Amazon-Web-Services

Exam Questions DOP-C02

AWS Certified DevOps Engineer - Professional

NEW QUESTION 1
A company runs applications in AWS accounts that are in an organization in AWS Organizations The applications use Amazon EC2 instances and Amazon S3. The company wants to detect potentially compromised EC2 instances suspicious network activity and unusual API activity in its existing AWS accounts and in any AWS accounts that the company creates in the future When the company detects one to these events the company wants to use an existing Amazon Simple Notification Service (Amazon SNS) topic to send a notification to its operational support team for investigation and remediation.
Which solution will meet these requirements in accordance with AWS best practices?

A. In the organization's management account configure an AWS account as the AmazonGuardDuty administrator accoun
B. In the GuardDuty administrator account add the company's existing AWS accounts to GuardDuty as members In the GuardDuty administrator account create an Amazon EventBridge rule with an event pattern to match GuardDuty events and to forward matching events to the SNS topic.
C. In the organization's management account configure Amazon GuardDuty to add newly created AWS accounts by invitation and to send invitations to the existing AWS accounts Create an AWS Cloud Formation stack set that accepts the GuardDuty invitation and creates an Amazon EventBridge rule Configure the rule with an event pattern to matc
D. GuardDuty events and to forward matching events to the SNS topi
E. Configure the Cloud Formation stack set to deploy into all AWS accounts in the organization.
F. In the organization's management accoun
G. create an AWS CloudTrail organization trail Activate the organization trail in all AWS accounts in the organizatio
H. Create an SCP that enables VPC Flow Logs in each account in the organizatio
I. Configure AWS Security Hub for the organization Create an Amazon EventBridge rule with an even pattern to match Security Hub events and to forward matching events to the SNS topic.
J. In the organization's management account configure an AWS account as the AWS CloudTrail administrator account in the CloudTrail administrator account create a CloudTrail organization trai
K. Add the company's existing AWS accounts to the organization trail Create an SCP that enables VPC Flow Logs in each account in the organizatio
L. Configure AWS Security Hub for the organizatio
M. Create an Amazon EventBridge rule with an event pattern to match Security Hub events and to forward matching events to the SNS topic.

Answer: B

Explanation:
It allows the company to detect potentially compromised EC2 instances, suspicious network activity, and unusual API activity in its existing AWS accounts and in any AWS accounts that the company creates in the future using Amazon GuardDuty. It also provides a solution for automatically adding future AWS accounts to GuardDuty by configuring GuardDuty to add newly created AWS accounts by invitation and to send invitations to the existing AWS accounts.

NEW QUESTION 2
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 Regio
B. Create an AWS CodeBuild project to perform a Git mirror operation of the primary Region's CodeCommit repository to the secondary Region's CodeCommit repositor
C. Create an AWS Lambda function that invokes the CodeBuild projec
D. Create an Amazon EventBridge rule that reacts to merge events in the primary Region's CodeCommit repositor
E. Configure the EventBridge rule to invoke the Lambda function.
F. Create an Amazon S3 bucket in the secondary Regio
G. 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 bucke
H. Create an AWS Lambda function that initiates the Fargate tas
I. Create an Amazon EventBridge rule that reacts to merge events in the CodeCommitrepositor
J. Configure the EventBridge rule to invoke the Lambda function.
K. Create an AWS CodeArtifact repository in the secondary Regio
L. Create an AWS CodePipeline pipeline that uses the primary Region's CodeCommit repository for the source actio
M. 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.
N. Create an AWS Cloud9 environment and a CodeCommit repository in the secondary Regio
O. Configure the primary Region's CodeCommit repository as a remote repository in the AWS Cloud9 environmen
P. Connect the secondary Region's CodeCommit repository to the AWS Cloud9 environment.

Answer: A

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 a CodeCommit repository in the secondary Region and setting up an AWS CodeBuild project to perform a Git mirror operation of the primary Region’s CodeCommit repository to the secondary Region’s repository. An AWS Lambda function is then created to invoke the CodeBuild project. Additionally, an Amazon EventBridge rule is 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.
References:
? AWS CodeCommit User Guide on resilience and disaster recovery1.
? AWS Documentation on monitoring CodeCommit events in Amazon EventBridge and Amazon CloudWatch Events2.

NEW QUESTION 3
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 resource
B. Deploy the template to all accounts from the organization's management account by using CloudFormation StackSet
C. Set the stack policy to deny Update:Delete actions.
D. Enable AWS Control Towe
E. Enroll the existing accounts in AWS Control Towe
F. Grant the individual account administrators access to CloudTrail and AWS Config.
G. Designate an AWS Config management accoun
H. Create AWS Config recorders in all accounts by using AWS CloudFormation StackSet
I. Deploy AWS Config rules to the organization by using the AWS Config management accoun
J. Create a CloudTrail organization trail in the organization's management accoun
K. Deny modification or deletion of the AWS Config recorders by using an SCP.
L. Create an AWS CloudFormation template that defines the standard account resource
M. 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.

Answer: D

NEW QUESTION 4
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 Region
B. Attach the policy to the root of the organization.
C. Configure AWS CloudTrail to send logs to Amazon CloudWatch Logs and enable it for all Region
D. Use a CloudWatch Logs metric filter to send an alert on any service activity in non-US Regions.
E. Use an AWS Lambda function that checks for AWS service activity and deploy it to all Region
F. Write an Amazon EventBridge rule that runs the Lambda function every hour, sending an alert if activity is found in a non-US Region.
G. 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.
H. Write an SCP using the aws: RequestedRegion condition key limiting access to US Region
I. Apply the policy to all users, groups, and roles

Answer: AB

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.
References:
? 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.

NEW QUESTION 5
An application runs on Amazon EC2 instances behind an Application Load Balancer (ALB). A DevOps engineer is using AWS CodeDeploy to release a new version. The deployment fails during the AllowTraffic lifecycle event, but a cause for the failure is not indicated in the deployment logs.
What would cause this?

A. The appspe
B. yml file contains an invalid script that runs in the AllowTraffic lifecycle hook.
C. The user who initiated the deployment does not have the necessary permissions tointeract with the ALB.
D. The health checks specified for the ALB target group are misconfigured.
E. The CodeDeploy agent was not installed in the EC2 instances that are pad of the ALB target group.

Answer: C

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

NEW QUESTION 6
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 metri
B. Use the recover action to stop and start the instanc
C. Use an S3 event notification to push the metadata to the instance when the instance is back up and running.
D. Configure AWS OpsWorks, and use the auto healing feature to stop and start the instanc
E. Use a lifecycle event in OpsWorks to pull the metadata from Amazon S3 and update it on the instance.
F. Use EC2 Auto Recovery to automatically stop and start the instance in case of a failur
G. Use an S3 event notification to push the metadata to the instance when the instance is back up and running.
H. Use AWS CloudFormation to create an EC2 instance that includes the UserData property for the EC2 resourc
I. Add a command in UserData to retrieve the application metadata from Amazon S3.

Answer: B

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

NEW QUESTION 7
A company uses Amazon S3 to store proprietary information. The development team creates buckets for new projects on a daily basis. The security team wants to ensure that all existing and future buckets have encryption logging and versioning enabled. Additionally, no buckets should ever be publicly read or write accessible.
What should a DevOps engineer do to meet these requirements?

A. Enable AWS CloudTrail and configure automatic remediation using AWS Lambda.
B. Enable AWS Conflg rules and configure automatic remediation using AWS Systems Manager documents.
C. Enable AWS Trusted Advisor and configure automatic remediation using Amazon EventBridge.
D. Enable AWS Systems Manager and configure automatic remediation using Systems Manager documents.

Answer: B

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

NEW QUESTION 8
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.

Answer: A

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

NEW QUESTION 9
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 EC
B. Amazon EC2, and Lambda as deploy providers.
C. Use AWS CodePipeline with AWS CodeDeploy as the deploy provider.
D. Use AWS CodePipeline with AWS Elastic Beanstalk as the deploy provider.
E. Use AWS CodeDeploy with GitHub integration to deploy the application.

Answer: B

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

NEW QUESTION 10
A company has an application and a CI/CD pipeline. The CI/CD pipeline consists of an AWS CodePipeline pipeline and an AWS CodeBuild project. The CodeBuild project runs tests against the application as part of the build process and outputs a test report. The company must keep the test reports for 90 days.
Which solution will meet these requirements?

A. Add a new stage in the CodePipeline pipeline after the stage that contains the CodeBuild projec
B. Create an Amazon S3 bucket to store the report
C. Configure an S3 deploy action type in the new CodePipeline stage with the appropriate path and format for the reports.
D. Add a report group in the CodeBuild project buildspec file with the appropriate path and format for the report
E. Create an Amazon S3 bucket to store the report
F. Configure an Amazon EventBridge rule that invokes an AWS Lambda function to copy the reports to the S3 bucket when a build is complete
G. Create an S3 Lifecycle rule to expire the objects after 90 days.
H. Add a new stage in the CodePipeline pipelin
I. Configure a test action type with the appropriate path and format for the report
J. Configure the report expiration time to be 90 days in the CodeBuild project buildspec file.
K. Add a report group in the CodeBuild project buildspec file with the appropriate path and format for the report
L. Create an Amazon S3 bucket to store the report
M. Configure the report group as an artifact in the CodeBuild project buildspec fil
N. Configure the S3 bucket as the artifact destinatio
O. Set the object expiration to 90 days.

