Amazon

Exam Questions AWS-Certified-Security-Specialty

Amazon AWS Certified Security - Specialty

NEW QUESTION 1
An IT department currently has a Java web application deployed on Apache Tomcat running on Amazon EC2 instances. All traffic to the EC2 instances is sent through an internet-facing Application Load Balancer (ALB)
The Security team has noticed during the past two days thousands of unusual read requests coming from hundreds of IP addresses. This is causing the Tomcat server to run out of threads and reject new connections
Which the SIMPLEST change that would address this server issue?

A. Create an Amazon CloudFront distribution and configure the ALB as the origin
B. Block the malicious IPs with a network access list (NACL).
C. Create an IAM Web Application Firewall (WAF). and attach it to the ALB
D. Map the application domain name to use Route 53

Answer: A

Explanation:
this is the simplest change that can address the server issue. CloudFront is a service that provides a global network of edge locations that cache and deliver web content. Creating a CloudFront distribution and configuring the ALB as the origin can help reduce the load on the Tomcat server by serving cached content to the end users. CloudFront can also provide protection against distributed denial-of-service (DDoS) attacks by filtering malicious traffic at the edge locations. The other options are either ineffective or complex for solving the server issue.

NEW QUESTION 2
A company manages three separate IAM accounts for its production, development, and test environments, Each Developer is assigned a unique IAM user under the development account. A new application hosted on an Amazon EC2 instance in the developer account requires read access to the archived documents stored in an Amazon S3 bucket in the production account.
How should access be granted?

A. Create an IAM role in the production account and allow EC2 instances in the development account to assume that role using the trust polic
B. Provide read access for the required S3 bucket to this role.
C. Use a custom identity broker to allow Developer IAM users to temporarily access the S3 bucket.
D. Create a temporary IAM user for the application to use in the production account.
E. Create a temporary IAM user in the production account and provide read access to Amazon S3.Generate the temporary IAM user's access key and secret key and store these on the EC2 instance used by the application in the development account.

Answer: A

Explanation:
https://IAM.amazon.com/premiumsupport/knowledge-center/cross-account-access-s3/

NEW QUESTION 3
A company discovers a billing anomaly in its AWS account. A security consultant investigates the anomaly and discovers that an employee who left the company 30 days ago still has access to the account.
The company has not monitored account activity in the past.
The security consultant needs to determine which resources have been deployed or reconfigured by the employee as quickly as possible.
Which solution will meet these requirements?

A. In AWS Cost Explorer, filter chart data to display results from the past 30 day
B. Export the results to a data tabl
C. Group the data table by re-source.
D. Use AWS Cost Anomaly Detection to create a cost monito
E. Access the detec-tion histor
F. Set the time frame to Last 30 day
G. In the search area, choose the service category.
H. In AWS CloudTrail, filter the event history to display results from the past 30 day
I. Create an Amazon Athena table that contains the dat
J. Parti-tion the table by event source.
K. Use AWS Audit Manager to create an assessment for the past 30 day
L. Apply a usage-based framework to the assessmen
M. Configure the assessment to as-sess by resource.

Answer: C

NEW QUESTION 4
A company wants to protect its website from man in-the-middle attacks by using Amazon CloudFront. Which solution will meet these requirements with the LEAST operational overhead?

A. Use the SimpleCORS managed response headers policy.
B. Use a Lambda@Edge function to add the Strict-Transport-Security response header.
C. Use the SecurityHeadersPolicy managed response headers policy.
D. Include the X-XSS-Protection header in a custom response headers policy.

Answer: C

Explanation:
https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/using-managed-response-headers-poli The SecurityHeadersPolicy is a managed policy provided by Amazon CloudFront that includes a set of recommended security headers to enhance the security of your website. These headers help protect against various types of attacks, including man-in-the-middle attacks. By applying the SecurityHeadersPolicy to your CloudFront distribution, the necessary security headers will be automatically added to the responses sent by CloudFront. This reduces operational overhead because you don't have to manually configure or manage the headers yourself.

NEW QUESTION 5
A company deploys a set of standard IAM roles in AWS accounts. The IAM roles are based on job functions within the company. To balance operational efficiency and security, a security engineer implemented AWS Organizations SCPs to restrict access to critical security services in all company accounts.
All of the company's accounts and OUs within AWS Organizations have a default FullAWSAccess SCP that is attached. The security engineer needs to ensure that no one can disable Amazon GuardDuty and AWS Security Hub. The security engineer also must not override other permissions that are granted by IAM policies that are defined in the accounts.
Which SCP should the security engineer attach to the root of the organization to meet these requirements?

A.
```
"Version": "2012-10-17",
"Statement"":[
  {
    "Effect": "Deny",
    "Action":[
      "guardduty:DeleteDetector",
      "guardduty:UpdateDetector",
      "securityhub:DisableSecurityHub"
    ],
    "Resource":[
    "*"
    ]
  }
 ]
}
```

B. A screenshot of a computer code Description automatically generated
```
{
 "Version": "2012-10-17",
 "Statement"":[
   {
      "Effect": "Deny",
      "Action":"*",
      "Resource": "*"
    },
    {
      "Effect": "Allow",
      "NotAction": [
      "guardduty:DeleteDetector",
      "guardduty:UpdateDetector",
      "securityhub:DisableSecurityHub"
    ],
    "Resource":[
    "*"
    ]
  }
 ]
}
```

C. A screenshot of a computer code Description automatically generated
```
{
 "Version": "2012-10-17",
 "Statement"":[
   {
      "Effect": "Allow",
      "Action":"*",
      "Resource": "*"
    },
    {
      "Effect": "Deny",
      "NotAction": [
      "guardduty:DeleteDetector",
      "guardduty:UpdateDetector",
      "securityhub:DisableSecurityHub"
    ],
    "Resource":[
    "*"
    ]
  }
 ]
}
```

D. A screenshot of a computer code Description automatically generated

```
{
 "Version": "2012-10-17",
 "Statement"":[
   {
       "Effect": "Allow",
        "NotAction": [
       "guardduty:DeleteDetector",
       "guardduty:UpdateDetector",
       "securityhub:DisableSecurityHub"
     ],
     "Resource":[
     "*"
     ]
   }
  ]
}
```

Answer: A

NEW QUESTION 6
A company needs to improve its ability to identify and prevent IAM policies that grant public access or cross-account access to resources. The company has implemented AWS Organizations and has started using AWS Identity and Access Management Access Analyzer to refine overly broad access to accounts in the organization.
A security engineer must automate a response in the company's organization for any newly created policies that are overly permissive. The automation must remediate external access and must notify the company's security team.
Which combination of steps should the security engineer take to meet these requirements? (Select THREE.)

A. Create an AWS Step Functions state machine that checks the resource type in the finding and adds an explicit Deny statement in the trust policy for the IAM rol
B. Configure the state machine to publish a notification to an Amazon SimpleNotification Service (Amazon SNS) topic.
C. Create an AWS Batch job that forwards any resource type findings to an AWS Lambda function.Configure the Lambda function to add an explicit Deny statement in the trust policy for the IAM rol
D. Configure the AWS Batch job to publish a notification to an Amazon Simple Notification Service (Amazon SNS) topic.
E. In Amazon EventBridge, create an event rule that matches active IAM Access Analyzer findings and invokes AWS Step Functions for resolution.
F. In Amazon CloudWatch, create a metric filter that matches active IAM Access Analyzer findings and invokes AWS Batch for resolution.
G. Create an Amazon Simple Queue Service (Amazon SQS) queu
H. Configure the queue to forward a notification to the security team that an external principal has been granted access to the specific IAM role and has been blocked.
I. Create an Amazon Simple Notification Service (Amazon SNS) topic for external or cross-account access notice
J. Subscribe the security team's email addresses to the topic.

Answer: ACF

Explanation:
The correct answer is A, C, and F.
To automate a response for any newly created policies that are overly permissive, the security engineer needs to use a combination of services that can monitor, analyze, remediate, and notify the security incidents.
Option A is correct because creating an AWS Step Functions state machine that checks the resource type in the finding and adds an explicit Deny statement in the trust policy for the IAM role is a valid way to remediate external access. AWS Step Functions is a service that allows you to coordinate multiple AWS services into serverless workflows. You can use Step Functions to invoke AWS Lambda functions, which can modify the IAM policies programmatically. You can also use Step Functions to publish a notification to an Amazon SNS topic, which can send messages to subscribers such as email addresses.
Option B is incorrect because creating an AWS Batch job that forwards any resource type findings to an AWS Lambda function is not a suitable way to automate a response. AWS Batch is a service that enables you to run batch computing workloads on AWS. Batch is designed for large-scale and long-running jobs that can benefit from parallelization and dynamic provisioning of compute resources. Batch is not intended for event-driven and real-time workflows that require immediate response.
Option C is correct because creating an Amazon EventBridge event rule that matches active IAM Access Analyzer findings and invokes AWS Step Functions for resolution is a valid way to monitor and analyze the security incidents. Amazon EventBridge is a serverless event bus service that allows you to connect your applications with data from various sources. EventBridge can use rules to match events and route them to targets for processing. You can use EventBridge to invoke AWS Step Functions state machines from the IAM Access Analyzer findings.
Option D is incorrect because creating an Amazon CloudWatch metric filter that matches active IAM Access Analyzer findings and invokes AWS Batch for resolution is not a suitable way to monitor and analyze the security incidents. Amazon CloudWatch is a service that provides monitoring and observability for your AWS resources and applications. CloudWatch can collect metrics, logs, and events from various sources and perform actions based on alarms or filters. However, CloudWatch cannot directly invoke AWS Batch jobs from the IAM Access Analyzer findings. You would need to use another service such as EventBridge or SNS to trigger the Batch job.
Option E is incorrect because creating an Amazon SQS queue that forwards a notification to the security team that an external principal has been granted access to the specific IAM role and has been blocked is not a valid way to notify the security incidents. Amazon SQS is a fully managed message queue service that enables you to decouple and scale microservices, distributed systems, and serverless applications. SQS can deliver messages to consumers that poll the queue for messages. However, SQS cannot directly forward a notification to the security team's email addresses. You would need to use another service such as SNS or SES to send email notifications.
Option F is correct because creating an Amazon SNS topic for external or cross-account access notices and subscribing the security team's email addresses to the topic is a valid way to notify the security incidents. Amazon SNS is a fully managed messaging service that enables you to decouple and scale microservices, distributed systems, and serverless applications. SNS can deliver messages to a variety of endpoints, such as email, SMS, or HTTP. You can use SNS to send email notifications to the security team when a critical security finding is detected.
References:

- AWS Step Functions
- AWS Batch
- Amazon EventBridge
- Amazon CloudWatch

- Amazon SQS
- Amazon SNS

NEW QUESTION 7
Company A has an AWS account that is named Account A. Company A recently acquired Company B, which has an AWS account that is named Account B. Company B stores its files in an Amazon S3 bucket.
The administrators need to give a user from Account A full access to the S3 bucket in Account B.
After the administrators adjust the IAM permissions for the user in AccountA to access the S3 bucket in Account B, the user still cannot access any files in the S3 bucket.
Which solution will resolve this issue?

A. In Account B, create a bucket ACL to allow the user from Account A to access the S3 bucket in Account B.
B. In Account B, create an object ACL to allow the user from Account A to access all the objects in the S3 bucket in Account B.
C. In Account B, create a bucket policy to allow the user from Account A to access the S3 bucket in Account B.
D. In Account B, create a user policy to allow the user from Account A to access the S3 bucket in Account B.

Answer: C

Explanation:
A bucket policy is a resource-based policy that defines permissions for a specific S3 bucket. It can be used to grant cross-account access to another AWS account or an IAM user or role in another account. A bucket policy can also specify which actions, resources, and conditions are allowed or denied.
A bucket ACL is an access control list that grants basic read or write permissions to predefined groups of users. It cannot be used to grant cross-account access to a specific IAM user or role in another account.
An object ACL is an access control list that grants basic read or write permissions to predefined groups of users for a specific object in an S3 bucket. It cannot be used to grant cross-account access to a specific IAM user or role in another account.
A user policy is an IAM policy that defines permissions for an IAM user or role in the same account. It cannot be used to grant cross-account access to another AWS account or an IAM user or role in another account.
For more information, see Provide cross-account access to objects in Amazon S3 buckets and Example 2: Bucket owner granting cross-account bucket permissions.

NEW QUESTION 8
A Security Engineer is asked to update an AWS CloudTrail log file prefix for an existing trail. When attempting to save the change in the CloudTrail console, the Security Engineer receives the following error message: `There is a problem with the bucket policy.` What will enable the Security Engineer to save the change?

