Amazon-Web-Services

Exam Questions SCS-C01

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 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 3
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 4
A security engineer recently rotated all IAM access keys in an AWS account. The security engineer then configured AWS Config and enabled the following AWS Config managed rules; mfa-enabled-for-iam-console-access, iam-user-mfa-enabled, access-key-rotated, and iam-user-unused-credentials-check.
The security engineer notices that all resources are displaying as noncompliant after the IAM GenerateCredentialReport API operation is invoked.
What could be the reason for the noncompliant status?

A. The IAM credential report was generated within the past 4 hours.
B. The security engineer does not have the GenerateCredentialReport permission.
C. The security engineer does not have the GetCredentialReport permission.
D. The AWS Config rules have a MaximumExecutionFrequency value of 24 hours.

Answer: D

Explanation:
The correct answer is D. The AWS Config rules have a MaximumExecutionFrequency value of 24 hours. According to the AWS documentation1, the MaximumExecutionFrequency parameter specifies the maximum frequency with which AWS Config runs evaluations for a rule. For AWS Config managed rules, this value can be one of the following:

- One_Hour
- Three_Hours
- Six_Hours
- Twelve_Hours
- TwentyFour_Hours

If the rule is triggered by configuration changes, it will still run evaluations when AWS Config delivers the configuration snapshot. However, if the rule is triggered periodically, it will not run evaluations more often than the specified frequency.
In this case, the security engineer enabled four AWS Config managed rules that are triggered periodically. Therefore, these rules will only run evaluations every 24 hours, regardless of when the IAM credential report is generated. This means that the resources will display as noncompliant until the next evaluation cycle, which could take up to 24 hours after the IAM access keys are rotated.
The other options are incorrect because:

- A. The IAM credential report can be generated at any time, but it will not affect the compliance status of the resources until the next evaluation cycle of the AWS Config rules.
- B. The security engineer was able to invoke the IAM GenerateCredentialReport API operation, which means they have the GenerateCredentialReport permission. This permission is required to generate a credential report that lists all IAM users in an AWS account and their credential status2.
- C. The security engineer does not need the GetCredentialReport permission to enable or evaluate AWS Config rules. This permission is required to retrieve a credential report that was previously generated by using the GenerateCredentialReport operation2.

References:
1: AWS::Config::ConfigRule - AWS CloudFormation 2: IAM: Generate and retrieve IAM credential reports

NEW QUESTION 5
A company uses AWS Organizations to manage several AWs accounts. The company processes a large volume of sensitive data. The company uses a serverless approach to microservices. The company stores all the data in either Amazon S3 or Amazon DynamoDB. The company reads the data by using either AWS lambda functions or container-based services that the company hosts on Amazon Elastic Kubernetes Service (Amazon EKS) on AWS Fargate.
The company must implement a solution to encrypt all the data at rest and enforce least privilege data access controls. The company creates an AWS Key Management Service (AWS KMS) customer managed key.
What should the company do next to meet these requirements?

A. Create a key policy that allows the kms:Decrypt action only for Amazon S3 and DynamoD
B. Create an SCP that denies the creation of S3 buckets and DynamoDB tables that are not encrypted with the key.
C. Create an 1AM policy that denies the kms:Decrypt action for the ke
D. Create a Lambda function than runs on a schedule to attach the policy to any new role
E. Create an AWS Config rule to send alerts for resources that are not encrypted with the key.
F. Create a key policy that allows the kms:Decrypt action only for Amazon S3, DynamoDB, Lambda, and Amazon EK
G. Create an SCP that denies the creation of S3 buckets and DynamoDB tables that are not encrypted with the key.
H. Create a key policy that allows the kms:Decrypt action only for Amazon S3, DynamoDB, Lambda, and Amazon EK
I. Create an AWS Config rule to send alerts for resources that are not encrypted with the key.

Answer: B

NEW QUESTION 6
A Security Engineer receives alerts that an Amazon EC2 instance on a public subnet is under an SFTP brute force attack from a specific IP address, which is a known malicious bot. What should the Security Engineer do to block the malicious bot?

A. Add a deny rule to the public VPC security group to block the malicious IP
B. Add the malicious IP to IAM WAF backhsted IPs
C. Configure Linux iptables or Windows Firewall to block any traffic from the malicious IP
D. Modify the hosted zone in Amazon Route 53 and create a DNS sinkhole for the malicious IP

Answer: D

Explanation:
what the Security Engineer should do to block the malicious bot. SFTP is a protocol that allows secure file transfer over SSH. EC2 is a service that provides virtual servers in the cloud. A public subnet is a subnet that has a route to an internet gateway, which allows it to communicate with the internet. A brute force attack is a type of attack that tries to guess passwords or keys by trying many possible combinations. A malicious bot is a software program that performs automated tasks for malicious purposes. Route 53 is a service that provides DNS resolution and domain name registration. A DNS sinkhole is a technique that redirects malicious or unwanted traffic to a different destination, such as a black hole server or a honeypot. By modifying the hosted zone in Route 53 and creating a DNS sinkhole for the malicious IP, the Security Engineer can block the malicious bot from reaching the EC2 instance on the public subnet. The other options are either ineffective or inappropriate for blocking the malicious bot.

