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 uses a single AWS account lo test applications on Amazon EC2 instances. The company has turned on AWS Config in the AWS account and has activated the restricted- ssh AWS Config managed rule.
The company needs an automated monitoring solution that will provide a customized notification in real time if any security group in the account is not compliant with the restricted-ssh rule. The customized notification must contain the name and ID of the noncompliant security group.
A DevOps engineer creates an Amazon Simple Notification Service (Amazon SNS) topic in the account and subscribes the appropriate personnel to the topic.
What should me DevOps engineer do next to meet these requirements?

A. Create an Amazon EventBridge rule that matches an AWS Config evaluation result of NON_COMPLIANT tor the restricted-ssh rul
B. Configure an input transformer for the EventBridge rule Configure the EventBridge rule to publish a notification to the SNS topic.
C. Configure AWS Config to send all evaluation results for the restricted-ssh rule to the SNS topi
D. Configure a filter policy on the SNS topic to send only notifications that contain the text of NON_COMPLIANT in the notification to subscribers.
E. Create an Amazon EventBridge rule that matches an AWS Config evaluation result of NON_COMPLIANT for the restricted-ssh rule Configure the EventBridge rule to invoke AWS Systems Manager Run Command on the SNS topic to customize a notification and to publish the notification to the SNS topic
F. Create an Amazon EventBridge rule that matches all AWS Config evaluation results of NON_COMPLIANT Configure an input transformer for the restricted-ssh rule Configure the EventBridge rule to publish a notification to the SNS topic.

Answer: A

Explanation:
Create an Amazon EventBridge (Amazon CloudWatch Events) rule that matches an AWS Config evaluation result of NON_COMPLIANT for the restricted-ssh rule. Configure an input transformer for the EventBridge (CloudWatch Events) rule. Configure the EventBridge (CloudWatch Events) rule to publish a notification to the SNS topic. This approach uses Amazon EventBridge (previously known as Amazon CloudWatch Events) to filter AWS Config evaluation results based on the restricted-ssh rule and its compliance status (NON_COMPLIANT). An input transformer can be used to customize the information contained in the notification, such as the name and ID of the noncompliant security group. The EventBridge (CloudWatch Events) rule can then be configured to publish a notification to the SNS topic, which will notify the appropriate personnel in real-time.

NEW QUESTION 3
A DevOps engineer is designing an application that integrates with a legacy REST API. The application has an AWS Lambda function that reads records from an Amazon Kinesis data stream. The Lambda function sends the records to the legacy REST API.
Approximately 10% of the records that the Lambda function sends from the Kinesis data stream have data errors and must be processed manually. The Lambda function event source configuration has an Amazon Simple Queue Service (Amazon SQS) dead-letter queue as an on-failure destination. The DevOps engineer has configured the Lambda function to process records in batches and has implemented retries in case of failure.
During testing the DevOps engineer notices that the dead-letter queue contains many records that have no data errors and that already have been processed by the legacy REST API. The DevOps engineer needs to configure the Lambda function's event source options to reduce the number of errorless records that are sent to the dead-letter queue.
Which solution will meet these requirements?

A. Increase the retry attempts
B. Configure the setting to split the batch when an error occurs
C. Increase the concurrent batches per shard
D. Decrease the maximum age of record

Answer: B

Explanation:

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

NEW QUESTION 4
A company requires an RPO of 2 hours and an RTO of 10 minutes for its data and application at all times. An application uses a MySQL database and Amazon EC2 web servers. The development team needs a strategy for failover and disaster recovery.
Which combination of deployment strategies will meet these requirements? (Select TWO.)

A. Create an Amazon Aurora cluster in one Availability Zone across multiple Regions as the data store Use Aurora's automatic recovery capabilities in the event of a disaster
B. Create an Amazon Aurora global database in two Regions as the data stor
C. In the event of a failure promote the secondary Region as the primary for the application.
D. Create an Amazon Aurora multi-master cluster across multiple Regions as the data stor
E. Use a Network Load Balancer to balance the database traffic in different Regions.
F. Set up the application in two Regions and use Amazon Route 53 failover-based routing that points to the Application Load Balancers in both Region
G. Use hearth checks to determine the availability in a given Regio
H. Use Auto Scaling groups in each Region to adjust capacity based on demand.
I. Set up the application m two Regions and use a multi-Region Auto Scaling group behind Application Load Balancers to manage the capacity based on deman
J. In the event of a disaster adjust the Auto Scaling group's desired instance count to increase baseline capacity in the failover Region.

Answer: BD

NEW QUESTION 5
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 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 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 8
A company is implementing a well-architected design for its globally accessible API stack. The design needs to ensure both high reliability and fast response times for users located in North America and Europe.
The API stack contains the following three tiers: Amazon API Gateway
AWS Lambda Amazon DynamoDB
Which solution will meet the requirements?

A. Configure Amazon Route 53 to point to API Gateway APIs in North America and Europe using health check
B. Configure the APIs to forward requests to a Lambda function in that Regio
C. Configure the Lambda functions to retrieve and update the data in a DynamoDB table in the same Region as the Lambda function.
D. Configure Amazon Route 53 to point to API Gateway APIs in North America and Europe using latency-based routing and health check
E. Configure the APIs to forward requests to a Lambda function in that Regio
F. Configure the Lambda functions to retrieve and update the data in a DynamoDB global table.
G. Configure Amazon Route 53 to point to API Gateway in North America, create a disaster recovery API in Europe, and configure both APIs to forward requests to the Lambda functions in that Regio
H. Retrieve the data from a DynamoDB global tabl
I. Deploy a Lambda function to check the North America API health every 5 minute
J. In the event of a failure, update Route 53 to point to the disaster recovery API.
K. Configure Amazon Route 53 to point to API Gateway API in North America using latency-based routin
L. Configure the API to forward requests to the Lambda function in the Region nearest to the use
M. Configure the Lambda function to retrieve and update the data in a DynamoDB table.

Answer: B

NEW QUESTION 9
A company has a data ingestion application that runs across multiple AWS accounts. The accounts are in an organization in AWS Organizations. The company needs to monitor the application and consolidate access to the application. Currently the company is running the application on Amazon EC2 instances from several Auto Scaling groups. The EC2 instances have no access to the internet because the data is sensitive Engineers have deployed the necessary VPC endpoints. The EC2 instances run a custom AMI that is built specifically tor the application.
To maintain and troubleshoot the application, system administrators need the ability to log in to the EC2 instances. This access must be automated and controlled centrally. The company's security team must receive a notification whenever the instances are accessed.
Which solution will meet these requirements?

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

Answer: C

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

NEW QUESTION 10
A company has a new AWS account that teams will use to deploy various applications. The teams will create many Amazon S3 buckets for application- specific purposes and to store AWS CloudTrail logs. The company has enabled Amazon Macie for the account.
A DevOps engineer needs to optimize the Macie costs for the account without compromising the account's functionality.
Which solutions will meet these requirements? (Select TWO.)

A. Exclude S3 buckets that contain CloudTrail logs from automated discovery.
B. Exclude S3 buckets that have public read access from automated discovery.
C. Configure scheduled daily discovery jobs for all S3 buckets in the account.
D. Configure discovery jobs to include S3 objects based on the last modified criterion.
E. Configure discovery jobs to include S3 objects that are tagged as production only.