A. Create a new trail with the updated log file prefix, and then delete the original trai
B. Update the existing bucket policy in the Amazon S3 console with the new log file prefix, and then update the log file prefix in the CloudTrail console.
C. Update the existing bucket policy in the Amazon S3 console to allow the Security Engineer's Principal to perform PutBucketPolicy, and then update the log file prefix in the CloudTrail console.
D. Update the existing bucket policy in the Amazon S3 console with the new log file prefix, and then update the log file prefix in the CloudTrail console.
E. Update the existing bucket policy in the Amazon S3 console to allow the Security Engineer's Principal to perform GetBucketPolicy, and then update the log file prefix in the CloudTrail console.

Answer: C

Explanation:
The correct answer is C. Update the existing bucket policy in the Amazon S3 console with the new log file prefix, and then update the log file prefix in the CloudTrail console.
According to the AWS documentation1, a bucket policy is a resource-based policy that you can use to grant access permissions to your Amazon S3 bucket and the objects in it. Only the bucket owner can associate a policy with a bucket. The permissions attached to the bucket apply to all of the objects in the bucket that are owned by the bucket owner.
When you create a trail in CloudTrail, you can specify an existing S3 bucket or create a new one to store your log files. CloudTrail automatically creates a bucket policy for your S3 bucket that grants CloudTrail write-only access to deliver log files to your bucket. The bucket policy also grants read-only access to AWS services that you can use to view and analyze your log data, such as Amazon Athena, Amazon CloudWatch Logs, and Amazon QuickSight.
If you want to update the log file prefix for an existing trail, you must also update the existing bucket policy in the S3 console with the new log file prefix. The log file prefix is part of the resource ARN that identifies the objects in your bucket that CloudTrail can access. If you don’t update the bucket policy with the new log file prefix, CloudTrail will not be able to deliver log files to your bucket, and you will receive an error message when you try to save the change in the CloudTrail console.
The other options are incorrect because:

- A. Creating a new trail with the updated log file prefix, and then deleting the original trail is not necessary and may cause data loss or inconsistency. You can simply update the existing trail and its associated bucket policy with the new log file prefix.
- B. Updating the existing bucket policy in the S3 console to allow the Security Engineer’s Principal to perform PutBucketPolicy is not relevant to this issue. The PutBucketPolicy action allows you to create or replace a policy on a bucket, but it does not affect CloudTrail’s ability to deliver log files to your bucket. You still need to update the existing bucket policy with the new log file prefix.
- D. Updating the existing bucket policy in the S3 console to allow the Security Engineer’s Principal to perform GetBucketPolicy is not relevant to this issue. The GetBucketPolicy action allows you to retrieve a policy on a bucket, but it does not affect CloudTrail’s ability to deliver log files to your bucket. You still need to update the existing bucket policy with the new log file prefix.

References:
1: Using bucket policies - Amazon Simple Storage Service

NEW QUESTION 9
A company wants to migrate its static primary domain website to AWS. The company hosts the website and DNS servers internally. The company wants the website to enforce SSL/TLS encryption block IP addresses from outside the United States (US), and take advantage of managed services whenever possible. Which solution will meet these requirements?

A. Migrate the website to Amazon S3 Import a public SSL certificate to an Application Loa
B. Balancer with rules to block traffic from outside the US Migrate DNS to Amazon Route 53.
C. Migrate the website to Amazon EC2 Import a public SSL certificate that is created by AWS Certificate Manager (ACM) to an Application Load Balancer with rules to block traffic from outside the US Update DNS accordingly.

D. Migrate the website to Amazon S3. Import a public SSL certificate to Amazon CloudFront Use AWS WAF rules to block traffic from outside the US Update DNS.accordingly
E. Migrate the website to Amazon S3 Import a public SSL certificate that is created by AWS Certificate Manager (ACM) to Amazo
F. CloudFront Configure CloudFront to block traffic from outside the U
G. Migrate DNS to Amazon Route 53.

Answer: D

Explanation:
To migrate the static website to AWS and meet the requirements, the following steps are required:

- Migrate the website to Amazon S3, which is a highly scalable and durable object storage service that can host static websites. To do this, create an S3 bucket with the same name as the domain name of the website, enable static website hosting for the bucket, upload the website files to the bucket, and configure the bucket policy to allow public read access to the objects. For more information, see Hosting a static website on Amazon S3.
- Import a public SSL certificate that is created by AWS Certificate Manager (ACM) to Amazon CloudFront, which is a global content delivery network (CDN) service that can improve the performance and security of web applications. To do this, request or import a public SSL certificate for the domain name of the website using ACM, create a CloudFront distribution with the S3 bucket as the origin, and associate the SSL certificate with the distribution. For more information, see Using alternate domain names and HTTPS.
- Configure CloudFront to block traffic from outside the US, which is one of the requirements. To do this, create a CloudFront web ACL using AWS WAF, which is a web application firewall service that lets you control access to your web applications. In the web ACL, create a rule that uses a geo match condition to block requests that originate from countries other than the US. Associate the web ACL with the CloudFront distribution. For more information, see How AWS WAF works with Amazon CloudFront features.
- Migrate DNS to Amazon Route 53, which is a highly available and scalable cloud DNS service that can route traffic to various AWS services. To do this, register or transfer your domain name to Route 53, create a hosted zone for your domain name, and create an alias record that points your domain name to your CloudFront distribution. For more information, see Routing traffic to an Amazon CloudFront web distribution by using your domain name.

The other options are incorrect because they either do not implement SSL/TLS encryption for the website (A), do not use managed services whenever possible (B), or do not block IP addresses from outside the US ©.Verified References:

- https://docs.aws.amazon.com/AmazonS3/latest/userguide/HostingWebsiteOnS3Setup.html
- https://docs.aws.amazon.com/waf/latest/developerguide/waf-cloudfront.html
- https://docs.aws.amazon.com/Route53/latest/DeveloperGuide/routing-to-cloudfront-distribution.html

NEW QUESTION 10
A company has an organization in AWS Organizations. The company wants to use AWS CloudFormation StackSets in the organization to deploy various AWS design patterns into environments. These patterns consist of Amazon EC2 instances, Elastic Load Balancing (ELB) load balancers, Amazon RDS databases, and Amazon Elastic Kubernetes Service (Amazon EKS) clusters or Amazon Elastic Container Service (Amazon ECS) clusters.
Currently, the company's developers can create their own CloudFormation stacks to increase the overall speed of delivery. A centralized CI/CD pipeline in a shared services AWS account deploys each CloudFormation stack.
The company's security team has already provided requirements for each service in accordance with internal standards. If there are any resources that do not comply with the internal standards, the security team must receive notification to take appropriate action. The security team must implement a notification solution that gives developers the ability to maintain the same overall delivery speed that they currently have.
Which solution will meet these requirements in the MOST operationally efficient way?

A. Create an Amazon Simple Notification Service (Amazon SNS) topi
B. Subscribe the security team's email addresses to the SNS topi
C. Create a custom AWS Lambda function that will run the aws cloudformation validate-template AWS CLI command on all CloudFormation templates before the build stage in the CI/CD pipelin
D. Configure the CI/CD pipeline to publish a notification to the SNS topic if any issues are found.
E. Create an Amazon Simple Notification Service (Amazon SNS) topi
F. Subscribe the security team's email addresses to the SNS topi
G. Create custom rules in CloudFormation Guard for each resource configuratio
H. In the ClICD pipeline, before the build stage, configure a Docker image to run the cfn-guard command on the CloudFormation templat
I. Configure the CI/CD pipeline to publish a notification to the SNS topic if any issues are found.
J. Create an Amazon Simple Notification Service (Amazon SNS) topic and an Am-azon Simple Queue Service (Amazon SQS) queu
K. Subscribe the security team's email addresses to the SNS topi
L. Create an Amazon S3 bucket in the shared services AWS accoun
M. Include an event notification to publish to the SQS queue when new objects are added to the S3 bucke
N. Require the de-velopers to put their CloudFormation templates in the S3 bucke
O. Launch EC2 instances that automatically scale based on the SQS queue dept
P. Con-figure the EC2 instances to use CloudFormation Guard to scan the templates and deploy the templates if there are no issue
Q. Configure the ClICD pipe-line to publish a notification to the SNS topic if any issues are found.
R. Create a centralized CloudFormation stack set that includes a standard set of resources that the developers can deploy in each AWS accoun
S. Configure each CloudFormation template to meet the security requirement
T. For any new resources or configurations, update the CloudFormation template and send the template to the security team for revie
. When the review is com-pleted, add the new CloudFormation stack to the repository for the devel-opers to use.

Answer: B

NEW QUESTION 10
A company wants to establish separate IAM Key Management Service (IAM KMS) keys to use for different IAM services. The company's security engineer created the following key policy lo allow the infrastructure deployment team to create encrypted Amazon Elastic Block Store (Amazon EBS) volumes by assuming the InfrastructureDeployment IAM role:

```
{
    "Version": "2012-10-17",
    "Id": "key-policy-ebs",
    "Statement": [
        {
            "Sid": "Enable IAM User Permissions",
            "Effect": "Allow",
            "Principal": {
                "AWS": "arn:aws:iam::123456789012:root"
            },
            "Action": "kms:*",
            "Resource": "*"
        },
        {
            "Sid": "Allow use of the key",
            "Effect": "Allow",
            "Principal": {
                "AWS": "arn:aws:iam::123456789012:role/aws-reserved/sso.amazonaws.com/InfrastructureDeployment"
            },
            "Action": [
                "kms:Encrypt",
                "kms:Decrypt",
                "kms:ReEncrypt*",
                "kms:GenerateDataKey*",
                "kms:DescribeKey",
                "kms:CreateGrant",
                "kms:ListGrants",
                "kms:RevokeGrant"
            ],
            "Resource": "*",
            "Condition": {
                "StringEquals": {
                    "kms:ViaService": "ec2.us-west-2.amazonaws.com"
                }
            }
        }
    ]
}
```

The security engineer recently discovered that IAM roles other than the InfrastructureDeployment role used this key (or other services. Which change to the policy should the security engineer make to resolve these issues?

A. In the statement block that contains the Sid "Allow use of the key", under the "Condition" block, change StringEquals to StringLike.
B. In the policy document, remove the statement Dlock that contains the Sid "Enable IAM User Permissions". Add key management policies to the KMS policy.
C. In the statement block that contains the Sid "Allow use of the Key", under the "Condition" block, change the Kms:ViaService value to ec2.us-east-1 .amazonIAM com.
D. In the policy document, add a new statement block that grants the kms:Disable' permission to the security engineer's IAM role.

Answer: C

Explanation:
To resolve the issues, the security engineer should make the following change to the policy:
> In the statement block that contains the Sid “Allow use of the key”, under the “Condition” block, change the Kms:ViaService value to ec2.us-east-1.amazonaws.com. This allows the security engineer to restrict the use of the key to only EC2 service in the us-east-1 region, and prevent other services from using the key.

NEW QUESTION 13
Auditors for a health care company have mandated that all data volumes be encrypted at rest Infrastructure is deployed mainly via IAM CloudFormation however third-party frameworks and manual deployment are required on some legacy systems
What is the BEST way to monitor, on a recurring basis, whether all EBS volumes are encrypted?

A. On a recurring basis, update an IAM user policies to require that EC2 instances are created with an encrypted volume
B. Configure an IAM Config rule lo run on a recurring basis 'or volume encryption
C. Set up Amazon Inspector rules tor volume encryption to run on a recurring schedule
D. Use CloudWatch Logs to determine whether instances were created with an encrypted volume

Answer: B

Explanation:
To support answer B, use the reference https://d1.IAMstatic.com/whitepapers/IAM-security-whitepaper.pdf "For example, IAM Config provides a managed IAM Config Rules to ensure that encryption is turned on for
all EBS volumes in your account."

NEW QUESTION 15
A company finds that one of its Amazon EC2 instances suddenly has a high CPU usage. The company does not know whether the EC2 instance is compromised or whether the operating system is performing background cleanup.
Which combination of steps should a security engineer take before investigating the issue? (Select THREE.)

A. Disable termination protection for the EC2 instance if termination protection has not been disabled.
B. Enable termination protection for the EC2 instance if termination protection has not been enabled.
C. Take snapshots of the Amazon Elastic Block Store (Amazon EBS) data volumes that are attached to the EC2 instance.
D. Remove all snapshots of the Amazon Elastic Block Store (Amazon EBS) data volumes that are attached to the EC2 instance.
E. Capture the EC2 instance metadata, and then tag the EC2 instance as under quarantine.