NEW QUESTION 7

A company uses AWS Organizations to manage a small number of AWS accounts. However, the company plans to add 1 000 more accounts soon. The company allows only a centralized security team to create IAM roles for all AWS accounts and teams. Application teams submit requests for IAM roles to the security team. The security team has a backlog of IAM role requests and cannot review and provision the IAM roles quickly.

The security team must create a process that will allow application teams to provision their own IAM roles. The process must also limit the scope of IAM roles and prevent privilege escalation.

Which solution will meet these requirements with the LEAST operational overhead?

A. Create an IAM group for each application tea
B. Associate policies with each IAM grou
C. Provision IAM users for each application team membe
D. Add the new IAM users to the appropriate IAM group by using role-based access control (RBAC).
E. Delegate application team leads to provision IAM rotes for each tea
F. Conduct a quarterly review of the IAM rotes the team leads have provisione
G. Ensure that the application team leads have the appropriate training to review IAM roles.
H. Put each AWS account in its own O
I. Add an SCP to each OU to grant access to only the AWS services that the teams plan to us
J. Include conditions tn the AWS account of each team.
K. Create an SCP and a permissions boundary for IAM role
L. Add the SCP to the root OU so that only roles that have the permissions boundary attached can create any new IAM roles.

Answer: D

Explanation:

To create a process that will allow application teams to provision their own IAM roles, while limiting the scope of IAM roles and preventing privilege escalation, the following steps are required:

- Create a service control policy (SCP) that defines the maximum permissions that can be granted to any IAM role in the organization. An SCP is a type of policy that you can use with AWS Organizations to manage permissions for all accounts in your organization. SCPs restrict permissions for entities in member accounts, including each AWS account root user, IAM users, and roles. For more information, see Service control policies overview.
- Create a permissions boundary for IAM roles that matches the SCP. A permissions boundary is an advanced feature for using a managed policy to set the maximum permissions that an identity-based policy can grant to an IAM entity. A permissions boundary allows an entity to perform only the actions that are allowed by both its identity-based policies and its permissions boundaries. For more information, see Permissions boundaries for IAM entities.
- Add the SCP to the root organizational unit (OU) so that it applies to all accounts in the organization.

This will ensure that no IAM role can exceed the permissions defined by the SCP, regardless of how it is created or modified.

- Instruct the application teams to attach the permissions boundary to any IAM role they create. This will prevent them from creating IAM roles that can escalate their own privileges or access resources they are not authorized to access.

This solution will meet the requirements with the least operational overhead, as it leverages AWS Organizations and IAM features to delegate and limit IAM role creation without requiring manual reviews or approvals.

The other options are incorrect because they either do not allow application teams to provision their own IAM roles (A), do not limit the scope of IAM roles or prevent privilege escalation (B), or do not take advantage of managed services whenever possible ©.

Verified References:

- https://docs.aws.amazon.com/IAM/latest/UserGuide/access_policies_boundaries.html

NEW QUESTION 8

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 bucket
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 9

A security engineer needs to implement a write-once-read-many (WORM) model for data that a company will store in Amazon S3 buckets. The company uses the S3 Standard storage class for all of its S3 buckets. The security engineer must en-sure that objects cannot be overwritten or deleted by any user, including the AWS account root user.

Which solution will meet these requirements?

A. Create new S3 buckets with S3 Object Lock enabled in compliance mod
B. Place objects in the S3 buckets.
C. Use S3 Glacier Vault Lock to attach a Vault Lock policy to new S3 bucket
D. Wait 24 hours to complete the Vault Lock proces
E. Place objects in the S3 buckets.
F. Create new S3 buckets with S3 Object Lock enabled in governance mod
G. Place objects in the S3 buckets.
H. Create new S3 buckets with S3 Object Lock enabled in governance mod
I. Add a legal hold to the S3 bucket
J. Place objects in the S3 buckets.

Answer: A

NEW QUESTION 10

A security engineer is designing an IAM policy to protect AWS API operations. The policy must enforce multi-factor authentication (MFA) for IAM users to access certain services in the AWS production account. Each session must remain valid for only 2 hours. The current version of the IAM policy is as follows:

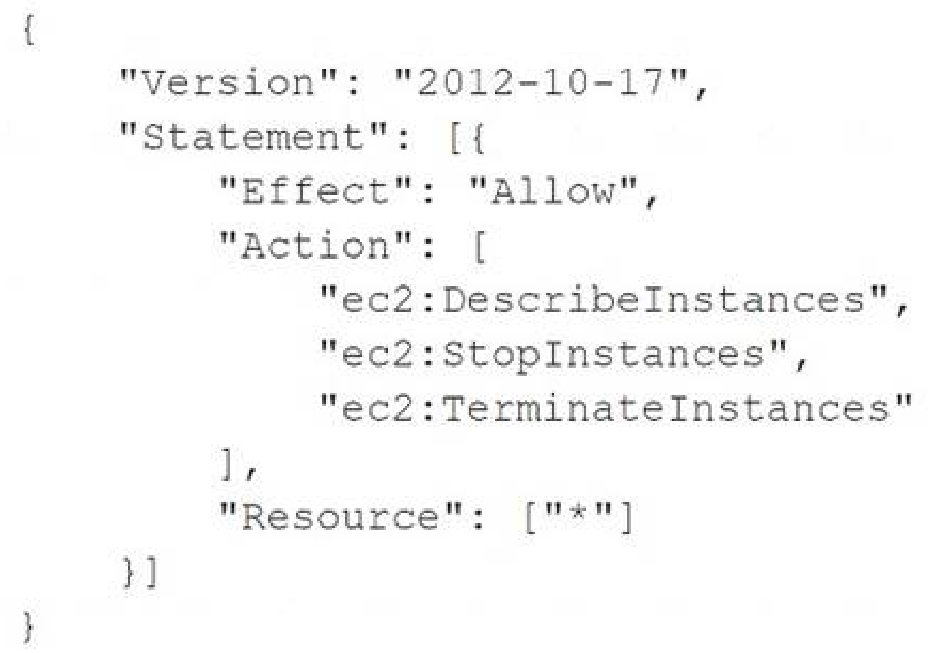

```
{
     "Version": "2012-10-17",
     "Statement": [{
         "Effect": "Allow",
         "Action": [
             "ec2:DescribeInstances",
             "ec2:StopInstances",
             "ec2:TerminateInstances"
         ],
         "Resource": ["*"]
     }]
}
```

Which combination of conditions must the security engineer add to the IAM policy to meet these requirements? (Select TWO.)

A. "Bool " : " aws : Multi FactorAuthPresent": "true" }
B. "B001 " : " aws : MultiFactorAuthPresent": "false" }
C. "NumericLessThan" : { " aws : Multi FactorAuthAge" : "7200"}
D. "NumericGreaterThan" : { " aws : MultiFactorAuthAge " : "7200"
E. "NumericLessThan" : { "MaxSessionDuration " : "7200"}

Answer: AC

Explanation:

The correct combination of conditions to add to the IAM policy is A and C. These conditions will ensure that IAM users must use MFA to access certain services in the AWS production account, and that each session will expire after 2 hours.

- Option A: “Bool” : { “aws:MultiFactorAuthPresent” : “true” } is a valid condition that checks if the principal (the IAM user) has authenticated with MFA before making the request. This condition will enforce MFA for the IAM users to access the specified services. This condition key is supported by all AWS services that support IAM policies1.
- Option B: “Bool” : { “aws:MultiFactorAuthPresent” : “false” } is the opposite of option A. This condition will allow access only if the principal has not authenticated with MFA, which is not the desired requirement. This condition key is supported by all AWS services that support IAM policies1.
- Option C: “NumericLessThan” : { “aws:MultiFactorAuthAge” : “7200” } is a valid condition that checks if the time since the principal authenticated with MFA is less than 7200 seconds (2 hours). This condition will enforce the session duration limit for the IAM users. This condition key is supported by all AWS services that

support IAM policies1.

- Option D: “NumericGreaterThan” : { “aws:MultiFactorAuthAge” : “7200” } is the opposite of option C. This condition will allow access only if the time since the principal authenticated with MFA is more than 7200 seconds (2 hours), which is not the desired requirement. This condition key is supported by all AWS services that support IAM policies1.
- Option E: “NumericLessThan” : { “MaxSessionDuration” : “7200” } is not a valid condition key.

MaxSessionDuration is a property of an IAM role, not a condition key. It specifies the maximum session duration (in seconds) for the role, which can be between 3600 and 43200 seconds (1 to 12 hours). This property can be set when creating or modifying a role, but it cannot be used as a condition in a policy2.

NEW QUESTION 10
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 15
A security engineer needs to build a solution to turn IAM CloudTrail back on in multiple IAM Regions in case it is ever turned off.
What is the MOST efficient way to implement this solution?

A. Use IAM Config with a managed rule to trigger the IAM-EnableCloudTrail remediation.
B. Create an Amazon EventBridge (Amazon CloudWatch Events) event with a cloudtrail.amazonIAM.com event source and a StartLogging event name to trigger an IAM Lambda function to call the StartLogging API.
C. Create an Amazon CloudWatch alarm with a cloudtrail.amazonIAM.com event source and a StopLogging event name to trigger an IAM Lambda function to call the StartLogging API.
D. Monitor IAM Trusted Advisor to ensure CloudTrail logging is enabled.

Answer: B

NEW QUESTION 16
A company recently had a security audit in which the auditors identified multiple potential threats. These potential threats can cause usage pattern changes such as DNS access peak, abnormal instance traffic, abnormal network interface traffic, and unusual Amazon S3 API calls. The threats can come from different sources and can occur at any time. The company needs to implement a solution to continuously monitor its system and identify all these incoming threats in near-real time.
Which solution will meet these requirements?

A. Enable AWS CloudTrail logs, VPC flow logs, and DNS log
B. Use Amazon CloudWatch Logs to manage these logs from a centralized account.
C. Enable AWS CloudTrail logs, VPC flow logs, and DNS log
D. Use Amazon Macie to monitor these logs from a centralized account.
E. Enable Amazon GuardDuty from a centralized accoun
F. Use GuardDuty to manage AWS CloudTrail logs, VPC flow logs, and DNS logs.
G. Enable Amazon Inspector from a centralized accoun
H. Use Amazon Inspector to manage AWS CloudTrail logs, VPC flow logs, and DNS logs.