Answer: B

Explanation:
The correct solution is to add a report group in the AWS CodeBuild project buildspec file with the appropriate path and format for the reports. Then, create an Amazon S3 bucket to store the reports. You should configure an Amazon EventBridge rule that invokes an AWS Lambda function to copy the reports to the S3 bucket when a build is completed. Finally, create an S3 Lifecycle rule to expire the objects after 90 days. This approach allows for the automated transfer of reports to long-term storage and ensures
they are retained for the required duration without manual intervention1. References:
? AWS CodeBuild User Guide on test reporting1.
? AWS CodeBuild User Guide on working with report groups2.
? AWS Documentation on using AWS CodePipeline with AWS CodeBuild3.

NEW QUESTION 10
A company has an organization in AWS Organizations. The organization includes workload accounts that contain enterprise applications. The company centrally manages users from an operations account. No users can be created in the workload accounts. The company recently added an operations team and must provide the operations team members with administrator access to each workload account.
Which combination of actions will provide this access? (Choose three.)

A. Create a SysAdmin role in the operations accoun
B. Attach the AdministratorAccess policy to the rol
C. Modify the trust relationship to allow the sts:AssumeRole action from the workload accounts.
D. Create a SysAdmin role in each workload accoun
E. Attach the AdministratorAccess policy to the rol
F. Modify the trust relationship to allow the sts:AssumeRole action from the operations account.
G. Create an Amazon Cognito identity pool in the operations accoun
H. Attach the SysAdmin role as an authenticated role.
I. In the operations account, create an IAM user for each operations team member.
J. In the operations account, create an IAM user group that is named SysAdmin
K. Add an IAM policy that allows the sts:AssumeRole action for the SysAdmin role in each workload accoun
L. Add all operations team members to the group.
M. Create an Amazon Cognito user pool in the operations accoun
N. Create an Amazon Cognito user for each operations team member.

Answer: BDE

Explanation:
https://docs.aws.amazon.com/IAM/latest/UserGuide/tutorial_cross-account- with-roles.html

NEW QUESTION 14
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 Complianc
B. Use calls to the put- compliance-items API action to scan and build a database of noncompliant EC2 instances based on their host placement configuratio
C. Use an Amazon DynamoDB table to store these instance IDs for fast acces
D. Generate a report through Systems Manager by calling the list-compliance-summaries API action.
E. Use custom Java code running on an EC2 instanc
F. Set up EC2 Auto Scaling for the instance depending on the number of instances to be checke
G. Send the list of noncompliant EC2 instance IDs to an Amazon SQS queu
H. Set up another worker instance to process instance IDs from the SQS queue and write them to Amazon DynamoD
I. 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.
J. Use AWS Confi
K. Identify all EC2 instances to be audited by enabling Config Recording on all Amazon EC2 resources for the regio
L. Create a custom AWS Config rule that triggers an AWS Lambda function by using the "config-rule-change-triggered" blueprint.Modify the LambdaevaluateCompliance () function to verify host placement to return a NON_COMPLIANT result if the instance is not running on an EC2 Dedicated Hos
M. Use the AWS Config report to address noncompliant instances.
N. Use AWS CloudTrai
O. Identify all EC2 instances to be audited by analyzing all calls to the EC2 RunCommand API actio
P. Invoke a AWS Lambda function that analyzes the host placement of the instanc
Q. Store the EC2 instance ID of noncompliant resources in an Amazon RDS for MySQL DB instanc
R. Generate a report by querying the RDS instance and exporting the query results to a CSV text file.

Answer: C

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.

NEW QUESTION 15
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 cal
B. Configure the rule to invoke an AWS Lambda function that enables trusted access to AWS Config for the organization.
C. In the organization's management account, create an AWS CloudFormation stack set to enable AWS Confi
D. Configure the stack set to deploy automatically when an account is created through Organizations.
E. In the organization's management account, create an SCP that allows the appropriate AWS Config API calls to enable AWS Confi
F. Apply the SCP to the root-level OU.
G. In the organization's management account, create an Amazon EventBridge rule that reacts to a CreateAccount API cal
H. Configure the rule to invoke an AWS Systems Manager Automation runbook to enable AWS Config for the account.

Answer: B

Explanation:
https://aws.amazon.com/about-aws/whats-new/2020/02/aws-cloudformation- stacksets-introduces-automatic-deployments-across-accounts-and-regions-through-aws- organizations/

NEW QUESTION 17
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 instance
B. Create route table entries for the IPv6 network, use EC2 instance types that support IPv6, and assign IPv6 addresses to each EC2 instance.
C. Assign each EC2 instance an IPv6 Elastic IP addres
D. Create a target group, and add the EC2 instances as target
E. Create a listener on port 443 of the ALB, and associate the target group with the ALB.
F. 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.
G. Add an IPv6 CIDR block to the VPC and subnets for the AL
H. Create a listener on port 443. and specify the dualstack IP address type on the AL
I. Create a target group, and add the EC2 instances as target
J. Associate the target group with the ALB.

Answer: D

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.

NEW QUESTION 22
A company deploys its corporate infrastructure on AWS across multiple AWS Regions and Availability Zones. The infrastructure is deployed on Amazon EC2 instances and connects with AWS loT Greengrass devices. The company deploys additional resources on on- premises servers that are located in the corporate headquarters.
The company wants to reduce the overhead involved in maintaining and updating its resources. The company's DevOps team plans to use AWS Systems Manager to implement automated management and application of patches. The DevOps team confirms that Systems Manager is available in the Regions that the resources are deployed m Systems Manager also is available in a Region near the corporate headquarters.
Which combination of steps must the DevOps team take to implement automated patch and configuration management across the company's EC2 instances loT devices and on- premises infrastructure? (Select THREE.)

A. Apply tags lo all the EC2 instance
B. AWS loT Greengrass devices, and on-premises server
C. Use Systems Manager Session Manager to push patches to all the tagged devices.
D. Use Systems Manager Run Command to schedule patching for the EC2 instances AWS loT Greengrass devices and on-premises servers.
E. Use Systems Manager Patch Manager to schedule patching loT the EC2 instances AWS loT Greengrass devices and on-premises servers as a Systems Manager maintenance window task.
F. Configure Amazon EventBridge to monitor Systems Manager Patch Manager for updates to patch baseline
G. Associate Systems Manager Run Command with the event lo initiate a patch action for all EC2 instances AWS loT Greengrass devices and on-premises servers.
H. Create an IAM instance profile for Systems Manager Attach the instance profile to all the EC2 instances in the AWS accoun
I. For the AWS loT Greengrass devices and on-premises servers create an IAM service role for Systems Manager.
J. Generate a managed-instance activation Use the Activation Code and Activation ID to install Systems Manager Agent (SSM Agent) on each server in the on-premises environment Update the AWS loT Greengrass IAM token exchange role Use the role to deploy SSM Agent on all the loT devices.

Answer: CEF

Explanation:
https://aws.amazon.com/blogs/mt/how-to-centrally-manage-aws-iot-greengrass-devices-using-aws-systems-manager/?force_isolation=true

NEW QUESTION 23
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 AR
B. create a new assertion safety rul
C. Apply the assertion safety rule to the two routing control
D. Configure the rule with the ATLEAST type with a threshold of 1.
E. In Route 53 ARC, create a new gating safety rul
F. Apply the assertion safety rule to the two routing control
G. Configure the rule with the OR type with a threshold of 1.
H. In Route 53 ARC, create a new resource se
I. Configure the resource set with an AWS: Route53: HealthCheck resource typ
J. Specify the ARNs of the two routing controls as the target resourc
K. Create a new readiness check for the resource set.
L. In Route 53 ARC, create a new resource se
M. Configure the resource set with an AWS: Route53RecoveryReadiness: DNSTargetResource resource typ
N. Add the domain names of the two Route 53 alias DNS records as the target resourc
O. Create a new readiness check for the resource set.

Answer: A

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. References:
? 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

NEW QUESTION 25
A DevOps engineer needs to back up sensitive Amazon S3 objects that are stored within an S3 bucket with a private bucket policy using S3 cross-Region replication functionality. The objects need to be copied to a target bucket in a different AWS Region and account.
Which combination of actions should be performed to enable this replication? (Choose three.)

A. Create a replication IAM role in the source account
B. Create a replication I AM role in the target account.
C. Add statements to the source bucket policy allowing the replication IAM role to replicate objects.
D. Add statements to the target bucket policy allowing the replication IAM role to replicate objects.
E. Create a replication rule in the source bucket to enable the replication.
F. Create a replication rule in the target bucket to enable the replication.