F. Immediately remove any entries in the EC2 instance metadata that contain sensitive information.

Answer: BCE

Explanation:
https://d1.awsstatic.com/WWPS/pdf/aws_security_incident_response.pdf

NEW QUESTION 19
A company uses Amazon RDS for MySQL as a database engine for its applications. A recent security audit revealed an RDS instance that is not compliant with company policy for encrypting data at rest. A security engineer at the company needs to ensure that all existing RDS databases are encrypted using server-side encryption and that any future deviations from the policy are detected.
Which combination of steps should the security engineer take to accomplish this? (Select TWO.)

A. Create an IAM Config rule to detect the creation of unencrypted RDS database
B. Create an Amazon EventBridge (Amazon CloudWatch Events) rule to trigger on the IAM Config rules compliance state change and use Amazon Simple Notification Service (Amazon SNS) to notify the security operations team.
C. Use IAM System Manager State Manager to detect RDS database encryption configuration drif
D. Create an Amazon EventBridge (Amazon CloudWatch Events) rule to track state changes and use Amazon Simple Notification Service (Amazon SNS) to notify the security operations team.
E. Create a read replica for the existing unencrypted RDS database and enable replica encryption in the proces
F. Once the replica becomes active, promote it into a standalone database instance and terminate the unencrypted database instance.
G. Take a snapshot of the unencrypted RDS databas
H. Copy the snapshot and enable snapshot encryption in the proces
I. Restore the database instance from the newly created encrypted snapsho
J. Terminate the unencrypted database instance.
K. Enable encryption for the identified unencrypted RDS instance by changing the configurations of the existing database

Answer: AD

NEW QUESTION 21
An ecommerce website was down for 1 hour following a DDoS attack Users were unable to connect to the website during the attack period. The ecommerce company's security team is worried about future potential attacks and wants to prepare for such events The company needs to minimize downtime in its response to similar attacks in the future.
Which steps would help achieve this9 (Select TWO)

A. Enable Amazon GuardDuty to automatically monitor for malicious activity and block unauthorized access.
B. Subscribe to IAM Shield Advanced and reach out to IAM Support in the event of an attack.
C. Use VPC Flow Logs to monitor network: traffic and an IAM Lambda function to automatically block an attacker's IP using security groups.
D. Set up an Amazon CloudWatch Events rule to monitor the IAM CloudTrail events in real time use IAM Config rules to audit the configuration, and use IAM Systems Manager for remediation.
E. Use IAM WAF to create rules to respond to such attacks

Answer: BE

Explanation:
To minimize downtime in response to DDoS attacks, the company should do the following:
- Subscribe to AWS Shield Advanced and reach out to AWS Support in the event of an attack. This provides access to 24x7 support from the AWS DDoS Response Team (DRT), as well as advanced detection and mitigation capabilities for network and application layer attacks.
- Use AWS WAF to create rules to respond to such attacks. This allows the company to filter web requests based on IP addresses, headers, body, or URI strings, and block malicious requests before they reach the web applications.

NEW QUESTION 23
A company has two AWS accounts. One account is for development workloads. The other account is for production workloads. For compliance reasons the production account contains all the AWS Key Management. Service (AWS KMS) keys that the company uses for encryption.
The company applies an IAM role to an AWS Lambda function in the development account to allow secure access to AWS resources. The Lambda function must access a specific KMS customer managed key that exists in the production account to encrypt the Lambda function's data.
Which combination of steps should a security engineer take to meet these requirements? (Select TWO.)

A. Configure the key policy for the customer managed key in the production account to allow access to the Lambda service.
B. Configure the key policy for the customer managed key in the production account to allow access to the IAM role of the Lambda function in the development account.
C. Configure a new IAM policy in the production account with permissions to use the customer managed ke
D. Apply the IAM policy to the IAM role that the Lambda function in the development account uses.
E. Configure a new key policy in the development account with permissions to use the customer managed ke
F. Apply the key policy to the IAM role that the Lambda function in the development account uses.
G. Configure the IAM role for the Lambda function in the development account by attaching an IAM policy that allows access to the customer managed key in the production account.

Answer: BE

Explanation:
To allow a Lambda function in one AWS account to access a KMS customer managed key in another AWS account, the following steps are required:
- Configure the key policy for the customer managed key in the production account to allow access to the IAM role of the Lambda function in the development account. A key policy is a resource-based policy that defines who can use or manage a KMS key. To grant cross-account access to a KMS key, you must specify the AWS account ID and the IAM role ARN of the external principal in the key policy statement. For more information, see Allowing users in other accounts to use a KMS key.
- Configure the IAM role for the Lambda function in the development account by attaching an IAM policy that allows access to the customer managed key in the production account. An IAM policy is an identity-based policy that defines what actions an IAM entity can perform on which resources. To allow an IAM role to use

a KMS key in another account, you must specify the KMS key ARN and the kms:Encrypt action (or any other action that requires access to the KMS key) in the IAM policy statement. For more information, see Using IAM policies with AWS KMS.
This solution will meet the requirements of allowing secure access to a KMS customer managed key across AWS accounts.
The other options are incorrect because they either do not grant cross-account access to the KMS key (A, C), or do not use a valid policy type for KMS keys (D).
Verified References:
https://docs.aws.amazon.com/kms/latest/developerguide/iam-policies.html

NEW QUESTION 27
A company's cloud operations team is responsible for building effective security for IAM cross-account access. The team asks a security engineer to help troubleshoot why some developers in the developer account (123456789012) in the developers group are not able to assume a cross-account role (ReadS3) into a production account (999999999999) to read the contents of an Amazon S3 bucket (productionapp). The two account policies are as follows:

Developer account 123456789012:
Developer group permissions:

```
{
    "Version": "2012-10-17",
    "Statement": {
        "Effect": "Allow",
        "Action": "sts:AssumeRole",
        "Resource": "arn:aws:iam::999999999999:role/ReadS3"
    }
}
```

Production account 999999999999:
Production account ReadS3 role policy:

```
{
    "Version": "2012-10-17",
    "Statement": [
        {
            "Effect": "Allow",
            "Action": "s3:ListAllMyBuckets",
            "Resource": "*"
        },
        {
            "Effect": "Allow",
            "Action": [
                "s3:ListBucket",
                "s3:GetBucketLocation"
            ],
```

Production account ReadS3 role policy - trust relationship:

```
{
    "Version": "2012-10-17",
    "Statement": [
        {
            "Effect": "Allow",
            "Principal": {
                "AWS": "arn:aws:iam::888888888888:root"
            },
            "Action": "sts:AssumeRole",
            "Condition": {
                "BoolIfExists": {
                    "aws:MultiFactorAuthPresent": "true"
                }
            }
        }
    ]
}
```

Which recommendations should the security engineer make to resolve this issue? (Select TWO.)

A. Ask the developers to change their password and use a different web browser.
B. Ensure that developers are using multi-factor authentication (MFA) when they log in to their developer account as the developer role.
C. Modify the production account ReadS3 role policy to allow the PutBucketPolicy action on the productionapp S3 bucket.
D. Update the trust relationship policy on the production account S3 role to allow the account number of the developer account.
E. Update the developer group permissions in the developer account to allow access to the productionapp S3 bucket.

Answer: AD

NEW QUESTION 29

There are currently multiple applications hosted in a VPC. During monitoring it has been noticed that multiple port scans are coming in from a specific IP Address block. The internal security team has requested that all offending IP Addresses be denied for the next 24 hours. Which of the following is the best method to quickly and temporarily deny access from the specified IP Address's.
Please select:

A. Create an AD policy to modify the Windows Firewall settings on all hosts in the VPC to deny access from the IP Address block.
B. Modify the Network ACLs associated with all public subnets in the VPC to deny access from the IP Address block.
C. Add a rule to all of the VPC Security Groups to deny access from the IP Address block.
D. Modify the Windows Firewall settings on all AMI'S that your organization uses in that VPC to deny access from the IP address block.

Answer: B

Explanation:
NACL acts as a firewall at the subnet level of the VPC and we can deny the offending IP address block at the subnet level using NACL rules to block the incoming traffic to the VPC instances. Since NACL rules are applied as per the Rule numbers make sure that this rule number should take precedence over other rule numbers if there are any such rules that will allow traffic from these IP ranges. The lowest rule number has more precedence over a rule that has a higher number.
The IAM Documentation mentions the following as a best practices for IAM users
For extra security, enable multi-factor authentication (MFA) for privileged IAM users (users who are allowed access to sensitive resources or APIs). With MFA, users have a device that generates a unique authentication code (a one-time password, or OTP). Users must provide both their normal credentials (like their user name and password) and the OTP. The MFA device can either be a special piece of hardware, or it can be a virtual device (for example, it can run in an app on a smartphone).
Options C is invalid because these options are not available Option D is invalid because there is not root access for users
For more information on IAM best practices, please visit the below URL: https://docs.IAM.amazon.com/IAM/latest/UserGuide/best-practices.html
The correct answer is: Modify the Network ACLs associated with all public subnets in the VPC to deny access from the IP Address block.
omit your Feedback/Queries to our Experts

NEW QUESTION 34
Amazon GuardDuty has detected communications to a known command and control endpoint from a company's Amazon EC2 instance. The instance was found to be running a vulnerable version of a common web framework. The company's security operations team wants to quickly identity other compute resources with the specific version of that framework installed.
Which approach should the team take to accomplish this task?

A. Scan all the EC2 instances for noncompliance with IAM Confi
B. Use Amazon Athena to query IAM CloudTrail logs for the framework installation
C. Scan all the EC2 instances with the Amazon Inspector Network Reachability rules package to identity instances running a web server with RecognizedPortWithListener findings
D. Scan all the EC2 instances with IAM Systems Manager to identify the vulnerable version of the web framework
E. Scan an the EC2 instances with IAM Resource Access Manager to identify the vulnerable version of the web framework

Answer: C

Explanation:
To quickly identify other compute resources with the specific version of the web framework installed, the team should do the following:
> Scan all the EC2 instances with AWS Systems Manager to identify the vulnerable version of the web framework. This allows the team to use AWS Systems Manager Inventory to collect and query
information about the software installed on their EC2 instances, and to filter the results by software name and version.

NEW QUESTION 35
A company has an application that uses dozens of Amazon DynamoDB tables to store data. Auditors find that the tables do not comply with the company's data protection policy.
The company's retention policy states that all data must be backed up twice each month: once at midnight on the 15th day of the month and again at midnight on the 25th day of the month. The company must retain the backups for 3 months.
Which combination of steps should a security engineer take to meet these re-quirements? (Select TWO.)

A. Use the DynamoDB on-demand backup capability to create a backup pla
B. Con-figure a lifecycle policy to expire backups after 3 months.
C. Use AWS DataSync to create a backup pla
D. Add a backup rule that includes a retention period of 3 months.
E. Use AVVS Backup to create a backup pla
F. Add a backup rule that includes a retention period of 3 months.
G. Set the backup frequency by using a cron schedule expressio
H. Assign each DynamoDB table to the backup plan.
I. Set the backup frequency by using a rate schedule expressio
J. Assign each DynamoDB table to the backup plan.

Answer: AD

NEW QUESTION 39
A company in France uses Amazon Cognito with the Cognito Hosted Ul as an identity broker for sign-in and sign-up processes. The company is marketing an application and expects that all the application's users will come from France.
When the company launches the application the company's security team observes fraudulent sign-ups for the application. Most of the fraudulent registrations are from users outside of France.
The security team needs a solution to perform custom validation at sign-up Based on the results of the validation the solution must accept or deny the registration request.
Which combination of steps will meet these requirements? (Select TWO.)

A. Create a pre sign-up AWS Lambda trigge
B. Associate the Amazon Cognito function with the Amazon Cognito user pool.
C. Use a geographic match rule statement to configure an AWS WAF web AC

D. Associate the web ACL with the Amazon Cognito user pool.
E. Configure an app client for the application's Amazon Cognito user poo
F. Use the app client ID to validate the requests in the hosted UI.
G. Update the application's Amazon Cognito user pool to configure a geographic restriction setting.
H. Use Amazon Cognito to configure a social identity provider (IdP) to validate the requests on the hosted UI.

Answer: B

Explanation:
https://docs.aws.amazon.com/cognito/latest/developerguide/user-pool-lambda-post-authentication.html

NEW QUESTION 43
A company is using Amazon Macie, AWS Firewall Manager, Amazon Inspector, and AWS Shield Advanced in its AWS account. The company wants to receive alerts if a DDoS attack occurs against the account.
Which solution will meet this requirement?

A. Use Macie to detect an active DDoS even
B. Create Amazon CloudWatch alarms that respond to Macie findings.
C. Use Amazon Inspector to review resources and to invoke Amazon CloudWatch alarms for any resources that are vulnerable to DDoS attacks.
D. Create an Amazon CloudWatch alarm that monitors Firewall Manager metrics for an active DDoS event.
E. Create an Amazon CloudWatch alarm that monitors Shield Advanced metrics for an active DDoS event.

