Exam Questions AWS-Certified-DevOps-Engineer-Professional

Amazon AWS Certified DevOps Engineer Professional

https://www.2passeasy.com/dumps/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 company has an application that runs on a fleet of Amazon EC2 instances. The application requires frequent restarts. The application logs contain error messages when a restart is required. The application logs are published to a log group in Amazon CloudWatch Logs.

An Amazon CloudWatch alarm notifies an application engineer through an Amazon Simple Notification Service (Amazon SNS) topic when the logs contain a large number of restart- related error messages. The application engineer manually restarts the application on the instances after the application engineer receives a notification from the SNS topic.

A DevOps engineer needs to implement a solution to automate the application restart on the instances without restarting the instances.

Which solution will meet these requirements in the MOST operationally efficient manner?

A. Configure an AWS Systems Manager Automation runbook that runs a script to restart the application on the instance
B. Configure the SNS topic to invoke the runbook.
C. Create an AWS Lambda function that restarts the application on the instance
D. Configure the Lambda function as an event destination of the SNS topic.
E. Configure an AWS Systems Manager Automation runbook that runs a script to restart the application on the instance
F. Create an AWS Lambda function to invoke the runboo
G. Configure the Lambda function as an event destination of the SNS topic.
H. Configure an AWS Systems Manager Automation runbook that runs a script to restart the application on the instance
I. Configure an Amazon EventBridge rule that reacts when the CloudWatch alarm enters ALARM stat
J. Specify the runbook as a target of the rule.

Answer: D

Explanation:
This solution meets the requirements in the most operationally efficient manner by automating the application restart process on the instances without restarting them. When the CloudWatch alarm enters the ALARM state, the EventBridge rule is triggered, which in turn invokes the Systems Manager Automation runbook that contains the script to restart the application on the instances.

NEW QUESTION 4
A company requires that its internally facing web application be highly available. The architecture is made up of one Amazon EC2 web server instance and one NAT instance that provides outbound internet access for updates and accessing public data.
Which combination of architecture adjustments should the company implement to achieve high availability? (Choose two.)

A. Add the NAT instance to an EC2 Auto Scaling group that spans multiple Availability Zone
B. Update the route tables.
C. Create additional EC2 instances spanning multiple Availability Zone
D. Add an Application Load Balancer to split the load between them.
E. Configure an Application Load Balancer in front of the EC2 instanc
F. Configure Amazon CloudWatch alarms to recover the EC2 instance upon host failure.
G. Replace the NAT instance with a NAT gateway in each Availability Zon
H. Update the route tables.
I. Replace the NAT instance with a NAT gateway that spans multiple Availability Zone
J. Update the route tables.

Answer: BD

Explanation:
https://docs.aws.amazon.com/vpc/latest/userguide/vpc-nat-gateway.html

NEW QUESTION 5
A company detects unusual login attempts in many of its AWS accounts. A DevOps engineer must implement a solution that sends a notification to the company's security team when multiple failed login attempts occur. The DevOps engineer has already created an Amazon Simple Notification Service (Amazon SNS) topic and has subscribed the security team to the SNS topic.
Which solution will provide the notification with the LEAST operational effort?

A. Configure AWS CloudTrail to send log management events to an Amazon CloudWatch Logs log grou
B. Create a CloudWatch Logs metric filter to match failed ConsoleLogin event
C. Create a CloudWatch alarm that is based on the metric filte
D. Configure an alarm action to send messages to the SNS topic.
E. Configure AWS CloudTrail to send log management events to an Amazon S3 bucke
F. Create an Amazon Athena query that returns a failure if the query finds failed logins in the logs in the S3 bucke
G. Create an Amazon EventBridge rule to periodically run the quer
H. Create a second EventBridge rule to detect when the query fails and to send a message to the SNS topic.
I. Configure AWS CloudTrail to send log data events to an Amazon CloudWatch Logs log grou
J. Create a CloudWatch logs metric filter to match failed Consolel_ogin event
K. Create a CloudWatch alarm that is based on the metric filte
L. Configure an alarm action to send messages to the SNS topic.
M. Configure AWS CloudTrail to send log data events to an Amazon S3 bucke
N. Configure an Amazon S3 event notification for the s3:ObjectCreated event typ
O. Filter the event type by ConsoleLogin failed event
P. Configure the event notification to forward to the SNS topic.

Answer: C

Explanation:
The correct answer is C. Configuring AWS CloudTrail to send log data events to an Amazon CloudWatch Logs log group and creating a CloudWatch logs metric filter to match failed ConsoleLogin events is the simplest and most efficient way to monitor and alert on failed login attempts. Creating a CloudWatch alarm that is based on the metric filter and configuring an alarm action to send messages to the SNS topic will ensure that the security team is notified when multiple failed login attempts occur. This solution requires the least operational effort compared to the other options.
Option A is incorrect because it involves configuring AWS CloudTrail to send log management events instead of log data events. Log management events are used to track changes to CloudTrail configuration, such as creating, updating, or deleting a trail. Log data events are used to track API activity in AWS accounts, such as login attempts. Therefore, option A will not capture the failed ConsoleLogin events.
Option B is incorrect because it involves creating an Amazon Athena query and two Amazon EventBridge rules to monitor and alert on failed login attempts. This is a more complex and costly solution than using CloudWatch logs and alarms. Moreover, option B relies on the query returning a failure, which may not happen if the query is executed successfully but does not find any failed logins.
Option D is incorrect because it involves configuring AWS CloudTrail to send log data events to an Amazon S3 bucket and configuring an Amazon S3 event notification for the s3:ObjectCreated event type. This solution will not work because the s3:ObjectCreated event type does not allow filtering by ConsoleLogin failed events. The event notification will be triggered for any object created in the S3 bucket, regardless of the event type. Therefore, option D will generate a lot of false positives and unnecessary notifications. References:
? AWS CloudTrail Log File Examples
? Creating CloudWatch Alarms for CloudTrail Events: Examples
? Monitoring CloudTrail Log Files with Amazon CloudWatch Logs

NEW QUESTION 6
A DevOps engineer used an AWS Cloud Formation custom resource to set up AD Connector. The AWS Lambda function ran and created AD Connector, but Cloud Formation is not transitioning from CREATE_IN_PROGRESS to CREATE_COMPLETE.
Which action should the engineer take to resolve this issue?

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

D. Ensure the Lambda function IAM role has ds ConnectDirectory permissions for the AWS account.

Answer: B

Explanation:

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

NEW QUESTION 7
A company has a legacy application A DevOps engineer needs to automate the process of building the deployable artifact for the legacy application. The solution must store the deployable artifact in an existing Amazon S3 bucket for future deployments to reference
Which solution will meet these requirements in the MOST operationally efficient way?

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

Answer: A

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

NEW QUESTION 8
A company has deployed an application in a production VPC in a single AWS account. The application is popular and is experiencing heavy usage. The company's security team wants to add additional security, such as AWS WAF, to the application deployment. However, the application's product manager is concerned about cost and does not want to approve the change unless the security team can prove that additional security is necessary.
The security team believes that some of the application's demand might come from users that have IP addresses that are on a deny list. The security team provides the deny list to a DevOps engineer. If any of the IP addresses on the deny list access the application, the security team wants to receive automated notification in near real time so that the security team can document that the application needs additional security. The DevOps engineer creates a VPC flow log for the production VPC.
Which set of additional steps should the DevOps engineer take to meet these requirements MOST cost-effectively?

