Amazon-Web-Services

Exam Questions DOP-C02

AWS Certified DevOps Engineer - Professional

NEW QUESTION 1
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 2
A DevOps engineer needs to apply a core set of security controls to an existing set of AWS accounts. The accounts are in an organization in AWS Organizations. Individual teams will administer individual accounts by using the AdministratorAccess AWS managed policy. For all accounts. AWS CloudTrail and AWS Config must be turned on in all available AWS Regions. Individual account administrators must not be able to edit or delete any of the baseline resources. However, individual account administrators must be able to edit or delete their own CloudTrail trails and AWS Config rules.
Which solution will meet these requirements in the MOST operationally efficient way?

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

Answer: D

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

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

Answer: B

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

NEW QUESTION 4
A company requires an RPO of 2 hours and an RTO of 10 minutes for its data and application at all times. An application uses a MySQL database and Amazon EC2 web servers. The development team needs a strategy for failover and disaster recovery.

Which combination of deployment strategies will meet these requirements? (Select TWO.)

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

Answer: BD

NEW QUESTION 5
A DevOps engineer is implementing governance controls for a company that requires its infrastructure to be housed within the United States. The engineer must restrict which AWS Regions can be used, and ensure an alert is sent as soon as possible if any activity outside the governance policy takes place. The controls should be automatically enabled on any new Region outside the United States (US).
Which combination of actions will meet these requirements? (Select TWO.)

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

Answer: AB

Explanation:
To implement governance controls that restrict AWS service usage to within the United States and ensure alerts for any activity outside the governance policy, the following actions will meet the requirements:
? A. Create an AWS Organizations SCP that denies access to all non-global services in non-US Regions. Attach the policy to the root of the organization. This action will effectively prevent users and roles in all accounts within the organization from accessing services in non-US Regions12.
? B. Configure AWS CloudTrail to send logs to Amazon CloudWatch Logs and enable it for all Regions. Use a CloudWatch Logs metric filter to send an alert on any service activity in non-US Regions. This action will allow monitoring of all AWS Regions and will trigger alerts if any activity is detected in non-US Regions, ensuring that the governance team is notified as soon as possible3.
References:
? AWS Documentation on Service Control Policies (SCPs) and how they can be used to manage permissions and restrict access based on Regions12.
? AWS Documentation on monitoring CloudTrail log files with Amazon CloudWatch Logs to set up alerts for specific activities3.

NEW QUESTION 6
A development team is using AWS CodeCommit to version control application code and AWS CodePipeline to orchestrate software deployments. The team has decided to use a remote main branch as the trigger for the pipeline to integrate code changes. A developer has pushed code changes to the CodeCommit repository, but noticed that the pipeline had no reaction, even after 10 minutes.
Which of the following actions should be taken to troubleshoot this issue?

A. Check that an Amazon EventBridge rule has been created for the main branch to trigger the pipeline.
B. Check that the CodePipeline service role has permission to access the CodeCommit repository.
C. Check that the developer's IAM role has permission to push to the CodeCommit repository.
D. Check to see if the pipeline failed to start because of CodeCommit errors in Amazon CloudWatch Logs.

Answer: A

Explanation:
When you create a pipeline from CodePipeline during the step-by-step it creates a CloudWatch Event rule for a given branch and repo
like this:
```
{
"source": [ "aws.codecommit"
],
"detail-type": [
"CodeCommit Repository State Change"
],
"resources": [
"arn:aws:codecommit:us-east-1:xxxxx:repo-name"
],
"detail": {
"event": [ "referenceCreated", "referenceUpdated"
],
"referenceType": [ "branch"
],
"referenceName": [ "master"
]
}
}
```
https://docs.aws.amazon.com/codepipeline/latest/userguide/pipelines-trigger-source-repo-changes-console.html

NEW QUESTION 7
A company wants to set up a continuous delivery pipeline. The company stores application code in a private GitHub repository. The company needs to deploy the application components to Amazon Elastic Container Service (Amazon ECS). Amazon EC2, and AWS Lambda. The pipeline must support manual approval actions.
Which solution will meet these requirements?

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

Answer: B

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

NEW QUESTION 8
A DevOps engineer has automated a web service deployment by using AWS CodePipeline with the following steps:
1) An AWS CodeBuild project compiles the deployment artifact and runs unit tests.
2) An AWS CodeDeploy deployment group deploys the web service to Amazon EC2 instances in the staging environment.
3) A CodeDeploy deployment group deploys the web service to EC2 instances in the production environment.
The quality assurance (QA) team requests permission to inspect the build artifact before the deployment to the production environment occurs. The QA team wants to run an internal penetration testing tool to conduct manual tests. The tool will be invoked by a REST API call.
Which combination of actions should the DevOps engineer take to fulfill this request? (Choose two.)

A. Insert a manual approval action between the test actions and deployment actions of the pipeline.
B. Modify the buildspec.yml file for the compilation stage to require manual approval before completion.
C. Update the CodeDeploy deployment groups so that they require manual approval to proceed.
D. Update the pipeline to directly call the REST API for the penetration testing tool.
E. Update the pipeline to invoke an AWS Lambda function that calls the REST API for thepenetration testing tool.