Answer: D

Explanation:
This answer is correct because AWS Shield Advanced is a service that provides comprehensive protection
against DDoS attacks of any size or duration. It also provides metrics and reports on the DDoS attack vectors, duration, and size. You can create an Amazon CloudWatch alarm that monitors Shield Advanced metrics such as DDoSAttackBitsPerSecond, DDoSAttackPacketsPerSecond, and DDoSAttackRequestsPerSecond to receive alerts if a DDoS attack occurs against your account.
For more information, see Monitoring AWS Shield Advanced with Amazon CloudWatch and AWS Shield Advanced metrics and alarms.

NEW QUESTION 47
A company has a relational database workload that runs on Amazon Aurora MySQL. According to new compliance standards the company must rotate all database credentials every 30 days. The company needs a solution that maximizes security and minimizes development effort.
Which solution will meet these requirements?

A. Store the database credentials in AWS Secrets Manage
B. Configure automatic credential rotation tor every 30 days.
C. Store the database credentials in AWS Systems Manager Parameter Stor
D. Create an AWS Lambda function to rotate the credentials every 30 days.
E. Store the database credentials in an environment file or in a configuration fil
F. Modify the credentials every 30 days.
G. Store the database credentials in an environment file or in a configuration fil
H. Create an AWS Lambda function to rotate the credentials every 30 days.

Answer: A

Explanation:
To rotate database credentials every 30 days, the most secure and efficient solution is to store the database credentials in AWS Secrets Manager and configure automatic credential rotation for every 30 days. Secrets Manager can handle the rotation of the credentials in both the secret and the database, and it can use AWS KMS to encrypt the credentials. Option B is incorrect because it requires creating a custom Lambda function to rotate the credentials, which is more effort than using Secrets Manager. Option C is incorrect because it stores the database credentials in an environment file or a configuration file, which is less secure than using Secrets Manager. Option D is incorrect because it combines the drawbacks of option B and option C. Verified References:
https://docs.aws.amazon.com/secretsmanager/latest/userguide/rotating-secrets.html
https://docs.aws.amazon.com/secretsmanager/latest/userguide/rotate-secrets_turn-on-for-other.html

NEW QUESTION 52
A company uses a third-party identity provider and SAML-based SSO for its AWS accounts. After the third-party identity provider renewed an expired signing certificate, users saw the following message when trying to log in:
Error: Response Signature Invalid (Service: AWSSecurityTokenService; Status Code: 400; Error Code: InvalidIdentityToken)
A security engineer needs to provide a solution that corrects the error and min-imizes operational overhead.
Which solution meets these requirements?

A. Upload the third-party signing certificate's new private key to the AWS identity provider entity defined in AWS Identity and Access Management (IAM) by using the AWS Management Console.
B. Sign the identity provider's metadata file with the new public ke
C. Upload the signature to the AWS identity provider entity defined in AWS Identity and Access Management (IAM) by using the AWS CU.
D. Download the updated SAML metadata file from the identity service provid-e
E. Update the file in the AWS identity provider entity defined in AWS Identity and Access Management (IAM) by using the AWS CLI.
F. Configure the AWS identity provider entity defined in AWS Identity and Ac-cess Management (IAM) to synchronously fetch the new public key by using the AWS Management Console.

Answer: C

Explanation:
This answer is correct because downloading the updated SAML metadata file from the identity service provider ensures that AWS has the latest information about the identity provider, including the new public key. Updating the file in the AWS identity provider entity defined in IAM by using the AWS CLI allows AWS to verify the signature of the SAML assertions sent by the identity provider. This solution also minimizes operational overhead because it can be automated with a script or a cron job.

NEW QUESTION 53
A company is using an AWS Key Management Service (AWS KMS) AWS owned key in its application to encrypt files in an AWS account The company's security team wants the ability to change to new key material for new files whenever a potential key breach occurs A security engineer must implement a solution that gives the security team the ability to change the key whenever the team wants to do so
Which solution will meet these requirements?

A. Create a new customer managed key Add a key rotation schedule to the key Invoke the key rotation schedule every time the security team requests a key change
B. Create a new AWS managed key Add a key rotation schedule to the key Invoke the key rotation schedule every time the security team requests a key change
C. Create a key alias Create a new customer managed key every time the security team requests a key change Associate the alias with the new key
D. Create a key alias Create a new AWS managed key every time the security team requests a key change Associate the alias with the new key

Answer: A

Explanation:
To meet the requirement of changing the key material for new files whenever a potential key breach occurs, the most appropriate solution would be to create a new customer managed key, add a key rotation schedule to the key, and invoke the key rotation schedule every time the security team requests a key change.
References: : Rotating AWS KMS keys - AWS Key Management Service

NEW QUESTION 58
Which of the following are valid configurations for using SSL certificates with Amazon CloudFront? (Select THREE)

A. Default AWS Certificate Manager certificate
B. Custom SSL certificate stored in AWS KMS
C. Default CloudFront certificate
D. Custom SSL certificate stored in AWS Certificate Manager
E. Default SSL certificate stored in AWS Secrets Manager
F. Custom SSL certificate stored in AWS IAM

Answer: ABC

Explanation:
The key length for an RSA certificate that you use with CloudFront is 2048 bits, even though ACM supports larger keys. If you use an imported certificate with CloudFront, your key length must be 1024 or 2048 bits and cannot exceed 2048 bits. You must import the certificate in the US East (N. Virginia) Region. You must have permission to use and import the SSL/TLS certificate
https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/cnames-and-https-requirements.html

NEW QUESTION 63
A security engineer must troubleshoot an administrator's inability to make an existing Amazon S3 bucket public in an account that is part of an organization n IAM Organizations. The administrator switched the role from the master account to a member account and then attempted to make one S3 bucket public. This action was immediately denied
Which actions should the security engineer take to troubleshoot the permissions issue? (Select TWO.)

A. Review the cross-account role permissions and the S3 bucket policy Verify that the Amazon S3 block public access option in the member account is deactivated.
B. Review the role permissions m the master account and ensure it has sufficient privileges to perform S3 operations
C. Filter IAM CloudTrail logs for the master account to find the original deny event and update the cross-account role m the member account accordingly Verify that the Amazon S3 block public access option in the master account is deactivated.
D. Evaluate the SCPs covering the member account and the permissions boundary of the role in the member account for missing permissions and explicit denies.
E. Ensure the S3 bucket policy explicitly allows the s3 PutBucketPublicAccess action for the role m the member account

Answer: DE

Explanation:
- A is incorrect because reviewing the cross-account role permissions and the S3 bucket policy is not enough to troubleshoot the permissions issue. You also need to verify that the Amazon S3 block public access option in the member account is deactivated, as well as the permissions boundary and the SCPs of the role in the member account.
- D is correct because evaluating the SCPs and the permissions boundary of the role in the member account can help you identify any missing permissions or explicit denies that could prevent the administrator from making the S3 bucket public.
- E is correct because ensuring that the S3 bucket policy explicitly allows the s3 PutBucketPublicAccess action for the role in the member account can help you override any block public access settings that could prevent the administrator from making the S3 bucket public.

NEW QUESTION 64
A company used a lift-and-shift approach to migrate from its on-premises data centers to the AWS Cloud. The company migrated on-premises VMS to Amazon EC2 in-stances. Now the company wants to replace some of components that are running on the EC2 instances with managed AWS services that provide similar functionality.
Initially, the company will transition from load balancer software that runs on EC2 instances to AWS Elastic Load Balancers. A security engineer must ensure that after this transition, all the load balancer logs are centralized and searchable for auditing. The security engineer must also ensure that metrics are generated to show which ciphers are in use.
Which solution will meet these requirements?

A. Create an Amazon CloudWatch Logs log grou
B. Configure the load balancers to send logs to the log grou
C. Use the CloudWatch Logs console to search the log
D. Create CloudWatch Logs filters on the logs for the required met-rics.
E. Create an Amazon S3 bucke
F. Configure the load balancers to send logs to the S3 bucke
G. Use Amazon Athena to search the logs that are in the S3 bucke

H. Create Amazon CloudWatch filters on the S3 log files for the re-quired metrics.
I. Create an Amazon S3 bucke
J. Configure the load balancers to send logs to the S3 bucke
K. Use Amazon Athena to search the logs that are in the S3 bucke
L. Create Athena queries for the required metric
M. Publish the metrics to Amazon CloudWatch.
N. Create an Amazon CloudWatch Logs log grou
O. Configure the load balancers to send logs to the log grou
P. Use the AWS Management Console to search the log
Q. Create Amazon Athena queries for the required metric
R. Publish the metrics to Amazon CloudWatch.

Answer: C

Explanation:

- Amazon S3 is a service that provides scalable, durable, and secure object storage. You can use Amazon S3 to store and retrieve any amount of data from anywhere on the web1
- AWS Elastic Load Balancing is a service that distributes incoming application or network traffic across multiple targets, such as EC2 instances, containers, or IP addresses. You can use Elastic Load Balancing to increase the availability and fault tolerance of your applications2
- Elastic Load Balancing supports access logging, which captures detailed information about requests sent to your load balancer. Each log contains information such as the time the request was received, the client's IP address, latencies, request paths, and server responses. You can use access logs to analyze traffic patterns and troubleshoot issues3
- You can configure your load balancer to store access logs in an Amazon S3 bucket that you specify.

You can also specify the interval for publishing the logs, which can be 5 or 60 minutes. The logs are stored in a hierarchical folder structure by load balancer name, IP address, year, month, day, and time.

- Amazon Athena is a service that allows you to analyze data in Amazon S3 using standard SQL. You can use Athena to run ad-hoc queries and get results in seconds. Athena is serverless, so there is no infrastructure to manage and you pay only for the queries that you run.
- You can use Athena to search the access logs that are stored in your S3 bucket. You can create a table in Athena that maps to your S3 bucket and then run SQL queries on the table. You can also use the Athena console or API to view and download the query results.
- You can also use Athena to create queries for the required metrics, such as the number of requests per cipher or protocol. You can then publish the metrics to Amazon CloudWatch, which is a service that monitors and manages your AWS resources and applications. You can use CloudWatch to collect and track metrics, create alarms, and automate actions based on the state of your resources.
- By using this solution, you can meet the requirements of ensuring that all the load balancer logs are centralized and searchable for auditing and that metrics are generated to show which ciphers are in use.

NEW QUESTION 66
Your CTO is very worried about the security of your IAM account. How best can you prevent hackers from completely hijacking your account?
Please select:

A. Use short but complex password on the root account and any administrators.
B. Use IAM IAM Geo-Lock and disallow anyone from logging in except for in your city.
C. Use MFA on all users and accounts, especially on the root account.
D. Don't write down or remember the root account password after creating the IAM account.

Answer: C

Explanation:
Multi-factor authentication can add one more layer of security to your IAM account Even when you go to your Security Credentials dashboard one of the items is to enable MFA on your root account
C:\Users\wk\Desktop\mudassar\Untitled.jpg

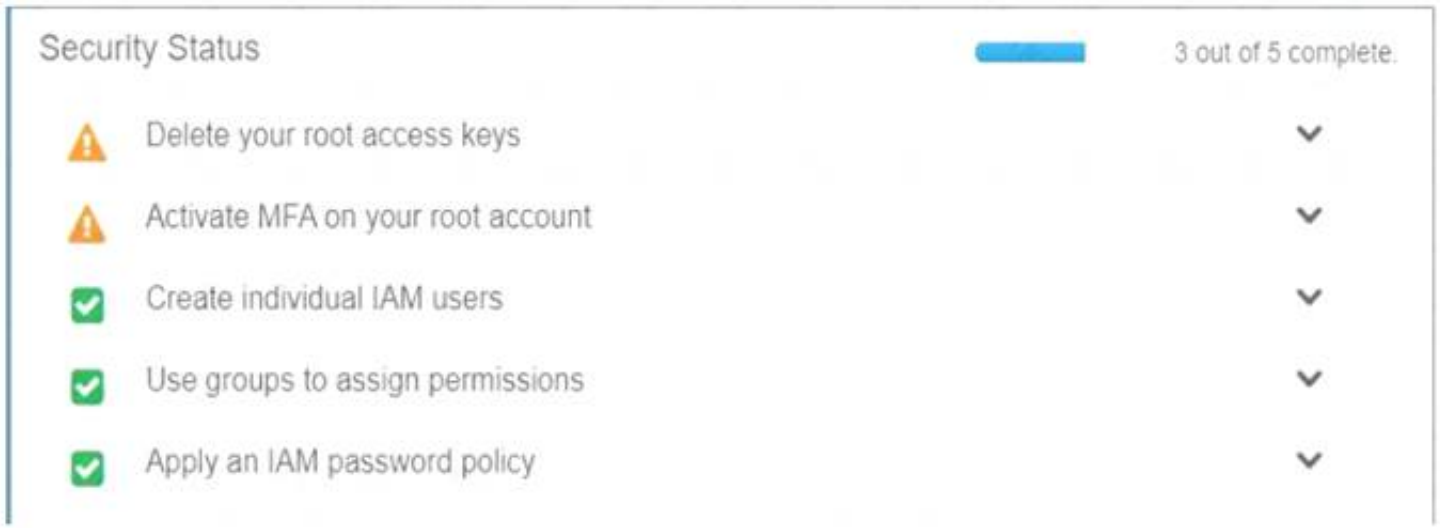

Option A is invalid because you need to have a good password policy Option B is invalid because there is no IAM Geo-Lock Option D is invalid because this is not a recommended practices For more information on MFA, please visit the below URL
http://docs.IAM.amazon.com/IAM/latest/UserGuide/id
credentials mfa.htmll
The correct answer is: Use MFA on all users and accounts, especially on the root account. Submit your Feedback/Queries to our Experts