A. Create a log group in Amazon CloudWatch Log
B. Configure the VPC flow log to capture accepted traffic and to send the data to the log grou
C. Create an Amazon CloudWatch metric filter for IP addresses on the deny lis
D. Create a CloudWatch alarm with the metric filter as inpu
E. Set the period to 5 minutes and the datapoints to alarm to 1. Use an Amazon Simple Notification Service (Amazon SNS) topic to send alarm notices to the security team.
F. Create an Amazon S3 bucket for log file
G. Configure the VPC flow log to capture all traffic and to send the data to the S3 bucke
H. Configure Amazon Athena to return all log files in the S3 bucket for IP addresses on the deny lis
I. Configure Amazon QuickSight to accept data from Athena and to publish the data as a dashboard that the security team can acces
J. Create a threshold alert of 1 for successful acces
K. Configure the alert to automatically notify the security team as frequently as possible when the alert threshold is met.
L. Create an Amazon S3 bucket for log file
M. Configure the VPC flow log to capture accepted traffic and to send the data to the S3 bucke
N. Configure an Amazon OpenSearch Service cluster and domain for the log file
O. Create an AWS Lambda function to retrieve the logs from the S3 bucket, format the logs, and load the logs into the OpenSearch Service cluste
P. Schedule the Lambda function to run every 5 minute
Q. Configure an alert and condition in OpenSearch Service to send alerts to the security team through an Amazon Simple Notification Service (Amazon SNS) topic when access from the IP addresses on the deny list is detected.
R. Create a log group in Amazon CloudWatch Log
S. Create an Amazon S3 bucket to hold query result
T. Configure the VPC flow log to capture all traffic and to send the data to the log grou
. Deploy an Amazon Athena CloudWatch connector in AWS Lambd
. Connect the connector to the log grou
. Configure Athena to periodically query for all accepted traffic from the IP addresses on the deny list and to store the results in the S3 bucke
. Configure an S3 event notification to automatically notify the security team through an Amazon Simple Notification Service (Amazon SNS) topic when new objects are added to the S3 bucket.

Answer: A

NEW QUESTION 9
A company builds a container image in an AWS CodeBuild project by running Docker commands. After the container image is built, the CodeBuild project uploads the container image to an Amazon S3 bucket. The CodeBuild project has an 1AM service role that has permissions to access the S3 bucket.
A DevOps engineer needs to replace the S3 bucket with an Amazon Elastic Container Registry (Amazon ECR) repository to store the container images. The DevOps engineer creates an ECR private image repository in the same AWS Region of the CodeBuild project. The DevOps engineer adjusts the 1AM service role with the permissions that are necessary to work with the new ECR repository. The DevOps engineer also places new repository information into the docker build command and the docker push command that are used in the buildspec.yml file.

When the CodeBuild project runs a build job, the job fails when the job tries to access the ECR repository.
Which solution will resolve the issue of failed access to the ECR repository?

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

Answer: A

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

NEW QUESTION 10
A company has an on-premises application that is written in Go. A DevOps engineer must move the application to AWS. The company's development team wants to enable blue/green deployments and perform A/B testing.
Which solution will meet these requirements?

A. Deploy the application on an Amazon EC2 instance, and create an AMI of the instanc
B. Use the AMI to create an automatic scaling launch configuration that is used in an Auto Scaling grou
C. Use Elastic Load Balancing to distribute traffi
D. When changes are made to the application, a new AMI will be created, which will initiate an EC2 instance refresh.
E. Use Amazon Lightsail to deploy the applicatio
F. Store the application in a zipped format in an Amazon S3 bucke
G. Use this zipped version to deploy new versions of the application to Lightsai
H. Use Lightsail deployment options to manage the deployment.
I. Use AWS CodeArtifact to store the application cod
J. Use AWS CodeDeploy to deploy the application to a fleet of Amazon EC2 instance
K. Use Elastic Load Balancing to distribute the traffic to the EC2 instance
L. When making changes to the application, upload a new version to CodeArtifact and create a new CodeDeploy deployment.
M. Use AWS Elastic Beanstalk to host the applicatio
N. Store a zipped version of the application in Amazon S3. Use that location to deploy new versions of the applicatio
O. Use Elastic Beanstalk to manage the deployment options.

Answer: D

Explanation:
 https://aws.amazon.com/quickstart/architecture/blue-green-deployment/

NEW QUESTION 10
A company is using AWS to run digital workloads. Each application team in the company has its own AWS account for application hosting. The accounts are consolidated in an organization in AWS Organizations.
The company wants to enforce security standards across the entire organization. To avoid noncompliance because of security misconfiguration, the company has enforced the use of AWS CloudFormation. A production support team can modify resources in the production environment by using the AWS Management Console to troubleshoot and resolve application-related issues.
A DevOps engineer must implement a solution to identify in near real time any AWS
service misconfiguration that results in noncompliance. The solution must automatically remediate the issue within 15 minutes of identification. The solution also must track noncompliant resources and events in a centralized dashboard with accurate timestamps.
Which solution will meet these requirements with the LEAST development overhead?

A. Use CloudFormation drift detection to identify noncompliant resource
B. Use drift detection events from CloudFormation to invoke an AWS Lambda function for remediatio
C. Configure the Lambda function to publish logs to an Amazon CloudWatch Logs log grou
D. Configure an Amazon CloudWatch dashboard to use the log group for tracking.
E. Turn on AWS CloudTrail in the AWS account
F. Analyze CloudTrail logs by using Amazon Athena to identify noncompliant resource
G. Use AWS Step Functions to track query results on Athena for drift detection and to invoke an AWS Lambda function for remediatio
H. For tracking, set up an Amazon QuickSight dashboard that uses Athena as the data source.
I. Turn on the configuration recorder in AWS Config in all the AWS accounts to identify noncompliant resource
J. Enable AWS Security Hub with the ~no-enable-default-standards option in all the AWS account
K. Set up AWS Config managed rules and custom rule
L. Set up automatic remediation by using AWS Config conformance pack
M. For tracking, set up a dashboard on Security Hub in a designated Security Hub administrator account.
N. Turn on AWS CloudTrail in the AWS account
O. Analyze CloudTrail logs by using Amazon CloudWatch Logs to identify noncompliant resource
P. Use CloudWatch Logs filters for drift detectio
Q. Use Amazon EventBridge to invoke the Lambda function for remediatio
R. Stream filtered CloudWatch logs to Amazon OpenSearch Servic
S. Set up a dashboard on OpenSearch Service for tracking.

Answer: C

Explanation:

The best solution is to use AWS Config and AWS Security Hub to identify and remediate noncompliant resources across multiple AWS accounts. AWS Config enables continuous monitoring of the configuration of AWS resources and evaluates them against desired configurations. AWS Config can also automatically remediate noncompliant resources by using conformance packs, which are a collection of AWS Config rules and remediation actions that can be deployed as a single entity. AWS Security Hub provides a comprehensive view of the security posture of AWS accounts and resources. AWS Security Hub can aggregate and normalize the findings from AWS Config and other AWS services, as well as from partner solutions. AWS Security Hub can also be used to create a dashboard for tracking noncompliant resources and events in a centralized location.

The other options are not optimal because they either require more development overhead, do not provide near real time detection and remediation, or do not provide a centralized dashboard for tracking.