Answer: C

Explanation:
Q: Which data sources does GuardDuty analyze? GuardDuty analyzes CloudTrail management event logs, CloudTrail S3 data event logs, VPC Flow Logs, DNS query logs, and Amazon EKS audit logs. GuardDuty can also scan EBS volume data for possible malware when GuardDuty Malware Protection is enabled and identifies suspicious behavior indicative of malicious software in EC2 instance or container workloads. The service is optimized to consume large data volumes for near real-time processing of security detections. GuardDuty gives you access to built-in detection techniques developed and optimized for the cloud, which are maintained and continuously improved upon by GuardDuty engineering.

NEW QUESTION 18
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 19
A company uses Amazon GuardDuty. The company's security team wants all High severity findings to automatically generate a ticket in a third-party ticketing system through email integration.
Which solution will meet this requirement?

A. Create a verified identity for the third-party ticketing email system in Amazon Simple Email Service (Amazon SES). Create an Amazon EventBridge rule that includes an event pattern that matches High severity GuardDuty finding
B. Specify the SES identity as the target for the EventBridge rule.
C. Create an Amazon Simple Notification Service (Amazon SNS) topi
D. Subscribe the third-party ticketing email system to the SNS topi
E. Create an Amazon EventBridge rule that includes an event pattern that matches High severity GuardDuty finding
F. Specify the SNS topic as the target for the EventBridge rule.
G. Use the GuardDuty CreateFilter API operation to build a filter in GuardDuty to monitor for High severity finding
H. Export the results of the filter to an Amazon Simple Notification Service (Amazon SNS) topi
I. Subscribe the third-party ticketing email system to the SNS topic.
J. Use the GuardDuty CreateFilter API operation to build a filter in GuardDuty to monitor for High severity finding
K. Create an Amazon Simple Notification Service (Amazon SNS) topi
L. Subscribe the third-party ticketing email system to the SNS topi
M. Create an Amazon EventBridge rule that includes an event pattern that matches GuardDuty findings that are selected by the filte
N. Specify the SNS topic as the target for the EventBridge rule.

Answer: B

Explanation:
The correct answer is B. Create an Amazon Simple Notification Service (Amazon SNS) topic. Subscribe the third-party ticketing email system to the SNS topic. Create an Amazon EventBridge rule that includes an event pattern that matches High severity GuardDuty findings. Specify the SNS topic as the target for the Event-Bridge rule.
According to the AWS documentation1, you can use Amazon EventBridge to create rules that match events from GuardDuty and route them to targets such as Amazon SNS topics. You can use event patterns to filter events based on criteria such as severity, type, or resource. For example, you can create a rule that matches only High severity findings and sends them to an SNS topic that is subscribed by a third-party ticketing email system. This way, you can automate the creation of tickets for High severity findings and notify the security team.

NEW QUESTION 21
A company has a single AWS account and uses an Amazon EC2 instance to test application code. The company recently discovered that the instance was compromised. The instance was serving up malware. The analysis of the instance showed that the instance was compromised 35 days ago.
A security engineer must implement a continuous monitoring solution that automatically notifies the company’s security team about compromised instances through an email distribution list for high severity findings. The security engineer must implement the solution as soon as possible.
Which combination of steps should the security engineer take to meet these requirements? (Choose three.)

A. Enable AWS Security Hub in the AWS account.
B. Enable Amazon GuardDuty in the AWS account.
C. Create an Amazon Simple Notification Service (Amazon SNS) topi
D. Subscribe the security team’s email distribution list to the topic.
E. Create an Amazon Simple Queue Service (Amazon SQS) queu
F. Subscribe the security team’s email distribution list to the queue.
G. Create an Amazon EventBridge (Amazon CloudWatch Events) rule for GuardDuty findings of high severit
H. Configure the rule to publish a message to the topic.
I. Create an Amazon EventBridge (Amazon CloudWatch Events) rule for Security Hub findings of high severit
J. Configure the rule to publish a message to the queue.

Answer: BCE

NEW QUESTION 26
A company wants to remove all SSH keys permanently from a specific subset of its Amazon Linux 2 Amazon EC2 instances that are using the same 1AM instance profile However three individuals who have IAM user accounts will need to access these instances by using an SSH session to perform critical duties
How can a security engineer provide the access to meet these requirements'?