NEW QUESTION 68
Your company is planning on using bastion hosts for administering the servers in IAM. Which of the following is the best description of a bastion host from a security perspective?
Please select:

A. A Bastion host should be on a private subnet and never a public subnet due to security concerns
B. A Bastion host sits on the outside of an internal network and is used as a gateway into the private network and is considered the critical strong point of the

network
C. Bastion hosts allow users to log in using RDP or SSH and use that session to S5H into internal network to access private subnet resources.
D. A Bastion host should maintain extremely tight security and monitoring as it is available to the public

Answer: C

Explanation:
A bastion host is a special purpose computer on a network specifically designed and configured to withstand attacks. The computer generally hosts a single application, for example a proxy server, and all other services are removed or limited to reduce the threat to the computer.
In IAM, A bastion host is kept on a public subnet. Users log on to the bastion host via SSH or RDP and then use that session to manage other hosts in the private subnets.
Options A and B are invalid because the bastion host needs to sit on the public network. Option D is invalid because bastion hosts are not used for monitoring For more information on bastion hosts, just browse to the below URL:
https://docsIAM.amazon.com/quickstart/latest/linux-bastion/architecture.htl
The correct answer is: Bastion hosts allow users to log in using RDP or SSH and use that session to SSH into internal network to access private subnet resources.
Submit your Feedback/Queries to our Experts

NEW QUESTION 72
A Network Load Balancer (NLB) target instance is not entering the InService state. A security engineer determines that health checks are failing.
Which factors could cause the health check failures? (Select THREE.)

A. The target instance's security group does not allow traffic from the NLB.
B. The target instance's security group is not attached to the NLB.
C. The NLB's security group is not attached to the target instance.
D. The target instance's subnet network ACL does not allow traffic from the NLB.
E. The target instance's security group is not using IP addresses to allow traffic from the NLB.
F. The target network ACL is not attached to the NLB.

Answer: ACD

NEW QUESTION 73
A security team is working on a solution that will use Amazon EventBridge (Amazon CloudWatch Events) to monitor new Amazon S3 objects. The solution will monitor for public access and for changes to any S3 bucket policy or setting that result in public access. The security team configures EventBridge to watch for specific API calls that are logged from AWS CloudTrail. EventBridge has an action to send an email notification through Amazon Simple Notification Service (Amazon SNS) to the security team immediately with details of the API call.
Specifically, the security team wants EventBridge to watch for the s3:PutObjectAcl, s3:DeleteBucketPolicy, and s3:PutBucketPolicy API invocation logs from CloudTrail. While developing the solution in a single account, the security team discovers that the s3:PutObjectAcl API call does not invoke an EventBridge event. However, the s3:DeleteBucketPolicy API call and the s3:PutBucketPolicy API call do invoke an event.
The security team has enabled CloudTrail for AWS management events with a basic configuration in the AWS Region in which EventBridge is being tested. Verification of the EventBridge event pattern indicates that the pattern is set up correctly. The security team must implement a solution so that the s3:PutObjectAcl API call will invoke an EventBridge event. The solution must not generate false notifications.
Which solution will meet these requirements?

A. Modify the EventBridge event pattern by selecting Amazon S3. Select All Events as the event type.
B. Modify the EventBridge event pattern by selecting Amazon S3. Select Bucket Level Operations as the event type.
C. Enable CloudTrail Insights to identify unusual API activity.
D. Enable CloudTrail to monitor data events for read and write operations to S3 buckets.

Answer: D

Explanation:
The correct answer is D. Enable CloudTrail to monitor data events for read and write operations to S3 buckets. According to the AWS documentation1, CloudTrail data events are the resource operations performed on or within a resource. These are also known as data plane operations. Data events are often high-volume activities. For example, Amazon S3 object-level API activity (such as GetObject, DeleteObject, and PutObject) is a data event.
By default, trails do not log data events. To record CloudTrail data events, you must explicitly add the
supported resources or resource types for which you want to collect activity. For more information, see Logging data events in the Amazon S3 User Guide2.
In this case, the security team wants EventBridge to watch for the s3:PutObjectAcl API invocation logs from CloudTrail. This API uses the acl subresource to set the access control list (ACL) permissions for a new or existing object in an S3 bucket3. This is a data event that affects the S3 object resource type. Therefore, the security team must enable CloudTrail to monitor data events for read and write operations to S3 buckets in order to invoke an EventBridge event for this API call.
The other options are incorrect because:

- A. Modifying the EventBridge event pattern by selecting Amazon S3 and All Events as the event type will not capture the s3:PutObjectAcl API call, because this is a data event and not a management event. Management events provide information about management operations that are performed on resources in your AWS account. These are also known as control plane operations4.
- B. Modifying the EventBridge event pattern by selecting Amazon S3 and Bucket Level Operations as the event type will not capture the s3:PutObjectAcl API call, because this is a data event that affects the S3 object resource type and not the S3 bucket resource type. Bucket level operations are management events that affect the configuration or metadata of an S3 bucket5.
- C. Enabling CloudTrail Insights to identify unusual API activity will not help the security team monitor new S3 objects or changes to any S3 bucket policy or setting that result in public access. CloudTrail Insights helps AWS users identify and respond to unusual activity associated with API calls and API error rates by continuously analyzing CloudTrail management events6. It does not analyze data events or generate EventBridge events.

References:
1: CloudTrail log event reference - AWS CloudTrail 2: Logging data events - AWS CloudTrail 3: PutObjectAcl - Amazon Simple Storage Service 4: [Logging management events - AWS CloudTrail] 5: [Amazon S3 Event Types - Amazon Simple Storage Service] 6: Logging Insights events for trails - AWS CloudTrail

NEW QUESTION 78
A security engineer is troubleshooting an AWS Lambda function that is named MyLambdaFunction. The function is encountering an error when the function attempts to read the objects in an Amazon S3 bucket that is named DOC-EXAMPLE-BUCKET. The S3 bucket has the following bucket policy:

```
{
    "Effect": "Allow",
    "Principal": {
        "Service": "lambda.amazonaws.com"
    },
    "Action": "s3:GetObject",
    "Resource": "arn:aws:s3:::DOC-EXAMPLE-BUCKET",
    "Condition": {
        "ArnLike": {
            "aws:SourceArn": "arn:aws:lambda:::function:MyLambdaFunction"
        }
    }
}
```

Which change should the security engineer make to the policy to ensure that the Lambda function can read the bucket objects?

A. Remove the Condition elemen
B. Change the Principal element to the following:{“AWS”: “arn "aws" ::: lambda ::: function:MyLambdaFunction”}
C. Change the Action element to the following: " s3:GetObject*"" s3:GetBucket*"
D. Change the Resource element to "arn:aws:s3:::DOC-EXAMPLE-BUCKET/*".
E. Change the Resource element to "arn:aws:lambda:::function:MyLambdaFunction". Change the Principal element to the following:{“Service”: “s3.amazonaws.com”}

Answer: C

Explanation:
The correct answer is C. Change the Resource element to “arn:aws:s3:::DOC-EXAMPLE-BUCKET/*”.
The reason is that the Resource element in the bucket policy specifies which objects in the bucket are affected by the policy. In this case, the policy only applies to the bucket itself, not the objects inside it. Therefore, the Lambda function cannot access the objects with the s3:GetObject permission. To fix this, the Resource element should include a wildcard (*) to match all objects in the bucket. This way, the policy grants the Lambda function permission to read any object in the bucket.
The other options are incorrect for the following reasons:

- A. Removing the Condition element would not help, because it only restricts access based on the source IP address of the request. The Principal element should not be changed to the Lambda function ARN, because it specifies who is allowed or denied access by the policy. The policy should allow access to any principal ("*") and rely on IAM roles or policies to control access to the Lambda function.
- B. Changing the Action element to include s3:GetBucket* would not help, because it would grant additional permissions that are not needed by the Lambda function, such as s3:GetBucketAcl or s3:GetBucketPolicy. The s3:GetObject* permission is sufficient for reading objects in the bucket.
- D. Changing the Resource element to the Lambda function ARN would not make sense, because it would mean that the policy applies to the Lambda function itself, not the bucket or its objects. The Principal element should not be changed to s3.amazonaws.com, because it would grant access to any AWS service that uses S3, not just Lambda.

NEW QUESTION 82
A developer is building a serverless application hosted on AWS that uses Amazon Redshift as a data store The application has separate modules for readwrite and read-only functionality The modules need their own database users for compliance reasons
Which combination of steps should a security engineer implement to grant appropriate access? (Select TWO.)

A. Configure cluster security groups for each application module to control access to database users that are required for read-only and readwrite
B. Configure a VPC endpoint for Amazon Redshift Configure an endpoint policy that maps database users to each application module, and allow access to the tables that are required for read-only and read/write
C. Configure an 1AM policy for each module Specify the ARN of an Amazon Redshift database user that allows the GetClusterCredentials API call
D. Create local database users for each module
E. Configure an 1AM policy for each module Specify the ARN of an 1AM user that allows the GetClusterCredentials API call

Answer: A

Explanation:
To grant appropriate access to separate modules for read-write and read-only functionality in a serverless application hosted on AWS that uses Amazon Redshift as a data store, a security engineer should configure cluster security groups for each application module to control access to database users that are required for read-only and readwrite, and configure an IAM policy for each module specifying the ARN of an IAM user that allows the GetClusterCredentials API call.
References: : Amazon Redshift - Amazon Web Services : Amazon Redshift - Amazon Web Services : Identity and Access Management - AWS Management Console : AWS Identity and Access Management - AWS Management Console

NEW QUESTION 84
A security engineer is checking an AWS CloudFormation template for vulnerabilities. The security engineer finds a parameter that has a default value that exposes an application's API key in plaintext. The parameter is referenced several times throughout the template. The security engineer must replace the parameter while maintaining the ability to reference the value in the template. Which solution will meet these requirements in the MOST secure way? {resolve:s3:MyBucketName:MyObjectName}}.

A. Store the API key value as a SecureString parameter in AWS Systems Manager Parameter Stor
B. In the template, replace all references to the value with {{resolve:ssm:MySSMParameterName:I}}.
C. Store the API key value in AWS Secrets Manage
D. In the template, replace all references to the value with { {resolve:secretsmanager:MySecretId:SecretString}}.
E. Store the API key value in Amazon DynamoD

F. In the template, replace all references to the value with{{resolve:dynamodb:MyTableName:MyPrimaryKey}}.
G. Store the API key value in a new Amazon S3 bucke
H. In the template, replace all references to the value with {

Answer: B

Explanation:
The correct answer is B. Store the API key value in AWS Secrets Manager. In the template, replace all references to the value with {{resolve:secretsmanager:MySecretId:SecretString}}.
This answer is correct because AWS Secrets Manager is a service that helps you protect secrets that are needed to access your applications, services, and IT resources. You can store and manage secrets such as database credentials, API keys, and other sensitive data in Secrets Manager. You can also use Secrets Manager to rotate, manage, and retrieve your secrets throughout their lifecycle1. Secrets Manager integrates with AWS CloudFormation, which allows you to reference secrets from your templates using the
{{resolve:secretsmanager:…}} syntax2. This way, you can avoid exposing your secrets in plaintext and still
use them in your resources.
The other options are incorrect because:

- A. Storing the API key value as a SecureString parameter in AWS Systems Manager Parameter Store is not a solution, because AWS CloudFormation does not support references to SecureString parameters. This means that you cannot use the {{resolve:ssm:…}} syntax to retrieve encrypted parameter values from Parameter Store3. You would have to use a custom resource or a Lambda function to decrypt the parameter value, which adds complexity and overhead to your template.
- C. Storing the API key value in Amazon DynamoDB is not a solution, because AWS CloudFormation does not support references to DynamoDB items. This means that you cannot use the
{{resolve:dynamodb:…}} syntax to retrieve item values from DynamoDB tables4. You would have to
use a custom resource or a Lambda function to query the DynamoDB table, which adds complexity and overhead to your template.
- D. Storing the API key value in a new Amazon S3 bucket is not a solution, because AWS CloudFormation does not support references to S3 objects. This means that you cannot use the
{{resolve:s3:…}} syntax to retrieve object values from S3 buckets5. You would have to use a custom resource or a Lambda function to download the object from S3, which adds complexity and overhead to your template.