Answer: ADE

Explanation:
S3 cross-Region replication (CRR) automatically replicates data between buckets across different AWS Regions. To enable CRR, you need to add a replication configuration to your source bucket that specifies the destination bucket, the IAM role, and the encryption type (optional). You also need to grant permissions to the IAM role to perform replication actions on both the source and destination buckets. Additionally, you can choose the destination storage class and enable additional replication options such as S3 Replication Time Control (S3 RTC) or S3 Batch Replication. https://medium.com/cloud-

techies/s3-same-region-replication-srr-and-cross-region-replication-crr-34d446806bab https://aws.amazon.com/getting-started/hands-on/replicate-data-using-amazon-s3- replication/ https://docs.aws.amazon.com/AmazonS3/latest/userguide/replication.html

NEW QUESTION 30
A DevOps team manages an API running on-premises that serves as a backend for an Amazon API Gateway endpoint. Customers have been complaining about high response latencies, which the development team has verified using the API Gateway latency metrics in Amazon CloudWatch. To identify the cause, the team needs to collect relevant data without introducing additional latency.
Which actions should be taken to accomplish this? (Choose two.)

A. Install the CloudWatch agent server side and configure the agent to upload relevant logs to CloudWatch.
B. Enable AWS X-Ray tracing in API Gateway, modify the application to capture request segments, and upload those segments to X-Ray during each request.
C. Enable AWS X-Ray tracing in API Gateway, modify the application to capture request segments, and use the X-Ray daemon to upload segments to X-Ray.
D. Modify the on-premises application to send log information back to API Gateway with each request.
E. Modify the on-premises application to calculate and upload statistical data relevant to the API service requests to CloudWatch metrics.

Answer: AC

Explanation:
https://docs.aws.amazon.com/AmazonCloudWatch/latest/monitoring/install-CloudWatch-Agent-on-premise.html
https://docs.aws.amazon.com/xray/latest/devguide/xray-api-sendingdata.html

NEW QUESTION 33
An Amazon EC2 instance is running in a VPC and needs to download an object from a restricted Amazon S3 bucket. When the DevOps engineer tries to download the object, an AccessDenied error is received,
What are the possible causes tor this error? (Select TWO,)

A. The 53 bucket default encryption is enabled.
B. There is an error in the S3 bucket policy.
C. The object has been moved to S3 Glacier.
D. There is an error in the IAM role configuration.
E. S3 Versioning is enabled.

Answer: BD

Explanation:
These are the possible causes for the AccessDenied error because they affect the permissions to access the S3 object from the EC2 instance. An S3 bucket policy is a resource-based policy that defines who can access the bucket and its objects, and what actions they can perform. An IAM role is an identity that can be assumed by an EC2 instance to grant it permissions to access AWS services and resources. If there is an error in the S3 bucket policy or the IAM role configuration, such as a missing or incorrect statement, condition, or principal, then the EC2 instance may not have the necessary permissions to download the object from the S3 bucket . https://docs.aws.amazon.com/AmazonS3/latest/userguide/example-bucket-policies.html
https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/iam-roles-for-amazon-ec2.html

NEW QUESTION 35
A DevOps team uses AWS CodePipeline, AWS CodeBuild, and AWS CodeDeploy to deploy an application. The application is a REST API that uses AWS Lambda functions and Amazon API Gateway Recent deployments have introduced errors that have affected many customers.
The DevOps team needs a solution that reverts to the most recent stable version of the application when an error is detected. The solution must affect the fewest customers possible.
Which solution Will meet these requirements With the MOST operational efficiency?

A. Set the deployment configuration in CodeDeploy to LambdaAllAtOnce Configure automatic rollbacks on the deployment group Create an Amazon CloudWatch alarm that detects HTTP Bad Gateway errors on API Gateway Configure the deployment group to roll back when the number of alarms meets the alarm threshold
B. Set the deployment configuration in CodeDeploy to LambdaCanary10Percent10Minute
C. Configure automatic rollbacks on the deployment group Create an Amazon CloudWatch alarm that detects HTTP Bad Gateway errors on API Gateway Configure the deployment group to roll back when the number of alarms meets the alarm threshold
D. Set the deployment configuration in CodeDeploy to LambdaAllAtOnce Configure manual rollbacks on the deployment grou
E. Create an Amazon Simple Notification Service (Amazon SNS) topc to send notifications every time a deployrnent fad
F. Configure the SNS topc to Invoke a new Lambda function that stops the current deployment and starts the most recent successful deployment
G. Set the deployment configuration in CodeDeploy to LambdaCanaryIOPercentIOMinutes Configure manual rollbacks on the deployment group Create a metric filter on an Amazon CloudWatch log group for API Gateway to monitor HTTP Bad Gateway error
H. Configure the metric filter to Invoke a new Lambda function that stops the current eployment and starts the most recent successful deployment

Answer: B

Explanation:
? Option A is incorrect because setting the deployment configuration to LambdaAllAtOnce means that the new version of the application will be deployed to all Lambda functions at once, affecting all customers. This does not meet the requirement of affecting the fewest customers possible. Moreover, configuring automatic rollbacks on the deployment group is not operationally efficient, as it requires manual intervention to fix the errors and redeploy the application.
? Option B is correct because setting the deployment configuration to LambdaCanary10Percent10Minutes means that the new version of the application will be deployed to 10 percent of the Lambda functions first, and then to the remaining 90 percent after 10 minutes. This minimizes the impact of errors on customers, as only 10 percent of them will be affected by a faulty deployment. Configuring automatic rollbacks on the deployment group also meets the requirement of reverting to the most recent stable version of the application when an error is detected. Creating a CloudWatch alarm that detects HTTP Bad Gateway errors on API Gateway is a valid way to monitor the health of the application and trigger a rollback if needed.
? Option C is incorrect because setting the deployment configuration to LambdaAllAtOnce means that the new version of the application will be deployed to all Lambda functions at once, affecting all customers. This does not meet the requirement of affecting the fewest customers possible. Moreover, configuring manual rollbacks on the deployment group is not operationally efficient, as it requires human intervention to stop the current deployment and start a new one. Creating an SNS topic to send notifications every time a deployment fails is not sufficient to detect errors in the application, as it does not monitor the API Gateway responses.
? Option D is incorrect because configuring manual rollbacks on the deployment
group is not operationally efficient, as it requires human intervention to stop the current deployment and start a new one. Creating a metric filter on a CloudWatch log group for API Gateway to monitor HTTP Bad Gateway errors is a valid way to monitor the health of the application, but invoking a new Lambda function to perform a rollback is unnecessary and complex, as CodeDeploy already provides automatic rollback functionality.

References:
? AWS CodeDeploy Deployment Configurations
? [AWS CodeDeploy Rollbacks]
? Amazon CloudWatch Alarms

NEW QUESTION 40
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 application
B. Put each application's code in a different branc
C. Merge the branches, and use AWS CodeBuild to build the application
D. Use AWS CodeDeploy to deploy the applications to one centralized application server.
E. Create one AWS CodeCommit repository for each of the application
F. Use AWS CodeBuild to build the applications one at a tim
G. Use AWS CodeDeploy to deploy the applications to one centralized application server.
H. Create one AWS CodeCommit repository for each of the application
I. Use AWS CodeBuild to build the applications one at a time and to create one AMI for each serve
J. Use AWS CloudFormation StackSets to automatically provision and decommission Amazon EC2 fleets by using these AMIs.
K. Create one AWS CodeCommit repository for each of the application
L. 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.

Answer: D

Explanation:
because of "as few maintenance tasks as possible on the underlying infrastructure". Fargate does that better than "one centralized application server"

NEW QUESTION 41
A company needs to implement failover for its application. The application includes an Amazon CloudFront distribution and a public Application Load Balancer (ALB) in an AWS Region. The company has configured the ALB as the default origin for the distribution.
After some recent application outages, the company wants a zero-second RTO. The company deploys the application to a secondary Region in a warm standby configuration. A DevOps engineer needs to automate the failover of the application to the secondary Region so that HTTP GET requests meet the desired R TO.
Which solution will meet these requirements?

A. Create a second CloudFront distribution that has the secondary ALB as the default origi
B. Create Amazon Route 53 alias records that have a failover policy and Evaluate Target Health set to Yes for both CloudFront distribution
C. Update the application to use the new record set.
D. Create a new origin on the distribution for the secondary AL
E. Create a new origin grou
F. Set the original ALB as the primary origi
G. Configure the origin group to fail over for HTTP 5xx status code
H. Update the default behavior to use the origin group.
I. Create Amazon Route 53 alias records that have a failover policy and Evaluate Target Health set to Yes for both ALB
J. Set the TTL of both records to
K. Update the distribution's origin to use the new record set.
L. Create a CloudFront function that detects HTTP 5xx status code
M. Configure the function to return a 307 Temporary Redirect error response to the secondary ALB if the function detects 5xx status code
N. Update the distribution's default behavior to send origin responses to the function.

