use the CloudFormation template from the us-east-1 output artifact.

Answer: AC

QUESTION 58
A company runs an application on one Amazon EC2 instance. Application metadata is stored in Amazon S3 and must be retrieved if the instance is restarted. The instance must restart or relaunch automatically if the instance becomes unresponsive.
Which solution will meet these requirements?

A. Create an Amazon CloudWatch alarm for the StatusCheckFailed metric. Use the recover action to stop and start the instance. Use an S3 event notification to push the metadata to the instance when the instance is back up and running.
B. Configure AWS OpsWorks, and use the auto healing feature to stop and start the instance. Use a lifecycle event in OpsWorks to pull the metadata from Amazon S3 and update it on the instance.
C. Use EC2 Auto Recovery to automatically stop and start the instance in case of a failure. Use an S3 event notification to push the metadata to the instance when the instance is back up and running.
D. Use AWS CloudFormation to create an EC2 instance that includes the UserData property for the EC2 resource. Add a command in UserData to retrieve the application metadata from Amazon S3.

Answer: B

QUESTION 59
A company has multiple AWS accounts. The company uses AWS IAM Identity Center (AWS Single Sign-On) that is

integrated with AWS Toolkit for Microsoft Azure DevOps. The attributes for access control feature is enabled in IAM Identity Center.
The attribute mapping list contains two entries. The department key is mapped to ${path:enterprise.department}. The costCenter key is mapped to ${path:enterprise.costCenter}.
All existing Amazon EC2 instances have a department tag that corresponds to three company departments (d1, d2, d3). A DevOps engineer must create policies based on the matching attributes. The policies must minimize administrative effort and must grant each Azure AD user access to only the EC2 instances that are tagged with the user's respective department name.
Which condition key should the DevOps engineer include in the custom permissions policies to meet these requirements?

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

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

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

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

Answer: C
Explanation:
https://aws.amazon.com/blogs/aws/new-attributes-based-access-control-with-aws-single-sign-on/

QUESTION 60
A company hosts a security auditing application in an AWS account. The auditing application uses an IAM role to access other AWS accounts. All the accounts are in the same organization in AWS Organizations.
A recent security audit revealed that users in the audited AWS accounts could modify or delete the auditing application's IAM role. The company needs to prevent any modification to the auditing application's IAM role by any entity other than a trusted administrator IAM role.
Which solution will meet these requirements?

A. Create an SCP that includes a Deny statement for changes to the auditing application's IAM role. Include a condition that allows the trusted administrator IAM role to make changes. Attach the SCP to the root of the organization.
B. Create an SCP that includes an Allow statement for changes to the auditing application's IAM role by the trusted administrator IAM role. Include a Deny statement for changes by all other IAM principals. Attach the SCP to the IAM service in each AWS account where the auditing application has an IAM role.
C. Create an IAM permissions boundary that includes a Deny statement for changes to the auditing application's IAM role. Include a condition that allows the trusted administrator IAM role to make changes. Attach the permissions boundary to the audited AWS accounts.
D. Create an IAM permissions boundary that includes a Deny statement for changes to the auditing

application's IAM role. Include a condition that allows the trusted administrator IAM role to make changes. Attach the permissions boundary to the auditing application's IAM role in the AWS accounts.

Answer: A

QUESTION 61
A company has an on-premises application that is written in Go. A DevOps engineer must move the application to AWS. The company's development team wants to enable blue/green deployments and perform A/B testing. Which solution will meet these requirements?

A. Deploy the application on an Amazon EC2 instance, and create an AMI of the instance. Use the AMI to create an automatic scaling launch configuration that is used in an Auto Scaling group. Use Elastic Load Balancing to distribute traffic. When changes are made to the application, a new AMI will be created, which will initiate an EC2 instance refresh.
B. Use Amazon Lightsail to deploy the application. Store the application in a zipped format in an Amazon S3 bucket. Use this zipped version to deploy new versions of the application to Lightsail. Use Lightsail deployment options to manage the deployment.
C. Use AWS CodeArtifact to store the application code. Use AWS CodeDeploy to deploy the application to a fleet of Amazon EC2 instances. Use Elastic Load Balancing to distribute the traffic to the EC2 instances. When making changes to the application, upload a new version to CodeArtifact and create a new CodeDeploy deployment.
D. Use AWS Elastic Beanstalk to host the application. Store a zipped version of the application in Amazon S3. Use that location to deploy new versions of the application. Use Elastic Beanstalk to manage the deployment options.

Answer: D