References:
1: What is AWS Secrets Manager? 2: Referencing AWS Secrets Manager secrets from Parameter Store parameters 3: Using dynamic references to specify template values 4: Amazon DynamoDB 5: Amazon Simple Storage Service (S3)

NEW QUESTION 89
A Security Engineer is troubleshooting an issue with a company's custom logging application. The application logs are written to an Amazon S3 bucket with event notifications enabled to send events lo an Amazon SNS topic. All logs are encrypted at rest using an IAM KMS CMK. The SNS topic is subscribed to an encrypted Amazon SQS queue. The logging application polls the queue for new messages that contain metadata about the S3 object. The application then reads the content of the object from the S3 bucket for indexing.
The Logging team reported that Amazon CloudWatch metrics for the number of messages sent or received is showing zero. No togs are being received.
What should the Security Engineer do to troubleshoot this issue?
A) Add the following statement to the IAM managed CMKs:

```
{
    "Sid": "Allow Amazon SNS to use this key",
    "Effect": "Allow",
    "Principal": {
        "Service": ["sns.amazonaws.com", "sqs.amazonaws.com", "s3.amazonaws.com"]
    },
    "Action": [
        "kms:Decrypt",
        "kms:GenerateDataKey*"
    ],
    "Resource": "*"
}
```

B)
Add the following statement to the CMK key policy:

```
{
    "Sid": "Allow Amazon SNS to use this key",
    "Effect": "Allow",
    "Principal": {
        "Service": "sns.amazonaws.com"
    },
    "Action": [
        "kms:Decrypt",
        "kms:GenerateDataKey*"
    ],
    "Resource": "*"
}
```

C)
Add the following statement to the CMK key policy:

```
{
    "Sid": "Allow Amazon SNS to use this key",
    "Effect": "Allow",
    "Principal": {
        "Service": "sqs.amazonaws.com"
    },
    "Action": [
        "kms:Decrypt",
        "kms:GenerateDataKey*"
    ],
    "Resource": "*"
}
```

D)
Add the following statement to the CMK key policy:

```
{
    "Sid": "Allow Amazon SNS to use this key",
    "Effect": "Allow",
    "Principal": {
        "Service": "s3.amazonaws.com"
    },
    "Action": [
        "kms:Decrypt",
        "kms:GenerateDataKey*"
    ],
    "Resource": "*"
}
```

A. Option A
B. Option B
C. Option C
D. Option D

Answer: D

NEW QUESTION 93
A company needs a security engineer to implement a scalable solution for multi-account authentication and authorization. The solution should not introduce additional user-managed architectural components. Native IAM features should be used as much as possible The security engineer has set up IAM Organizations w1th all features activated and IAM SSO enabled.
Which additional steps should the security engineer take to complete the task?

A. Use AD Connector to create users and groups for all employees that require access to IAM accounts.Assign AD Connector groups to IAM accounts and link to the IAM roles in accordance with the employees'job functions and access requirements Instruct employees to access IAM accounts by using the IAM Directory Service user portal.
B. Use an IAM SSO default directory to create users and groups for all employees that require access to IAM account
C. Assign groups to IAM accounts and link to permission sets in accordance with the employees'job functions and access requirement
D. Instruct employees to access IAM accounts by using the IAM SSO user portal.
E. Use an IAM SSO default directory to create users and groups for all employees that require access to IAM account
F. Link IAM SSO groups to the IAM users present in all accounts to inherit existing permission
G. Instruct employees to access IAM accounts by using the IAM SSO user portal.
H. Use IAM Directory Service tor Microsoft Active Directory to create users and groups for all employees that require access to IAM accounts Enable IAM Management Console access in the created directory and specify IAM SSO as a source cl information tor integrated accounts and permission set
I. Instruct employees to access IAM accounts by using the IAM Directory Service user portal.

Answer: B

NEW QUESTION 97
A security engineer receives a notice from the AWS Abuse team about suspicious activity from a Linux-based Amazon EC2 instance that uses Amazon Elastic Block Store (Amazon EBS>-based storage The instance is making connections to known malicious addresses
The instance is in a development account within a VPC that is in the us-east-1 Region The VPC contains an internet gateway and has a subnet in us-east-1a and us-easMb Each subnet is associate with a route table that uses the internet gateway as a default route Each subnet also uses the default network ACL The suspicious EC2 instance runs within the us-east-1 b subnet. During an initial investigation a security engineer discovers that the suspicious instance is the only instance that runs in the subnet
Which response will immediately mitigate the attack and help investigate the root cause?

A. Log in to the suspicious instance and use the netstat command to identify remote connections Use the IP addresses from these remote connections to create deny rules in the security group of the instance Install diagnostic tools on the instance for investigation Update the outbound network ACL for the subnet inus-east-lb to explicitly deny all connections as the first rule during the investigation of the instance
B. Update the outbound network ACL for the subnet in us-east-1b to explicitly deny all connections as the first rule Replace the security group with a new security group that allows connections only from a diagnostics security group Update the outbound network ACL for the us-east-1b subnet to remove the deny all rule Launch a new EC2 instance that has diagnostic tools Assign the new security group to the new EC2 instance Use the new EC2 instance to investigate the suspicious instance
C. Ensure that the Amazon Elastic Block Store (Amazon EBS) volumes that are attached to the suspicious EC2 instance will not delete upon termination Terminate the instance Launch a new EC2 instance inus-east-1a that has diagnostic tools Mount the EBS volumes from the terminated instance for investigation
D. Create an AWS WAF web ACL that denies traffic to and from the suspicious instance Attach the AWS WAF web ACL to the instance to mitigate the attack Log in to the instance and install diagnostic tools to investigate the instance

Answer: B

Explanation:
This option suggests updating the outbound network ACL for the subnet in us-east-1b to explicitly deny all connections as the first rule, replacing the security group with a new one that only allows connections from a diagnostics security group, and launching a new EC2 instance with diagnostic tools to investigate the suspicious instance. This option will immediately mitigate the attack and provide the necessary tools for investigation.

NEW QUESTION 98
A company is using IAM Secrets Manager to store secrets for its production Amazon RDS database. The Security Officer has asked that secrets be rotated every 3 months. Which solution would allow the company to securely rotate the secrets? (Select TWO.)

A. Place the RDS instance in a public subnet and an IAM Lambda function outside the VP
B. Schedule the Lambda function to run every 3 months to rotate the secrets.
C. Place the RDS instance in a private subnet and an IAM Lambda function inside the VPC in the private subne
D. Configure the private subnet to use a NAT gatewa
E. Schedule the Lambda function to run every 3 months to rotate the secrets.
F. Place the RDS instance in a private subnet and an IAM Lambda function outside the VP
G. Configure the private subnet to use an internet gatewa
H. Schedule the Lambda function to run every 3 months lo rotate the secrets.
I. Place the RDS instance in a private subnet and an IAM Lambda function inside the VPC in the private subne
J. Schedule the Lambda function to run quarterly to rotate the secrets.
K. Place the RDS instance in a private subnet and an IAM Lambda function inside the VPC in the private subne
L. Configure a Secrets Manager interface endpoin
M. Schedule the Lambda function to run every 3 months to rotate the secrets.

Answer: BE

Explanation:
these are the solutions that can securely rotate the secrets for the production RDS database using Secrets Manager. Secrets Manager is a service that helps you manage secrets such as database credentials, API keys, and passwords. You can use Secrets Manager to rotate secrets automatically by using a Lambda function that runs on a schedule. The Lambda function needs to have access to both the RDS instance and the Secrets Manager service. Option B places the RDS instance in a private subnet and the Lambda function in the same VPC in another private subnet. The private subnet with the Lambda function needs to use a NAT gateway to access Secrets Manager over the internet. Option E places the RDS instance and the Lambda function in the same private subnet and configures a Secrets Manager interface endpoint, which is a private connection between the VPC and Secrets Manager. The other options are either insecure or incorrect for rotating secrets using Secrets Manager.

NEW QUESTION 102
A company is running an Amazon RDS for MySQL DB instance in a VPC. The VPC must not send or receive network traffic through the internet.
A security engineer wants to use AWS Secrets Manager to rotate the DB instance credentials automatically. Because of a security policy, the security engineer cannot use the standard AWS Lambda function that Secrets Manager provides to rotate the credentials.
The security engineer deploys a custom Lambda function in the VPC. The custom Lambda function will be responsible for rotating the secret in Secrets Manager.
The security engineer edits the DB instance's security group to allow connections from this function. When the function is invoked, the function cannot communicate with Secrets Manager to rotate the secret properly.
What should the security engineer do so that the function can rotate the secret?

A. Add an egress-only internet gateway to the VP
B. Allow only the Lambda function's subnet to route traffic through the egress-only internet gateway.
C. Add a NAT gateway to the VP
D. Configure only the Lambda function's subnet with a default route through the NAT gateway.
E. Configure a VPC peering connection to the default VPC for Secrets Manage
F. Configure the Lambda function's subnet to use the peering connection for routes.
G. Configure a Secrets Manager interface VPC endpoin
H. Include the Lambda function's private subnet during the configuration process.

Answer: D

Explanation:
You can establish a private connection between your VPC and Secrets Manager by creating an interface VPC endpoint. Interface endpoints are powered by AWS PrivateLink, a technology that enables you to privately access Secrets Manager APIs without an internet gateway, NAT device, VPN connection, or AWS Direct Connect connection. Reference:
https://docs.aws.amazon.com/secretsmanager/latest/userguide/vpc-endpoint-overview.html
The correct answer is D. Configure a Secrets Manager interface VPC endpoint. Include the Lambda function’s private subnet during the configuration process.
A Secrets Manager interface VPC endpoint is a private connection between the VPC and Secrets Manager that does not require an internet gateway, NAT device, VPN connection, or AWS Direct Connect connection1. By configuring a Secrets Manager interface VPC endpoint, the security engineer can enable the custom Lambda function to communicate with Secrets Manager without sending or receiving network traffic through the internet. The security engineer must include the Lambda function’s private subnet during the configuration process to allow the function to use the endpoint2.
The other options are incorrect for the following reasons:

- A. An egress-only internet gateway is a VPC component that allows outbound communication over IPv6 from instances in the VPC to the internet, and prevents the internet from initiating an IPv6 connection with the instances3. However, this option does not meet the requirement that the VPC must not send or receive network traffic through the internet. Moreover, an egress-only internet gateway is for use with IPv6 traffic only, and Secrets Manager does not support IPv6 addresses2.
- B. A NAT gateway is a VPC component that enables instances in a private subnet to connect to the internet or other AWS services, but prevents the internet from initiating connections with those instances4. However, this option does not meet the requirement that the VPC must not send or receive network traffic through the internet. Additionally, a NAT gateway requires an elastic IP address, which is a public IPv4 address4.
- C. A VPC peering connection is a networking connection between two VPCs that enables you to route traffic between them using private IPv4 addresses or IPv6 addresses5. However, this option does not work because Secrets Manager does not have a default VPC that can be peered with. Furthermore, a VPC peering connection does not provide a private connection to Secrets Manager APIs without an internet gateway or other devices2.

NEW QUESTION 103
A company is using AWS Organizations to implement a multi-account strategy. The company does not have on-premises infrastructure. All workloads run on AWS. The company currently has eight member accounts. The company anticipates that it will have no more than 20 AWS accounts total at any time.

The company issues a new security policy that contains the following requirements:
• No AWS account should use a VPC within the AWS account for workloads.
• The company should use a centrally managed VPC that all AWS accounts can access to launch workloads in subnets.
• No AWS account should be able to modify another AWS account's application resources within the centrally managed VPC.
• The centrally managed VPC should reside in an existing AWS account that is named Account-A within an organization.
The company uses an AWS CloudFormation template to create a VPC that contains multiple subnets in Account-A. This template exports the subnet IDs through the CloudFormation Outputs section.
Which solution will complete the security setup to meet these requirements?