Answer: AE

NEW QUESTION 9
A company manages multiple AWS accounts by using AWS Organizations with OUS for the different business divisions, The company is updating their corporate network to use new IP address ranges. The company has 10 Amazon S3 buckets in different AWS accounts. The S3 buckets store reports for the different divisions. The S3 bucket configurations allow only private corporate network IP addresses to access the S3 buckets.
A DevOps engineer needs to change the range of IP addresses that have permission to access the contents of the S3 buckets The DevOps engineer also needs to revoke the permissions of two OUS in the company
Which solution will meet these requirements?

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

Answer: C

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

? S3 Bucket Policies
? Service Control Policies
? Permissions Boundaries

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

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

Answer: BDE

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

NEW QUESTION 10
A company has enabled all features for its organization in AWS Organizations. The organization contains 10 AWS accounts. The company has turned on AWS CloudTrail in all the accounts. The company expects the number of AWS accounts in the organization to increase to 500 during the next year. The company plans to use multiple OUs for these accounts.
The company has enabled AWS Config in each existing AWS account in the organization.
A DevOps engineer must implement a solution that enables AWS Config automatically for all future AWS accounts that are created in the organization.
Which solution will meet this requirement?

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

Answer: B

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

NEW QUESTION 11
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 13
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 14
A DevOps engineer at a company is supporting an AWS environment in which all users use AWS IAM Identity Center (AWS Single Sign-On). The company wants to immediately disable credentials of any new IAM user and wants the security team to receive a notification.
Which combination of steps should the DevOps engineer take to meet these requirements? (Choose three.)

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

Answer: ACE

NEW QUESTION 16
A company is launching an application. The application must use only approved AWS services. The account that runs the application was created less than 1 year ago and is assigned to an AWS Organizations OU.
The company needs to create a new Organizations account structure. The account structure must have an appropriate SCP that supports the use of only services that are currently active in the AWS account.
The company will use AWS Identity and Access Management (IAM) Access Analyzer in the solution.
Which solution will meet these requirements?

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

Answer: A

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

organization's root is not a good option because it might affect other accounts and OUs in the organization that have different service requirements or approvals.
Creating an SCP that allows the services that IAM Access Analyzer identifies and attaching it to the management account is not a valid option because SCPs cannot be attached directly to accounts, only to OUs or roots.
References:
? 1: Using AWS Identity and Access Management Access Analyzer - AWS Identity and Access Management
? 2: Generate a policy based on access activity - AWS Identity and Access Management
? 3: Organizing your accounts into OUs - AWS Organizations
? 4: Service control policies - AWS Organizations
? 5: How SCPs work - AWS Organizations

NEW QUESTION 18
A company uses an Amazon API Gateway regional REST API to host its application API. The REST API has a custom domain. The REST API's default endpoint is deactivated.
The company's internal teams consume the API. The company wants to use mutual TLS between the API and the internal teams as an additional layer of authentication.
Which combination of steps will meet these requirements? (Select TWO.)

A. Use AWS Certificate Manager (ACM) to create a private certificate authority (CA). Provision a client certificate that is signed by the private CA.
B. Provision a client certificate that is signed by a public certificate authority (CA). Import the certificate into AWS Certificate Manager (ACM).
C. Upload the provisioned client certificate to an Amazon S3 bucke
D. Configure the API Gateway mutual TLS to use the client certificate that is stored in the S3 bucket as the trust store.
E. Upload the provisioned client certificate private key to an Amazon S3 bucke
F. Configure the API Gateway mutual TLS to use the private key that is stored in the S3 bucket as the trust store.
G. Upload the root private certificate authority (CA) certificate to an Amazon S3 bucke
H. Configure the API Gateway mutual TLS to use the private CA certificate that is stored in the S3 bucket as the trust store.

Answer: AE

Explanation:
Mutual TLS (mTLS) authentication requires two-way authentication between the client and the server. For Amazon API Gateway, you can enable mTLS for a custom domain name, which requires clients to present X.509 certificates to verify their identity to access your API. To set up mTLS, you would typically use AWS Certificate Manager (ACM) to create a private certificate authority (CA) and provision a client certificate signed by this private
CA. The root CA certificate is then uploaded to an Amazon S3 bucket and configured in API Gateway as the trust store12.
References:
? Introducing mutual TLS authentication for Amazon API Gateway1.
? Configuring mutual TLS authentication for a REST API2.
? AWS Private Certificate Authority details3.
? AWS Certificate Manager Private Certificate Authority updates4.

NEW QUESTION 20
A company manages AWS accounts for application teams in AWS Control Tower. Individual application teams are responsible for securing their respective AWS accounts.
A DevOps engineer needs to enable Amazon GuardDuty for all AWS accounts in which the application teams have not already enabled GuardDuty. The DevOps engineer is using AWS CloudFormation StackSets from the AWS Control Tower management account.
How should the DevOps engineer configure the CloudFormation template to prevent failure during the StackSets deployment?

A. Create a CloudFormation custom resource that invokes an AWS Lambda functio
B. Configure the Lambda function to conditionally enable GuardDuty if GuardDuty is not already enabled in the accounts.
C. Use the Conditions section of the CloudFormation template to enable GuardDuty in accounts where GuardDuty is not already enabled.
D. Use the CloudFormation F
E. GetAtt intrinsic function to check whether GuardDuty is already enabled If GuardDuty is not already enabled use the Resources section of the CloudFormation template to enable GuardDuty.
F. Manually discover the list of AWS account IDs where GuardDuty is not enabled Use the CloudFormation Fn: ImportValue intrinsic function to import the list of account IDs into the CloudFormation template to skip deployment for the listed AWS accounts.