Option A is not optimal because CloudFormation drift detection is not a near real time solution. Drift detection has to be manually initiated on each stack or resource, or scheduled using a cron expression. Drift detection also does not provide remediation

actions, so a custom Lambda function has to be developed and invoked. CloudWatch Logs and dashboard can be used for tracking, but they do not provide a comprehensive view of the security posture of the AWS accounts and resources.

Option B is not optimal because CloudTrail logs analysis using Athena is not a near real time solution. Athena queries have to be manually run or scheduled using a cron expression. Athena also does not provide remediation actions, so a custom Lambda function has to be developed and invoked. Step Functions can be used to orchestrate the query and remediation workflow, but it adds more complexity and cost. QuickSight dashboard can be used for tracking, but it does not provide a comprehensive view of the security posture of the AWS accounts and resources.

Option D is not optimal because CloudTrail logs analysis using CloudWatch Logs is not a near real time solution. CloudWatch Logs filters have to be manually created or updated for each resource type and configuration change. CloudWatch Logs also does not provide remediation actions, so a custom Lambda function has to be developed and invoked. EventBridge can be used to trigger the Lambda function, but it adds more complexity and cost. OpenSearch Service dashboard can be used for tracking, but it does not provide a comprehensive view of the security posture of the AWS accounts and resources. References:
? AWS Config conformance packs
? Introducing AWS Config conformance packs
? Managing conformance packs across all accounts in your organization

NEW QUESTION 15
A company's security team requires that all external Application Load Balancers (ALBs) and Amazon API Gateway APIs are associated with AWS WAF web ACLs. The company
has hundreds of AWS accounts, all of which are included in a single organization in AWS Organizations. The company has configured AWS Config for the organization. During an audit, the company finds some externally facing ALBs that are not associated with AWS WAF web ACLs.
Which combination of steps should a DevOps engineer take to prevent future violations? (Choose two.)

A. Delegate AWS Firewall Manager to a security account.
B. Delegate Amazon GuardDuty to a security account.
C. Create an AWS Firewall Manager policy to attach AWS WAF web ACLs to any newly created ALBs and API Gateway APIs.
D. Create an Amazon GuardDuty policy to attach AWS WAF web ACLs to any newly created ALBs and API Gateway APIs.
E. Configure an AWS Config managed rule to attach AWS WAF web ACLs to any newly created ALBs and API Gateway APIs.

Answer: AC

Explanation:
 If instead you want to automatically apply the policy to existing in-scope resources, choose Auto remediate any noncompliant resources. This option creates a web ACL in each applicable account within the AWS organization and associates the web ACL with the resources in the accounts. When you choose Auto remediate any noncompliant resources, you can also choose to remove existing web ACL associations from in-scope resources, for the web ACLs that aren't managed by another active Firewall Manager policy. If you choose this option, Firewall Manager first associates the policy's web ACL with the resources, and then removes the prior associations. If a resource has an association with another web ACL that's managed by a different active Firewall Manager policy, this choice doesn't affect that association.

NEW QUESTION 18
A company has developed an AWS Lambda function that handles orders received through an API. The company is using AWS CodeDeploy to deploy the Lambda function as the final stage of a CI/CD pipeline.
A DevOps engineer has noticed there are intermittent failures of the ordering API for a few seconds after deployment. After some investigation the DevOps engineer believes the failures are due to database changes not having fully propagated before the Lambda function is invoked
How should the DevOps engineer overcome this?

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

Answer: A

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

NEW QUESTION 20
A company is divided into teams Each team has an AWS account and all the accounts are in an organization in AWS Organizations. Each team must retain full administrative rights to its AWS account. Each team also must be allowed to access only AWS services that the company approves for use AWS services must gam approval through a request and approval process.
How should a DevOps engineer configure the accounts to meet these requirements?

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

E. Create an SCP that allows access to only approved AWS service
F. Attach the SCP to the root OU of the organizatio
G. Remove the FullAWSAccess SCP from the root OU of the organization.

Answer: C

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

NEW QUESTION 24
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 26
A business has an application that consists of five independent AWS Lambda functions.
The DevOps engineer has built a CI/CD pipeline using AWS CodePipeline and AWS CodeBuild that builds tests packages and deploys each Lambda function in sequence. The pipeline uses an Amazon EventBridge rule to ensure the pipeline starts as quickly as possible after a change is made to the application source code.
After working with the pipeline for a few months the DevOps engineer has noticed the pipeline takes too long to complete.
What should the DevOps engineer implement to BEST improve the speed of the pipeline?

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

Answer: C

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

NEW QUESTION 29
A company updated the AWS Cloud Formation template for a critical business application. The stack update process failed due to an error in the updated template and AWS CloudFormation automatically began the stack rollback process Later a DevOps engineer discovered that the application was still unavailable and that the stack was in the UPDATE_ROLLBACK_FAILED state.
Which combination of actions should the DevOps engineer perform so that the stack rollback can complete successfully? (Select TWO.)

A. Attach the AWSC loud Formation FullAccess IAM policy to the AWS CtoudFormation role.
B. Automatically recover the stack resources by using AWS CloudFormation drift detection.
C. Issue a ContinueUpdateRollback command from the AWS CloudFormation console or the AWS CLI.
D. Manually adjust the resources to match the expectations of the stack.
E. Update the existing AWS CloudFormation stack by using the original template.

Answer: CD

Explanation:
https://docs.aws.amazon.com/cli/latest/reference/cloudformation/continue- update-rollback.html For a specified stack that is in the UPDATE_ROLLBACK_FAILED state, continues rolling it back to the UPDATE_ROLLBACK_COMPLETE state. Depending on the cause of the failure, you can manually fix the error and continue the rollback. By continuing the rollback, you can return your stack to a working state (the UPDATE_ROLLBACK_COMPLETE state), and then try to update the stack again.

NEW QUESTION 31
A global company manages multiple AWS accounts by using AWS Control Tower. The company hosts internal applications and public applications.

Each application team in the company has its own AWS account for application hosting. The accounts are consolidated in an organization in AWS Organizations. One of the AWS Control Tower member accounts serves as a centralized DevOps account with CI/CD pipelines that application teams use to deploy applications to their respective target AWS accounts. An 1AM role for deployment exists in the centralized DevOps account.

An application team is attempting to deploy its application to an Amazon Elastic Kubernetes Service (Amazon EKS) cluster in an application AWS account. An 1AM role for deployment exists in the application AWS account. The deployment is through an AWS CodeBuild project that is set up in the centralized DevOps account. The CodeBuild project uses an 1AM service role for CodeBuild. The deployment is failing with an Unauthorized error during attempts to connect to the cross-account EKS cluster from CodeBuild.

Which solution will resolve this error?