Answer: B

Explanation:
To implement failover for the application to the secondary Region so that HTTP GET requests meet the desired RTO, the DevOps engineer should use the following solution:
? Create a new origin on the distribution for the secondary ALB. A CloudFront origin
is the source of the content that CloudFront delivers to viewers. By creating a new origin for the secondary ALB, the DevOps engineer can configure CloudFront to route traffic to the secondary Region when the primary Region is unavailable1
? Create a new origin group. Set the original ALB as the primary origin. Configure
the origin group to fail over for HTTP 5xx status codes. An origin group is a logical grouping of two origins: a primary origin and a secondary origin. By creating an origin group, the DevOps engineer can specify which origin CloudFront should use as a fallback when the primary origin fails. The DevOps engineer can also define which HTTP status codes should trigger a failover from the primary origin to the secondary origin. By setting the original ALB as the primary origin and configuring the origin group to fail over for HTTP 5xx status codes, the DevOps engineer can ensure that CloudFront will switch to the secondary ALB when the primary ALB returns server errors2
? Update the default behavior to use the origin group. A behavior is a set of rules
that CloudFront applies when it receives requests for specific URLs or file types. The default behavior applies to all requests that do not match any other behaviors. By updating the default behavior to use the origin group, the DevOps engineer can enable failover routing for all requests that are sent to the distribution3
This solution will meet the requirements because it will automate the failover of the
application to the secondary Region with zero-second RTO. When CloudFront receives an HTTP GET request, it will first try to route it to the primary ALB in the primary Region. If the primary ALB is healthy and returns a successful response, CloudFront will deliver it to the viewer. If the primary ALB is unhealthy or returns an HTTP 5xx status code, CloudFront will automatically route the request to the secondary ALB in the secondary Region and deliver its response to the viewer.
The other options are not correct because they either do not provide zero-second RTO or do not work as expected. Creating a second CloudFront distribution that has the secondary ALB as the default origin and creating Amazon Route 53 alias records that have a failover policy is not a good option because it will introduce additional latency and complexity to the solution. Route 53 health checks and DNS propagation can take several minutes or longer, which means that viewers might experience delays or errors when accessing the application during a failover event. Creating Amazon Route 53 alias records that have a failover policy and

Evaluate Target Health set to Yes for both ALBs and setting the TTL of both records to O is not a valid option because it will not work with CloudFront distributions. Route 53 does not support health checks for alias records that point to CloudFront distributions, so it cannot detect if an ALB behind a distribution is healthy or not.
Creating a CloudFront function that detects HTTP 5xx status codes and returns a 307 Temporary Redirect error response to the secondary ALB is not a valid option because it will not provide zero-second RTO. A 307 Temporary Redirect error response tells viewers to retry their requests with a different URL, which means that viewers will have to make an additional request and wait for another response from CloudFront before reaching the secondary ALB.
References:
? 1: Adding, Editing, and Deleting Origins - Amazon CloudFront
? 2: Configuring Origin Failover - Amazon CloudFront
? 3: Creating or Updating a Cache Behavior - Amazon CloudFront

NEW QUESTION 46
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 ne
B. Network Load Balancer (NLB) and EC2 instances that use the shared private subnets Use the NLB DNS names for communication between microservices.
C. 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.
D. 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.
E. 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 organizatio
F. In each of the microservice VPC
G. 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.

Answer: B

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.

NEW QUESTION 49
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 accoun
B. 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.
C. Set up AWS Config in the accoun
D. Use a managed rule that returns a compliance failure for EC2::Volume resources that do not have a Backup Frequency tag applie
E. Configure a remediation action that uses a custom AWS Systems Manager Automation runbook to apply the Backup_Frequency tag with a value of weekly.
F. Turn on AWS CloudTrail in the accoun
G. Create an Amazon EventBridge rule that reacts to EBS CreateVolume event
H. Configure a custom AWS Systems Manager Automation runbook to apply the Backup_Frequency tag with a value of weekl
I. Specify the runbook as the target of the rule.
J. Turn on AWS CloudTrail in the accoun
K. Create an Amazon EventBridge rule that reacts to EBS CreateVolume events or EBS ModifyVolume event
L. Configure a custom AWS Systems Manager Automation runbook to apply the Backup_Frequency tag with a value of weekl
M. Specify the runbook as the target of the rule.

Answer: B

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.

NEW QUESTION 54
A company's security policies require the use of security hardened AMIS in production environments. A DevOps engineer has used EC2 Image Builder to create a pipeline that builds the AMIs on a recurring schedule.
The DevOps engineer needs to update the launch templates of the companys Auto Scaling groups. The Auto Scaling groups must use the newest AMIS during the launch of Amazon EC2 instances.
Which solution will meet these requirements with the MOST operational efficiency?

A. Configure an Amazon EventBridge rule to receive new AMI events from Image Builde
B. Target an AWS Systems Manager Run Command document that updates the launch templates of the Auto Scaling groups with the newest AMI ID.
C. Configure an Amazon EventBridge rule to receive new AMI events from Image Builde
D. Target an AWS Lambda function that updates the launch templates of the Auto Scaling groups with the newest AMI ID.
E. Configure the launch template to use a value from AWS Systems Manager Parameter Store for the AMI I
F. Configure the Image Builder pipeline to update the Parameter Store value with the newest AMI ID.
G. Configure the Image Builder distribution settings to update the launch templates with the newest AMI I
H. Configure the Auto Scaling groups to use the newest version of the launch template.

Answer: C

Explanation:
? The most operationally efficient solution is to use AWS Systems Manager Parameter Store1 to store the AMI ID and reference it in the launch template2. This way, the launch template does not need to be updated every time a new AMI is created by Image Builder. Instead, the Image Builder pipeline can update the Parameter Store value with the newest AMI ID3, and the Auto Scaling group can launch instances using the latest value from Parameter Store.
? The other solutions require updating the launch template or creating a new version of it every time a new AMI is created, which adds complexity and overhead. Additionally, using EventBridge rules and Lambda functions or Run Command documents introduces additional dependencies and potential points of failure.
References: 1: AWS Systems Manager Parameter Store 2: Using AWS Systems Manager parameters instead of AMI IDs in launch templates 3: Update an SSM parameter with
Image Builder

NEW QUESTION 56
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.

Answer: BCE

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.

NEW QUESTION 60
A company deploys updates to its Amazon API Gateway API several times a week by using an AWS CodePipeline pipeline. As part of the update process the company exports the JavaScript SDK for the API from the API. Gateway console and uploads the SDK to an Amazon S3 bucket
The company has configured an Amazon CloudFront distribution that uses the S3 bucket as an origin Web client then download the SDK by using the CloudFront distribution's endpoint. A DevOps engineer needs to implement a solution to make the new SDK available automatically during new API deployments.
Which solution will meet these requirements?

A. Create a CodePipeline action immediately after the deployment stage of the AP
B. Configure the action to invoke an AWS Lambda functio
C. Configure the Lambda function to download the SDK from API Gateway, upload the SDK to the S3 bucket and create a CloudFront invalidation for the SDK path.
D. Create a CodePipeline action immediately after the deployment stage of the API Configure the action to use the CodePipelme integration with AP
E. Gateway to export the SDK to Amazon S3 Create another action that uses the CodePipeline integration with Amazon S3 to invalidate the cache for the SDK path.
F. Create an Amazon EventBridge rule that reacts to UpdateStage events from aws apigateway Configure the rule to invoke an AWS Lambda function to download the SDK from API Gateway upload the SDK to the S3 bucket and call the CloudFront API to create an invalidation for the SDK path.
G. Create an Amazon EventBridge rule that reacts to Creat
H. Deployment events from aws apigatewa
I. Configure the rule to invoke an AWS Lambda function to download the SDK from AP
J. Gateway upload the SDK to the S3 bucket and call the S3 API to invalidate the cache for the SDK path.

Answer: A

Explanation:
This solution would allow the company to automate the process of updating the SDK and making it available to web clients. By adding a CodePipeline action

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

NEW QUESTION 63
A company runs its container workloads in AWS App Runner. A DevOps engineer manages the company's container repository in Amazon Elastic Container Registry (Amazon ECR).
The DevOps engineer must implement a solution that continuously monitors the container repository. The solution must create a new container image when the solution detects an operating system vulnerability or language package vulnerability.
Which solution will meet these requirements?