A. Assign an 1AM policy to the instance profile to allow the EC2 instances to be managed by AWS Systems Manager Provide the 1AM user accounts with permission to use Systems Manager Remove the SSH keys from the EC2 instances Use Systems Manager Inventory to select the EC2 instance and connect
B. Assign an 1AM policy to the 1AM user accounts to provide permission to use AWS Systems Manager Run Command Remove the SSH keys from the EC2 instances Use Run Command to open an SSH connection to the EC2 instance
C. Assign an 1AM policy to the instance profile to allow the EC2 instances to be managed by AWS Systems Manager Provide the 1AM user accounts with permission to use Systems Manager Remove the SSH keys from the EC2 instances Use Systems Manager Session Manager to select the EC2 instance and connect
D. Assign an 1AM policy to the 1AM user accounts to provide permission to use the EC2 service in the AWS Management Console Remove the SSH keys from the EC2 instances Connect to the EC2 instance as the ec2-user through the AWS Management Console's EC2 SSH client method

Answer: C

Explanation:
To provide access to the three individuals who have IAM user accounts to access the Amazon Linux 2 Amazon EC2 instances that are using the same IAM instance profile, the most appropriate solution would be to assign an IAM policy to the instance profile to allow the EC2 instances to be managed by AWS Systems Manager, provide the IAM user accounts with permission to use Systems Manager, remove the SSH keys from the EC2 instances, and use Systems Manager Session Manager to select the EC2 instance and connect.
References: : AWS Systems Manager Session Manager - AWS Systems Manager : AWS Systems Manage AWS Management Console : AWS Identity and Access Management - AWS Management Console : Am Elastic Compute Cloud - Amazon Web Services : Amazon Linux 2 - Amazon Web Services : AWS Syst Manager - AWS Management Console : AWS Systems Manager - AWS Management Console : AWS Systems Manager - AWS Management Console

NEW QUESTION 27
A company has recently recovered from a security incident that required the restoration of Amazon EC2 instances from snapshots.
After performing a gap analysis of its disaster recovery procedures and backup strategies, the company is concerned that, next time, it will not be able to recover the EC2 instances if the AWS account was compromised and Amazon EBS snapshots were deleted.
All EBS snapshots are encrypted using an AWS KMS CMK. Which solution would solve this problem?

A. Create a new Amazon S3 bucke
B. Use EBS lifecycle policies to move EBS snapshots to the new S3 bucke
C. Move snapshots to Amazon S3 Glacier using lifecycle policies, and apply Glacier Vault Lock policies to prevent deletion.
D. Use AWS Systems Manager to distribute a configuration that performs local backups of all attached disks to Amazon S3.
E. Create a new AWS account with limited privilege
F. Allow the new account to access the AWS KMS key used to encrypt the EBS snapshots, and copy the encrypted snapshots to the new account on a recurring basis.
G. Use AWS Backup to copy EBS snapshots to Amazon S3.

Answer: C

Explanation:
This answer is correct because creating a new AWS account with limited privileges would provide an isolated and secure backup destination for the EBS snapshots. Allowing the new account to access the AWS KMS key used to encrypt the EBS snapshots would enable cross-account snapshot sharing without requiring re-encryption. Copying the encrypted snapshots to the new account on a recurring basis would ensure that the backups are up-to-date and consistent.

NEW QUESTION 28
Your CTO thinks your IAM account was hacked. What is the only way to know for certain if there was unauthorized access and what they did, assuming your hackers are very sophisticated IAM engineers and doing everything they can to cover their tracks?
Please select:

A. Use CloudTrail Log File Integrity Validation.
B. Use IAM Config SNS Subscriptions and process events in real time.
C. Use CloudTrail backed up to IAM S3 and Glacier.
D. Use IAM Config Timeline forensics.

Answer: A

Explanation:
The IAM Documentation mentions the following
To determine whether a log file was modified, deleted, or unchanged after CloudTrail delivered it you can use CloudTrail log file integrity validation. This feature is built using industry standard algorithms: SHA-256 for hashing and SHA-256 with RSA for digital signing. This makes it computationally infeasible to modify, delete or forge CloudTrail log files without detection. You can use the IAM CLI to validate the files in the location where CloudTrail delivered them
Validated log files are invaluable in security and forensic investigations. For example, a validated log file enables you to assert positively that the log file itself has not changed, or that particular user credentials performed specific API activity. The CloudTrail log file integrity validation process also lets you know if a log file has been deleted or changed, or assert positively that no log files were delivered to your account during a given period of time.
Options B.C and D is invalid because you need to check for log File Integrity Validation for cloudtrail logs For more information on Cloudtrail log file validation, please visit the below URL: http://docs.IAM.amazon.com/IAMcloudtrail/latest/userguide/cloudtrail-log-file-validation-intro.html
The correct answer is: Use CloudTrail Log File Integrity Validation. omit your Feedback/Queries to our Expert

NEW QUESTION 33
A company uses an external identity provider to allow federation into different IAM accounts. A security engineer for the company needs to identify the federated user that terminated a production Amazon EC2 instance a week ago.
What is the FASTEST way for the security engineer to identify the federated user?

A. Review the IAM CloudTrail event history logs in an Amazon S3 bucket and look for the TerminateInstances event to identify the federated user from the role session name.
B. Filter the IAM CloudTrail event history for the TerminateInstances event and identify the assumed IAM rol
C. Review the AssumeRoleWithSAML event call in CloudTrail to identify the corresponding username.
D. Search the IAM CloudTrail logs for the TerminateInstances event and note the event tim
E. Review the IAM Access Advisor tab for all federated role
F. The last accessed time should match the time when the instance was terminated.
G. Use Amazon Athena to run a SQL query on the IAM CloudTrail logs stored in an Amazon S3 bucket and filter on the TerminateInstances even
H. Identify the corresponding role and run another query to filter the AssumeRoleWithWebIdentity event for the user name.