A. Configure the application account's deployment 1AM role to have a trust relationship with the centralized DevOps accoun
B. Configure the trust relationship to allow the sts:AssumeRole actio
C. Configure the application account's deployment 1AM role to have the required access to the EKS cluste
D. Configure the EKS cluster aws-auth ConfigMap to map the role to the appropriate system permissions.
E. Configure the centralized DevOps account's deployment I AM role to have a trust relationship with the application accoun
F. Configure the trust relationship to allow the sts:AssumeRole actio
G. Configure the centralized DevOps account's deployment 1AM role to allow the required access to CodeBuild.
H. Configure the centralized DevOps account's deployment 1AM role to have a trust relationship with the application accoun
I. Configure the trust relationship to allow the sts:AssumeRoleWithSAML actio
J. Configure the centralized DevOps account's deployment 1AM role to allow the required access to CodeBuild.
K. Configure the application account's deployment 1AM role to have a trust relationship with the AWS Control Tower management accoun
L. Configure the trust relationship to allow the sts:AssumeRole actio
M. Configure the application account's deployment 1AM role to have the required access to the EKS cluste
N. Configure the EKS cluster aws-auth ConfigMap to map the role to the appropriate system permissions.

Answer: A

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

NEW QUESTION 34
A company must encrypt all AMIs that the company shares across accounts. A DevOps engineer has access to a source account where an unencrypted custom AMI has been built. The DevOps engineer also has access to a target account where an Amazon EC2 Auto Scaling group will launch EC2 instances from the AMI. The DevOps engineer must share the AMI with the target account.
The company has created an AWS Key Management Service (AWS KMS) key in the source account.
Which additional steps should the DevOps engineer perform to meet the requirements? (Choose three.)

A. In the source account, copy the unencrypted AMI to an encrypted AM
B. Specify the KMS key in the copy action.
C. In the source account, copy the unencrypted AMI to an encrypted AM
D. Specify the default Amazon Elastic Block Store (Amazon EBS) encryption key in the copy action.
E. In the source account, create a KMS grant that delegates permissions to the Auto Scaling group service-linked role in the target account.
F. In the source account, modify the key policy to give the target account permissions to create a gran
G. In the target account, create a KMS grant that delegates permissions to the Auto Scaling group service-linked role.
H. In the source account, share the unencrypted AMI with the target account.
I. In the source account, share the encrypted AMI with the target account.

Answer: ADF

Explanation:
The Auto Scaling group service-linked role must have a specific grant in the source account in order to decrypt the encrypted AMI. This is because the service-linked role does not have permissions to assume the default IAM role in the source account. The following steps are required to meet the requirements:
? In the source account, copy the unencrypted AMI to an encrypted AMI. Specify the KMS key in the copy action.
? In the source account, create a KMS grant that delegates permissions to the Auto Scaling group service-linked role in the target account.
? In the source account, share the encrypted AMI with the target account.
? In the target account, attach the KMS grant to the Auto Scaling group service- linked role.
The first three steps are the same as the steps that I described earlier. The fourth step is required to grant the Auto Scaling group service-linked role permissions to decrypt the AMI
in the target account.

NEW QUESTION 35
A company uses AWS CodeArtifact to centrally store Python packages. The CodeArtifact repository is configured with the following repository policy.

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

A development team is building a new project in an account that is in an organization in AWS Organizations. The development team wants to use a Python library that has already been stored in the CodeArtifact repository in the organization. The development team uses AWS CodePipeline and AWS CodeBuild to build the new application. The CodeBuild job that the development team uses to build the application is configured to run in a VPC Because of compliance requirements the VPC has no internet connectivity.

The development team creates the VPC endpoints for CodeArtifact and updates the CodeBuild buildspec yaml file. However, the development team cannot download the Python library from the repository.

Which combination of steps should a DevOps engineer take so that the development team can use Code Artifact? (Select TWO.)

A. Create an Amazon S3 gateway endpoint Update the route tables for the subnets thatare running the CodeBuild job.
B. Update the repository policy's Principal statement to include the ARN of the role that the CodeBuild project uses.
C. Share the CodeArtifact repository with the organization by using AWS Resource Access Manager (AWS RAM).
D. Update the role that the CodeBuild project uses so that the role has sufficient permissions to use the CodeArtifact repository.
E. Specify the account that hosts the repository as the delegated administrator for CodeArtifact in the organization.

Answer: AD

Explanation:
 "AWS CodeArtifact operates in multiple Availability Zones and stores artifact data and metadata in Amazon S3 and Amazon DynamoDB. Your encrypted data is redundantly stored across multiple facilities and multiple devices in each facility, making it highly available and highly durable."
https://aws.amazon.com/codeartifact/features/ With no internet connectivity, a gateway endpoint becomes necessary to access S3.

NEW QUESTION 36
A highly regulated company has a policy that DevOps engineers should not log in to their Amazon EC2 instances except in emergencies. It a DevOps engineer does log in the security team must be notified within 15 minutes of the occurrence.
Which solution will meet these requirements'?

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

Answer: B

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

NEW QUESTION 41
A company is hosting a web application in an AWS Region. For disaster recovery purposes, a second region is being used as a standby. Disaster recovery requirements state that session data must be replicated between regions in near-real time and 1% of requests should route to the secondary region to continuously

verify system functionality. Additionally, if there is a disruption in service in the main region, traffic should be automatically routed to the secondary region, and the secondary region must be able to
scale up to handle all traffic.
How should a DevOps engineer meet these requirements?

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

Answer: D

NEW QUESTION 42
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 47
A company needs to ensure that flow logs remain configured for all existing and new VPCs in its AWS account. The company uses an AWS CloudFormation stack to manage its VPCs. The company needs a solution that will work for any VPCs that any IAM user creates.
Which solution will meet these requirements?

A. Add the resource to the CloudFormation stack that creates the VPCs.
B. Create an organization in AWS Organization
C. Add the company's AWS account to the organizatio
D. Create an SCP to prevent users from modifying VPC flow logs.
E. Turn on AWS Confi
F. Create an AWS Config rule to check whether VPC flow logs are turned o
G. Configure automatic remediation to turn on VPC flow logs.
H. Create an IAM policy to deny the use of API calls for VPC flow log
I. Attach the IAM policy to all IAM users.

Answer: C

Explanation:
To meet the requirements of ensuring that flow logs remain configured for all existing and new VPCs in the AWS account, the company should use AWS Config and automatic remediation. AWS Config is a service that enables customers to assess, audit, and evaluate the configurations of their AWS resources. AWS Config continuously monitors and records the configuration changes of the AWS resources and evaluates them against desired configurations. Customers can use AWS Config rules to define the desired configuration state of their AWS resources and trigger actions when a resource configuration violates a rule.
One of the AWS Config rules that customers can use is vpc-flow-logs-enabled, which checks whether VPC flow logs are enabled for all VPCs in an AWS account. Customers can also configure automatic remediation for this rule, which means that AWS Config will automatically enable VPC flow logs for any VPCs that do not have them enabled. Customers can specify the destination (CloudWatch Logs or S3) and the traffic type (all, accept, or reject) for the flow logs as remediation parameters. By using AWS Config and automatic remediation, the company can ensure that flow logs remain configured for all existing and new VPCs in its AWS account, regardless of who creates them or how they are created.
The other options are not correct because they do not meet the requirements or follow best practices. Adding the resource to the CloudFormation stack that creates the VPCs is not a sufficient solution because it will only work for VPCs that are created by using the CloudFormation stack. It will not work for VPCs that are created by using other methods, such as the console or the API. Creating an organization in AWS Organizations and creating an SCP to prevent users from modifying VPC flow logs is not a good solution because it will not ensure that flow logs are enabled for all VPCs in the first place. It will only prevent users from disabling or changing flow logs after they are enabled. Creating an IAM policy to deny the use of API calls for VPC flow logs and attaching it to all IAM users is not a valid solution because it will prevent users from enabling or disabling flow logs at all.
It will also not work for VPCs that are created by using other methods, such as the console or CloudFormation.
References:

? 1: AWS::EC2::FlowLog - AWS CloudFormation
? 2: Amazon VPC Flow Logs extends CloudFormation Support to custom format subscriptions, 1-minute aggregation intervals and tagging
? 3: Logging IP traffic using VPC Flow Logs - Amazon Virtual Private Cloud
? : About AWS Config - AWS Config
? : vpc-flow-logs-enabled - AWS Config
? : Remediate Noncompliant Resources with AWS Config Rules - AWS Config

NEW QUESTION 48
A company is implementing AWS CodePipeline to automate its testing process The company wants to be notified when the execution state fails and used the following custom event pattern in Amazon EventBridge:

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

Which type of events will match this event pattern?

