1) A company controls the source code for an application in AWS CodeCommit. The company is creating a CI/CD pipeline for the application by using AWS CodePipeline. The pipeline must start automatically when changes occur to the main branch of the CodeCommit repository. Changes occur frequently every day, so the pipeline must be as responsive as possible.

What should a DevOps engineer do to meet these requirements?

A) Configure the pipeline to periodically check the repository's main branch for changes. Start the pipeline when changes are detected.
B) Configure an Amazon EventBridge (Amazon CloudWatch Events) rule to detect changes to the repository's main branch. Configure the pipeline to start in response to the changes.
C) Configure the repository to periodically run an AWS Lambda function. Configure the function to check the repository's main branch and to start the pipeline when the function detects changes.
D) Configure the repository to publish a notification to an Amazon Simple Notification Service (Amazon SNS) topic when changes occur to the repository's main branch. Subscribe the pipeline to the SNS topic.

2) A DevOps team has an application that stores critical company assets in an existing Amazon S3 bucket. The team uses a single AWS Region. A new company policy requires the team to deploy the application to multiple Regions. The assets must always be accessible. Users must use the same endpoint to access the assets.

Which combination of steps should the team take to meet these requirements in the MOST operationally efficient way? (Select THREE.)

A) Use AWS CloudFormation StackSets to create a new S3 bucket that has versioning enabled in each required Region. Copy the assets from the existing S3 bucket to the new S3 buckets. Create an AWS Lambda function to copy files that are added to the new S3 bucket in the primary Region to the additional Regions.
B) Use AWS CloudFormation StackSets to create a new S3 bucket that has versioning enabled in each required Region. Create multiple S3 replication rules on the new S3 bucket in the primary Region to replicate all its contents to the additional Regions. Copy the assets from the existing S3 bucket to the new S3 bucket in the primary Region.
C) Create an Amazon CloudFront distribution. Configure new origins for each S3 bucket. Create an origin group that contains all the newly created origins. Update the default behavior of the distribution to use the new origin group.
D) Create an Amazon CloudFront distribution. Configure new origins for each S3 bucket. Create a Lambda@Edge function to validate the availability of the origin and to route the viewer request to an available nearby origin.
E) Create an Amazon Route 53 alias record. Configure a failover routing policy that uses the newly created S3 buckets as a target.
F) Create an Amazon Route 53 alias record. Configure a simple routing policy that uses the Amazon CloudFront distribution as a target.

3) A company is using AWS CodeBuild to build an application. Company policy requires all build artifacts to be encrypted at rest. The company must limit access to the artifacts to IAM users in an operations IAM group that have permission to assume an operations IAM role.

Which solution will meet these requirements?

A) Add a post-build command to the CodeBuild build specification to push build objects to an Amazon S3 bucket. Set a bucket policy that prevents upload to the bucket unless the request includes the x-amz-server-side-encryption header. Add a Deny statement for all actions with a NotPrincipal element that references the operations IAM group.
B) Add a post-build command to the CodeBuild build specification to push build objects to an Amazon S3 bucket. Configure an S3 event notification to invoke an AWS Lambda function to get the object, encrypt the object, and put the object back into the S3 bucket with a tag key of Encrypted and a tag value of True. Set a bucket policy with a Deny statement for all actions with a NotPrincipal element that references the operations IAM group. Include in the policy a Condition element that references the Encrypted tag.
C) Add a post-build command to the CodeBuild build specification to push build objects to an Amazon S3 bucket that has S3 default encryption enabled. Set a bucket policy that contains a Deny statement for all actions with a NotPrincipal element that references the operations IAM role.
D) Add a post-build command to the CodeBuild build specification to call the AWS Key Management Service (AWS KMS) Encrypt API operation and pass the artifact to AWS KMS for encryption with a specified KMS key. Push the encrypted artifact to an Amazon S3 bucket. Set up the operations IAM group as the only user for the specified KMS key.

4) A DevOps engineer needs to implement a blue/green deployment process for an application on AWS. The DevOps engineer must gradually shift the traffic between the environments.