A. Use EC2 Image Builder to create a container image pipelin
B. Use Amazon ECR as the target repositor
C. Turn on enhanced scanning on the ECR repositor
D. Create an Amazon EventBridge rule to capture an Inspector2 finding even
E. Use the event to invoke the image pipelin
F. Re-upload the container to the repository.
G. Use EC2 Image Builder to create a container image pipelin
H. Use Amazon ECR as the target repositor
I. Enable Amazon GuardDuty Malware Protection on the container workloa
J. Create an Amazon EventBridge rule to capture a GuardDuty finding even
K. Use the event to invoke the image pipeline.
L. Create an AWS CodeBuild project to create a container imag
M. Use Amazon ECR as the target repositor
N. Turn on basic scanning on the repositor
O. Create an Amazon EventBridge rule to capture an ECR image action even
P. Use the event to invoke the CodeBuild projec
Q. Re-upload the container to the repository.
R. Create an AWS CodeBuild project to create a container imag
S. Use Amazon ECR as the target repositor
T. Configure AWS Systems Manager Compliance to scan all managed node
. Create an Amazon EventBridge rule to capture a configuration compliance state change even
. Use the event to invoke the CodeBuild project.

Answer: A

Explanation:
The solution that meets the requirements is to use EC2 Image Builder to create a container image pipeline, use Amazon ECR as the target repository, turn on enhanced scanning on the ECR repository, create an Amazon EventBridge rule to capture an Inspector2 finding event, and use the event to invoke the image pipeline. Re-upload the container to the repository.
This solution will continuously monitor the container repository for vulnerabilities using enhanced scanning, which is a feature of Amazon ECR that provides detailed information and guidance on how to fix security issues found in your container images. Enhanced scanning uses Inspector2, a security assessment service that integrates with Amazon ECR and generates findings for any vulnerabilities detected in your images. You can use Amazon EventBridge to create a rule that triggers an action when an Inspector2 finding event occurs. The action can be to invoke an EC2 Image Builder pipeline, which is a
service that automates the creation of container images. The pipeline can use the latest patches and updates to build a new container image and upload it to the same ECR repository, replacing the vulnerable image.
The other options are not correct because they do not meet all the requirements or use services that are not relevant for the scenario.
Option B is not correct because it uses Amazon GuardDuty Malware Protection, which is a feature of GuardDuty that detects malicious activity and unauthorized behavior on your AWS accounts and resources. GuardDuty does not scan container images for vulnerabilities, nor does it integrate with Amazon ECR or EC2 Image Builder.
Option C is not correct because it uses basic scanning on the ECR repository, which only provides a summary of the vulnerabilities found in your container images. Basic scanning does not use Inspector2 or generate findings that can be captured by Amazon EventBridge. Moreover, basic scanning does not provide guidance on how to fix the vulnerabilities.
Option D is not correct because it uses AWS Systems Manager Compliance, which is a feature of Systems Manager that helps you monitor and manage the compliance status of your AWS resources based on AWS Config rules and AWS Security Hub standards. Systems Manager Compliance does not scan container images for vulnerabilities, nor does it integrate with Amazon ECR or EC2 Image Builder.

NEW QUESTION 66
A company hosts applications in its AWS account Each application logs to an individual Amazon CloudWatch log group. The company's CloudWatch costs for ingestion are increasing
A DevOps engineer needs to Identify which applications are the source of the increased logging costs.
Which solution Will meet these requirements?

A. Use CloudWatch metrics to create a custom expression that Identifies the CloudWatch log groups that have the most data being written to them.
B. Use CloudWatch Logs Insights to create a set of queries for the application log groups to Identify the number of logs written for a period of time
C. Use AWS Cost Explorer to generate a cost report that details the cost for CloudWatch usage
D. Use AWS CloudTrail to filter for CreateLogStream events for each application

Answer: C

Explanation:
The correct answer is C.
A comprehensive and detailed explanation is:
? Option A is incorrect because using CloudWatch metrics to create a custom expression that identifies the CloudWatch log groups that have the most data being written to them is not a valid solution. CloudWatch metrics do not provide information about the size or volume of data being ingested by CloudWatch logs. CloudWatch metrics only provide information about the number of events, bytes, and errors that occur within a log group or stream. Moreover, creating a custom expression with CloudWatch metrics would require using the search_web tool, which is not necessary for this use case.
? Option B is incorrect because using CloudWatch Logs Insights to create a set of queries for the application log groups to identify the number of logs written for a period of time is not a valid solution. CloudWatch Logs Insights can help analyze and filter log events based on patterns and expressions, but it does not provide information about the cost or billing of CloudWatch logs. CloudWatch Logs Insights also charges based on the amount of data scanned by each query, which could increase the logging costs further.

? Option C is correct because using AWS Cost Explorer to generate a cost report that details the cost for CloudWatch usage is a valid solution. AWS Cost Explorer is a tool that helps visualize, understand, and manage AWS costs and usage over time. AWS Cost Explorer can generate custom reports that show the breakdown of costs by service, region, account, tag, or any other dimension. AWS Cost Explorer can also filter and group costs by usage type, which can help identify the specific CloudWatch log groups that are the source of the increased logging costs.
? Option D is incorrect because using AWS CloudTrail to filter for CreateLogStream events for each application is not a valid solution. AWS CloudTrail is a service that records API calls and account activity for AWS services, including CloudWatch logs. However, AWS CloudTrail does not provide information about the cost or billing of CloudWatch logs. Filtering for CreateLogStream events would only show when a new log stream was created within a log group, but not how much data was ingested or stored by that log stream.
References:
? CloudWatch Metrics
? CloudWatch Logs Insights
? AWS Cost Explorer
? AWS CloudTrail

NEW QUESTION 71
A company has a guideline that every Amazon EC2 instance must be launched from an AMI that the company's security team produces Every month the security team sends an email message with the latest approved AMIs to all the development teams.
The development teams use AWS CloudFormation to deploy their applications. When developers launch a new service they have to search their email for the latest AMIs that the security department sent. A DevOps engineer wants to automate the process that the security team uses to provide the AMI IDs to the development teams.
What is the MOST scalable solution that meets these requirements?

A. Direct the security team to use CloudFormation to create new versions of the AMIs and to list! the AMI ARNs in an encrypted Amazon S3 object as part of the stack's Outputs Section Instruct the developers to use a cross-stack reference to load the encrypted S3 object and obtain the most recent AMI ARNs.
B. Direct the security team to use a CloudFormation stack to create an AWS CodePipeline pipeline that builds new AMIs and places the latest AMI ARNs in an encrypted Amazon S3 object as part of the pipeline output Instruct the developers to use a cross-stack reference within their own CloudFormation template to obtain the S3 object location and the most recent AMI ARNs.
C. Direct the security team to use Amazon EC2 Image Builder to create new AMIs and to place the AMI ARNs as parameters in AWS Systems Manager Parameter Store Instruct the developers to specify a parameter of type SSM in their CloudFormation stack to obtain the most recent AMI ARNs from Parameter Store.
D. Direct the security team to use Amazon EC2 Image Builder to create new AMIs and to create an Amazon Simple Notification Service (Amazon SNS) topic so that every development team can receive notification
E. When the development teams receive a notification instruct them to write an AWS Lambda function that will update their CloudFormation stack with the most recent AMI ARNs.

Answer: C

Explanation:
https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/dynamic-references.html

NEW QUESTION 72
A company is hosting a web application in an AWS Region. For disaster recovery purposes, a second region is being used as a standby. Disaster recovery requirements state that session data must be replicated between regions in near-real time and 1% of requests should route to the secondary region to continuously verify system functionality. Additionally, if there is a disruption in service in the main region, traffic should be automatically routed to the secondary region, and the secondary region must be able to
scale up to handle all traffic.
How should a DevOps engineer meet these requirements?

A. In both regions, deploy the application on AWS Elastic Beanstalk and use Amazon DynamoDB global tables for session dat
B. Use an Amazon Route 53 weighted routing policy with health checks to distribute the traffic across the regions.
C. In both regions, launch the application in Auto Scaling groups and use DynamoDB for session dat
D. Use a Route 53 failover routing policy with health checks to distribute the traffic across the regions.
E. In both regions, deploy the application in AWS Lambda, exposed by Amazon API Gateway, and use Amazon RDS for PostgreSQL with cross-region replication for session dat
F. Deploy the web application with client-side logic to call the API Gateway directly.
G. In both regions, launch the application in Auto Scaling groups and use DynamoDB global tables for session dat
H. Enable an Amazon CloudFront weighted distribution across region
I. Point the Amazon Route 53 DNS record at the CloudFront distribution.

Answer: D

NEW QUESTION 75
A company uses AWS CodeArtifact to centrally store Python packages. The CodeArtifact repository is configured with the following repository policy.

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

A development team is building a new project in an account that is in an organization in AWS Organizations. The development team wants to use a Python library that has already been stored in the CodeArtifact repository in the organization. The development team uses AWS CodePipeline and AWS CodeBuild to build the new application. The CodeBuild job that the development team uses to build the application is configured to run in a VPC Because of compliance requirements the VPC has no internet connectivity.
The development team creates the VPC endpoints for CodeArtifact and updates the CodeBuild buildspec yaml file. However, the development team cannot download the Python library from the repository.
Which combination of steps should a DevOps engineer take so that the development team can use Code Artifact? (Select TWO.)

A. Create an Amazon S3 gateway endpoint Update the route tables for the subnets thatare 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.