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

Answer: B

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

NEW QUESTION 52
To run an application, a DevOps engineer launches an Amazon EC2 instance with public IP addresses in a public subnet. A user data script obtains the application artifacts and installs them on the instances upon launch. A change to the security classification of the application now requires the instances to run with no access to the internet. While the instances launch successfully and show as healthy, the application does not seem to be installed.
Which of the following should successfully install the application while complying with the new rule?

A. Launch the instances in a public subnet with Elastic IP addresses attache
B. Once the application is installed and running, run a script to disassociate the Elastic IP addresses afterwards.
C. Set up a NAT gatewa
D. Deploy the EC2 instances to a private subne
E. Update the private subnet's route table to use the NAT gateway as the default route.
F. Publish the application artifacts to an Amazon S3 bucket and create a VPC endpoint for S3. Assign an IAM instance profile to the EC2 instances so they can read the application artifacts from the S3 bucket.
G. Create a security group for the application instances and allow only outbound traffic to the artifact repositor
H. Remove the security group rule once the install is complete.

Answer: C

Explanation:
 EC2 instances running in private subnets of a VPC can now have controlled access to S3 buckets, objects, and API functions that are in the same region as the VPC. You can use an S3 bucket policy to indicate which VPCs and which VPC Endpoints have access to your S3 buckets 1-
https://aws.amazon.com/pt/blogs/aws/new-vpc-endpoint-for- amazon-s3/

NEW QUESTION 56
A company hosts applications in its AWS account Each application logs to an individual Amazon CloudWatch log group. The company's CloudWatch costs for ingestion are increasing
A DevOps engineer needs to Identify which applications are the source of the increased logging costs.
Which solution Will meet these requirements?

A. Use CloudWatch metrics to create a custom expression that Identifies the CloudWatch log groups that have the most data being written to them.

B. Use CloudWatch Logs Insights to create a set of queries for the application log groups to Identify the number of logs written for a period of time
C. Use AWS Cost Explorer to generate a cost report that details the cost for CloudWatch usage
D. Use AWS CloudTrail to filter for CreateLogStream events for each application

Answer: C

Explanation:
The correct answer is C.
A comprehensive and detailed explanation is:
? Option A is incorrect because using CloudWatch metrics to create a custom expression that identifies the CloudWatch log groups that have the most data being written to them is not a valid solution. CloudWatch metrics do not provide information about the size or volume of data being ingested by CloudWatch logs. CloudWatch metrics only provide information about the number of events, bytes, and errors that occur within a log group or stream. Moreover, creating a custom expression with CloudWatch metrics would require using the search_web tool, which is not necessary for this use case.
? Option B is incorrect because using CloudWatch Logs Insights to create a set of queries for the application log groups to identify the number of logs written for a period of time is not a valid solution. CloudWatch Logs Insights can help analyze and filter log events based on patterns and expressions, but it does not provide information about the cost or billing of CloudWatch logs. CloudWatch Logs Insights also charges based on the amount of data scanned by each query, which could increase the logging costs further.
? Option C is correct because using AWS Cost Explorer to generate a cost report that details the cost for CloudWatch usage is a valid solution. AWS Cost Explorer is a tool that helps visualize, understand, and manage AWS costs and usage over time. AWS Cost Explorer can generate custom reports that show the breakdown of costs by service, region, account, tag, or any other dimension. AWS Cost Explorer can also filter and group costs by usage type, which can help identify the specific CloudWatch log groups that are the source of the increased logging costs.
? Option D is incorrect because using AWS CloudTrail to filter for CreateLogStream events for each application is not a valid solution. AWS CloudTrail is a service that records API calls and account activity for AWS services, including CloudWatch logs. However, AWS CloudTrail does not provide information about the cost or billing of CloudWatch logs. Filtering for CreateLogStream events would only show when a new log stream was created within a log group, but not how much data was ingested or stored by that log stream.
References:
? CloudWatch Metrics
? CloudWatch Logs Insights
? AWS Cost Explorer
? AWS CloudTrail

NEW QUESTION 61
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 63
A security review has identified that an AWS CodeBuild project is downloading a database population script from an Amazon S3 bucket using an unauthenticated request. The security team does not allow unauthenticated requests to S3 buckets for this project.
How can this issue be corrected in the MOST secure manner?

A. Add the bucket name to the AllowedBuckets section of the CodeBuild project setting
B. Update the build spec to use the AWS CLI to download the database population script.
C. Modify the S3 bucket settings to enable HTTPS basic authentication and specify a toke
D. Update the build spec to use cURL to pass the token and download the database population script.
E. Remove unauthenticated access from the S3 bucket with a bucket polic
F. Modify the service role for the CodeBuild project to include Amazon S3 acces
G. Use the AWS CLI to download the database population script.
H. Remove unauthenticated access from the S3 bucket with a bucket polic
I. Use the AWS CLI to download the database population script using an IAM access key and a secret access key.

Answer: C

Explanation:
A bucket policy is a resource-based policy that defines who can access a specific S3 bucket and what actions they can perform on it. By removing unauthenticated access from the bucket policy, you can prevent anyone without valid credentials from accessing the bucket. A service role is an IAM role that allows an AWS service, such as CodeBuild, to perform actions on your behalf. By modifying the service role for the CodeBuild project to include Amazon S3 access, you can grant the project permission to read and write objects in the S3 bucket. The AWS CLI is a command-line tool that allows you to interact with AWS services, such as S3, using commands in your terminal. By using the AWS CLI to download the database population script, you can leverage the service role credentials and encryption to secure the data transfer.
For more information, you can refer to these web pages:
? [Using bucket policies and user policies - Amazon Simple Storage Service]
? [Create a service role for CodeBuild - AWS CodeBuild]
? [AWS Command Line Interface]

NEW QUESTION 66
An IT team has built an AWS CloudFormation template so others in the company can quickly and reliably deploy and terminate an application. The template creates an Amazon EC2 instance with a user data script to install the application and an Amazon S3 bucket that the application uses to serve static webpages while it is running.
All resources should be removed when the CloudFormation stack is deleted. However, the team observes that CloudFormation reports an error during stack deletion, and the S3 bucket created by the stack is not deleted.
How can the team resolve the error in the MOST efficient manner to ensure that all resources are deleted without errors?