The application runs on Amazon EC2 instances behind an Application Load Balancer (ALB). The instances run in an Amazon EC2 Auto Scaling group. The application stores data on an Amazon RDS Multi-AZ DB instance. Amazon Route 53 provides external DNS.

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

A) Create a second Auto Scaling group behind the same ALB.
B) Create a second Auto Scaling group behind a second ALB.
C) In Route 53, create a second alias record that points to the new environment. Use a failover routing policy to choose between the two records.
D) In Route 53, create a second alias record that points to the new environment. Use a weighted routing policy to choose between the two records.
E) Configure the new EC2 instances to use the primary RDS DB instance.
F) Configure the new EC2 instances to use the standby RDS DB instance.

5) A company runs an application on Amazon EC2 instances that use the latest version of the Amazon Linux 2 AMI. When server administrators apply new security patches, the server administrators manually remove affected instances from service, patch the instances, and place the instances back into service.

A new security policy requires the company to apply security patches within 7 days after patches are released. The company's security team must verify that all the EC2 instances are compliant with this policy. The patching must occur during a time that has the least impact on users.

Which solution will automate compliance with these requirements?

A) Configure an AWS CodeBuild project to download and apply patches to all the instances over SSH. Use an Amazon EventBridge (Amazon CloudWatch Events) scheduled rule to run the CodeBuild project during a maintenance window.
B) Use AWS Systems Manager Patch Manager to create a patch baseline. Create a script on the EC2 instances to use the AWS CLI to pull the latest patches from Patch Manager. Create a cron job to schedule the script to run during a maintenance window.
C) Create a script to apply any available security patches. Create a cron job to schedule the script to run during a maintenance window. Install the script and cron job on the application AMI. Redeploy the application.
D) Enlist all the EC2 instances in an AWS Systems Manager Patch Manager patch group. Use Patch Manager to create a patch baseline. Configure a maintenance window to apply the patch baseline.

6) A company uses AWS CloudTrail on all its AWS accounts and sends all trails to a centralized Amazon S3 bucket. The company sends specified events to a third-party logging tool by using S3 event notifications and an AWS Lambda function.

The company has hired a security services provider to set up a security operations center. The security services provider wants to receive the CloudTrail logs through an Amazon Simple Queue Service (Amazon SQS) queue.

The company must continue to use S3 event notifications and the Lambda function to send events to the third-party logging tool.

What is the MOST operationally efficient way to meet these requirements?

A) Add an additional notification to the S3 bucket for all CreateObject events to send all objects to the SQS queue.
B) Replace the existing S3 event notification destination with an Amazon Simple Notification Service (Amazon SNS) topic. Subscribe the Lambda function and the SQS queue to the topic.
C) Replace the existing S3 event notification destination with an Amazon Kinesis data stream. Create consumers for the Lambda function and the SQS queue.
D) Configure the trail to send logs to Amazon CloudWatch Logs. Subscribe the SQS queue to the CloudWatch Logs log group.

7) A company is reviewing its AWS account security policies. The company has staff members in different countries and wants to monitor its AWS accounts for unusual behavior that is associated with an IAM identity.

The company wants to send a notification to any staff member for whom unusual activity is detected. The company also wants to send a notification to the user's team leader. An external messaging platform will send the notifications. The platform requires a target user-id for each recipient.

The company already has an API on AWS that the company can use to return the user-id of the staff member and the team leader from IAM user names. The company manages its AWS accounts by using AWS Organizations.

Which solution will meet these requirements?

A) Designate an account in the organization as the Amazon GuardDuty administrator. Add the company's AWS accounts as GuardDuty member accounts that are associated with the GuardDuty administrator account. Create an AWS Lambda function to perform the user-id lookup and to send notifications to the external messaging platform. Create an Amazon EventBridge (Amazon CloudWatch Events) rule in the GuardDuty administrator account to match the Impact:IAMUser/AnomalousBehavior notification type and invoke the Lambda function.