A. Use a CloudFormation template in the member accounts to launch workload
B. Configure the template to use the Fn::ImportValue function to obtain the subnet ID values.
C. Use a transit gateway in the VPC within Account-
D. Configure the member accounts to use the transit gateway to access the subnets in Account-A to launch workloads.
E. Use AWS Resource Access Manager (AWS RAM) to share Account-A's VPC subnets with the remaining member account
F. Configure the member accounts to use the shared subnets to launch workloads.
G. Create a peering connection between Account-A and the remaining member account
H. Configure the member accounts to use the subnets in Account-A through the VPC peering connection to launch workloads.

Answer: C

Explanation:
The correct answer is C. Use AWS Resource Access Manager (AWS RAM) to share Account-A's VPC
subnets with the remaining member accounts. Configure the member accounts to use the shared subnets to launch workloads.
This answer is correct because AWS RAM is a service that helps you securely share your AWS resources across AWS accounts, within your organization or organizational units (OUs), and with IAM roles and users for supported resource types1. One of the supported resource types is VPC subnets2, which means you can share the subnets in Account-A's VPC with the other member accounts using AWS RAM. This way, you can meet the requirements of using a centrally managed VPC, avoiding duplicate VPCs in each account, and launching workloads in shared subnets. You can also control the access to the shared subnets by using IAM policies and resource-based policies3, which can prevent one account from modifying another account's resources.
The other options are incorrect because:

- A. Using a CloudFormation template in the member accounts to launch workloads and using the Fn::ImportValue function to obtain the subnet ID values is not a solution, because Fn::ImportValue can only import values that have been exported by another stack within the same region4. This means that you cannot use Fn::ImportValue to reference the subnet IDs that are exported by Account-A's CloudFormation template, unless all the member accounts are in the same region as Account-A. This option also does not avoid creating duplicate VPCs in each account, which is one of the requirements.
- B. Using a transit gateway in the VPC within Account-A and configuring the member accounts to use the transit gateway to access the subnets in Account-A to launch workloads is not a solution, because a transit gateway does not allow you to launch workloads in another account's subnets. A transit gateway is a network transit hub that enables you to route traffic between your VPCs and on-premises networks5, but it does not enable you to share subnets across accounts.
- D. Creating a peering connection between Account-A and the remaining member accounts and configuring the member accounts to use the subnets in Account-A through the VPC peering connection to launch workloads is not a solution, because a VPC peering connection does not allow you to launch workloads in another account's subnets. A VPC peering connection is a networking connection between two VPCs that enables you to route traffic between them privately6, but it does not enable you to share subnets across accounts.

References:
1: What is AWS Resource Access Manager? 2: Shareable AWS resources 3: Managing permissions for shared resources 4: Fn::ImportValue 5: What is a transit gateway? 6: What is VPC peering?

NEW QUESTION 105
A company is using Amazon Route 53 Resolver for its hybrid DNS infrastructure. The company has set up Route 53 Resolver forwarding rules for authoritative domains that are hosted on on-premises DNS servers.
A new security mandate requires the company to implement a solution to log and query DNS traffic that goes to the on-premises DNS servers. The logs must show details of the source IP address of the instance from which the query originated. The logs also must show the DNS name that was requested in Route 53 Resolver.
Which solution will meet these requirements?

A. Use VPC Traffic Mirrorin
B. Configure all relevant elastic network interfaces as the traffic source, include amazon-dns in the mirror filter, and set Amazon CloudWatch Logs as the mirror targe
C. Use CloudWatch Insights on the mirror session logs to run queries on the source IP address and DNS name.
D. Configure VPC flow logs on all relevant VPC
E. Send the logs to an Amazon S3 bucke
F. Use Amazon Athena to run SQL queries on the source IP address and DNS name.
G. Configure Route 53 Resolver query logging on all relevant VPC
H. Send the logs to Amazon CloudWatch Log
I. Use CloudWatch Insights to run queries on the source IP address and DNS name.
J. Modify the Route 53 Resolver rules on the authoritative domains that forward to the on-premises DNS server
K. Send the logs to an Amazon S3 bucke
L. Use Amazon Athena to run SQL queries on the source IP address and DNS name.

Answer: C

Explanation:
The correct answer is C. Configure Route 53 Resolver query logging on all relevant VPCs. Send the logs to Amazon CloudWatch Logs. Use CloudWatch Insights to run queries on the source IP address and DNS name.
According to the AWS documentation1, Route 53 Resolver query logging lets you log the DNS queries that Route 53 Resolver handles for your VPCs. You can send the logs to CloudWatch Logs, Amazon S3, or Kinesis Data Firehose. The logs include information such as the following:

- The AWS Region where the VPC was created
- The ID of the VPC that the query originated from
- The IP address of the instance that the query originated from
- The instance ID of the resource that the query originated from
- The date and time that the query was first made
- The DNS name requested (such as prod.example.com)
-

The DNS record type (such as A or AAAA)

- The DNS response code, such as NoError or ServFail
- The DNS response data, such as the IP address that is returned in response to the DNS query

You can use CloudWatch Insights to run queries on your log data and analyze the results using graphs and statistics2. You can filter and aggregate the log data based on any field, and use operators and functions to perform calculations and transformations. For example, you can use CloudWatch Insights to find out how many queries were made for a specific domain name, or which instances made the most queries.
Therefore, this solution meets the requirements of logging and querying DNS traffic that goes to the
on-premises DNS servers, showing details of the source IP address of the instance from which the query originated, and the DNS name that was requested in Route 53 Resolver.
The other options are incorrect because:

- A. Using VPC Traffic Mirroring would not capture the DNS queries that go to the on-premises DNS servers, because Traffic Mirroring only copies network traffic from an elastic network interface of an EC2 instance to a target for analysis3. Traffic Mirroring does not include traffic that goes through a Route 53 Resolver outbound endpoint, which is used to forward queries to on-premises DNS servers4. Therefore, this solution would not meet the requirements.
- B. Configuring VPC flow logs on all relevant VPCs would not capture the DNS name that was requested in Route 53 Resolver, because flow logs only record information about the IP traffic going to and from network interfaces in a VPC5. Flow logs do not include any information about the content or payload of a packet, such as a DNS query or response. Therefore, this solution would not meet the requirements.
- D. Modifying the Route 53 Resolver rules on the authoritative domains that forward to the on-premises DNS servers would not enable logging of DNS queries, because Resolver rules only specify how to forward queries for specified domain names to your network6. Resolver rules do not have any logging functionality by themselves. Therefore, this solution would not meet the requirements. References:

1: Resolver query logging - Amazon Route 53 2: Analyzing log data with CloudWatch Logs Insights - Amazon CloudWatch 3: What is Traffic Mirroring? - Amazon Virtual Private Cloud 4: Outbound Resolver endpoints - Amazon Route 53 5: Logging IP traffic using VPC Flow Logs - Amazon Virtual Private Cloud 6: Managing forwarding rules - Amazon Route 53

NEW QUESTION 109
There is a requirement for a company to transfer large amounts of data between IAM and an on-premise location. There is an additional requirement for low latency and high consistency traffic to IAM. Given these requirements how would you design a hybrid architecture? Choose the correct answer from the options below
Please select:

A. Provision a Direct Connect connection to an IAM region using a Direct Connect partner.
B. Create a VPN tunnel for private connectivity, which increases network consistency and reduces latency.
C. Create an iPSec tunnel for private connectivity, which increases network consistency and reduces latency.
D. Create a VPC peering connection between IAM and the Customer gateway.

Answer: A

Explanation:
IAM Direct Connect makes it easy to establish a dedicated network connection from your premises to IAM. Using IAM Direct Connect you can establish private connectivity between IAM and your datacenter, office, or colocation environment which in many cases can reduce your network costs, increase bandwidth throughput and provide a more consistent network experience than Internet-based connections.
Options B and C are invalid because these options will not reduce network latency Options D is invalid because this is only used to connect 2 VPC's
For more information on IAM direct connect, just browse to the below URL: https://IAM.amazon.com/directconnect
The correct answer is: Provision a Direct Connect connection to an IAM region using a Direct Connect partner. omit your Feedback/Queries to our Experts

NEW QUESTION 114
A company is using IAM Organizations to develop a multi-account secure networking strategy. The company plans to use separate centrally managed accounts for shared services, auditing, and security inspection. The company plans to provide dozens of additional accounts to application owners for production and development environments.
Company security policy requires that all internet traffic be routed through a centrally managed security inspection layer in the security inspection account. A security engineer must recommend a solution that minimizes administrative overhead and complexity.
Which solution meets these requirements?

A. Use IAM Control Towe
B. Modify the default Account Factory networking template to automatically associate new accounts with a centrally managed VPC through a VPC peering connection and to create a default route to the VPC peer in the default route tabl
C. Create an SCP that denies the CreateInternetGateway actio
D. Attach the SCP to all accounts except the security inspection account.
E. Create a centrally managed VPC in the security inspection accoun
F. Establish VPC peering connections between the security inspection account and other account
G. Instruct account owners to create default routes in their account route tables that point to the VPC pee
H. Create an SCP that denies theAttach InternetGateway actio
I. Attach the SCP to all accounts except the security inspection account.
J. Use IAM Control Towe
K. Modify the default Account Factory networking template to automatically associate new accounts with a centrally managed transitgateway and to create a default route to the transit gateway in the default route tabl
L. Create an SCP that denies the AttachInternetGateway actio
M. Attach the SCP to all accounts except the security inspection account.
N. Enable IAM Resource Access Manager (IAM RAM) for IAM Organization
O. Create a shared transit gateway, and make it available by using an IAM RAM resource shar
P. Create an SCP that denies the CreateInternetGateway actio
Q. Attach the SCP to all accounts except the security inspection accoun
R. Create routes in the route tables of all accounts that point to the shared transit gateway.

Answer: C

NEW QUESTION 118
A security engineer needs to configure an Amazon S3 bucket policy to restrict access to an S3 bucket that is named DOC-EXAMPLE-BUCKET. The policy must allow access to only DOC-EXAMPLE-BUCKET from only the following endpoint: vpce-1a2b3c4d. The policy must deny all access to DOC-EXAMPLE-BUCKET if

the specified endpoint is not used.
Which bucket policy statement meets these requirements?

A. A computer code with black text Description automatically generated

```
"Statement": [
    {
        "Sid": "Access-to-specific-VPCE-only",
        "Principal": "*",
        "Action": "s3:*",
        "Effect": "Allow",
        "Resource": ["arn:aws:s3:::DOC-EXAMPLE-BUCKET",
                     "arn:aws:s3:::DOC-EXAMPLE-BUCKET/*"],
        "Condition": {
            "StringNotEquals": {
                "aws:sourceVpce": "vpce-1a2b3c4d"
            }
        }
    }
]
```

B. A computer code with black text Description automatically generated

```
"Statement": [
    {
        "Sid": "Access-to-specific-VPCE-only",
        "Principal": "*",
        "Action": "s3:*",
        "Effect": "Deny",
        "Resource": ["arn:aws:s3:::DOC-EXAMPLE-BUCKET",
                     "arn:aws:s3:::DOC-EXAMPLE-BUCKET/*"],
        "Condition": {
            "StringNotEquals": {
                "aws:sourceVpce": "vpce-1a2b3c4d"
            }
        }
    }
]
```

C. A computer code with black text Description automatically generated

```
"Statement": [
    {
        "Sid": "Access-to-specific-VPCE-only",
        "Principal": "*",
        "Action": "s3:*",
        "Effect": "Deny",
        "Resource": ["arn:aws:s3:::DOC-EXAMPLE-BUCKET",
                     "arn:aws:s3:::DOC-EXAMPLE-BUCKET/*"],
        "Condition": {
            "StringEquals": {
                "aws:sourceVpce": "vpce-1a2b3c4d"
            }
        }
    }
]
```

D. A computer code with black text Description automatically generated

```
"Statement": [
    {
        "Sid": "Access-to-specific-VPCE-only",
        "Principal": "*",
        "Action": "s3:*",
        "Effect": "Allow",
        "Resource": ["arn:aws:s3:::DOC-EXAMPLE-BUCKET",
                     "arn:aws:s3:::DOC-EXAMPLE-BUCKET/*"],
        "Condition": {
            "StringEquals": {
                "aws:sourceVpce": "vpce-1a2b3c4d"
            }
        }
    }
]
```

Answer: B

Explanation:
https://docs.aws.amazon.com/AmazonS3/latest/userguide/example-bucket-policies-vpc-endpoint.html

NEW QUESTION 123
A company is building a data processing application that uses AWS Lambda functions The application's Lambda functions need to communicate with an Amazon RDS OB instance that is deployed within a VPC in the same AWS account
Which solution meets these requirements in the MOST secure way?