A. Add a DelelionPolicy attribute to the S3 bucket resource, with the value Delete forcing the bucket to be removed when the stack is deleted.
B. Add a custom resource with an AWS Lambda function with the DependsOn attribute specifying the S3 bucket, and an IAM rol
C. Write the Lambda function to delete all objects from the bucket when RequestType is Delete.
D. Identify the resource that was not delete
E. Manually empty the S3 bucket and then delete it.
F. Replace the EC2 and S3 bucket resources with a single AWS OpsWorks Stacks resourc
G. Define a custom recipe for the stack to create and delete the EC2 instance and the S3 bucket.

Answer: B

Explanation:
https://aws.amazon.com/premiumsupport/knowledge-center/cloudformation-s3-custom-resources/

NEW QUESTION 70
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 74
A company has multiple development groups working in a single shared AWS account. The Senior Manager of the groups wants to be alerted via a third-party API call when the creation of resources approaches the service limits for the account.
Which solution will accomplish this with the LEAST amount of development effort?

A. Create an Amazon CloudWatch Event rule that runs periodically and targets an AWS Lambda functio
B. Within the Lambda function, evaluate the current state of the AWS environment and compare deployed resource values to resource limits on the accoun
C. Notify the Senior Manager if the account is approaching a service limit.
D. Deploy an AWS Lambda function that refreshes AWS Trusted Advisor checks, and configure an Amazon CloudWatch Events rule to run the Lambda function periodicall
E. Create another CloudWatch Events rule with an event pattern matching Trusted Advisor events and a target Lambda functio
F. In the target Lambda function, notify the Senior Manager.
G. Deploy an AWS Lambda function that refreshes AWS Personal Health Dashboard checks, and configure an Amazon CloudWatch Events rule to run the Lambda function periodicall
H. Create another CloudWatch Events rule with an event pattern matching Personal Health Dashboard events and a target Lambda functio
I. In the target Lambda function, notify the Senior Manager.
J. Add an AWS Config custom rule that runs periodically, checks the AWS service limit status, and streams notifications to an Amazon SNS topi
K. Deploy an AWS Lambda function that notifies the Senior Manager, and subscribe the Lambda function to the SNS topic.

Answer: B

Explanation:
To meet the requirements, the company needs to create a solution that alerts the Senior Manager when the creation of resources approaches the service limits for the account with the least amount of development effort. The company can use AWS Trusted Advisor, which is a service that provides best practice recommendations for cost optimization, performance, security, and service limits. The company can deploy an AWS Lambda function that refreshes Trusted Advisor checks, and configure an Amazon CloudWatch Events rule to run the Lambda function periodically. This will ensure that Trusted Advisor checks are up to date and reflect the current state of the account. The company can then create another CloudWatch Events rule with an event pattern matching Trusted Advisor events and a target Lambda function. The event pattern can filter for events related to service limit checks and their status. The target Lambda function can notify the Senior Manager via a third-party API call if the event indicates that the account is approaching or exceeding a service limit.

NEW QUESTION 79
A company's application development team uses Linux-based Amazon EC2 instances as bastion hosts. Inbound SSH access to the bastion hosts is restricted to specific IP addresses, as defined in the associated security groups. The company's security team wants to receive a notification if the security group rules are modified to allow SSH access from any IP address.
What should a DevOps engineer do to meet this requirement?

A. Create an Amazon EventBridge rule with a source of aws.cloudtrail and the event name AuthorizeSecurityGroupIngres
B. Define an Amazon Simple Notification Service (Amazon SNS) topic as the target.
C. Enable Amazon GuardDuty and check the findings for security groups in AWS Security Hu
D. Configure an Amazon EventBridge rule with a custom pattern that matches GuardDuty events with an output of NON_COMPLIAN
E. Define an Amazon Simple Notification Service (Amazon SNS) topic as the target.
F. Create an AWS Config rule by using the restricted-ssh managed rule to check whether security groups disallow unrestricted incoming SSH traffi
G. Configure automatic remediation to publish a message to an Amazon Simple Notification Service (Amazon SNS) topic.
H. Enable Amazon Inspecto
I. Include the Common Vulnerabilities and Exposures-1.1 rules package to check the security groups that are associated with the bastion host
J. Configure Amazon Inspector to publish a message to an Amazon Simple Notification Service (Amazon SNS) topic.

Answer: A

Explanation:
https://aws.amazon.com/premiumsupport/knowledge-center/monitor-security-group-changes-ec2/

NEW QUESTION 84
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 85
A DevOps team uses AWS CodePipeline, AWS CodeBuild, and AWS CodeDeploy to deploy an application. The application is a REST API that uses AWS Lambda functions and Amazon API Gateway Recent deployments have introduced errors that have affected many customers.
The DevOps team needs a solution that reverts to the most recent stable version of the application when an error is detected. The solution must affect the fewest customers possible.
Which solution Will meet these requirements With the MOST operational efficiency?

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

Answer: B

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

NEW QUESTION 87
An Amazon EC2 instance is running in a VPC and needs to download an object from a restricted Amazon S3 bucket. When the DevOps engineer tries to download the object, an AccessDenied error is received,
What are the possible causes tor this error? (Select TWO,)

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

Answer: BD

Explanation:
These are the possible causes for the AccessDenied error because they affect the permissions to access the S3 object from the EC2 instance. An S3 bucket policy is a resource-based policy that defines who can access the bucket and its objects, and what actions they can perform. An IAM role is an identity that can be assumed by an EC2 instance to grant it permissions to access AWS services and resources. If there is an error in the S3 bucket policy or the IAM role configuration, such as a missing or incorrect statement, condition, or principal, then the EC2 instance may not have the necessary permissions to download the object from the S3 bucket . https://docs.aws.amazon.com/AmazonS3/latest/userguide/example-bucket-policies.html
https://docs.aws.amazon.com/AWSEC2/latest/2UserGuide/iam-roles-for-amazon-ec2.html

NEW QUESTION 88
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 89
A company's application is currently deployed to a single AWS Region. Recently, the company opened a new office on a different continent. The users in the new office are experiencing high latency. The company's application runs on Amazon EC2 instances behind an Application Load Balancer (ALB) and uses Amazon DynamoDB as the database layer. The instances run in an EC2 Auto Scaling group across multiple Availability Zones. A DevOps engineer is tasked with minimizing application response times and improving availability for users in both Regions.
Which combination of actions should be taken to address the latency issues? (Choose three.)

A. Create a new DynamoDB table in the new Region with cross-Region replication enabled.
B. Create new ALB and Auto Scaling group global resources and configure the new ALB to direct traffic to the new Auto Scaling group.
C. Create new ALB and Auto Scaling group resources in the new Region and configure the new ALB to direct traffic to the new Auto Scaling group.
D. Create Amazon Route 53 records, health checks, and latency-based routing policies to route to the ALB.
E. Create Amazon Route 53 aliases, health checks, and failover routing policies to route to the ALB.
F. Convert the DynamoDB table to a global table.

Answer: CDF

Explanation:
 C. Create new ALB and Auto Scaling group resources in the new Region and configure the new ALB to direct traffic to the new Auto Scaling group. This will allow users in the new Region to access the application with lower latency by reducing the network hops between the user and the application servers.
* D. Create Amazon Route 53 records, health checks, and latency-based routing policies to route to the ALB. This will enable Route 53 to route user traffic to the nearest healthy ALB, based on the latency between the user and the ALBs.
* F. Convert the DynamoDB table to a global table. This will enable reads and writes to the table in both Regions with low latency, improving the overall response time of the application

NEW QUESTION 90
A company has multiple accounts in an organization in AWS Organizations. The company's SecOps team needs to receive an Amazon Simple Notification Service (Amazon SNS) notification if any account in the organization turns off the Block Public Access feature on an Amazon S3 bucket. A DevOps engineer must implement this change without affecting the operation of any AWS accounts. The implementation must ensure that individual member accounts in the organization cannot turn off the notification.
Which solution will meet these requirements?