Answer: A

Explanation:
 This solution will meet the requirements because it will use a CloudFormation custom resource to execute custom logic during the stack set operation. A custom resource is a resource that you define in your template and that is associated with an AWS Lambda function. The Lambda function runs whenever the custom resource is created, updated, or deleted, and can perform any actions that are supported by the AWS SDK. In this case, the Lambda function can use the GuardDuty API to check whether GuardDuty is already enabled in each target account, and if not, enable it. This way, the DevOps engineer can avoid deploying the stack set to accounts that already have GuardDuty enabled, and prevent failure during the deployment.

NEW QUESTION 23
A DevOps engineer is creating an AWS CloudFormation template to deploy a web service. The web service will run on Amazon EC2 instances in a private subnet behind an Application Load Balancer (ALB). The DevOps engineer must ensure that the service can accept requests from clients that have IPv6 addresses.
What should the DevOps engineer do with the CloudFormation template so that IPv6 clients can access the web service?

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

Answer: D

Explanation:
it involves adding an IPv6 CIDR block to the VPC and subnets for the ALB and specifying the dualstack IP address type on the ALB listener. This allows the ALB to listen on both IPv4 and IPv6 addresses, and forward requests to the EC2 instances that are added as targets to the target group associated with the ALB.

NEW QUESTION 26
A company has deployed a critical application in two AWS Regions. The application uses an Application Load Balancer (ALB) in both Regions. The company has Amazon Route 53 alias DNS records for both ALBs.
The company uses Amazon Route 53 Application Recovery Controller to ensure that the application can fail over between the two Regions. The Route 53 ARC configuration includes a routing control for both Regions. The company uses Route 53 ARC to perform quarterly disaster recovery (DR) tests.
During the most recent DR test, a DevOps engineer accidentally turned off both routing controls. The company needs to ensure that at least one routing control is turned on at all times.
Which solution will meet these requirements?

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

Answer: A

Explanation:
The correct solution is to create a new assertion safety rule in Route 53 ARC and apply it to the two routing controls. An assertion safety rule is a type of safety rule that ensures that a minimum number of routing controls are always enabled. The ATLEAST type of assertion safety rule specifies the minimum number of routing controls that must be enabled for the rule to evaluate as healthy. By setting the threshold to 1, the rule ensures that at least one routing control is always turned on. This prevents the scenario where both routing controls are accidentally turned off and the application becomes unavailable in both Regions.
The other solutions are incorrect because they do not use safety rules to prevent both routing controls from being turned off. A gating safety rule is a type of safety rule that prevents routing control state changes that violate the rule logic. The OR type of gating safety rule specifies that one or more routing controls must be enabled for the rule to evaluate as healthy. However, this rule does not prevent a user from turning off both routing controls manually. A resource set is a collection of resources that are tested for readiness by Route 53 ARC. A readiness check is a test that verifies that all the resources in a resource set are operational. However, these concepts are not related to routing control states or safety rules. Therefore, creating a new resource set and a new readiness check will not ensure that at least one routing control is turned on at all times. References:
? Routing control in Amazon Route 53 Application Recovery Controller
? Viewing and updating routing control states in Route 53 ARC
? Creating a control panel in Route 53 ARC
? Creating safety rules in Route 53 ARC

NEW QUESTION 29
A DevOps team manages an API running on-premises that serves as a backend for an Amazon API Gateway endpoint. Customers have been complaining about high response latencies, which the development team has verified using the API Gateway latency metrics in Amazon CloudWatch. To identify the cause, the team needs to collect relevant data without introducing additional latency.
Which actions should be taken to accomplish this? (Choose two.)

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

Answer: AC

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

NEW QUESTION 33
A company has multiple AWS accounts. The company uses AWS IAM Identity Center (AWS Single Sign-On) that is integrated with AWS Toolkit for Microsoft Azure DevOps. The attributes for access control feature is enabled in IAM Identity Center.
The attribute mapping list contains two entries. The department key is mapped to
${path:enterprise.department}. The costCenter key is mapped to
${path:enterprise.costCenter}.
All existing Amazon EC2 instances have a department tag that corresponds to three company departments (d1, d2, d3). A DevOps engineer must create policies based on the matching attributes. The policies must minimize administrative effort and must grant each Azure AD user access to only the EC2 instances that are tagged with the user's respective department name.
Which condition key should the DevOps engineer include in the custom permissions policies to meet these requirements?
A.

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

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

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

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

A.

Answer: C

Explanation:
https://docs.aws.amazon.com/singlesignon/latest/userguide/configure- abac.html

NEW QUESTION 34
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 35
A company has migrated its container-based applications to Amazon EKS and want to establish automated email notifications. The notifications sent to each email address are for specific activities related to EKS components. The solution will include Amazon SNS topics and an AWS Lambda function to evaluate incoming log events and publish messages to the correct SNS topic.
Which logging solution will support these requirements?

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

Answer: A

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

NEW QUESTION 38
A DevOps team is merging code revisions for an application that uses an Amazon RDS Multi-AZ DB cluster for its production database. The DevOps team uses continuous integration to periodically verify that the application works. The DevOps team needs to test the changes before the changes are deployed to the production database.
Which solution will meet these requirements'?