A. Configure the DB instance to allow public access Update the DB instance security group to allow access from the Lambda public address space for the AWS Region

B. Deploy the Lambda functions inside the VPC Attach a network ACL to the Lambda subnet Provide outbound rule access to the VPC CIDR range only Update the DB instance security group to allow traffic from 0 0 0 0/0
C. Deploy the Lambda functions inside the VPC Attach a security group to the Lambda functions Provide outbound rule access to the VPC CIDR range only Update the DB instance security group to allow traffic from the Lambda security group
D. Peer the Lambda default VPC with the VPC that hosts the DB instance to allow direct network access without the need for security groups

Answer: C

Explanation:
The AWS documentation states that you can deploy the Lambda functions inside the VPC and attach a security group to the Lambda functions. You can then provide outbound rule access to the VPC CIDR range only and update the DB instance security group to allow traffic from the Lambda security group. This method is the most secure way to meet the requirements.
References: : AWS Lambda Developer Guide

NEW QUESTION 128
A company has an AWS account that includes an Amazon S3 bucket. The S3 bucket uses server-side encryption with AWS KMS keys (SSE-KMS) to encrypt all the objects at rest by using a customer managed key. The S3 bucket does not have a bucket policy.
An IAM role in the same account has an IAM policy that allows s3 List* and s3 Get' permissions for the S3 bucket. When the IAM role attempts to access an object in the S3 bucket the role receives an access denied message.
Why does the IAM rote not have access to the objects that are in the S3 bucket?

A. The IAM rote does not have permission to use the KMS CreateKey operation.
B. The S3 bucket lacks a policy that allows access to the customer managed key that encrypts the objects.
C. The IAM rote does not have permission to use the customer managed key that encrypts the objects that are in the S3 bucket.
D. The ACL of the S3 objects does not allow read access for the objects when the objects ace encrypted at rest.

Answer: C

Explanation:
When using server-side encryption with AWS KMS keys (SSE-KMS), the requester must have both Amazon S3 permissions and AWS KMS permissions to access the objects. The Amazon S3 permissions are for the bucket and object operations, such as s3:ListBucket and s3:GetObject. The AWS KMS permissions are for the key operations, such as kms:GenerateDataKey and kms:Decrypt. In this case, the IAM role has the necessary Amazon S3 permissions, but not the AWS KMS permissions to use the customer managed key that encrypts the objects. Therefore, the IAM role receives an access denied message when trying to access the objects. Verified References:

- https://docs.aws.amazon.com/AmazonS3/latest/userguide/troubleshoot-403-errors.html
- https://repost.aws/knowledge-center/s3-access-denied-error-kms
- https://repost.aws/knowledge-center/cross-account-access-denied-error-s3

NEW QUESTION 133
A company is using AWS Organizations to manage multiple AWS accounts for its hu-man resources, finance, software development, and production departments. All the company's developers are part of the software development AWS account.
The company discovers that developers have launched Amazon EC2 instances that were preconfigured with software that the company has not approved for use.
The company wants to implement a solution to ensure that developers can launch EC2 instances with only approved software applications and only in the software de-velopment AWS account.
Which solution will meet these requirements?

A. In the software development account, create AMIS of preconfigured instanc-es that include only approved softwar
B. Include the AMI IDs in the condi-tion section of an AWS CloudFormation template to launch the appropriate AMI based on the AWS Regio
C. Provide the developers with theCloudFor-mation template to launch EC2 instances in the software development ac-count.
D. Create an Amazon EventBridge rule that runs when any EC2 RunInstances API event occurs in the software development accoun
E. Specify AWS Systems Man-ager Run Command as a target of the rul
F. Configure Run Command to run a script that will install all approved software onto the instances that the developers launch.
G. Use an AWS Service Catalog portfolio that contains EC2 products with ap-propriate AMIS that include only approved softwar
H. Grant the developers permission to portfolio access only the Service Catalog to launch a prod-uct in the software development account.
I. In the management account, create AMIS of preconfigured instances that in-clude only approved softwar
J. Use AWS CloudFormation StackSets to launch the AMIS across any AWS account in the organizatio
K. Grant the developers permission to launch the stack sets within the management account.

Answer: C

NEW QUESTION 135
A development team is using an IAM Key Management Service (IAM KMS) CMK to try to encrypt and decrypt a secure string parameter from IAM Systems Manager Parameter Store. However, the development team receives an error message on each attempt.
Which issues that are related to the CMK could be reasons for the error? (Select TWO.)

A. The CMK that is used in the attempt does not exist.
B. The CMK that is used in the attempt needs to be rotated.
C. The CMK that is used in the attempt is using the CMK's key ID instead of the CMK ARN.
D. The CMK that is used in the attempt is not enabled.
E. The CMK that is used in the attempt is using an alias.

Answer: AD

NEW QUESTION 140
A developer at a company uses an SSH key to access multiple Amazon EC2 instances. The company discovers that the SSH key has been posted on a public GitHub repository. A security engineer verifies that the key has not been used recently.
How should the security engineer prevent unauthorized access to the EC2 instances?

A. Delete the key pair from the EC2 consol
B. Create a new key pair.
C. Use the ModifylnstanceAttribute API operation to change the key on any EC2 instance that is using the key.
D. Restrict SSH access in the security group to only known corporate IP addresses.
E. Update the key pair in any AMI that is used to launch the EC2 instance
F. Restart the EC2 instances.

Answer: C

Explanation:
To prevent unauthorized access to the EC2 instances, the security engineer should do the following:
> Restrict SSH access in the security group to only known corporate IP addresses. This allows the security engineer to use a virtual firewall that controls inbound and outbound traffic for their EC2 instances, and limit SSH access to only trusted sources.

NEW QUESTION 142
A company's IAM account consists of approximately 300 IAM users. Now there is a mandate that an access change is required for 100 IAM users to have unlimited privileges to S3.As a system administrator, how can you implement this effectively so that there is no need to apply the policy at the individual user level? Please select:

A. Create a new role and add each user to the IAM role
B. Use the IAM groups and add users, based upon their role, to different groups and apply the policy to group
C. Create a policy and apply it to multiple users using a JSON script
D. Create an S3 bucket policy with unlimited access which includes each user's IAM account ID

Answer: B

Explanation:
Option A is incorrect since you don't add a user to the IAM Role Option C is incorrect since you don't assign multiple users to a policy Option D is incorrect since this is not an ideal approach
An IAM group is used to collectively manage users who need the same set of permissions. By having groups, it becomes easier to manage permissions. So if you change the permissions on the group scale, it will affect all the users in that group
For more information on IAM Groups, just browse to the below URL: https://docs.IAM.amazon.com/IAM/latest/UserGuide/id_eroups.html
The correct answer is: Use the IAM groups and add users, based upon their role, to different groups and apply the policy to group
Submit your Feedback/Queries to our Experts

NEW QUESTION 146
A company uses Amazon EC2 Linux instances in the AWS Cloud. A member of the company's security team recently received a report about common vulnerability identifiers on the instances.
A security engineer needs to verify patching and perform remediation if the instances do not have the correct patches installed. The security engineer must determine which EC2 instances are at risk and must implement a solution to automatically update those instances with the applicable patches.
What should the security engineer do to meet these requirements?

A. Use AWS Systems Manager Patch Manager to view vulnerability identifiers for missing patches on the instance
B. Use Patch Manager also to automate the patching process.
C. Use AWS Shield Advanced to view vulnerability identifiers for missing patches on the instance
D. Use AWS Systems Manager Patch Manager to automate the patching process.
E. Use Amazon GuardDuty to view vulnerability identifiers for missing patches on the instance
F. Use Amazon Inspector to automate the patching process.
G. Use Amazon Inspector to view vulnerability identifiers for missing patches on the instance
H. Use Amazon Inspector also to automate the patching process.

Answer: A

Explanation:
https://aws.amazon.com/about-aws/whats-new/2020/10/now-use-aws-systems-manager-to-view-vulnerability-id

NEW QUESTION 148
A company is evaluating the use of AWS Systems Manager Session Manager to gam access to the company's Amazon EC2 instances. However, until the company implements the change, the company must protect the key file for the EC2 instances from read and write operations by any other users.
When a security administrator tries to connect to a critical EC2 Linux instance during an emergency, the security administrator receives the following error. "Error Unprotected private key file - Permissions for' ssh/my_private_key pern' are too open".
Which command should the security administrator use to modify the private key Me permissions to resolve this error?

A. chmod 0040 ssh/my_private_key pern
B. chmod 0400 ssh/my_private_key pern
C. chmod 0004 ssh/my_private_key pern
D. chmod 0777 ssh/my_private_key pern

Answer: B

Explanation:
The error message indicates that the private key file permissions are too open, meaning that other users can read or write to the file. This is a security risk, as the private key should be accessible only by the owner of the file. To fix this error, the security administrator should use the chmod command to change the permissions of the private key file to 0400, which means that only the owner can read the file and no one else can read or write to it.
The chmod command takes a numeric argument that represents the permissions for the owner, group, and others in octal notation. Each digit corresponds to a set of permissions: read (4), write (2), and execute (1). The digits are added together to get the final permissions for each category. For example, 0400 means that the owner has read permission (4) and no other permissions (0), and the group and others have no permissions at all (0).
The other options are incorrect because they either do not change the permissions at all (D), or they give too much or too little permissions to the owner, group, or others (A, C).

Verified References:

- https://superuser.com/questions/215504/permissions-on-private-key-in-ssh-folder
- https://www.baeldung.com/linux/ssh-key-permissions

NEW QUESTION 149
A company uses several AWS CloudFormation stacks to handle the deployment of a suite of applications. The leader of the company's application development team notices that the stack deployments fail with permission errors when some team members try to deploy the stacks. However, other team members can deploy the stacks successfully.
The team members access the account by assuming a role that has a specific set of permissions that are necessary for the job responsibilities of the team members. All team members have permissions to perform operations on the stacks.
Which combination of steps will ensure consistent deployment of the stacks MOST securely? (Select THREE.)

A. Create a service role that has a composite principal that contains each service that needs the necessary permission
B. Configure the role to allow the sts:AssumeRole action.
C. Create a service role that has cloudformation.amazonaws.com as the service principa
D. Configure the role to allow the sts:AssumeRole action.
E. For each required set of permissions, add a separate policy to the role to allow those permission
F. Add the ARN of each CloudFormation stack in the resource field of each policy.
G. For each required set of permissions, add a separate policy to the role to allow those permission
H. Add the ARN of each service that needs the per-missions in the resource field of the corresponding policy.
I. Update each stack to use the service role.
J. Add a policy to each member role to allow the iam:PassRole actio
K. Set the policy's resource field to the ARN of the service role.

Answer: BDF

NEW QUESTION 154
A company's policy requires that all API keys be encrypted and stored separately from source code in a centralized security account. This security account is managed by the company's security team However, an audit revealed that an API key is steed with the source code of an IAM Lambda function m an IAM CodeCommit repository in the DevOps account
How should the security learn securely store the API key?

A. Create a CodeCommit repository in the security account using IAM Key Management Service (IAMKMS) tor encryption Require the development team to migrate the Lambda source code to this repository
B. Store the API key in an Amazon S3 bucket in the security account using server-side encryption with Amazon S3 managed encryption keys (SSE-S3) to encrypt the key Create a resigned URL tor the S3 ke
C. and specify the URL m a Lambda environmental variable in the IAM CloudFormation template Update the Lambda function code to retrieve the key using the URL and call the API
D. Create a secret in IAM Secrets Manager in the security account to store the API key using IAM Key Management Service (IAM KMS) tor encryption Grant access to the IAM role used by the Lambda function so that the function can retrieve the key from Secrets Manager and call the API
E. Create an encrypted environment variable for the Lambda function to store the API key using IAM Key Management Service (IAM KMS) tor encryption Grant access to the IAM role used by the Lambda function so that the function can decrypt the key at runtime

Answer: C

Explanation:
To securely store the API key, the security team should do the following:

- Create a secret in AWS Secrets Manager in the security account to store the API key using AWS Key Management Service (AWS KMS) for encryption. This allows the security team to encrypt and manage the API key centrally, and to configure automatic rotation schedules for it.
- Grant access to the IAM role used by the Lambda function so that the function can retrieve the key from Secrets Manager and call the API. This allows the security team to avoid storing the API key with the source code, and to use IAM policies to control access to the secret.

NEW QUESTION 158
.....

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

AWS-Certified-Security-Specialty Practice Exam Features:

* AWS-Certified-Security-Specialty Questions and Answers Updated Frequently

* AWS-Certified-Security-Specialty Practice Questions Verified by Expert Senior Certified Staff

* AWS-Certified-Security-Specialty Most Realistic Questions that Guarantee you a Pass on Your FirstTry

* AWS-Certified-Security-Specialty Practice Test Questions in Multiple Choice Formats and Updatesfor 1 Year