Answer: AD

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

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 12
A company uses AWS CodePipeline pipelines to automate releases of its application A typical pipeline consists of three stages build, test, and deployment. The company has been using a separate AWS CodeBuild project to run scripts for each stage. However, the company now wants to use AWS CodeDeploy to handle the deployment stage of the pipelines.
The company has packaged the application as an RPM package and must deploy the application to a fleet of Amazon EC2 instances. The EC2 instances are in an EC2 Auto Scaling group and are launched from a common AMI.
Which combination of steps should a DevOps engineer perform to meet these requirements? (Choose two.)

A. Create a new version of the common AMI with the CodeDeploy agent installe
B. Update the IAM role of the EC2 instances to allow access to CodeDeploy.
C. Create a new version of the common AMI with the CodeDeploy agent installe
D. Create an AppSpec file that contains application deployment scripts and grants access to CodeDeploy.
E. Create an application in CodeDeplo
F. Configure an in-place deployment typ
G. Specify the Auto Scaling group as the deployment targe
H. Add a step to the CodePipeline pipeline to use EC2 Image Builder to create a new AM
I. Configure CodeDeploy to deploy the newly created AMI.
J. Create an application in CodeDeplo
K. Configure an in-place deployment typ
L. Specify the Auto Scaling group as the deployment targe
M. Update the CodePipeline pipeline to use the CodeDeploy action to deploy the application.
N. Create an application in CodeDeplo
O. Configure an in-place deployment typ
P. Specify the EC2 instances that are launched from the common AMI as the deployment targe
Q. Update the CodePipeline pipeline to use the CodeDeploy action to deploy the application.

Answer: AD

Explanation:
https://docs.aws.amazon.com/codedeploy/latest/userguide/integrations-aws-auto-scaling.html

NEW QUESTION 17
A company manages multiple AWS accounts by using AWS Organizations with OUS for the different business divisions, The company is updating their corporate network to use new IP address ranges. The company has 10 Amazon S3 buckets in different AWS accounts. The S3 buckets store reports for the different divisions. The S3 bucket configurations allow only private corporate network IP addresses to access the S3 buckets.
A DevOps engineer needs to change the range of IP addresses that have permission to access the contents of the S3 buckets The DevOps engineer also needs to revoke the permissions of two OUS in the company
Which solution will meet these requirements?

A. Create a new SCP that has two statements, one that allows access to the new range of IP addresses for all the S3 buckets and one that demes access to the

old range of IP addresses for all the S3 bucket
B. Set a permissions boundary for the OrganzauonAccountAccessRole role In the two OUS to deny access to the S3 buckets.
C. Create a new SCP that has a statement that allows only the new range of IP addresses to access the S3 bucket
D. Create another SCP that denies access to the S3 bucket
E. Attach the second SCP to the two OUS
F. On all the S3 buckets, configure resource-based policies that allow only the new range of IP addresses to access the S3 bucket
G. Create a new SCP that denies access to the S3 bucket
H. Attach the SCP to the two OUs.
I. On all the S3 buckets, configure resource-based policies that allow only the new range of IP addresses to access the S3 bucket
J. Set a permissions boundary for the OrganizationAccountAccessRole role in the two OUS to deny access to the S3 buckets.

Answer: C

Explanation:
The correct answer is C.
A comprehensive and detailed explanation is:
? Option A is incorrect because creating a new SCP that has two statements, one that allows access to the new range of IP addresses for all the S3 buckets and one that denies access to the old range of IP addresses for all the S3 buckets, is not a valid solution. SCPs are not resource-based policies, and they cannot specify the S3 buckets or the IP addresses as resources or conditions. SCPs can only control the actions that can be performed by the principals in the organization, not the access to specific resources. Moreover, setting a permissions boundary for the OrganizationAccountAccessRole role in the two OUs to deny access to the S3 buckets is not sufficient to revoke the permissions of the two OUs, as there might be other roles or users in those OUs that can still access the S3 buckets.
? Option B is incorrect because creating a new SCP that has a statement that allows
only the new range of IP addresses to access the S3 buckets is not a valid solution, for the same reason as option A. SCPs are not resource-based policies, and they cannot specify the S3 buckets or the IP addresses as resources or conditions. Creating another SCP that denies access to the S3 buckets and attaching it to the two OUs is also not a valid solution, as SCPs cannot specify the S3 buckets as resources either.
? Option C is correct because it meets both requirements of changing the range of IP addresses that have permission to access the contents of the S3 buckets and revoking the permissions of two OUs in the company. On all the S3 buckets, configuring resource-based policies that allow only the new range of IP addresses to access the S3 buckets is a valid way to update the IP address ranges, as resource-based policies can specify both resources and conditions. Creating a new SCP that denies access to the S3 buckets and attaching it to the two OUs is also a valid way to revoke the permissions of those OUs, as SCPs can deny actions such as s3:PutObject or s3:GetObject on any resource.
? Option D is incorrect because setting a permissions boundary for the OrganizationAccountAccessRole role in the two OUs to deny access to the S3 buckets is not sufficient to revoke the permissions of the two OUs, as there might be other roles or users in those OUs that can still access the S3 buckets. A permissions boundary is a policy that defines the maximum permissions that an IAM entity can have. However, it does not revoke any existing permissions that are granted by other policies.
References:
? AWS Organizations
? S3 Bucket Policies
? Service Control Policies
? Permissions Boundaries

NEW QUESTION 21
A DevOps engineer is deploying a new version of a company's application in an AWS CodeDeploy deployment group associated with its Amazon EC2 instances. After some time, the deployment fails. The engineer realizes that all the events associated with the specific deployment ID are in a Skipped status and code was not deployed in the instances associated with the deployment group.
What are valid reasons for this failure? (Select TWO.).

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

Answer: AD

Explanation:
https://docs.aws.amazon.com/codedeploy/latest/userguide/troubleshooting- deployments.html#troubleshooting-skipped-lifecycle-events

NEW QUESTION 26
A company runs a workload on Amazon EC2 instances. The company needs a control that requires the use of Instance Metadata Service Version 2 (IMDSv2) on all EC2 instances in the AWS account. If an EC2 instance does not prevent the use of Instance Metadata Service Version 1 (IMDSv1), the EC2 instance must be terminated.
Which solution will meet these requirements?

A. Set up AWS Config in the accoun
B. Use a managed rule to check EC2 instance
C. Configure the rule to remediate the findings by using AWS Systems Manager Automation to terminate the instance.
D. Create a permissions boundary that prevents the ec2:RunInstance action if the ec2:MetadataHttpTokens condition key is not set to a value of require
E. Attach the permissions boundary to the IAM role that was used to launch the instance.
F. Set up Amazon Inspector in the accoun
G. Configure Amazon Inspector to activate deep inspection for EC2 instance
H. Create an Amazon EventBridge rule for an Inspector2 findin
I. Set an AWS Lambda function as the target to terminate the instance.
J. Create an Amazon EventBridge rule for the EC2 instance launch successful even
K. Send the event to an AWS Lambda function to inspect the EC2 metadata and to terminate the instance.

Answer: B

Explanation:

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

NEW QUESTION 30
A DevOps engineer at a company is supporting an AWS environment in which all users use AWS IAM Identity Center (AWS Single Sign-On). The company wants to immediately disable credentials of any new IAM user and wants the security team to receive a notification.
Which combination of steps should the DevOps engineer take to meet these requirements? (Choose three.)