Answer: B

Explanation:
The fastest way to identify the federated user who terminated a production Amazon EC2 instance is to filter the IAM CloudTrail event history for the TerminateInstances event and identify the assumed IAM role. Then, review the AssumeRoleWithSAML event call in CloudTrail to identify the corresponding username. This method does not require any additional tools or queries, and it directly links the IAM role with the federated user.
Option A is incorrect because the role session name may not be the same as the federated user name, and it may not be unique or descriptive enough to identify the user.

Option C is incorrect because the IAM Access Advisor tab only shows when a role was last accessed, not by whom or for what purpose. It also does not show the specific time of access, only the date.
Option D is incorrect because using Amazon Athena to run SQL queries on the IAM CloudTrail logs is not the fastest way to identify the federated user, as it requires creating a table schema and running multiple queries. It also assumes that the federation is done using web identity providers, not SAML providers, as indicated by the AssumeRoleWithWebIdentity event.
References:

- AWS Identity and Access Management
- Logging AWS STS API Calls with AWS CloudTrail
- [Using Amazon Athena to Query S3 Data for CloudTrail Analysis]

NEW QUESTION 36
A startup company is using a single AWS account that has resources in a single AWS Region. A security engineer configures an AWS Cloud Trail trail in the same Region to deliver log files to an Amazon S3 bucket by using the AWS CLI.
Because of expansion, the company adds resources in multiple Regions. The secu-rity engineer notices that the logs from the new Regions are not reaching the S3 bucket.
What should the security engineer do to fix this issue with the LEAST amount of operational overhead?

A. Create a new CloudTrail trai
B. Select the new Regions where the company added resources.
C. Change the S3 bucket to receive notifications to track all actions from all Regions.
D. Create a new CloudTrail trail that applies to all Regions.
E. Change the existing CloudTrail trail so that it applies to all Regions.

Answer: D

Explanation:
The correct answer is D. Change the existing CloudTrail trail so that it applies to all Regions.
According to the AWS documentation1, you can configure CloudTrail to deliver log files from multiple Regions to a single S3 bucket for a single account. To change an existing single-Region trail to log in all Regions, you must use the AWS CLI and add the --is-multi-region-trail option to the update-trail command2. This will ensure that you log global service events and capture all management event activity in your account.
Option A is incorrect because creating a new CloudTrail trail for each Region will incur additional costs and increase operational overhead. Option B is incorrect because changing the S3 bucket to receive notifications will not affect the delivery of log files from other Regions. Option C is incorrect because creating a new CloudTrail trail that applies to all Regions will result in duplicate log files for the original Region and also incur additional costs.

NEW QUESTION 38
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 43

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:

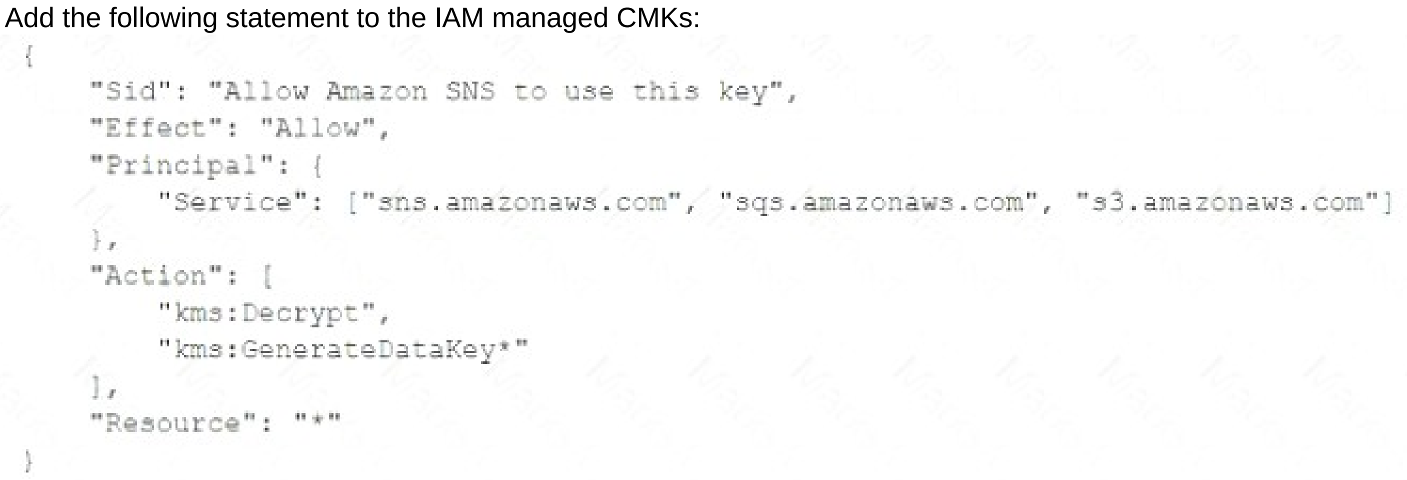

```
{
    "Sid": "Allow Amazon SNS to use this key",
    "Effect": "Allow",
    "Principal": {
        "Service": ["sns.amazonaws.com", "sqs.amazonaws.com", "s3.amazonaws.com"]
    },
    "Action": [
        "kms:Decrypt",
        "kms:GenerateDataKey*"
    ],
    "Resource": "*"
}
```