A. Designate an account to be the delegated Amazon GuardDuty administrator accoun
B. Turn on GuardDuty for all accounts across the organizatio
C. In the GuardDuty administrator account, create an SNS topi
D. Subscribe the SecOps team's email address to the SNS topi
E. In the same account, create an Amazon EventBridge rule that uses an event pattern for GuardDuty findings and a target of the SNS topic.
F. Create an AWS CloudFormation template that creates an SNS topic and subscribes the SecOps team's email address to the SNS topi
G. In the template, include an Amazon EventBridge rule that uses an event pattern of CloudTrail activity for s3:PutBucketPublicAccessBlock and a target of the SNS topi
H. Deploy the stack to every account in the organization by using CloudFormation StackSets.
I. Turn on AWS Config across the organizatio
J. In the delegated administrator account, create an SNS topi
K. Subscribe the SecOps team's email address to the SNS topi
L. Deploy a conformance pack that uses the s3-bucket-level-public-access-prohibited AWS Config managed rule in each account and uses an AWS Systems Manager document to publish an event to the SNS topic to notify the SecOps team.
M. Turn on Amazon Inspector across the organizatio
N. In the Amazon Inspector delegated administrator account, create an SNS topi
O. Subscribe the SecOps team's email address to the SNS topi
P. In the same account, create an Amazon EventBridge rule that uses an event pattern for public network exposure of the S3 bucket and publishes an event to the SNS topic to notify the SecOps team.

Answer: C

Explanation:
 Amazon GuardDuty is primarily on threat detection and response, not configuration monitoring A conformance pack is a collection of AWS Config rules and remediation actions that can be easily deployed as a single entity in an account and a Region or across an organization in AWS Organizations. https://docs.aws.amazon.com/config/latest/developerguide/conformance-packs.html https://docs.aws.amazon.com/config/latest/developerguide/s3-account-level-public-access- blocks.html

NEW QUESTION 93
A company sells products through an ecommerce web application The company wants a dashboard that shows a pie chart of product transaction details. The company wants to integrate the dashboard With the company's existing Amazon CloudWatch dashboards
Which solution Will meet these requirements With the MOST operational efficiency?

A. Update the ecommerce application to emit a JSON object to a CloudWatch log group for each processed transactio
B. Use CloudWatch Logs Insights to query the log group and to visualize the results in a pie chart format Attach the results to the desired CloudWatch dashboard.
C. Update the ecommerce application to emit a JSON object to an Amazon S3 bucket for each processed transactio
D. Use Amazon Athena to query the S3 bucket and to visualize the results In a Pie chart forma

E. Export the results from Athena Attach the results to the desired CloudWatch dashboard
F. Update the ecommerce application to use AWS X-Ray for instrumentatio
G. Create a new X-Ray subsegment Add an annotation for each processed transactio
H. Use X-Ray traces to query the data and to visualize the results in a pie chart format Attach the results to the desired CloudWatch dashboard
I. Update the ecommerce application to emit a JSON object to a CloudWatch log group for each processed transaction_ Create an AWS Lambda function to aggregate and write the results to Amazon DynamoD
J. Create a Lambda subscription filter for the log fil
K. Attach the results to the desired CloudWatch dashboard.

Answer: A

Explanation:
The correct answer is A.
A comprehensive and detailed explanation is:
? Option A is correct because it meets the requirements with the most operational efficiency. Updating the ecommerce application to emit a JSON object to a CloudWatch log group for each processed transaction is a simple and cost- effective way to collect the data needed for the dashboard. Using CloudWatch Logs Insights to query the log group and to visualize the results in a pie chart format is also a convenient and integrated solution that leverages the existing CloudWatch dashboards. Attaching the results to the desired CloudWatch dashboard is straightforward and does not require any additional steps or services.
? Option B is incorrect because it introduces unnecessary complexity and cost.
Updating the ecommerce application to emit a JSON object to an Amazon S3 bucket for each processed transaction is a valid way to store the data, but it requires creating and managing an S3 bucket and its permissions. Using Amazon Athena to query the S3 bucket and to visualize the results in a pie chart format is also a valid way to analyze the data, but it incurs charges based on the amount of
data scanned by each query. Exporting the results from Athena and attaching them to the desired CloudWatch dashboard is also an extra step that adds more overhead and latency.
? Option C is incorrect because it uses AWS X-Ray for an inappropriate purpose.
Updating the ecommerce application to use AWS X-Ray for instrumentation is a good practice for monitoring and tracing distributed applications, but it is not designed for aggregating product transaction details. Creating a new X-Ray subsegment and adding an annotation for each processed transaction is possible, but it would clutter the X-Ray service map and make it harder to debug performance issues. Using X-Ray traces to query the data and to visualize the results in a pie chart format is also possible, but it would require custom code and logic that are not supported by X-Ray natively. Attaching the results to the desired CloudWatch dashboard is also not supported by X-Ray directly, and would require additional steps or services.
? Option D is incorrect because it introduces unnecessary complexity and cost.
Updating the ecommerce application to emit a JSON object to a CloudWatch log group for each processed transaction is a simple and cost-effective way to collect the data needed for the dashboard, as in option A. However, creating an AWS Lambda function to aggregate and write the results to Amazon DynamoDB is redundant, as CloudWatch Logs Insights can already perform aggregation queries on log data. Creating a Lambda subscription filter for the log file is also redundant, as CloudWatch Logs Insights can already access log data directly. Attaching the results to the desired CloudWatch dashboard would also require additional steps or services, as DynamoDB does not support native integration with CloudWatch dashboards.
References:
? CloudWatch Logs Insights
? Amazon Athena
? AWS X-Ray
? AWS Lambda
? Amazon DynamoDB

NEW QUESTION 98
A company has multiple member accounts that are part of an organization in AWS Organizations. The security team needs to review every Amazon EC2 security group and their inbound and outbound rules. The security team wants to programmatically retrieve this information from the member accounts using an AWS Lambda function in the management account of the organization.
Which combination of access changes will meet these requirements? (Choose three.)

A. Create a trust relationship that allows users in the member accounts to assume the management account IAM role.
B. Create a trust relationship that allows users in the management account to assume the IAM roles of the member accounts.
C. Create an IAM role in each member account that has access to the AmazonEC2ReadOnlyAccess managed policy.
D. Create an I AM role in each member account to allow the sts:AssumeRole action against the management account IAM role's ARN.
E. Create an I AM role in the management account that allows the sts:AssumeRole action against the member account IAM role's ARN.
F. Create an IAM role in the management account that has access to the AmazonEC2ReadOnlyAccess managed policy.