A. Create an Amazon EventBridge rule that reacts to an IAM CreateUser API call in AWS CloudTrail.
B. Create an Amazon EventBridge rule that reacts to an IAM GetLoginProfile API call in AWS CloudTrail.
C. Create an AWS Lambda function that is a target of the EventBridge rul
D. Configure the Lambda function to disable any access keys and delete the login profiles that are associated with the IAM user.
E. Create an AWS Lambda function that is a target of the EventBridge rul
F. Configure the Lambda function to delete the login profiles that are associated with the IAM user.
G. Create an Amazon Simple Notification Service (Amazon SNS) topic that is a target of the EventBridge rul
H. Subscribe the security team's group email address to the topic.
I. Create an Amazon Simple Queue Service (Amazon SQS) queue that is a target of the Lambda functio
J. Subscribe the security team's group email address to the queue.

Answer: ACE

NEW QUESTION 33
A DevOps engineer has developed an AWS Lambda function The Lambda function starts an AWS CloudFormation drift detection operation on all supported resources for a specific CloudFormation stack The Lambda function then exits Its invocation The DevOps engineer has created an Amazon EventBrdge scheduled rule that Invokes the Lambda function every hour. An Amazon Simple Notification Service (Amazon SNS) topic already exists In the AWS account. The DevOps engineer has subscribed to the SNS topic to receive notifications
The DevOps engineer needs to receive a notification as soon as possible when drift is detected in this specific stack configuration.
Which solution Will meet these requirements?

A. Configure the existing EventBridge rule to also target the SNS topic Configure an SNS subscription filter policy to match the Cloud Formation stac
B. Attach the subscription filter policy to the SNS tomc
C. Create a second Lambda function to query the CloudFormation API for the drift detection results for the stack Configure the second Lambda function to publish a message to the SNS topic If drift ts detected Adjust the existing EventBridge rule to also target the second Lambda function
D. Configure Amazon GuardDuty in the account with drift detection for all CloudFormation stack
E. Create a second EventBndge rule that reacts to the GuardDuty drift detection event finding for the specific CloudFormation stac
F. Configure the SNS topic as a target of the second EventBridge rule.
G. Configure AWS Config in the accoun
H. Use the cloudformation-stack-drift-detection- check managed rul
I. Create a second EventBndge rule that reacts to a compliance change event for the CloudFormaUon stac
J. Configure the SNS topc as a target of the second EventBridge rule.

Answer: D

Explanation:
A comprehensive and detailed explanation is:
? Option A is incorrect because EventBridge rules cannot filter events based on the message body or attributes of the target service. Therefore, configuring an SNS subscription filter policy to match the CloudFormation stack will not work. The SNS topic will receive all events from the EventBridge rule, regardless of the stack name or drift status.
? Option B is incorrect because it introduces unnecessary complexity and cost.
Creating a second Lambda function to query the CloudFormation API for the drift detection results is redundant, since CloudFormation already publishes drift detection events to EventBridge. Moreover, invoking two Lambda functions every hour will incur more charges than invoking one.
? Option C is incorrect because GuardDuty does not provide drift detection for CloudFormation stacks. GuardDuty is a threat detection service that monitors for malicious activity and unauthorized behavior in AWS accounts and workloads. It does not monitor or report on configuration changes or drifts in CloudFormation stacks.
? Option D is correct because it leverages AWS Config and its managed rule for drift detection. AWS Config is a service that enables you to assess, audit, and evaluate the configurations of your AWS resources. It can detect configuration changes and drifts in CloudFormation stacks using the cloudformation-stack-drift-detection- check managed rule. This rule triggers an AWS Config event when a stack drifts from its expected template configuration. By creating a second EventBridge rule that reacts to this event for the specific stack, the DevOps engineer can configure the SNS topic as a target and receive a notification as soon as possible when drift is detected.
References:
? AWS Config
? Amazon SNS subscription filter policies
? Amazon EventBridge rules

NEW QUESTION 38
A company is launching an application. The application must use only approved AWS services. The account that runs the application was created less than 1 year ago and is assigned to an AWS Organizations OU.
The company needs to create a new Organizations account structure. The account structure must have an appropriate SCP that supports the use of only services that are currently active in the AWS account.
The company will use AWS Identity and Access Management (IAM) Access Analyzer in the solution.
Which solution will meet these requirements?

A. Create an SCP that allows the services that IAM Access Analyzer identifie
B. Create an OU for the accoun
C. Move the account into the new O
D. Attach the new SCP to the new O
E. Detach the default FullAWSAccess SCP from the new OU.

F. Create an SCP that denies the services that IAM Access Analyzer identifie
G. Create an OU for the accoun
H. Move the account into the new OI
I. Attach the new SCP to the new OU.
J. Create an SCP that allows the services that IAM Access Analyzer identifie
K. Attach the new SCP to the organization's root.
L. Create an SCP that allows the services that IAM Access Analyzer identifie
M. Create an OU for the accoun
N. Move the account into the new O
O. Attach the new SCP to the management accoun
P. Detach the default FullAWSAccess SCP from the new OU.

Answer: A

Explanation:
To meet the requirements of creating a new Organizations account structure with an appropriate SCP that supports the use of only services that are currently active in the AWS account, the company should use the following solution:
? Create an SCP that allows the services that IAM Access Analyzer identifies. IAM Access Analyzer is a service that helps identify potential resource-access risks by analyzing resource-based policies in the AWS environment. IAM Access Analyzer can also generate IAM policies based on access activity in the AWS CloudTrail logs. By using IAM Access Analyzer, the company can create an SCP that grants only the permissions that are required for the application to run, and denies all other services. This way, the company can enforce the use of only approved AWS services and reduce the risk of unauthorized access12
? Create an OU for the account. Move the account into the new OU. An OU is a container for accounts within an organization that enables you to group accounts that have similar business or security requirements. By creating an OU for the account, the company can apply policies and manage settings for the account as a group. The company should move the account into the new OU to make it subject to the policies attached to the OU3
? Attach the new SCP to the new OU. Detach the default FullAWSAccess SCP from the new OU. An SCP is a type of policy that specifies the maximum permissions for an organization or organizational unit (OU). By attaching the new SCP to the new OU, the company can restrict the services that are available to all accounts in that OU, including the account that runs the application. The company should also detach the default FullAWSAccess SCP from the new OU, because this policy allows all actions on all AWS services and might override or conflict with the new SCP45
The other options are not correct because they do not meet the requirements or follow best practices. Creating an SCP that denies the services that IAM Access Analyzer identifies is not a good option because it might not cover all possible services that are not approved or required for the application. A deny policy is also more difficult to maintain and update than an allow policy. Creating an SCP that allows the services that IAM Access Analyzer identifies and attaching it to the organization’s root is not a good option because it might affect other accounts and OUs in the organization that have different service requirements or approvals. Creating an SCP that allows the services that IAM Access Analyzer identifies and attaching it to the management account is not a valid option because SCPs cannot be attached directly to accounts, only to OUs or roots.
References:
? 1: Using AWS Identity and Access Management Access Analyzer - AWS Identity and Access Management
? 2: Generate a policy based on access activity - AWS Identity and Access Management
? 3: Organizing your accounts into OUs - AWS Organizations
? 4: Service control policies - AWS Organizations
? 5: How SCPs work - AWS Organizations

NEW QUESTION 43
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 47
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 50
A video-sharing company stores its videos in Amazon S3. The company has observed a sudden increase in video access requests, but the company does not know which videos are most popular. The company needs to identify the general access pattern for the video files. This pattern includes the number of users who access a certain file on a given day, as well as the numb A DevOps engineer manages a large commercial website that runs on Amazon EC2 The website uses Amazon Kinesis Data Streams to collect and process web togs The DevOps engineer manages the Kinesis consumer application, which also runs on Amazon EC2 Sudden increases of data cause the Kinesis consumer application to (all behind and the Kinesis data streams drop records before the records can be processed The DevOps engineer must implement a solution to improve stream handling
Which solution meets these requirements with the MOST operational efficiency" er of pull requests for certain files.
How can the company meet these requirements with the LEAST amount of effort?