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

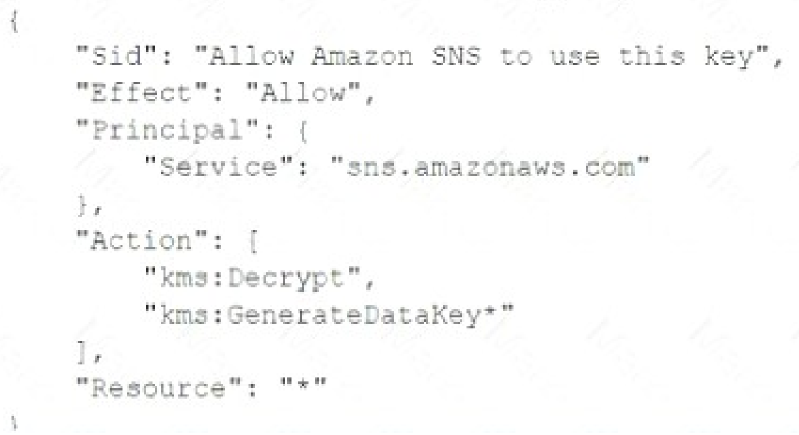

```
{
    "Sid": "Allow Amazon SNS to use this key",
    "Effect": "Allow",
    "Principal": {
        "Service": "sns.amazonaws.com"
    },
    "Action": [
        "kms:Decrypt",
        "kms:GenerateDataKey*"
    ],
    "Resource": "*"
}
```

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

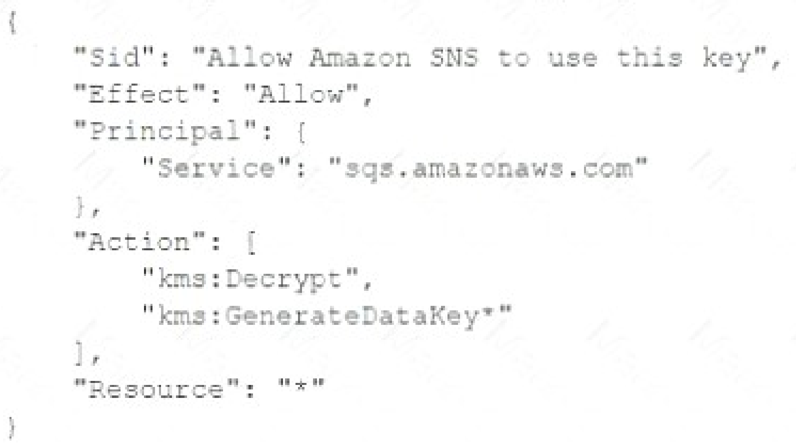

```
{
    "Sid": "Allow Amazon SNS to use this key",
    "Effect": "Allow",
    "Principal": {
        "Service": "sqs.amazonaws.com"
    },
    "Action": [
        "kms:Decrypt",
        "kms:GenerateDataKey*"
    ],
    "Resource": "*"
}
```

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

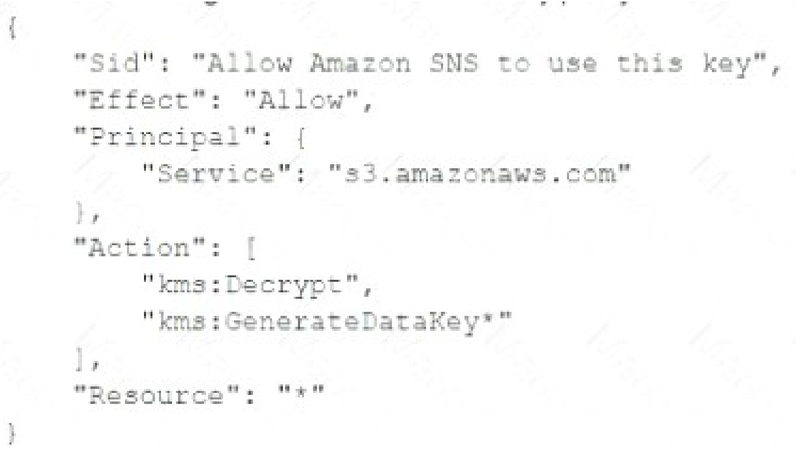

```
{
    "Sid": "Allow Amazon SNS to use this key",
    "Effect": "Allow",
    "Principal": {
        "Service": "s3.amazonaws.com"
    },
    "Action": [
        "kms:Decrypt",
        "kms:GenerateDataKey*"
    ],
    "Resource": "*"
}
```

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

Answer: D

NEW QUESTION 47
A Security Engineer is building a Java application that is running on Amazon EC2. The application communicates with an Amazon RDS instance and authenticates with a user name and password.
Which combination of steps can the Engineer take to protect the credentials and minimize downtime when the credentials are rotated? (Choose two.)