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

Answer: A

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

NEW QUESTION 42
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 45
A company's DevOps engineer is creating an AWS Lambda function to process notifications from an Amazon Simple Notification Service (Amazon SNS) topic. The Lambda function will process the notification messages and will write the contents of the notification messages to an Amazon RDS Multi-AZ DB instance.
During testing a database administrator accidentally shut down the DB instance. While the database was down the company lost several of the SNS notification messages that were delivered during that time.
The DevOps engineer needs to prevent the loss of notification messages in the future Which solutions will meet this requirement? (Select TWO.)

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

Answer: CD

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

NEW QUESTION 47
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 52
A company's application teams use AWS CodeCommit repositories for their applications.
The application teams have repositories in multiple AWS accounts. All accounts are in an organization in AWS Organizations.
Each application team uses AWS IAM Identity Center (AWS Single Sign-On) configured with an external IdP to assume a developer IAM role. The developer role allows the application teams to use Git to work with the code in the repositories.
A security audit reveals that the application teams can modify the main branch in any repository. A DevOps engineer must implement a solution that allows the application teams to modify the main branch of only the repositories that they manage.
Which combination of steps will meet these requirements? (Select THREE.)

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

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

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

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

Answer: ADF

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

NEW QUESTION 54
A company requires its developers to tag all Amazon Elastic Block Store (Amazon EBS) volumes in an account to indicate a desired backup frequency. This requirement Includes EBS volumes that do not require backups. The company uses custom tags named Backup_Frequency that have values of none, dally, or weekly that correspond to the desired backup frequency. An audit finds that developers are occasionally not tagging the EBS volumes.
A DevOps engineer needs to ensure that all EBS volumes always have the Backup_Frequency tag so that the company can perform backups at least weekly unless a different value is specified.
Which solution will meet these requirements?

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

Answer: B

Explanation:
The following are the steps that the DevOps engineer should take to ensure that all EBS volumes always have the Backup_Frequency tag so that the company can perform backups at least weekly unless a different value is specified:
? Set up AWS Config in the account.
? Use a managed rule that returns a compliance failure for EC2::Volume resources that do not have a Backup Frequency tag applied.
? Configure a remediation action that uses a custom AWS Systems Manager Automation runbook to apply the Backup_Frequency tag with a value of weekly.
The managed rule AWS::Config::EBSVolumesWithoutBackupTag will return a compliance failure for any EBS volume that does not have the Backup_Frequency tag applied. The remediation action will then use the Systems Manager Automation runbook to apply the Backup_Frequency tag with a value of weekly to the EBS volume.

NEW QUESTION 56
An IT team has built an AWS CloudFormation template so others in the company can quickly and reliably deploy and terminate an application. The template creates an Amazon EC2 instance with a user data script to install the application and an Amazon S3 bucket that the application uses to serve static webpages while it is running.
All resources should be removed when the CloudFormation stack is deleted. However, the team observes that CloudFormation reports an error during stack deletion, and the S3 bucket created by the stack is not deleted.
How can the team resolve the error in the MOST efficient manner to ensure that all resources are deleted without errors?

A. Add a DelelionPolicy attribute to the S3 bucket resource, with the value Delete forcing the bucket to be removed when the stack is deleted.
B. Add a custom resource with an AWS Lambda function with the DependsOn attribute specifying the S3 bucket, and an IAM 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 61
A company deploys updates to its Amazon API Gateway API several times a week by using an AWS CodePipeline pipeline. As part of the update process the company exports the JavaScript SDK for the API from the API. Gateway console and uploads the SDK to an Amazon S3 bucket
The company has configured an Amazon CloudFront distribution that uses the S3 bucket as an origin Web client then download the SDK by using the CloudFront distribution's endpoint. A DevOps engineer needs to implement a solution to make the new SDK available automatically during new API deployments.
Which solution will meet these requirements?

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

Answer: A

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

NEW QUESTION 64
A company uses AWS Organizations to manage its AWS accounts. The company has a root OU that has a child OU. The root OU has an SCP that allows all actions on all resources. The child OU has an SCP that allows all actions for Amazon DynamoDB and AWS Lambda, and denies all other actions.
The company has an AWS account that is named vendor-data in the child OU. A DevOps engineer has an 1AM user that is attached to the AdministratorAccess 1AM policy in the vendor-data account. The DevOps engineer attempts to launch an Amazon EC2 instance in the vendor-data account but receives an access denied error.
Which change should the DevOps engineer make to launch the EC2 instance in the vendor-data account?

A. Attach the AmazonEC2FullAccess 1AM policy to the 1AM user.
B. Create a new SCP that allows all actions for Amazon EC2. Attach the SCP to the vendor-data account.
C. Update the SCP in the child OU to allow all actions for Amazon EC2.
D. Create a new SCP that allows all actions for Amazon EC2. Attach the SCP to the root OU.

Answer: C

Explanation:
The correct answer is C. Updating the SCP in the child OU to allow all actions for Amazon EC2 will enable the DevOps engineer to launch the EC2 instance in the vendor-data account. SCPs are applied to OUs and accounts in a hierarchical manner, meaning that the SCPs attached to the parent OU are inherited by the child OU and accounts. Therefore, the SCP in the child OU overrides the SCP in the root OU and denies all actions except for DynamoDB and Lambda. By adding EC2 to the allowed actions in the child OU's SCP, the DevOps engineer can access EC2 resources in the vendor-data account.
Option A is incorrect because attaching the AmazonEC2FullAccess IAM policy to the IAM user will not grant the user access to EC2 resources. IAM policies are evaluated after SCPs, so even if the IAM policy allows EC2 actions, the SCP will still deny them.
Option B is incorrect because creating a new SCP that allows all actions for EC2 and attaching it to the vendor-data account will not work. SCPs are not cumulative, meaning that only one SCP is applied to an account at a time. The SCP attached to the account will be the SCP attached to the OU that contains the account. Therefore, option B will not change the SCP that is applied to the vendor-data account.
Option D is incorrect because creating a new SCP that allows all actions for EC2 and attaching it to the root OU will not work. As explained earlier, the SCP in the child OU overrides the SCP in the root OU and denies all actions except for DynamoDB and Lambda. Therefore, option D will not affect the SCP that is applied to the vendor-data account.