B) Designate an account in the organization as the Amazon Detective administrator. Add the company's AWS accounts as Detective member accounts that are associated with the Detective administrator account. Create an AWS Lambda function to perform the user-id lookup and to send notifications to the external messaging platform. Create an Amazon EventBridge (Amazon CloudWatch Events) rule in the Detective administrator account to match the Impact:IAMUser/AnomalousBehavior notification type and invoke the Lambda function.

C) Designate an account in the organization as the Amazon GuardDuty administrator. Add the company's AWS accounts as GuardDuty member accounts that are associated with the GuardDuty administrator account. Create an AWS Lambda function to perform the user-id lookup and to send notifications to the external messaging platform. Create an Amazon Simple Notification Service (Amazon SNS) topic in the GuardDuty administrator account to match the Impact:IAMUser/AnomalousBehavior notification type and invoke the Lambda function.

D) Designate an account in the organization as the Amazon Detective administrator. Add the company's AWS accounts as Detective member accounts that are associated with the Detective administrator account. Create an AWS Lambda function to perform the user-id lookup and to send notifications to the external messaging platform. Create an Amazon Simple Notification Service (Amazon SNS) topic in the Detective administrator account to match the Impact:IAMUser/AnomalousBehavior notification type and invoke the Lambda function.

8) A development team is designing an application that has a large customer base spread across three AWS Regions. The application will use an Amazon DynamoDB table that must be available in all three Regions to deliver low-latency data access. When the table is updated in one Region, the changes must seamlessly propagate to the other Regions.

How should a DevOps engineer configure the table to meet these requirements with the LEAST operational overhead?

A) Create a DynamoDB table in each of the three Regions. Give each table the same name.
B) Configure three DynamoDB tables in each of the three Regions. Use the AWS SDK for DynamoDB to synchronize data changes among the tables.
C) Configure a multi-Region, multi-active DynamoDB global table that includes the three Regions.
D) Use DynamoDB global tables to configure a primary table in one Region and a read replica in each of the other Regions.

9) A company has a legacy API that runs on a fleet of Amazon EC2 instances behind a public Application Load Balancer (ALB). The ALB has access logging enabled and stores the access logs in Amazon S3. The API is available through the hostname api.example.com. The company uses Amazon Route 53 to manage the hostname.

Developers have rebuilt five of the API endpoints by using a different AWS Lambda function for each endpoint. A DevOps engineer wants to test the new versions of the Lambda functions with a limited number of random customers. To ensure compatibility with an existing log processing service, the test must not affect the ALB access logs.

How should the DevOps engineer perform the test to meet these requirements?

A) Add the five Lambda functions as targets to the existing target group for the EC2 instances. Set the weight in the target group of each Lambda function target to be less than the EC2 instance targets. Amend the default rule on the ALB to enable target group-level stickiness.
B) Create a single target group that includes all the Lambda functions as individual targets. On the ALB, create a new listener rule that includes a host header condition that matches the API endpoint's hostname. Add the target group to the listener rule. Specify a lower weight for the new target group than the weight of the default rule's target group.
C) Create a new ALB and a new target group for each Lambda function. Create a new listener rule that includes a host header condition that matches each of the endpoints and forwards traffic to the target groups. Create a new Route 53 alias record with a weight of 10. Update the existing Route 53 record for the api.example.com hostname with a weight of 90.
D) Create a new target group for each Lambda function. On the ALB, create new listener rules that include a path condition that matches each of the different endpoints. Set the rules to be weighted between the Lambda function target group for that endpoint and the instance-based target group.

10) A DevOps engineer is managing a legacy application on AWS. The application is a monolithic Windows program that runs on a single Amazon EC2 instance. The source code for the application is not available, so the application cannot be modified.

The application has a memory leak and malfunctions when memory utilization on the EC2 instance increases to more than 90%. The DevOps engineer has configured the unified Amazon CloudWatch agent on the EC2 instance to collect the operation system's memory utilization metrics.

The DevOps engineer needs to implement a solution to prevent the application from malfunctioning.

Which combination of steps will meet these requirements with the MOST operational efficiency? (Select TWO.)