Answer: BCE

Explanation:
 https://aws.amazon.com/premiumsupport/knowledge-center/lambda-function-assume-iam-role/ https://kreuzwerker.de/post/aws-multi-account-setups-reloaded

NEW QUESTION 100
A company's developers use Amazon EC2 instances as remote workstations. The company is concerned that users can create or modify EC2 security groups to allow unrestricted inbound access.
A DevOps engineer needs to develop a solution to detect when users create unrestricted security group rules. The solution must detect changes to security group rules in near real time, remove unrestricted rules, and send email notifications to the security team. The DevOps engineer has created an AWS Lambda function that checks for security group ID from input, removes rules that grant unrestricted access, and sends notifications through Amazon Simple Notification Service (Amazon SNS).
What should the DevOps engineer do next to meet the requirements?

A. Configure the Lambda function to be invoked by the SNS topi
B. Create an AWS CloudTrail subscription for the SNS topi
C. Configure a subscription filter for security group modification events.
D. Create an Amazon EventBridge scheduled rule to invoke the Lambda functio
E. Define a schedule pattern that runs the Lambda function every hour.
F. Create an Amazon EventBridge event rule that has the default event bus as the sourc
G. Define the rule's event pattern to match EC2 security group creation and modification event
H. Configure the rule to invoke the Lambda function.
I. Create an Amazon EventBridge custom event bus that subscribes to events from all AWS service
J. Configure the Lambda function to be invoked by the custom event bus.

Answer: C

Explanation:
To meet the requirements, the DevOps engineer should create an Amazon EventBridge event rule that has the default event bus as the source. The rule's event pattern should match EC2 security group creation and modification events, and it should be configured to invoke the Lambda function. This solution will allow for near real-time detection of security group rule changes and will trigger the Lambda function to remove any unrestricted rules and send email notifications to the security team. https://repost.aws/knowledge-center/monitor-security-group-changes-ec2

NEW QUESTION 101
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 102
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 103
A company uses AWS Organizations and AWS Control Tower to manage all the company's AWS accounts. The company uses the Enterprise Support plan.
A DevOps engineer is using Account Factory for Terraform (AFT) to provision new accounts. When new accounts are provisioned, the DevOps engineer notices that the support plan for the new accounts is set to the Basic Support plan. The DevOps engineer needs to implement a solution to provision the new accounts with the Enterprise Support plan.
Which solution will meet these requirements?

A. Use an AWS Config conformance pack to deploy the account-part-of-organizations AWS Config rule and to automatically remediate any noncompliant accounts.
B. Create an AWS Lambda function to create a ticket for AWS Support to add the account to the Enterprise Support pla
C. Grant the Lambda function the support:ResolveCase permission.
D. Add an additional value to the control_tower_parameters input to set the AWSEnterpriseSupport parameter as the organization's management account number.
E. Set the aft_feature_enterprise_support feature flag to True in the AFT deployment input configuratio
F. Redeploy AFT and apply the changes.

Answer: D

Explanation:
AWS Organizations is a service that helps to manage multiple AWS accounts. AWS Control Tower is a service that makes it easy to set up and govern secure, compliant multi-account AWS environments. Account Factory for Terraform (AFT) is an AWS Control Tower feature that provisions new accounts using Terraform templates. To provision new accounts with the Enterprise Support plan, the DevOps engineer can set the aft_feature_enterprise_support feature flag to True in the AFT deployment input configuration. This flag enables the Enterprise Support plan for newly provisioned accounts.
https://docs.aws.amazon.com/controltower/latest/userguide/aft-feature-options.html

NEW QUESTION 108
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 113
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 114
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 116
A company uses AWS CloudFormation stacks to deploy updates to its application. The stacks consist of different resources. The resources include AWS Auto Scaling groups, Amazon EC2 instances, Application Load Balancers (ALBs), and other resources that are necessary to launch and maintain independent stacks. Changes to application resources outside of CloudFormation stack updates are not allowed.
The company recently attempted to update the application stack by using the AWS CLI. The stack failed to update and produced the following error message:
"ERROR: both the deployment and the CloudFormation stack rollback failed. The deployment failed because the following resource(s) failed to update: [AutoScalingGroup]."
The stack remains in a status of UPDATE_ROLLBACK_FAILED. * Which solution will resolve this issue?

A. Update the subnet mappings that are configured for the ALB
B. Run the aws cloudformation update-stack-set AWS CLI command.
C. Update the 1AM role by providing the necessary permissions to update the stac

D. Run the aws cloudformation continue-update-rollback AWS CLI command.
E. Submit a request for a quota increase for the number of EC2 instances for the accoun
F. Run the aws cloudformation cancel-update-stack AWS CLI command.
G. Delete the Auto Scaling group resourc
H. Run the aws cloudformation rollback-stack AWS CLI command.

Answer: B

Explanation:
https://repost.aws/knowledge-center/cloudformation-update-rollback-failed If your stack is stuck in the UPDATE_ROLLBACK_FAILED state after a failed update, then the only actions that you can perform on the stack are the ContinueUpdateRollback or DeleteStack operations.

NEW QUESTION 121
A company is using AWS CodePipeline to deploy an application. According to a new guideline, a member of the company's security team must sign off on any application changes before the changes are deployed into production. The approval must be recorded and retained.
Which combination of actions will meet these requirements? (Select TWO.)

A. Configure CodePipeline to write actions to Amazon CloudWatch Logs.
B. Configure CodePipeline to write actions to an Amazon S3 bucket at the end of each pipeline stage.
C. Create an AWS CloudTrail trail to deliver logs to Amazon S3.
D. Create a CodePipeline custom action to invoke an AWS Lambda function for approva
E. Create a policy that gives the security team access to manage CodePipeline custom actions.
F. Create a CodePipeline manual approval action before the deployment ste
G. Create a policy that grants the security team access to approve manual approval stages.

Answer: CE

Explanation:
To meet the new guideline for application deployment, the company can use a combination of AWS CodePipeline and AWS CloudTrail. A manual approval action in CodePipeline allows the security team to review and approve changes before they are deployed. This action can be configured to pause the pipeline until approval is granted, ensuring that no changes move to production without the necessary sign-off. Additionally, by creating an AWS CloudTrail trail, all actions taken within CodePipeline, including approvals, are recorded and delivered to an Amazon S3 bucket. This provides an audit trail that can be retained for compliance and review purposes.
References:
? AWS CodePipeline's manual approval action provides a way to ensure that a member of the security team can review and approve changes before they are deployed1.
? AWS CloudTrail integration with CodePipeline allows for the recording and retention of all pipeline actions, including approvals, which can be stored in Amazon S3 for record-keeping2.

NEW QUESTION 123
......

THANKS FOR TRYING THE DEMO OF OUR PRODUCT

Visit Our Site to Purchase the Full Set of Actual AWS-Certified-DevOps-Engineer-Professional Exam Questions With Answers.

We Also Provide Practice Exam Software That Simulates Real Exam Environment And Has Many Self-Assessment Features. Order the AWS-Certified-DevOps-Engineer-Professional Product From:

https://www.2passeasy.com/dumps/AWS-Certified-DevOps-Engineer-Professional/

Money Back Guarantee

AWS-Certified-DevOps-Engineer-Professional Practice Exam Features:

* AWS-Certified-DevOps-Engineer-Professional Questions and Answers Updated Frequently

* AWS-Certified-DevOps-Engineer-Professional Practice Questions Verified by Expert Senior Certified Staff

* AWS-Certified-DevOps-Engineer-Professional Most Realistic Questions that Guarantee you a Pass on Your FirstTry

* AWS-Certified-DevOps-Engineer-Professional Practice Test Questions in Multiple Choice Formats and Updatesfor 1 Year