A. Activate S3 server access loggin
B. Import the access logs into an Amazon Aurora databas
C. Use an Aurora SQL query to analyze the access patterns.
D. Activate S3 server access loggin
E. Use Amazon Athena to create an external table with the log file
F. Use Athena to create a SQL query to analyze the access patterns.
G. Invoke an AWS Lambda function for every S3 object access even
H. Configure the Lambda function to write the file access information, such as use
I. S3 bucket, and file key, to an Amazon Aurora databas
J. Use an Aurora SQL query to analyze the access patterns.
K. Record an Amazon CloudWatch Logs log message for every S3 object access even
L. Configure a CloudWatch Logs log stream to write the file access information, such as user, S3 bucket, and file key, to an Amazon Kinesis Data Analytics for SQL applicatio
M. Perform a sliding window analysis.

Answer: B

Explanation:
Activating S3 server access logging and using Amazon Athena to create an external table with the log files is the easiest and most cost-effective way to analyze access patterns. This option requires minimal setup and allows for quick analysis of the access
patterns with SQL queries. Additionally, Amazon Athena scales automatically to match the query load, so there is no need for additional infrastructure provisioning or management.

NEW QUESTION 55
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 56
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 60
A company has migrated its container-based applications to Amazon EKS and want to establish automated email notifications. The notifications sent to each email address are for specific activities related to EKS components. The solution will include Amazon SNS topics and an AWS Lambda function to evaluate incoming log events and publish messages to the correct SNS topic.
Which logging solution will support these requirements?

A. Enable Amazon CloudWatch Logs to log the EKS component
B. Create a CloudWatch subscription filter for each component with Lambda as the subscription feed destination.
C. Enable Amazon CloudWatch Logs to log the EKS component
D. Create CloudWatch Logs Insights queries linked to Amazon EventBridge events that invoke Lambda.
E. Enable Amazon S3 logging for the EKS component
F. Configure an Amazon CloudWatch subscription filter for each component with Lambda as the subscription feed destination.
G. Enable Amazon S3 logging for the EKS component
H. Configure S3 PUT Object event notifications with AWS Lambda as the destination.

Answer: A

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

NEW QUESTION 64
A company recently migrated its legacy application from on-premises to AWS. The application is hosted on Amazon EC2 instances behind an Application Load Balancer which is behind Amazon API Gateway. The company wants to ensure users experience minimal disruptions during any deployment of a new version of the application. The company also wants to ensure it can quickly roll back updates if there is an issue.
Which solution will meet these requirements with MINIMAL changes to the application?

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

Answer: A

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

NEW QUESTION 69
A DevOps team is merging code revisions for an application that uses an Amazon RDS Multi-AZ DB cluster for its production database. The DevOps team uses continuous integration to periodically verify that the application works. The DevOps team needs to test the changes before the changes are deployed to the production database.
Which solution will meet these requirements'?

A. Use a buildspec file in AWS CodeBuild to restore the DB cluster from a snapshot of the production database run integration tests, and drop the restored database after verification.
B. Deploy the application to productio
C. Configure an audit log of data control language (DCL) operations to capture database activities to perform if verification fails.
D. Create a snapshot of the DB duster before deploying the application Use the Update requires Replacement property on the DB instance in AWS CloudFormation to deploy the application and apply the changes.
E. Ensure that the DB cluster is a Multi-AZ deploymen
F. Deploy the application with the update
G. Fail over to the standby instance if verification fails.

Answer: A

Explanation:
This solution will meet the requirements because it will create a temporary copy of the production database using a snapshot, run the integration tests on the copy, and delete the copy after the tests are done. This way, the production database will not be affected by the code revisions, and the DevOps team can test the changes before deploying them to production. A buildspec file is a YAML file that contains the commands and settings that CodeBuild uses to run a build1. The buildspec file can specify the steps to restore the DB cluster from a snapshot, run the integration tests, and drop the restored database2

NEW QUESTION 74
A company's development team uses AVMS Cloud Formation to deploy its application resources The team must use for an changes to the environment The team cannot use AWS Management Console or the AWS CLI to make manual changes directly.
The team uses a developer IAM role to access the environment The role is configured with the Admnistratoraccess managed policy. The company has created a new Cloudformationdeployment IAM role that has the following policy.

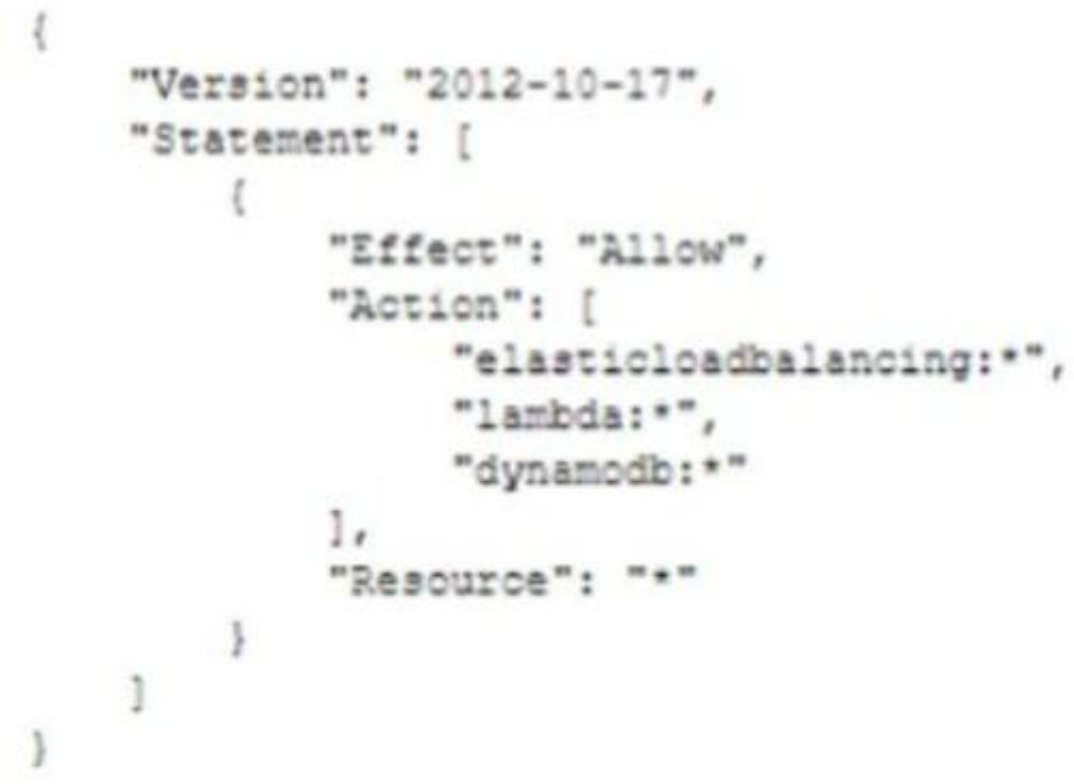

```
{
    "Version": "2012-10-17",
    "Statement": [
        {
            "Effect": "Allow",
            "Action": [
                "elasticloadbalancing:*",
                "lambda:*",
                "dynamodb:*"
            ],
            "Resource": "*"
        }
    ]
}
```

The company wants ensure that only CloudFormation can use the new role. The development team cannot make any manual changes to the deployed resources. Which combination of steps meet these requirements? (Select THREE.)