NEW QUESTION 65

A company runs its container workloads in AWS App Runner. A DevOps engineer manages the company's container repository in Amazon Elastic Container Registry (Amazon ECR).

The DevOps engineer must implement a solution that continuously monitors the container repository. The solution must create a new container image when the solution detects an operating system vulnerability or language package vulnerability.

Which solution will meet these requirements?

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

Answer: A

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

NEW QUESTION 68
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 72
A DevOps engineer is working on a data archival project that requires the migration of on- premises data to an Amazon S3 bucket. The DevOps engineer develops a script that incrementally archives on-premises data that is older than 1 month to Amazon S3. Data that is transferred to Amazon S3 is deleted from the on-premises location The script uses the S3 PutObject operation.
During a code review the DevOps engineer notices that the script does not verity whether the data was successfully copied to Amazon S3. The DevOps engineer must update the script to ensure that data is not corrupted during transmission. The script must use MD5 checksums to verify data integrity before the on-premises data is deleted.
Which solutions for the script will meet these requirements'? (Select TWO.)

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

Answer: BD

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

NEW QUESTION 77
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 78
A rapidly growing company wants to scale for developer demand for AWS development environments. Development environments are created manually in the AWS Management Console. The networking team uses AWS CloudFormation to manage the networking infrastructure, exporting stack output values for the Amazon VPC and all subnets. The development environments have common standards, such as Application Load Balancers, Amazon EC2 Auto Scaling groups, security groups, and Amazon DynamoDB tables.
To keep up with demand, the DevOps engineer wants to automate the creation of development environments. Because the infrastructure required to support the application is expected to grow, there must be a way to easily update the deployed infrastructure. CloudFormation will be used to create a template for the

development environments.
Which approach will meet these requirements and quickly provide consistent AWS environments for developers?

A. Use Fn::ImportValue intrinsic functions in the Resources section of the template to retrieve Virtual Private Cloud (VPC) and subnet value
B. Use CloudFormation StackSets for the development environments, using the Count input parameter to indicate the number of environments neede
C. Use the UpdateStackSet command to update existing development environments.
D. Use nested stacks to define common infrastructure component
E. To access the exported values, use TemplateURL to reference the networking team's templat
F. To retrieve Virtual Private Cloud (VPC) and subnet values, use Fn::ImportValue intrinsic functions in the Parameters section of the root templat
G. Use the CreateChangeSet and ExecuteChangeSet commands to update existing development environments.
H. Use nested stacks to define common infrastructure component
I. Use Fn::ImportValue intrinsic functions with the resources of the nested stack to retrieve Virtual Private Cloud (VPC) and subnet value
J. Use the CreateChangeSet and ExecuteChangeSet commands to update existing development environments.
K. Use Fn::ImportValue intrinsic functions in the Parameters section of the root template to retrieve Virtual Private Cloud (VPC) and subnet value
L. Define the development resources in the order they need to be created in the CloudFormation nested stack
M. Use the CreateChangeSe
N. and ExecuteChangeSet commands to update existing development environments.

Answer: C

Explanation:
 https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/intrinsic-function- reference-importvalue.html
https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/intrinsic-function-reference-importvalue.html CF of network exports the VPC, subnet or needed information CF of application imports the above information to its stack and UpdateChangeSet/ ExecuteChangeSet

NEW QUESTION 83
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 86
An ecommerce company has chosen AWS to host its new platform. The company's DevOps team has started building an AWS Control Tower landing zone. The DevOps team has set the identity store within AWS IAM Identity Center (AWS Single Sign-On) to external identity provider (IdP) and has configured SAML 2.0. The DevOps team wants a robust permission model that applies the principle of least privilege. The model must allow the team to build and manage only the team's own resources.
Which combination of steps will meet these requirements? (Choose three.)

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

Answer: BCF

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

NEW QUESTION 88
An application running on a set of Amazon EC2 instances in an Auto Scaling group requires a configuration file to operate. The instances are created and

maintained with AWS CloudFormation. A DevOps engineer wants the instances to have the latest configuration file when launched and wants changes to the configuration file to be reflected on all the instances with a minimal delay when the CloudFormation template is updated. Company policy requires that application configuration files be maintained along with AWS infrastructure configuration files m source control.
Which solution will accomplish this?

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

Answer: D

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

NEW QUESTION 91
A company has a single AWS account that runs hundreds of Amazon EC2 instances in a single AWS Region. New EC2 instances are launched and terminated each hour in the account. The account also includes existing EC2 instances that have been running for longer than a week.
The company's security policy requires all running EC2 instances to use an EC2 instance profile. If an EC2 instance does not have an instance profile attached, the EC2 instance must use a default instance profile that has no IAM permissions assigned.
A DevOps engineer reviews the account and discovers EC2 instances that are running without an instance profile. During the review, the DevOps engineer also observes that new EC2 instances are being launched without an instance profile.
Which solution will ensure that an instance profile is attached to all existing and future EC2 instances in the Region?