Answer: AD

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.

NEW QUESTION 78
A DevOps engineer has implemented a CI/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 metri
C. 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
D. 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.
E. 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.

Answer: A

Explanation:

https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/aws-attribute-updatepolicy.html

NEW QUESTION 79
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.

Answer: BD

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

NEW QUESTION 82
An application running on a set of Amazon EC2 instances in an Auto Scaling group requires a configuration file to operate. The instances are created and maintained with AWS CloudFormation. A DevOps engineer wants the instances to have the latest configuration file when launched and wants changes to the configuration file to be reflected on all the instances with a minimal delay when the CloudFormation template is updated. Company policy requires that application configuration files be maintained along with AWS infrastructure configuration files m source control.
Which solution will accomplish this?

A. In the CloudFormaiion template add an AWS Config rul
B. Place the configuration file content in the rule's InputParameters property and set the Scope property to the EC2 Auto Scaling grou
C. Add an AWS Systems Manager Resource Data Sync resource to the template to poll for updates to the configuration.
D. In the CloudFormation template add an EC2 launch template resourc
E. Place the configuration file content in the launch templat
F. Configure the cfn-mit script to run when the instance is launched and configure the cfn-hup script to poll for updates to the configuration.
G. In the CloudFormation template add an EC2 launch template resourc
H. Place the configuration file content in the launch templat
I. Add an AWS Systems Manager Resource Data Sync resource to the template to poll for updates to the configuration.
J. In the CloudFormation template add CloudFormation imt metadat
K. Place the configuration file content m the metadat
L. Configure the cfn-init script to run when the instance is launched and configure the cfn-hup script to poll for updates to the configuration.

Answer: D

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

NEW QUESTION 83
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 call
B. Configure the rule to invoke an AWS Lambda function to attach the default instance profile to the EC2 instances.
C. Configure the ec2-instance-profile-attached AWS Config managed rule with a trigger type of configuration change
D. Configure an automatic remediation action that invokes an AWS Systems Manager Automation runbook to attach the default instance profile to the EC2 instances.
E. Configure an Amazon EventBridge rule that reacts to EC2 StartInstances API call
F. Configure the rule to invoke an AWS Systems Manager Automation runbook to attach the default instance profile to the EC2 instances.
G. Configure the iam-role-managed-policy-check AWS Config managed rule with a trigger type of configuration change
H. Configure an automatic remediation action that invokes an AWS Lambda function to attach the default instance profile to the EC2 instances.

Answer: B

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

NEW QUESTION 88
A DevOps engineer is architecting a continuous development strategy for a company's software as a service (SaaS) web application running on AWS. For application and security reasons users subscribing to this application are distributed across multiple. Application Load Balancers (ALBs) each of which has a dedicated Auto Scaling group and fleet of Amazon EC2 instances The application does not require a build stage and when it is committed to AWS CodeCommit, the application must trigger a simultaneous deployment to all ALBs Auto Scaling groups and EC2 fleets.
Which architecture will meet these requirements with the LEAST amount of configuration?

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

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

Answer: C

Explanation:
https://docs.aws.amazon.com/codedeploy/latest/userguide/deployment- groups.html

NEW QUESTION 90
A company uses AWS Secrets Manager to store a set of sensitive API keys that an AWS Lambda function uses. When the Lambda function is invoked, the Lambda function retrieves the API keys and makes an API call to an external service. The Secrets Manager secret is encrypted with the default AWS Key Management Service (AWS KMS) key.
A DevOps engineer needs to update the infrastructure to ensure that only the Lambda function's execution role can access the values in Secrets Manager. The solution must apply the principle of least privilege.
Which combination of steps will meet these requirements? (Select TWO.)

A. Update the default KMS key for Secrets Manager to allow only the Lambda function's execution role to decrypt.
B. Create a KMS customer managed key that trusts Secrets Manager and allows the Lambda function's execution role to decryp
C. Update Secrets Manager to use the new customer managed key.
D. Create a KMS customer managed key that trusts Secrets Manager and allows the account's :root principal to decryp
E. Update Secrets Manager to use the new customer managed key.
F. Ensure that the Lambda function's execution role has the KMS permissions scoped on the resource leve
G. Configure the permissions so that the KMS key can encrypt the Secrets Manager secret.
H. Remove all KMS permissions from the Lambda function's execution role.

Answer: BD

Explanation:
The requirement is to update the infrastructure to ensure that only the Lambda function’s execution role can access the values in Secrets Manager. The solution must apply the principle of least privilege, which means granting the minimum permissions necessary to perform a task.
To do this, the DevOps engineer needs to use the following steps:
? Create a KMS customer managed key that trusts Secrets Manager and allows the Lambda function’s execution role to decrypt. A customer managed key is a symmetric encryption key that is fully managed by the customer. The customer can define the key policy, which specifies who can use and manage the key. By creating a customer managed key, the DevOps engineer can restrict the decryption permission to only the Lambda function’s execution role, and prevent other principals from accessing the secret values. The customer managed key also needs to trust Secrets Manager, which means allowing Secrets Manager to use the key to encrypt and decrypt secrets on behalf of the customer.
? Update Secrets Manager to use the new customer managed key. Secrets Manager allows customers to choose which KMS key to use for encrypting each secret. By default, Secrets Manager uses the default KMS key for Secrets Manager, which is a service-managed key that is shared by all customers in the same AWS Region. By updating Secrets Manager to use the new customer managed key, the DevOps engineer can ensure that only the Lambda function’s execution role can decrypt the secret values using that key.
? Ensure that the Lambda function’s execution role has the KMS permissions scoped on the resource level. The Lambda function’s execution role is an IAM role that grants permissions to the Lambda function to access AWS services and resources. The role needs to have KMS permissions to use the customer managed key for decryption. However, to apply the principle of least privilege, the role should have the permissions scoped on the resource level, which means specifying the ARN of the customer managed key as a condition in the IAM policy statement. This way, the role can only use that specific key and not any other KMS keys in the account.

NEW QUESTION 95
A DevOps engineer manages a web application that runs on Amazon EC2 instances behind an Application Load Balancer (ALB). The instances run in an EC2 Auto Scaling group across multiple Availability Zones. The engineer needs to implement a deployment strategy that:
Launches a second fleet of instances with the same capacity as the original fleet. Maintains the original fleet unchanged while the second fleet is launched. Transitions traffic to the second fleet when the second fleet is fully deployed. Terminates the original fleet automatically 1 hour after transition.
Which solution will satisfy these requirements?

A. Use an AWS CloudFormation template with a retention policy for the ALB set to 1 hou
B. Update the Amazon Route 53 record to reflect the new ALB.
C. Use two AWS Elastic Beanstalk environments to perform a blue/green deployment from the original environment to the new on
D. Create an application version lifecycle policy to terminate the original environment in 1 hour.
E. Use AWS CodeDeploy with a deployment group configured with a blue/green deployment configuration Select the option Terminate the original instances in the deployment group with a waiting period of 1 hour.
F. Use AWS Elastic Beanstalk with the configuration set to Immutabl
G. Create an.ebextension using the Resources key that sets the deletion policy of the ALB to 1 hour, and deploy the application.

Answer: C

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

NEW QUESTION 99
A business has an application that consists of five independent AWS Lambda functions.
The DevOps engineer has built a CI/CD pipeline using AWS CodePipeline and AWS CodeBuild that builds tests packages and deploys each Lambda function in sequence. The pipeline uses an Amazon EventBridge rule to ensure the pipeline starts as quickly as possible after a change is made to the application source code.
After working with the pipeline for a few months the DevOps engineer has noticed the pipeline takes too long to complete.
What should the DevOps engineer implement to BEST improve the speed of the pipeline?

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

B. Create a custom CodeBuild execution environment that includes a symmetricmultiprocessing configuration to run the builds in parallel.
C. Modify the CodePipeline configuration to run actions for each Lambda function in parallel by specifying the same runorder.
D. Modify each CodeBuild protect to run within a VPC and use dedicated instances to increase throughput.

Answer: C

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

NEW QUESTION 102
A company is divided into teams Each team has an AWS account and all the accounts are in an organization in AWS Organizations. Each team must retain full administrative rights to its AWS account. Each team also must be allowed to access only AWS services that the company approves for use AWS services must gam approval through a request and approval process.
How should a DevOps engineer configure the accounts to meet these requirements?

A. Use AWS CloudFormation StackSets to provision IAM policies in each account to deny access to restricted AWS service
B. In each account configure AWS Config rules that ensure that the policies are attached to IAM principals in the account.
C. Use AWS Control Tower to provision the accounts into OUs within the organization Configure AWS Control Tower to enable AWS IAM identity Center (AWS Single Sign-On). Configure 1AM Identity Center to provide administrative access Include deny policies on user roles for restricted AWS services.
D. Place all the accounts under a new top-level OU within the organization Create an SCP that denies access to restricted AWS services Attach the SCP to the OU.
E. Create an SCP that allows access to only approved AWS service
F. Attach the SCP to the root OU of the organizatio
G. Remove the FullAWSAccess SCP from the root OU of the organization.