A. Remove the AdministratorAccess polic
B. Assign the ReadOnlyAccess managed IAM policy to the developer rol
C. Instruct the developers to use the CloudFormationDeployment role as a CloudFormation service role when the developers deploy new stacks.
D. Update the trust of CloudFormationDeployment role to allow the developer IAM role to assume the CloudFormationDepoyment role.
E. Configure the IAM to be to get and pass the CloudFormationDeployment role if cloudformation actions for resources,
F. Update the trust Of the CloudFormationDepoyment role to anow the cloudformation.amazonaws.com AWS principal to perform the iam:AssumeR01e action
G. Remove me Administratoraccess polic
H. Assign the ReadOnly/Access managed IAM policy to the developer role Instruct the developers to assume the CloudFormatondeployment role when the developers new stacks
I. Add an IAM policy to CloudFormationDeplyment to allow cloudformation * on an Add a policy that allows the iam.PassR01e action for ARN of if iam PassedT0Service equal cloudformation.amazonaws.com

Answer: ADF

Explanation:
A comprehensive and detailed explanation is:
? Option A is correct because removing the AdministratorAccess policy and assigning the ReadOnlyAccess managed IAM policy to the developer role is a valid way to prevent the developers from making any manual changes to the deployed resources. The AdministratorAccess policy grants full access to all AWS resources and actions, which is not necessary for the developers. The ReadOnlyAccess policy grants read-only access to most AWS resources and actions, which is sufficient for the developers to view the status of their stacks. Instructing the developers to use the CloudFormationDeployment role as a CloudFormation service role when they deploy new stacks is also a valid way to ensure that only CloudFormation can use the new role. A CloudFormation service role is an IAM role that allows CloudFormation to make calls to resources in a stack on behalf of the user1. The user can specify a service role when they create or update a stack, and CloudFormation will use that role's credentials for all operations that are performed on that stack1.
? Option B is incorrect because updating the trust of CloudFormationDeployment role to allow the developer IAM role to assume the CloudFormationDeployment role is not a valid solution. This would allow the developers to manually assume the CloudFormationDeployment role and perform actions on the deployed resources, which is not what the company wants. The trust of CloudFormationDeployment role should only allow the cloudformation.amazonaws.com AWS principal to assume the role, as in option D.
? Option C is incorrect because configuring the IAM user to be able to get and pass the CloudFormationDeployment role if cloudformation actions for resources is not a valid solution. This would allow the developers to manually pass the CloudFormationDeployment role to other services or resources, which is not what the company wants. The IAM user should only be able to pass the CloudFormationDeployment role as a service role when they create or update a stack with CloudFormation, as in option A.
? Option D is correct because updating the trust of CloudFormationDeployment role
to allow the cloudformation.amazonaws.com AWS principal to perform the iam:AssumeRole action is a valid solution. This allows CloudFormation to assume the CloudFormationDeployment role and access resources in other services on behalf of the user2. The trust policy of an IAM role defines which entities can assume the role2. By specifying cloudformation.amazonaws.com as the principal, you grant permission only to CloudFormation to assume this role.
? Option E is incorrect because instructing the developers to assume the
CloudFormationDeployment role when they deploy new stacks is not a valid solution. This would allow the developers to manually assume the CloudFormationDeployment role and perform actions on the deployed resources, which is not what the company wants. The developers should only use the CloudFormationDeployment role as a service role when they deploy new stacks with CloudFormation, as in option A.
? Option F is correct because adding an IAM policy to CloudFormationDeployment
that allows cloudformation:* on all resources and adding a policy that allows the iam:PassRole action for ARN of CloudFormationDeployment if iam:PassedToService equals cloudformation.amazonaws.com are valid solutions. The first policy grants permission for CloudFormationDeployment to perform any action with any resource using cloudformation.amazonaws.com as a service principal3. The second policy grants permission for passing this role only if it is passed by cloudformation.amazonaws.com as a service principal4. This ensures that only CloudFormation can use this role.
References:
? 1: AWS CloudFormation service roles
? 2: How to use trust policies with IAM roles
? 3: AWS::IAM::Policy
? 4: IAM: Pass an IAM role to a specific AWS service

NEW QUESTION 76
A company's DevOps engineer is creating an AWS Lambda function to process notifications from an Amazon Simple Notification Service (Amazon SNS) topic. The Lambda function will process the notification messages and will write the contents of the notification messages to an Amazon RDS Multi-AZ DB instance. During testing a database administrator accidentally shut down the DB instance. While the database was down the company lost several of the SNS notification messages that were delivered during that time.

The DevOps engineer needs to prevent the loss of notification messages in the future Which solutions will meet this requirement? (Select TWO.)

A. Replace the RDS Multi-AZ DB instance with an Amazon DynamoDB table.
B. Configure an Amazon Simple Queue Service (Amazon SQS) queue as a destination of the Lambda function.
C. Configure an Amazon Simple Queue Service (Amazon SQS> dead-letter queue for the SNS topic.
D. Subscribe an Amazon Simple Queue Service (Amazon SQS) queue to the SNS topic Configure the Lambda function to process messages from the SQS queue.
E. Replace the SNS topic with an Amazon EventBridge event bus Configure an EventBridge rule on the new event bus to invoke the Lambda function for each event.

Answer: CD

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

NEW QUESTION 79
A company provides an application to customers. The application has an Amazon API Gateway REST API that invokes an AWS Lambda function. On initialization, the Lambda function loads a large amount of data from an Amazon DynamoDB table. The data load process results in long cold-start times of 8-10 seconds. The DynamoDB table has DynamoDB Accelerator (DAX) configured.
Customers report that the application intermittently takes a long time to respond to requests. The application receives thousands of requests throughout the day. In the middle of the day, the application experiences 10 times more requests than at any other time of the day. Near the end of the day, the application's request volume decreases to 10% of its normal total.
A DevOps engineer needs to reduce the latency of the Lambda function at all times of the day.
Which solution will meet these requirements?

A. Configure provisioned concurrency on the Lambda function with a concurrency value of 1. Delete the DAX cluster for the DynamoDB table.
B. Configure reserved concurrency on the Lambda function with a concurrency value of 0.
C. Configure provisioned concurrency on the Lambda functio
D. Configure AWS Application Auto Scaling on the Lambda function with provisioned concurrency values set to a minimum of 1 and a maximum of 100.
E. Configure reserved concurrency on the Lambda functio
F. Configure AWS Application Auto Scaling on the API Gateway API with a reserved concurrency maximum value of 100.

Answer: C

Explanation:
The following are the steps that the DevOps engineer should take to reduce the latency of the Lambda function at all times of the day:
? Configure provisioned concurrency on the Lambda function.
? Configure AWS Application Auto Scaling on the Lambda function with provisioned concurrency values set to a minimum of 1 and a maximum of 100.
The provisioned concurrency setting ensures that there is always a minimum number of Lambda function instances available to handle requests. The Application Auto Scaling setting will automatically scale the number of Lambda function instances up or down based on the demand for the application.
This solution will ensure that the Lambda function is able to handle the increased load during the middle of the day, while also keeping the cold-start latency low.
The following are the reasons why the other options are not correct:
? Option A is incorrect because it will not reduce the cold-start latency of the Lambda function.
? Option B is incorrect because it will not scale the number of Lambda function instances up or down based on demand.
? Option D is incorrect because it will only configure reserved concurrency on the API Gateway API, which will not affect the Lambda function.