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

Answer: B

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

NEW QUESTION 96
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 99

A company uses AWS Organizations to manage multiple accounts. Information security policies require that all unencrypted Amazon EBS volumes be marked as non-compliant. A DevOps engineer needs to automatically deploy the solution and ensure that this compliance check is always present.
Which solution will accomplish this?

A. Create an AWS CloudFormation template that defines an AWS Inspector rule to check whether EBS encryption is enable
B. Save the template to an Amazon S3 bucket that has been shared with all accounts within the compan
C. Update the account creation script pointing to the CloudFormation template in Amazon S3.
D. Create an AWS Config organizational rule to check whether EBS encryption is enabled and deploy the rule using the AWS CL
E. Create and apply an SCP to prohibit stopping and deleting AWS Config across the organization.
F. Create an SCP in Organization
G. Set the policy to prevent the launch of Amazon EC2 instances without encryption on the EBS volumes using a conditional expressio
H. Apply the SCP to all AWS account
I. Use Amazon Athena to analyze the AWS CloudTrail output, looking for events that deny an ec2: RunInstances action.
J. Deploy an IAM role to all accounts from a single trusted accoun
K. Build a pipeline withAWS CodePipeline with a stage in AWS Lambda to assume the IAM role, and list all EBS volumes in the accoun
L. Publish a report to Amazon S3.

Answer: B

Explanation:
 https://docs.aws.amazon.com/config/latest/developerguide/ec2-ebs-encryption-by-default.html

NEW QUESTION 100
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 102
A company is divided into teams Each team has an AWS account and all the accounts are in an organization in AWS Organizations. Each team must retain full administrative rights to its AWS account. Each team also must be allowed to access only AWS services that the company approves for use AWS services must gam approval through a request and approval process.
How should a DevOps engineer configure the accounts to meet these requirements?

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

Answer: C

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

NEW QUESTION 107
A company is hosting a static website from an Amazon S3 bucket. The website is available to customers at example.com. The company uses an Amazon Route 53 weighted routing policy with a TTL of 1 day. The company has decided to replace the existing static website with a dynamic web application. The dynamic web application uses an Application Load Balancer (ALB) in front of a fleet of Amazon EC2 instances.
On the day of production launch to customers, the company creates an additional Route 53 weighted DNS record entry that points to the ALB with a weight of 255 and a TTL of 1 hour. Two days later, a DevOps engineer notices that the previous static website is displayed sometimes when customers navigate to example.com.
How can the DevOps engineer ensure that the company serves only dynamic content for example.com?

A. Delete all objects, including previous versions, from the S3 bucket that contains the static website content.
B. Update the weighted DNS record entry that points to the S3 bucke

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

Answer: D

NEW QUESTION 110
A company uses AWS Directory Service for Microsoft Active Directory as its identity provider (IdP). The company requires all infrastructure to be defined and deployed by AWS CloudFormation.
A DevOps engineer needs to create a fleet of Windows-based Amazon EC2 instances to host an application. The DevOps engineer has created a CloudFormation template that contains an EC2 launch template, IAM role, EC2 security group, and EC2 Auto Scaling group. The DevOps engineer must implement a solution that joins all EC2 instances to the domain of the AWS Managed Microsoft AD directory.
Which solution will meet these requirements with the MOST operational efficiency?

A. In the CloudFormation template, create an AWS::SSM::Document resource that joins the EC2 instance to the AWS Managed Microsoft AD domain by using the parameters for the existing director
B. Update the launch template to include the SSMAssociation property to use the new SSM documen
C. Attach the AmazonSSMManagedInstanceCore and AmazonSSMDirectoryServiceAccess AWS managed policies to the IAM role that the EC2 instances use.
D. In the CloudFormation template, update the launch template to include specific tags that propagate on launc
E. Create an AWS::SSM::Association resource to associate the AWS- JoinDirectoryServiceDomain Automation runbook with the EC2 instances that have the specified tag
F. Define the required parameters to join the AWS Managed Microsoft AD director
G. Attach the AmazonSSMManagedInstanceCore and AmazonSSMDirectoryServiceAccess AWS managed policies to the IAM role that the EC2 instances use.
H. Store the existing AWS Managed Microsoft AD domain connection details in AWS Secrets Manage
I. In the CloudFormation template, create an AWS::SSM::Association resource to associate the AWS-CreateManagedWindowsInstanceWithApproval Automation runbook with the EC2 Auto Scaling grou
J. Pass the ARNs for the parameters from Secrets Manager to join the domai
K. Attach the AmazonSSMDirectoryServiceAccess and SecretsManagerReadWrite AWS managed policies to the IAM role that the EC2 instances use.
L. Store the existing AWS Managed Microsoft AD domain administrator credentials in AWS Secrets Manage
M. In the CloudFormation template, update the EC2 launch template to include user dat
N. Configure the user data to pull the administrator credentials from Secrets Manager and to join the AWS Managed Microsoft AD domai
O. Attach the AmazonSSMManagedInstanceCore and SecretsManagerReadWrite AWS managed policies to the IAM role that the EC2 instances use.