Answer: C

Explanation:
https://docs.aws.amazon.com/vpc/latest/userguide/managed-prefix-lists.html A managed prefix list is a set of one or more CIDR blocks. You can use prefix lists to make it easier to configure and maintain your security groups and route tables. https://docs.aws.amazon.com/vpc/latest/userguide/sharing-managed-prefix-lists.html With AWS Resource Access Manager (AWS RAM), the owner of a prefix list can share a prefix list with the following: Specific AWS accounts inside or outside of its organization in AWS Organizations An organizational unit inside its organization in AWS Organizations An entire organization in AWS Organizations

NEW QUESTION 107
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 bucke
C. Apply a weight of 0. Specify the domain reset option to propagate changes immediately.
D. Configure webpage redirect requests on the S3 bucket with a hostname that redirects to the ALB.
E. Remove the weighted DNS record entry that points to the S3 bucket from the example.com hosted zon
F. Wait for DNS propagation to become complete.

Answer: D

NEW QUESTION 112
A company's application uses a fleet of Amazon EC2 On-Demand Instances to analyze and process data. The EC2 instances are in an Auto Scaling group. The Auto Scaling group is a target group for an Application Load Balancer (ALB). The application analyzes critical data that cannot tolerate interruption. The application also analyzes noncritical data that can withstand interruption.
The critical data analysis requires quick scalability in response to real-time application demand. The noncritical data analysis involves memory consumption. A DevOps engineer must implement a solution that reduces scale-out latency for the critical data. The solution also must process the noncritical data.
Which combination of steps will meet these requirements? (Select TWO.)

A. For the critical data, modify the existing Auto Scaling grou
B. Create a warm pool instance in the stopped stat
C. Define the warm pool siz
D. Create a new version of the launch template that has detailed monitoring enable
E. use Spot Instances.
F. For the critical data, modify the existing Auto Scaling grou
G. Create a warm pool instance in the stopped stat
H. Define the warm pool siz
I. Create a new version of the launch template that has detailed monitoring enable
J. Use On-Demand Instances.
K. For the critical dat
L. modify the existing Auto Scaling grou
M. Create a lifecycle hook to ensure that bootstrap scripts are completed successfull
N. Ensure that the application on the instances is ready to accept traffic before the instances are registere
O. Create a new version of the launch template that has detailed monitoring enabled.
P. For the noncritical data, create a second Auto Scaling group that uses a launch templat
Q. Configure the launch template to install the unified Amazon CloudWatch agent and to configure the CloudWatch agent with a custom memory utilization metri

R. Use Spot Instance
S. Add the new Auto Scaling group as the target group for the AL
T. Modify the application to use two target groups for critical data and noncritical data.
. For the noncritical data, create a second Auto Scaling grou
. Choose the predefined memory utilization metric type for the target tracking scaling polic
. Use Spot Instance
. Add the new Auto Scaling group as the target group for the AL
. Modify the application to use two target groups for critical data and noncritical data.

Answer: BD

Explanation:
? For the critical data, using a warm pool1 can reduce the scale-out latency by having pre-initialized EC2 instances ready to serve the application traffic. Using On-Demand Instances can ensure that the instances are always available and not interrupted by Spot interruptions2.
? For the noncritical data, using a second Auto Scaling group with Spot Instances can reduce the cost and leverage the unused capacity of EC23. Using a launch template with the CloudWatch agent4 can enable the collection of memory utilization metrics, which can be used to scale the group based on the memory demand. Adding the second group as a target group for the ALB and modifying the application to use two target groups can enable routing the traffic based on the data type.
References: 1: Warm pools for Amazon EC2 Auto Scaling 2: Amazon EC2 On-Demand Capacity Reservations 3: Amazon EC2 Spot Instances 4: Metrics collected by the CloudWatch agent

NEW QUESTION 113
A company has developed an AWS Lambda function that handles orders received through an API. The company is using AWS CodeDeploy to deploy the Lambda function as the final stage of a CI/CD pipeline.
A DevOps engineer has noticed there are intermittent failures of the ordering API for a few seconds after deployment. After some investigation the DevOps engineer believes the failures are due to database changes not having fully propagated before the Lambda function is invoked
How should the DevOps engineer overcome this?

A. Add a BeforeAllowTraffic hook to the AppSpec file that tests and waits for any necessary database changes before traffic can flow to the new version of the Lambda function.
B. Add an AfterAllowTraffic hook to the AppSpec file that forces traffic to wait for any pending database changes before allowing the new version of the Lambda function to respond.
C. Add a BeforeAllowTraffic hook to the AppSpec file that tests and waits for any necessary database changes before deploying the new version of the Lambda function.
D. Add a validateService hook to the AppSpec file that inspects incoming traffic and rejects the payload if dependent services such as the database are not yet ready.

Answer: A

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

NEW QUESTION 117
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

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

NEW QUESTION 118
A DevOps engineer is setting up a container-based architecture. The engineer has decided to use AWS CloudFormation to automatically provision an Amazon ECS cluster and an Amazon EC2 Auto Scaling group to launch the EC2 container instances. After successfully creating the CloudFormation stack, the engineer noticed that, even though the ECS cluster and the EC2 instances were created successfully and the stack finished the creation, the EC2 instances were associating with a different cluster.
How should the DevOps engineer update the CloudFormation template to resolve this issue?

A. Reference the EC2 instances in the AWS: ECS: Cluster resource and reference the ECS cluster in the AWS: ECS: Service resource.

B. Reference the ECS cluster in the AWS: AutoScaling: LaunchConfiguration resource of the UserData property.
C. Reference the ECS cluster in the AWS:EC2: lnstance resource of the UserData property.
D. Reference the ECS cluster in the AWS: CloudFormation: CustomResource resource to trigger an AWS Lambda function that registers the EC2 instances with the appropriate ECS cluster.

Answer: B

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

NEW QUESTION 120
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 stat
B. Configure an autoscaling:EC2_INSTANCE_LAUNCHING lifecycle hook on the Auto Scaling grou
C. 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 grou
E. Configure an autoscaling:EC2_INSTANCE_LAUNCHING lifecycle hook on the Auto Scaling grou
F. Modify the application to complete the lifecycle hook when the application is ready to serve requests.
G. Configure a warm pool for the Auto Scaling group with warmed EC2 instances in the Running stat
H. Configure an autoscaling:EC2_INSTANCE_LAUNCHING lifecycle hook on the Auto Scaling grou
I. Modify the application to complete the lifecycle hook when the application is ready to serve requests.
J. Increase the maximum instance count of the Auto Scaling grou
K. Configure an autoscaling:EC2_INSTANCE_LAUNCHING lifecycle hook on the Auto Scaling grou
L. 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

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.

NEW QUESTION 122
A Company uses AWS CodeCommit for source code control. Developers apply their changes to various feature branches and create pull requests to move those changes to the main branch when the changes are ready for production.
The developers should not be able to push changes directly to the main branch. The company applied the AWSCodeCommitPowerUser managed policy to the developers' IAM role, and now these developers can push changes to the main branch directly on every repository in the AWS account.
What should the company do to restrict the developers' ability to push changes to the main branch directly?

A. Create an additional policy to include a Deny rule for the GitPush and PutFile action
B. Include a restriction for the specific restriction for the specific repositories in the policy repositories in the policy statement with a condition that references the main branch.A Create an additional policy to include a Deny rule for the GitPush and PutFile actions Include a restriction for the specific repositories in the policy statement with a condition that references the main branch
C. Remove the IAM policy, and add an AWSCodeCommitReadOnly managed polic
D. Add an Allow rule for the GitPush and PutFile actions for the specific repositories in the policy statement with a condition that references the mam branch.
E. Modify the IAM policy Include a Deny rule for the GitPush and PutFile actions for the specific repositories in the policy statement with a condition that references the main branch.
F. Create an additional policy to include an Allow rule for the GitPush and PutFile action
G. Include a restriction for the specific repositories in the policy statement with a condition that references the feature branches.

Answer: A

Explanation:
By default, the AWSCodeCommitPowerUser managed policy allows users to push changes to any branch in any repository in the AWS account. To restrict the developers' ability to push changes to the main branch directly, an additional policy is needed that explicitly denies these actions for the main branch.
The Deny rule should be included in a policy statement that targets the specific repositories and includes a condition that references the main branch. The policy statement should look something like this:

```
{
"Effect": "Deny", "Action": [ "codecommit:GitPush", "codecommit:PutFile"
],
"Resource": "arn:aws:codecommit:<region>:<account-id>:<repository-name>", "Condition": {
"StringEqualsIfExists": { "codecommit:References": [ "refs/heads/main"
]
}
}
```

NEW QUESTION 127
A DevOps engineer is planning to deploy a Ruby-based application to production. The application needs to interact with an Amazon RDS for MySQL database and

should have automatic scaling and high availability. The stored data in the database is critical and should persist regardless of the state of the application stack. The DevOps engineer needs to set up an automated deployment strategy for the application with automatic rollbacks. The solution also must alert the application team when a deployment fails.
Which combination of steps will meet these requirements? (Select THREE.)