NEW QUESTION 82
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 85
A company's application teams use AWS CodeCommit repositories for their applications.
The application teams have repositories in multiple AWS accounts. All accounts are in an organization in AWS Organizations.
Each application team uses AWS IAM Identity Center (AWS Single Sign-On) configured with an external IdP to assume a developer IAM role. The developer role allows the application teams to use Git to work with the code in the repositories.
A security audit reveals that the application teams can modify the main branch in any repository. A DevOps engineer must implement a solution that allows the application teams to modify the main branch of only the repositories that they manage.
Which combination of steps will meet these requirements? (Select THREE.)

A. Update the SAML assertion to pass the user's team nam
B. Update the IAM role's trust policy to add an access-team session tag that has the team name.
C. Create an approval rule template for each team in the Organizations management accoun
D. Associate the template with all the repositorie
E. Add the developer role ARN as an approver.
F. Create an approval rule template for each accoun
G. Associate the template with all repositorie
H. Add the "aws:ResourceTag/access-team":"$;{aws:PrincipalTag/access- team}" condition to the approval rule template.
I. For each CodeCommit repository, add an access-team tag that has the value set to the name of the associated team.
J. Attach an SCP to the account
K. Include the following statement:

```
{
    "Effect": "Deny",
    "Action": [
        "codecommit:GitPush",
        "codecommit:PutFile",
        "codecommit:Merge*"
    ],
    "Resource": "*",
    "Condition": {
        "StringEqualsIfExists": {
            "codecommit:References": ["refs/heads/main"]
        },
        "StringNotEquals": {
            "aws:ResourceTag/access-team": "$ ;{aws:PrincipalTag/access-team}"
        },
        "Null": {
            "codecommit:References": "false"
        }
    }
}
```

L. Create an IAM permissions boundary in each accoun
M. Include the following statement:

```
{
    "Effect": "Allow",
    "Action": [
        "codecommit:GitPush",
        "codecommit:PutFile",
        "codecommit:Merge*"
    ],
    "Resource": "*",
    "Condition": {
        "StringEqualsIfExists": {
            "codecommit:References": ["refs/heads/main"]
        },
        "StringNotEquals": {
            "aws:ResourceTag/access-team": "$ ;{aws:PrincipalTag/access-team}"
        },
        "Null": {
            "codecommit:References": "false"
        }
    }
}
```

Answer: ADF

Explanation:
Short Explanation: To meet the requirements, the DevOps engineer should update the SAML assertion to pass the user's team name, update the IAM role's trust policy to add an access-team session tag that has the team name, create an IAM permissions boundary in each account, and for each CodeCommit repository, add an access-team tag that has the value set to the name of the associated team.
References:
? Updating the SAML assertion to pass the user's team name allows the DevOps engineer to use IAM tags to identify which team a user belongs to. This can help enforce fine-grained access control based on the user's team membership1.
? Updating the IAM role's trust policy to add an access-team session tag that has the team name allows the DevOps engineer to use IAM condition keys to restrict access based on the session tag value2. For example, the DevOps engineer can use the aws:PrincipalTag condition key to match the access-team tag of the user with the access-team tag of the repository3.
? Creating an IAM permissions boundary in each account allows the DevOps engineer to set the maximum permissions that an identity-based policy can grant to an IAM entity. An entity's permissions boundary allows it to perform only the actions that are allowed by both its identity-based policies and its permissions boundaries4. For example, the DevOps engineer can use a permissions boundary policy to limit the actions that a user can perform on CodeCommit repositories based on their access-team tag5.
? For each CodeCommit repository, adding an access-team tag that has the value set to the name of the associated team allows the DevOps engineer to use resource tags to identify which team manages a repository. This can help enforce fine-grained access control based on the resource tag value6.
? The other options are incorrect because:

NEW QUESTION 86
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 87
A DevOps engineer needs to configure a blue green deployment for an existing three-tier application. The application runs on Amazon EC2 instances and uses an Amazon RDS database The EC2 instances run behind an Application Load Balancer (ALB) and are in an Auto Scaling group.
The DevOps engineer has created a launch template and an Auto Scaling group for the blue environment. The DevOps engineer also has created a launch template and an Auto Scaling group for the green environment. Each Auto Scaling group deploys to a matching blue or green target group. The target group also specifies which software blue or green gets loaded on the EC2 instances. The ALB can be configured to send traffic to the blue environments target group or the

green environments target group. An Amazon Route 53 record for www example com points to the ALB.
The deployment must move traffic all at once between the software on the blue environment's EC2 instances to the newly deployed software on the green environments EC2 instances
What should the DevOps engineer do to meet these requirements?

A. Start a rolling restart to the Auto Scaling group tor the green environment to deploy the new software on the green environment's EC2 instances When the rolling restart is complete, use an AWS CLI command to update the ALB to send traffic to the green environment's target group.
B. Use an AWS CLI command to update the ALB to send traffic to the green environment's target grou
C. Then start a rolling restart of the Auto Scaling group for the green environment to deploy the new software on the green environment's EC2 instances.
D. Update the launch template to deploy the green environment's software on the blue environment's EC2 instances Keep the target groups and Auto Scaling groups unchanged in both environments Perform a rolling restart of the blue environment's EC2 instances.
E. Start a rolling restart of the Auto Scaling group for the green environment to deploy the new software on the green environment's EC2 instances When the rolling restart is complete, update the Route 53 DNS to point to the green environments endpoint on the ALB.

Answer: A

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

NEW QUESTION 92
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 95
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 97
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 101
A company has an application that is using a MySQL-compatible Amazon Aurora Multi-AZ DB cluster as the database. A cross-Region read replica has been created for disaster recovery purposes. A DevOps engineer wants to automate the promotion of the replica so it becomes the primary database instance in the event of a failure.
Which solution will accomplish this?

A. Configure a latency-based Amazon Route 53 CNAME with health checks so it points to both the primary and replica endpoint
B. Subscribe an Amazon SNS topic to Amazon RDS failure notifications from AWS CloudTrail and use that topic to invoke an AWS Lambda function that will promote the replica instance as the primary.
C. Create an Aurora custom endpoint to point to the primary database instanc
D. Configure the application to use this endpoin
E. Configure AWS CloudTrail to run an AWS Lambda function to promote the replica instance and modify the custom endpoint to point to the newly promoted instance.
F. Create an AWS Lambda function to modify the application's AWS CloudFormation template to promote the replica, apply the template to update the stack, and point the application to the newly promoted instanc
G. Create an Amazon CloudWatch alarm to invoke this Lambda function after the failure event occurs.
H. Store the Aurora endpoint in AWS Systems Manager Parameter Stor
I. Create an Amazon EventBridge event that detects the database failure and runs an AWS Lambda function to promote the replica instance and update the endpoint URL stored in AWS Systems Manager Parameter Stor
J. Code the application to reload the endpoint from Parameter Store if a database connection fails.

Answer: D

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

NEW QUESTION 102
A DevOps engineer is working on a data archival project that requires the migration of on- premises data to an Amazon S3 bucket. The DevOps engineer develops a script that incrementally archives on-premises data that is older than 1 month to Amazon S3. Data that is transferred to Amazon S3 is deleted from the on-premises location The script uses the S3 PutObject operation.
During a code review the DevOps engineer notices that the script does not verity whether the data was successfully copied to Amazon S3. The DevOps engineer must update the script to ensure that data is not corrupted during transmission. The script must use MD5 checksums to verify data integrity before the on-premises data is deleted.
Which solutions for the script will meet these requirements'? (Select TWO.)