A) Create an Amazon EventBridge (Amazon CloudWatch Events) rule that publishes to an Amazon Simple Notification Service (Amazon SNS) topic when memory utilization increases to more than 80%.
B) Create a metric filter on memory utilization in Amazon CloudWatch Logs. Create a CloudWatch alarm on the memory utilization filter. Configure the alarm to publish to an Amazon Simple Notification Service (Amazon SNS) topic when the memory utilization increases to more than 80%.
C) Create a CloudWatch alarm on the memory utilization metric. Configure the alarm to publish to an Amazon Simple Notification Service (Amazon SNS) topic when the memory utilization increases to more than 80%.
D) Configure an AWS Lambda function to restart the application by using AWS Systems Manager Run Command. Subscribe the Lambda function to the Amazon Simple Notification Service (Amazon SNS) topic.
E) Configure the EC2 instance to run a script that restarts the application. Subscribe the EC2 instance to the Amazon Simple Notification Service (Amazon SNS) topic.

Answers

1) B – Option B is the recommended solution and is the most responsive of the given options. The change will directly produce the event, and the event will directly start the pipeline. The periodic checks in option A will work, but they will not launch the pipeline until the next periodic check occurs. Option C is not a feature that AWS CodeCommit supports. Option D is not a valid method to start the pipeline.

2) B, C, F – There are three parts to this question. Part 1 is how to deploy the workload to multiple AWS Regions. Part 2 is how to provide access from a single endpoint for multiple deployments. Part 3 is how to handle failover events by using DNS.

Part 1: For option B, AWS CloudFormation StackSets provides an operationally efficient multi-Region deployment strategy for the Region-specific Amazon S3 buckets. S3 replication copies new and existing objects in the primary Region to multiple deployment Regions. Option A is incorrect because it is less operationally efficient. The AWS Lambda function is unnecessary because S3 replication can provide the appropriate functionality without custom code.

Part 2: For option C, you can use an Amazon CloudFront distribution to make a single endpoint available to resolve to multiple origins. You can configure CloudFront custom origins to create high availability origin failover that requires a shorter connection timeout, fewer connection attempts, or both. Option D is feasible but is less operationally efficient because it involves custom code within the Lambda@Edge function. The custom code is unnecessary because of native handling within the origin configurations.

Part 3: For option F, because the CloudFront origin configurations are handling the failover, Route 53 is providing a simple routing policy user-friendly domain name to the CloudFront distribution. Option E is incorrect because there are not multiple records to benefit from failover routing.

3) C – Amazon S3 default encryption ensures that the artifacts are encrypted at rest. The Deny statement with the NotPrincipal element set to the operations IAM role will deny access to the S3 bucket except for requests that use the role. The scenario implies that the operations role has a permissions policy that allows access to the bucket.

Options A and B are incorrect because the bucket policy is referencing the IAM group and not the role. Option A is also incorrect because AWS recommends the use of default encryption over a bucket policy to enforce encryption. Option B also allows artifacts to be stored at rest briefly without encryption. Option D is incorrect because AWS Key Management Service (AWS KMS) Encrypt API operations would be useful for encryption of plaintext values such as a password, but not for encryption of a build artifact file, archive, or object.

4) B, D, E – A blue/green deployment consists of two separate environments. The blue environment contains Amazon EC2 instances in an Auto Scaling group that run the current production version of the application. The green environment contains EC2 instances in another Auto Scaling group that run the new version of the application. Each Auto Scaling group is behind its own Application Load Balancer (ALB), so you can configure two alias records as endpoints in Amazon Route 53 and use a weighted routing policy to gradually shift traffic from the ALB for the blue environment to the ALB for the green environment. Unless schema changes are necessary for the new release, it is best to point both environments to the same database so that the data remains consistent during the cutover.

Option A is incorrect because you need two ALBs as endpoints so that you can use Route 53 to gradually shift the traffic. Option C is incorrect because a failover routing policy sends all traffic to a single endpoint unless a health check detects a failure. Therefore, this option cannot gradually shift the traffic. Option F is incorrect because the standby instance in an Amazon RDS Multi-AZ DB instance is a hot standby and is not available for reads or writes.