A. Deploy the application on AWS Elastic Beanstal
B. Deploy an Amazon RDS for MySQL DB instance as part of the Elastic Beanstalk configuration.
C. Deploy the application on AWS Elastic Beanstal
D. Deploy a separate Amazon RDS for MySQL DB instance outside of Elastic Beanstalk.
E. Configure a notification email address that alerts the application team in the AWS Elastic Beanstalk configuration.
F. Configure an Amazon EventBridge rule to monitor AWS Health event
G. Use an Amazon Simple Notification Service (Amazon SNS) topic as a target to alert the application team.
H. Use the immutable deployment method to deploy new application versions.
I. Use the rolling deployment method to deploy new application versions.

Answer: BDE

Explanation:
For deploying a Ruby-based application with requirements for interaction with an Amazon RDS for MySQL database, automatic scaling, high availability, and data persistence, the following steps will meet the requirements:
? B. Deploy the application on AWS Elastic Beanstalk. Deploy a separate Amazon
RDS for MySQL DB instance outside of Elastic Beanstalk. This approach ensures that the database persists independently of the Elastic Beanstalk environment, which can be torn down and recreated without affecting the database123.
? E. Use the immutable deployment method to deploy new application
versions. Immutable deployments provide a zero-downtime deployment method that ensures that if any part of the deployment process fails, the environment is rolled back to the original state automatically4.
? D. Configure an Amazon EventBridge rule to monitor AWS Health events. Use an
Amazon Simple Notification Service (Amazon SNS) topic as a target to alert the application team. This setup allows for automated monitoring and alerting of the application team in case of deployment failures or other health events56.
References:
? AWS Elastic Beanstalk documentation on deploying Ruby applications1.
? AWS documentation on application auto-scaling7.
? AWS documentation on automated deployment strategies with automatic rollbacks and alerts456.

NEW QUESTION 128
A company builds a container image in an AWS CodeBuild project by running Docker commands. After the container image is built, the CodeBuild project uploads the container image to an Amazon S3 bucket. The CodeBuild project has an 1AM service role that has permissions to access the S3 bucket.
A DevOps engineer needs to replace the S3 bucket with an Amazon Elastic Container Registry (Amazon ECR) repository to store the container images. The DevOps engineer creates an ECR private image repository in the same AWS Region of the CodeBuild project. The DevOps engineer adjusts the 1AM service role with the permissions that are necessary to work with the new ECR repository. The DevOps engineer also places new repository information into the docker build command and the docker push command that are used in the buildspec.yml file.
When the CodeBuild project runs a build job, the job fails when the job tries to access the ECR repository.
Which solution will resolve the issue of failed access to the ECR repository?

A. Update the buildspec.yml file to log in to the ECR repository by using the aws ecr get- login-password AWS CLI command to obtain an authentication toke
B. Update the docker login command to use the authentication token to access the ECR repository.
C. Add an environment variable of type SECRETS_MANAGER to the CodeBuild projec
D. In the environment variable, include the ARN of the CodeBuild project's IAM service rol
E. Update the buildspec.yml file to use the new environment variable to log in with the docker login command to access the ECR repository.
F. Update the ECR repository to be a public image repositor
G. Add an ECR repository policy that allows the 1AM service role to have access.
H. Update the buildspec.yml file to use the AWS CLI to assume the 1AM service role for ECR operation
I. Add an ECR repository policy that allows the 1AM service role to have access.

Answer: A

Explanation:
(A) When Docker communicates with an Amazon Elastic Container Registry (ECR) repository, it requires authentication. You can authenticate your Docker client to the Amazon ECR registry with the help of the AWS CLI (Command Line Interface). Specifically, you can use the "aws ecr get-login-password" command to get an authorization token and then use Docker's "docker login" command with that token to authenticate to the registry. You would need to perform these steps in your buildspec.yml file before attempting to push or pull images from/to the ECR repository.

NEW QUESTION 133
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.

Answer: A

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.

NEW QUESTION 138
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.

Answer: B

Explanation:

Reference: https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/crpg-ref- responses.html

NEW QUESTION 142
A company needs a strategy for failover and disaster recovery of its data and application. The application uses a MySQL database and Amazon EC2 instances. The company requires a maximum RPO of 2 hours and a maximum RTO of 10 minutes for its data and application at all times.
Which combination of deployment strategies will meet these requirements? (Select TWO.)

A. Create an Amazon Aurora Single-AZ cluster in multiple AWS Regions as the data stor
B. Use Aurora's automatic recovery capabilities in the event of a disaster.
C. Create an Amazon Aurora global database in two AWS Regions as the data stor
D. In the event of a failure, promote the secondary Region to the primary for the applicatio
E. Update the application to use the Aurora cluster endpoint in the secondary Region.
F. Create an Amazon Aurora cluster in multiple AWS Regions as the data stor
G. Use a Network Load Balancer to balance the database traffic in different Regions.
H. Set up the application in two AWS Region
I. Use Amazon Route 53 failover routing that points to Application Load Balancers in both Region
J. Use health checks and Auto Scaling groups in each Region.
K. Set up the application in two AWS Region
L. Configure AWS Global Accelerator to point to Application Load Balancers (ALBs) in both Region
M. Add both ALBs to a single endpoint grou
N. Use health checks and Auto Scaling groups in each Region.

Answer: BE

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

NEW QUESTION 147
A company is launching an application that stores raw data in an Amazon S3 bucket. Three applications need to access the data to generate reports. The data must be redacted differently for each application before
the applications can access the data.
Which solution will meet these requirements?

A. Create an S3 bucket for each applicatio
B. Configure S3 Same-Region Replication (SRR) from the raw data's S3 bucket to each application's S3 bucke
C. Configure each application to consume data from its own S3 bucket.
D. Create an Amazon Kinesis data strea
E. Create an AWS Lambda function that isinvoked by object creation events in the raw data's S3 bucke
F. Program the Lambda function to redact data for each applicatio

G. Publish the data on the Kinesis data strea
H. Configure each application to consume data from the Kinesis data stream.
I. For each application, create an S3 access point that uses the raw data's S3 bucket as the destinatio
J. Create an AWS Lambda function that is invoked by object creation events in the raw data's S3 bucke
K. Program the Lambda function to redact data for each applicatio
L. Store the data in each application's S3 access poin
M. Configure each application to consume data from its own S3 access point.
N. Create an S3 access point that uses the raw data's S3 bucket as the destinatio
O. For each application, create an S3 Object Lambda access point that uses the S3 access poin
P. Configure the AWS Lambda function for each S3 Object Lambda access point to redact data when objects are retrieve
Q. Configure each application to consume data from its own S3 Object Lambda access point.

Answer: D

Explanation:
? The best solution is to use S3 Object Lambda1, which allows you to add your own code to S3 GET, LIST, and HEAD requests to modify and process data as it is returned to an application2. This way, you can redact the data differently for each application without creating and storing multiple copies of the data or running proxies.
? The other solutions are less efficient or scalable because they require replicating the data to multiple buckets, streaming the data through Kinesis, or storing the data in S3 access points.
References: 1: Amazon S3 Features | Object Lambda | AWS 2: Transforming objects with S3 Object Lambda - Amazon Simple Storage Service

NEW QUESTION 152
A company is testing a web application that runs on Amazon EC2 instances behind an Application Load Balancer. The instances run in an Auto Scaling group across multiple Availability Zones. The company uses a blue green deployment process with immutable instances when deploying new software.
During testing users are being automatically logged out of the application at random times. Testers also report that when a new version of the application is deployed all users are logged out. The development team needs a solution to ensure users remain logged m across scaling events and application deployments.
What is the MOST operationally efficient way to ensure users remain logged in?

A. Enable smart sessions on the load balancer and modify the application to check tor an existing session.
B. Enable session sharing on the toad balancer and modify the application to read from the session store.
C. Store user session information in an Amazon S3 bucket and modify the application to read session information from the bucket.
D. Modify the application to store user session information in an Amazon ElastiCache cluster.

Answer: D

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

NEW QUESTION 157
......

Thank You for Trying Our Product

We offer two products:

1st - We have Practice Tests Software with Actual Exam Questions

2nd - Questons and Answers in PDF Format

DOP-C02 Practice Exam Features:

* DOP-C02 Questions and Answers Updated Frequently

* DOP-C02 Practice Questions Verified by Expert Senior Certified Staff

* DOP-C02 Most Realistic Questions that Guarantee you a Pass on Your FirstTry

* DOP-C02 Practice Test Questions in Multiple Choice Formats and Updatesfor 1 Year

100% Actual & Verified — Instant Download, Please Click
[Order The DOP-C02 Practice Test Here]