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

Answer: BD

Explanation:
https://docs.aws.amazon.com/AmazonS3/latest/userguide/checking-object- integrity.html

NEW QUESTION 104
A company wants to deploy a workload on several hundred Amazon EC2 instances. The company will provision the EC2 instances in an Auto Scaling group by using a launch template.
The workload will pull files from an Amazon S3 bucket, process the data, and put the results into a different S3 bucket. The EC2 instances must have least-privilege permissions and must use temporary security credentials.
Which combination of steps will meet these requirements? (Select TWO.)

A. Create an IAM role that has the appropriate permissions for S3 bucket
B. Add the IAM role to an instance profile.
C. Update the launch template to include the IAM instance profile.
D. Create an IAM user that has the appropriate permissions for Amazon S3. Generate a secret key and token.
E. Create a trust anchor and profil
F. Attach the IAM role to the profile.
G. Update the launch templat
H. Modify the user data to use the new secret key and token.

Answer: AB

Explanation:
To meet the requirements of deploying a workload on several hundred EC2 instances with least-privilege permissions and temporary security credentials, the company should use an IAM role and an instance profile. An IAM role is a way to grant permissions to an entity that you trust, such as an EC2 instance. An instance profile is a container for an IAM role that you can use to pass role information to an EC2 instance when the instance starts. By using an IAM role and an instance profile, the EC2 instances can automatically receive temporary security credentials from the AWS Security Token Service (STS) and use them to access the S3 buckets. This way, the company does not need to manage or rotate any long-term credentials, such as IAM users or access keys.
To use an IAM role and an instance profile, the company should create an IAM role that has the appropriate permissions for S3 buckets. The permissions should allow the EC2 instances to read from the source S3 bucket and write to the destination S3 bucket. The company should also create a trust policy for the IAM role that specifies that EC2 is allowed to assume the role. Then, the company should add the IAM role to an instance profile. An instance profile can have only one IAM role, so the company does not need to create
multiple roles or profiles for this scenario.
Next, the company should update the launch template to include the IAM instance profile. A launch template is a way to save launch parameters for EC2 instances, such as the instance type, security group, user data, and IAM instance profile. By using a launch template, the company can ensure that all EC2 instances in the Auto Scaling group have consistent configuration and permissions. The company should specify the name or ARN of the IAM instance profile in the launch template. This way, when the Auto Scaling group launches new EC2 instances based on the launch template, they will automatically receive the IAM role and its permissions through the instance profile.
The other options are not correct because they do not meet the requirements or follow best practices. Creating an IAM user and generating a secret key and token is not a good option because it involves managing long-term credentials that need to be rotated regularly. Moreover, embedding credentials in user data is not secure because user data is visible to anyone who can describe the EC2 instance. Creating a trust anchor and profile is not a valid option because trust anchors are used for certificate-based authentication, not for IAM roles or instance profiles. Modifying user data to use a new secret key and token is also not a good option because it requires updating user data every time the credentials change, which is not scalable or efficient.
References:
? 1: AWS Certified DevOps Engineer - Professional Certification | AWS Certification
| AWS
? 2: DevOps Resources - Amazon Web Services (AWS)
? 3: Exam Readiness: AWS Certified DevOps Engineer - Professional
? : IAM Roles for Amazon EC2 - AWS Identity and Access Management
? : Working with Instance Profiles - AWS Identity and Access Management
? : Launching an Instance Using a Launch Template - Amazon Elastic Compute Cloud
? : Temporary Security Credentials - AWS Identity and Access Management

NEW QUESTION 106
A company hosts a security auditing application in an AWS account. The auditing application uses an IAM role to access other AWS accounts. All the accounts are in the same organization in AWS Organizations.

A recent security audit revealed that users in the audited AWS accounts could modify or delete the auditing application's IAM role. The company needs to prevent any modification to the auditing application's IAM role by any entity other than a trusted administrator IAM role.
Which solution will meet these requirements?

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

Answer: A

Explanation:
https://docs.aws.amazon.com/organizations/latest/userguide/orgs_manage_policies_scps. html?icmpid=docs_orgs_console
SCPs (Service Control Policies) are the best way to restrict permissions at the organizational level, which in this case would be used to restrict modifications to the IAM role used by the auditing application, while still allowing trusted administrators to make changes to it. Options C and D are not as effective because IAM permission boundaries are applied to IAM entities (users, groups, and roles), not the account itself, and must be applied to all IAM entities in the account.

NEW QUESTION 111
A company is performing vulnerability scanning for all Amazon EC2 instances across many accounts. The accounts are in an organization in AWS Organizations. Each account's VPCs are attached to a shared transit gateway. The VPCs send traffic to the internet through a central egress VPC. The company has enabled Amazon Inspector in a delegated administrator account and has enabled scanning for all member accounts.
A DevOps engineer discovers that some EC2 instances are listed in the "not scanning" tab in Amazon Inspector.
Which combination of actions should the DevOps engineer take to resolve this issue? (Choose three.)

A. Verify that AWS Systems Manager Agent is installed and is running on the EC2 instances that Amazon Inspector is not scanning.
B. Associate the target EC2 instances with security groups that allow outbound communication on port 443 to the AWS Systems Manager service endpoint.
C. Grant inspector: StartAssessmentRun permissions to the IAM role that the DevOps engineer is using.
D. Configure EC2 Instance Connect for the EC2 instances that Amazon Inspector is not scanning.
E. Associate the target EC2 instances with instance profiles that grant permissions to communicate with AWS Systems Manager.
F. Create a managed-instance activatio
G. Use the Activation Code and the Activation ID to register the EC2 instances.

Answer: ABE

Explanation:
https://docs.aws.amazon.com/inspector/latest/user/scanning-ec2.html

NEW QUESTION 114
A company uses an organization in AWS Organizations that has all features enabled. The company uses AWS Backup in a primary account and uses an AWS Key Management Service (AWS KMS) key to encrypt the backups.
The company needs to automate a cross-account backup of the resources that AWS Backup backs up in the primary account. The company configures cross-account backup in the Organizations management account. The company creates a new AWS account in the organization and configures an AWS Backup backup vault in the new account. The company creates a KMS key in the new account to encrypt the backups. Finally, the company configures a new backup plan in the primary account. The destination for the new backup plan is the backup vault in the new account.
When the AWS Backup job in the primary account is invoked, the job creates backups in the primary account. However, the backups are not copied to the new account's backup vault.
Which combination of steps must the company take so that backups can be copied to the new account's backup vault? (Select TWO.)

A. Edit the backup vault access policy in the new account to allow access to the primary account.
B. Edit the backup vault access policy in the primary account to allow access to the new account.
C. Edit the backup vault access policy in the primary account to allow access to the KMS key in the new account.
D. Edit the key policy of the KMS key in the primary account to share the key with the new account.
E. Edit the key policy of the KMS key in the new account to share the key with the primary account.

Answer: AE

Explanation:
To enable cross-account backup, the company needs to grant permissions to both the backup vault and the KMS key in the destination account. The backup vault access policy in the destination account must allow the primary account to copy backups into the vault. The key policy of the KMS key in the destination account must allow the primary account to use the key to encrypt and decrypt the backups. These steps are described in the AWS documentation12. Therefore, the correct answer is A and E.
References:
? 1: Creating backup copies across AWS accounts - AWS Backup
? 2: Using AWS Backup with AWS Organizations - AWS Backup

NEW QUESTION 116
An ecommerce company has chosen AWS to host its new platform. The company's DevOps team has started building an AWS Control Tower landing zone. The DevOps team has set the identity store within AWS IAM Identity Center (AWS Single Sign-On) to external identity provider (IdP) and has configured SAML 2.0.
The DevOps team wants a robust permission model that applies the principle of least privilege. The model must allow the team to build and manage only the team's own resources.
Which combination of steps will meet these requirements? (Choose three.)