Answer: B

Explanation:
 To meet the requirements, the DevOps engineer needs to create a solution that joins all EC2 instances to the domain of the AWS Managed Microsoft AD directory with the most operational efficiency. The DevOps engineer can use AWS Systems Manager Automation to automate the domain join process using an existing runbook called AWS- JoinDirectoryServiceDomain. This runbook can join Windows instances to an AWS Managed Microsoft AD or Simple AD directory by using PowerShell commands. The DevOps engineer can create an AWS::SSM::Association resource in the CloudFormation template to associate the runbook with the EC2 instances that have specific tags. The tags can be defined in the launch template and propagated on launch to the EC2 instances. The DevOps engineer can also define the required parameters for the runbook, such as the directory ID, directory name, and organizational unit. The DevOps engineer can attach the AmazonSSMManagedInstanceCore and AmazonSSMDirectoryServiceAccess AWS managed policies to the IAM role that the EC2 instances use. These policies grant the necessary permissions for Systems Manager and Directory Service operations.

NEW QUESTION 112
A company plans to use Amazon CloudWatch to monitor its Amazon EC2 instances. The company needs to stop EC2 instances when the average of the NetworkPacketsIn metric is less than 5 for at least 3 hours in a 12-hour time window. The company must evaluate the metric every hour. The EC2 instances must continue to run if there is missing data for the NetworkPacketsIn metric during the evaluation period.
A DevOps engineer creates a CloudWatch alarm for the NetworkPacketsIn metric. The DevOps engineer configures a threshold value of 5 and an evaluation period of 1 hour.
Which set of additional actions should the DevOps engineer take to meet these requirements?

A. Configure the Datapoints to Alarm value to be 3 out of 12. Configure the alarm to treat missing data as breaching the threshol
B. Add an AWS Systems Manager action to stop the instance when the alarm enters the ALARM state.
C. Configure the Datapoints to Alarm value to be 3 out of 12. Configure the alarm to treat missing data as not breaching the threshol
D. Add an EC2 action to stop the instance when the alarm enters the ALARM state.
E. Configure the Datapoints to Alarm value to be 9 out of 12. Configure the alarm to treat missing data as breaching the threshol
F. Add an EC2 action to stop the instance when the alarm enters the ALARM state.
G. Configure the Datapoints to Alarm value to be 9 out of 12. Configure the alarm to treat missing data as not breaching the threshol
H. Add an AWS Systems Manager action to stop the instance when the alarm enters the ALARM state.

Answer: B

Explanation:
To meet the requirements, the DevOps engineer needs to configure the CloudWatch alarm to stop the EC2 instances when the average of the NetworkPacketsIn metric is less than 5 for at least 3 hours in a 12-hour time window. This means that the alarm should trigger when 3 out of 12 datapoints are below the threshold of 5. The alarm should also treat missing data as not breaching the threshold, so that the EC2 instances continue to run if there is no data for the metric during the evaluation period. The DevOps engineer can add an EC2 action to stop the instance when the alarm enters the ALARM state, which is a built-in action type for CloudWatch alarms.

NEW QUESTION 115
An AWS CodePipeline pipeline has implemented a code release process. The pipeline is integrated with AWS CodeDeploy to deploy versions of an application to multiple Amazon EC2 instances for each CodePipeline stage.
During a recent deployment the pipeline failed due to a CodeDeploy issue. The DevOps team wants to improve monitoring and notifications during deployment to decrease resolution times.
What should the DevOps engineer do to create notifications. When issues are discovered?

A. Implement Amazon CloudWatch Logs for CodePipeline and CodeDeploy create an AWS Config rule to evaluate code deployment issues, and create an

Amazon Simple Notification Service (Amazon SNS) topic to notify stakeholders of deployment issues.
B. Implement Amazon EventBridge for CodePipeline and CodeDeploy create an AWS Lambda function to evaluate code deployment issues, and create an Amazon Simple Notification Service (Amazon SNS) topic to notify stakeholders of deployment issues.
C. Implement AWS CloudTrail to record CodePipeline and CodeDeploy API call information create an AWS Lambda function to evaluate code deployment issues and create an Amazon Simple Notification Service (Amazon SNS) topic to notify stakeholders of deployment issues.
D. Implement Amazon EventBridge for CodePipeline and CodeDeploy create an Amazo
E. Inspector assessment target to evaluate code deployment issues and create an Amazon Simpl
F. Notification Service (Amazon SNS) topic to notify stakeholders of deployment issues.

Answer: B

Explanation:
 AWS CloudWatch Events can be used to monitor events across different AWS resources, and a CloudWatch Event Rule can be created to trigger an AWS Lambda function when a deployment issue is detected in the pipeline. The Lambda function can then evaluate the issue and send a notification to the appropriate stakeholders through an Amazon SNS topic. This approach allows for real-time notifications and faster resolution times.

NEW QUESTION 117
A Company uses AWS CodeCommit for source code control. Developers apply their changes to various feature branches and create pull requests to move those changes to the main branch when the changes are ready for production.
The developers should not be able to push changes directly to the main branch. The company applied the AWSCodeCommitPowerUser managed policy to the developers' IAM role, and now these developers can push changes to the main branch directly on every repository in the AWS account.
What should the company do to restrict the developers' ability to push changes to the main branch directly?