5) D – Patch Manager, a capability of AWS Systems Manager, will automatically apply security patches during a maintenance window according to a list of approved patches that you define in a patch baseline. The company's security team can view the patch compliance of the instances in the Systems Manager console or pull a summary by using the AWS CLI.

Option A is incorrect because AWS CodeBuild builds your source code into artifacts. CodeBuild does not deploy patches to instances. Option B is incorrect because you do not need to schedule the Amazon Linux 2 preinstalled Systems Manager Agent (SSM Agent) to pull the patches. You only need to associate the patching configuration with a Systems Manager maintenance window. Option C is incorrect because it does not include a way for the security team to verify patch compliance. This option also includes a single point of failure in the cron job.

6) B – To implement a fanout messaging scenario for Amazon S3 event notifications of one event to many consumers, you can move the S3 event notification destination to an Amazon Simple Notification Service (Amazon SNS) topic. You can subscribe multiple consumers, such as the AWS Lambda function and the Amazon Simple Queue Service (Amazon SQS) queue, to the topic without changing the Lambda function code.

Option A is not possible and will result in a "Configuration is ambiguously defined" error because of overlapping notification event prefixes and suffixes. Option C is incorrect because Amazon Kinesis Data Streams is not a valid S3 event notification destination. Option D is incorrect because it is an incomplete solution. You cannot subscribe an SQS queue directly to an Amazon CloudWatch Logs log group.

7) A – Amazon GuardDuty findings are published to Amazon EventBridge (Amazon CloudWatch Events) as events that can invoke an AWS Lambda function target. Option B is incorrect because Amazon Detective will not by itself detect unusual activity. Detective provides analysis information related to a given finding. Option C is incorrect because Amazon Simple Notification Service (Amazon SNS) can filter messages by attributes and not by message contents. An EventBridge (CloudWatch Events) rule would be required to publish to the SNS topic. Option D is incorrect because of the same reasons that options B and C are incorrect.

8) C – Amazon DynamoDB global tables start as single-Region tables that you can make available for multi-Region and multi-active workloads. Global tables provide Region-specific workloads with low-latency data access without requiring you to configure or manage a replication solution.

Option A is incorrect because the use of a separate table in each Region would require an additional replication solution. Option B is incorrect because the development and management of a synchronization process across the tables would be unnecessary operational overhead. Option D is incorrect because global tables are multi-Region, multi-active tables that do not have read replicas.

9) D – This scenario is similar to a blue/green deployment and a canary deployment. Only the existing Application Load Balancer (ALB) is required for this solution. Target groups support a single AWS Lambda function as a registered target. Therefore, this solution requires five target groups, one for each endpoint. With each endpoint having its own path, new path conditions are needed in the listener rules to facilitate the weighted distribution of requests across the existing EC2 target group and the new Lambda function target groups.

Options A and B are incorrect because you cannot register multiple Lambda functions to a single target group. Option A is also incorrect because target group-level stickiness would negate the benefit of weighted routing for limited testing. Option B is also incorrect because weighted rules are assigned at the individual rule level and are not evaluated across multiple rules. Option C is incorrect because it would affect the ALB access logs by generating different access logs based on the new load balancer ID. Additionally, listener rules that include a host header condition would not be effective for URI level testing.

10) C, D – Option C is correct because the unified Amazon CloudWatch agent publishes system-level metrics as CloudWatch metrics that can be used directly for alarms. Option D is correct because an AWS Lambda function that is subscribed to an Amazon Simple Notification Service (Amazon SNS) topic can send commands to Amazon EC2 instances by using the AWS API or SDK to initiate AWS Systems Manager Run Command.

Options A and B are incorrect because they are not as operationally efficient as the correct answers. In option A, an Amazon EventBridge (Amazon CloudWatch Events) rule is not necessary because the CloudWatch alarm can directly notify Amazon SNS. Option B is incorrect because the published system-level CloudWatch metrics negate the need for a CloudWatch Logs based metric filter to generate metrics for an alarm. Option E is incorrect because you cannot subscribe an EC2 instance to an SNS topic.