A. Create IAM policies that include the required permission
B. Include the aws:PrincipalTag condition key.
C. Create permission set
D. Attach an inline policy that includes the required permissions and uses the aws:PrincipalTag condition key to scope the permissions.
E. Create a group in the Id
F. Place users in the grou
G. Assign the group to accounts and the permission sets in IAM Identity Center.
H. Create a group in the Id
I. Place users in the grou
J. Assign the group to OUs and IAM policies.
K. Enable attributes for access control in IAM Identity Cente
L. Apply tags to user
M. Map the tags as key-value pairs.
N. Enable attributes for access control in IAM Identity Cente
O. Map attributes from the IdP as key-value pairs.

Answer: BCF

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

NEW QUESTION 120
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 121
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 124
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 128
An online retail company based in the United States plans to expand its operations to Europe and Asia in the next six months. Its product currently runs on Amazon EC2 instances behind an Application Load Balancer. The instances run in an Amazon EC2 Auto Scaling group across multiple Availability Zones. All data is stored in an Amazon Aurora database instance.
When the product is deployed in multiple regions, the company wants a single product catalog across all regions, but for compliance purposes, its customer information and purchases must be kept in each region.
How should the company meet these requirements with the LEAST amount of application changes?

A. Use Amazon Redshift for the product catalog and Amazon DynamoDB tables for the customer information and purchases.
B. Use Amazon DynamoDB global tables for the product catalog and regional tables for the customer information and purchases.
C. Use Aurora with read replicas for the product catalog and additional local Aurora instances in each region for the customer information and purchases.
D. Use Aurora for the product catalog and Amazon DynamoDB global tables for the customer information and purchases.

Answer: C

NEW QUESTION 129
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 131
A DevOps engineer is building a multistage pipeline with AWS CodePipeline to build, verify, stage, test, and deploy an application. A manual approval stage is required between the test stage and the deploy stage. The development team uses a custom chat tool with webhook support that requires near-real-time notifications.
How should the DevOps engineer configure status updates for pipeline activity and approval requests to post to the chat tool?

A. Create an Amazon CloudWatch Logs subscription that filters on CodePipeline Pipeline Execution State Chang
B. Publish subscription events to an Amazon Simple Notification Service (Amazon SNS) topi
C. Subscribe the chat webhook URL to the SNS topic, and complete the subscription validation.
D. Create an AWS Lambda function that is invoked by AWS CloudTrail event
E. When a CodePipeline Pipeline Execution State Change event is detected, send the event details to the chat webhook URL.
F. Create an Amazon EventBridge rule that filters on CodePipeline Pipeline Execution State Chang
G. Publish the events to an Amazon Simple Notification Service (Amazon SNS) topi
H. Create an AWS Lambda function that sends event details to the chat webhook UR
I. Subscribe the function to the SNS topic.
J. Modify the pipeline code to send the event details to the chat webhook URL at the end of each stag
K. Parameterize the URL so that each pipeline can send to a different URL based on the pipeline environment.

Answer: C

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

NEW QUESTION 135
A company uses AWS Storage Gateway in file gateway mode in front of an Amazon S3 bucket that is used by multiple resources. In the morning when business begins, users do not see the objects processed by a third party the previous evening. When a DevOps engineer looks directly at the S3 bucket, the data is there, but it is missing in Storage Gateway.
Which solution ensures that all the updated third-party files are available in the morning?

A. Configure a nightly Amazon EventBridge event to invoke an AWS Lambda function to run the RefreshCache command for Storage Gateway.
B. Instruct the third party to put data into the S3 bucket using AWS Transfer for SFTP.
C. Modify Storage Gateway to run in volume gateway mode.
D. Use S3 Same-Region Replication to replicate any changes made directly in the S3 bucket to Storage Gateway.

Answer: A

Explanation:
https://docs.aws.amazon.com/storagegateway/latest/APIReference/API_RefreshCache.ht ml " It only updates the cached inventory to reflect changes in the inventory of the objects in the S3 bucket. This operation is only supported in the S3 File Gateway types."

NEW QUESTION 138
A company needs to implement failover for its application. The application includes an Amazon CloudFront distribution and a public Application Load Balancer (ALB) in an AWS Region. The company has configured the ALB as the default origin for the distribution.
After some recent application outages, the company wants a zero-second RTO. The company deploys the application to a secondary Region in a warm standby configuration. A DevOps engineer needs to automate the failover of the application to the secondary Region so that HTTP GET requests meet the desired RTO.
Which solution will meet these requirements?

A. Create a second CloudFront distribution that has the secondary ALB as the default 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 TargetHealth set to Yes for both ALB
J. Set the TTL of both records to 0. Update the distribution's origin to use the new record set.
K. Create a CloudFront function that detects HTTP 5xx status code
L. Configure the function to return a 307 Temporary Redirect error response to the secondary ALB if the function detects 5xx status code
M. Update the distribution's default behavior to send origin responses to the function.

Answer: B

Explanation:
The best solution to implement failover for the application is to use CloudFront origin groups. Origin groups allow CloudFront to automatically switch to a secondary origin when the primary origin is unavailable or returns specific HTTP status codes that indicate a failure1. This way, CloudFront can serve the requests from the secondary ALB in the secondary Region without any delay or redirection. To set up origin groups, the DevOps engineer needs to create a new origin on the distribution for the secondary ALB, create a new origin group with the original ALB as the primary origin and the secondary ALB as the secondary origin, and configure the origin group to fail over for HTTP 5xx status
codes. Then, the DevOps engineer needs to update the default behavior to use the origin group instead of the single origin2.
The other options are not as effective or efficient as the solution in option B. Option A is not suitable because creating a second CloudFront distribution will increase the complexity and cost of the application. Moreover, using Route 53 alias records with a failover policy will introduce some delay in detecting and switching to the secondary CloudFront distribution, which may not meet the zero-second RTO requirement. Option C is not feasible because CloudFront does not support using Route 53 alias records as origins3. Option D is not advisable because using a CloudFront function to redirect the requests to the secondary ALB will add an extra round-trip and latency to the failover process, which may also not meet the zero-second RTO requirement.
References:
? 1: Optimizing high availability with CloudFront origin failover - Amazon CloudFront
? 2: Creating an origin group - Amazon CloudFront
? 3: Values That You Specify When You Create or Update a Web Distribution - Amazon CloudFront

NEW QUESTION 140
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 141
A company builds an application that uses an Application Load Balancer in front of Amazon EC2 instances that are in an Auto Scaling group. The application is stateless. The Auto Scaling group uses a custom AMI that is fully prebuilt. The EC2 instances do not have a custom bootstrapping process.
The AMI that the Auto Scaling group uses was recently deleted. The Auto Scaling group's scaling activities show failures because the AMI ID does not exist.
Which combination of steps should a DevOps engineer take to meet these requirements? (Select THREE.)

A. Create a new launch template that uses the new AMI.
B. Update the Auto Scaling group to use the new launch template.
C. Reduce the Auto Scaling group's desired capacity to O.
D. Increase the Auto Scaling group's desired capacity by I.
E. Create a new AMI from a running EC2 instance in the Auto Scaling group.
F. Create a new AMI by copying the most recent public AMI of the operating system that the EC2 instances use.

Answer: ABF

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

NEW QUESTION 145
.....

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