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

Answer: A

Explanation:
 By default, the AWSCodeCommitPowerUser managed policy allows users to push changes to any branch in any repository in the AWS account. To restrict the developers' ability to push changes to the main branch directly, an additional policy is needed that explicitly denies these actions for the main branch.
The Deny rule should be included in a policy statement that targets the specific repositories and includes a condition that references the main branch. The policy statement should look something like this:
{
"Effect": "Deny", "Action": ["codecommit:GitPush", "codecommit:PutFile"
],
"Resource": "arn:aws:codecommit:<region>:<account-id>:<repository-name>", "Condition": {
"StringEqualsIfExists": { "codecommit:References": ["refs/heads/main"
]
}
}
}

NEW QUESTION 118
A company has configured an Amazon S3 event source on an AWS Lambda function The company needs the Lambda function to run when a new object is created or an existing object IS modified In a particular S3 bucket The Lambda function will use the S3 bucket name and the S3 object key of the incoming event to read the contents of the created or modified S3 object The Lambda function will parse the contents and save the parsed contents to an Amazon DynamoDB table.
The Lambda function's execution role has permissions to read from the S3 bucket and to write to the DynamoDB table, During testing, a DevOps engineer discovers that the Lambda
function does not run when objects are added to the S3 bucket or when existing objects are modified.
Which solution will resolve this problem?

A. Increase the memory of the Lambda function to give the function the ability to process large files from the S3 bucket.
B. Create a resource policy on the Lambda function to grant Amazon S3 the permission to invoke the Lambda function for the S3 bucket
C. Configure an Amazon Simple Queue Service (Amazon SQS) queue as an OnFailure destination for the Lambda function
D. Provision space in the /tmp folder of the Lambda function to give the function the ability to process large files from the S3 bucket

Answer: B

Explanation:
? Option A is incorrect because increasing the memory of the Lambda function does not address the root cause of the problem, which is that the Lambda function is not triggered by the S3 event source. Increasing the memory of the Lambda function might improve its performance or reduce its execution time, but it does not affect its invocation. Moreover, increasing the memory of the Lambda function might incur higher costs, as Lambda charges based on the amount of memory allocated to the function.
? Option B is correct because creating a resource policy on the Lambda function to grant Amazon S3 the permission to invoke the Lambda function for the S3 bucket is a necessary step to configure an S3 event source. A resource policy is a JSON document that defines who can access a Lambda resource and under what conditions. By granting Amazon S3 permission to invoke the Lambda function, the company ensures that the Lambda function runs when a new object is created or an existing object is modified in the S3 bucket1.
? Option C is incorrect because configuring an Amazon Simple Queue Service (Amazon SQS) queue as an On-Failure destination for the Lambda function does not help with triggering the Lambda function. An On-Failure destination is a feature that allows Lambda to send events to another service, such as SQS or Amazon Simple Notification Service (Amazon SNS), when a function invocation fails. However, this feature only applies to asynchronous invocations, and S3 event sources use synchronous invocations. Therefore, configuring an SQS queue as an On-Failure destination would have no effect on the problem.
? Option D is incorrect because provisioning space in the /tmp folder of the Lambda function does not address the root cause of the problem, which is that the Lambda function is not triggered by the S3 event source. Provisioning space in the /tmp folder of the Lambda function might help with processing large files from

the S3 bucket, as it provides temporary storage for up to 512 MB of data. However, it does not affect the invocation of the Lambda function.
References:
? Using AWS Lambda with Amazon S3
? Lambda resource access permissions
? AWS Lambda destinations
? [AWS Lambda file system]

NEW QUESTION 121
A company runs an application on Amazon EC2 instances. The company uses a series of AWS CloudFormation stacks to define the application resources. A developer performs updates by building and testing the application on a laptop and then uploading the build output and CloudFormation stack templates to Amazon S3. The developer's peers review the changes before the developer performs the CloudFormation stack update and installs a new version of the application onto the EC2 instances.
The deployment process is prone to errors and is time-consuming when the developer updates each EC2 instance with the new application. The company wants to automate as much of the application deployment process as possible while retaining a final manual approval step before the modification of the application or resources.
The company already has moved the source code for the application and the CloudFormation templates to AWS CodeCommit. The company also has created an AWS CodeBuild project to build and test the application.
Which combination of steps will meet the company's requirements? (Choose two.)

A. Create an application group and a deployment group in AWS CodeDeplo
B. Install the CodeDeploy agent on the EC2 instances.
C. Create an application revision and a deployment group in AWS CodeDeplo
D. Create an environment in CodeDeplo
E. Register the EC2 instances to the CodeDeploy environment.
F. Use AWS CodePipeline to invoke the CodeBuild job, run the CloudFormation update, and pause for a manual approval ste
G. After approval, start the AWS CodeDeploy deployment.
H. Use AWS CodePipeline to invoke the CodeBuild job, create CloudFormation change sets for each of the application stacks, and pause for a manual approval ste
I. After approval, run the CloudFormation change sets and start the AWS CodeDeploy deployment.
J. Use AWS CodePipeline to invoke the CodeBuild job, create CloudFormation change sets for each of the application stacks, and pause for a manual approval ste
K. After approval, start the AWS CodeDeploy deployment.

Answer: AD

Explanation:
 A- https://docs.aws.amazon.com/codedeploy/latest/userguide/codedeploy- agent.html D - This option correctly utilizes AWS CodePipeline to invoke the CodeBuild job and create CloudFormation change sets. It adds a manual approval step before executing the change sets and starting the AWS CodeDeploy deployment. This ensures that the deployment process is automated while retaining the final manual approval step.

NEW QUESTION 123
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 128

......

Thank You for Trying Our Product

We offer two products:

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

2nd - Questons and Answers in PDF Format

DOP-C02 Practice Exam Features:

* DOP-C02 Questions and Answers Updated Frequently

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

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

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

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