A. Have a Database Administrator encrypt the credentials and store the ciphertext in Amazon S3. Grant permission to the instance role associated with the EC2

instance to read the object and decrypt the ciphertext.
B. Configure a scheduled job that updates the credential in AWS Systems Manager Parameter Store and notifies the Engineer that the application needs to be restarted.
C. Configure automatic rotation of credentials in AWS Secrets Manager.
D. Store the credential in an encrypted string parameter in AWS Systems Manager Parameter Stor
E. Grant permission to the instance role associated with the EC2 instance to access the parameter and the AWS KMS key that is used to encrypt it.
F. Configure the Java application to catch a connection failure and make a call to AWS Secrets Manager to retrieve updated credentials when the password is rotate
G. Grant permission to the instance role associated with the EC2 instance to access Secrets Manager.

Answer: CE

Explanation:
AWS Secrets Manager is a service that helps you manage, retrieve, and rotate secrets such as database credentials, API keys, and other sensitive information. By configuring automatic rotation of credentials in AWS Secrets Manager, you can ensure that your secrets are changed regularly and securely, without requiring manual intervention or application downtime. You can also specify the rotation frequency and the rotation function that performs the logic of changing the credentials on the database and updating the secret in Secrets Manager1.
* E. Configure the Java application to catch a connection failure and make a call to AWS Secrets Manager to retrieve updated credentials when the password is rotated. Grant permission to the instance role associated with the EC2 instance to access Secrets Manager.
By configuring the Java application to catch a connection failure and make a call to AWS Secrets Manager to retrieve updated credentials, you can avoid hard-coding the credentials in your application code or configuration files. This way, your application can dynamically obtain the latest credentials from Secrets Manager whenever the password is rotated, without needing to restart or redeploy the application. To enable this, you need to grant permission to the instance role associated with the EC2 instance to access Secrets Manager using IAM policies2. You can also use the AWS SDK for Java to integrate your application with Secrets Manager3.

NEW QUESTION 51
A company deploys a distributed web application on a fleet of Amazon EC2 instances. The fleet is behind an Application Load Balancer (ALB) that will be configured to terminate the TLS connection. All TLS traffic to the ALB must stay secure, even if the certificate private key is compromised.
How can a security engineer meet this requirement?

A. Create an HTTPS listener that uses a certificate that is managed by IAM Certificate Manager (ACM).
B. Create an HTTPS listener that uses a security policy that uses a cipher suite with perfect toward secrecy (PFS).
C. Create an HTTPS listener that uses the Server Order Preference security feature.
D. Create a TCP listener that uses a custom security policy that allows only cipher suites with perfect forward secrecy (PFS).

Answer: A

NEW QUESTION 52
A team is using AWS Secrets Manager to store an application database password. Only a limited number of IAM principals within the account can have access to the secret. The principals who require access to the secret change frequently. A security engineer must create a solution that maximizes flexibility and scalability.
Which solution will meet these requirements?

A. Use a role-based approach by creating an IAM role with an inline permissions policy that allows access to the secre
B. Update the IAM principals in the role trust policy as required.
C. Deploy a VPC endpoint for Secrets Manage
D. Create and attach an endpoint policy that specifies the IAM principals that are allowed to access the secre
E. Update the list of IAM principals as required.
F. Use a tag-based approach by attaching a resource policy to the secre
G. Apply tags to the secret and the IAM principal
H. Use the aws:PrincipalTag and aws:ResourceTag IAM condition keys to control access.
I. Use a deny-by-default approach by using IAM policies to deny access to the secret explicitl
J. Attach the policies to an IAM grou
K. Add all IAM principals to the IAM grou
L. Remove principals from the group when they need acces
M. Add the principals to the group again when access is no longer allowed.

Answer: C

NEW QUESTION 57
A company is using AWS Organizations to manage multiple accounts. The company needs to allow an IAM user to use a role to access resources that are in another organization's AWS account.
Which combination of steps must the company perform to meet this requirement? (Select TWO.)

A. Create an identity policy that allows the sts: AssumeRole action in the AWS account that contains the resource
B. Attach the identity policy to the IAM user.
C. Ensure that the sts: AssumeRole action is allowed by the SCPs of the organization that owns the resources that the IAM user needs to access.
D. Create a role in the AWS account that contains the resource
E. Create an entry in the role's trust policy that allows the IAM user to assume the rol
F. Attach the trust policy to the role.
G. Establish a trust relationship between the IAM user and the AWS account that contains the resources.
H. Create a role in the IAM user's AWS accoun
I. Create an identity policy that allows the sts: AssumeRole actio
J. Attach the identity policy to the role.

Answer: BC

Explanation:
To allow cross-account access to resources using IAM roles, the following steps are required:

- Create a role in the AWS account that contains the resources (the trusting account) and specify the AWS account that contains the IAM user (the trusted account) as a trusted entity in the role's trust policy. This allows users from the trusted account to assume the role and access resources in the trusting account.

- Ensure that the IAM user has permission to assume the role in their own AWS account. This can be done by creating an identity policy that allows the sts:AssumeRole action and attaching it to the IAM user or their group.
- Ensure that there are no service control policies (SCPs) in the organization that owns the resources that deny or restrict access to the sts:AssumeRole action or the role itself. SCPs are applied to all accounts in an organization and can override any permissions granted by IAM policies.

Verified References:

- https://repost.aws/knowledge-center/cross-account-access-iam
- https://docs.aws.amazon.com/organizations/latest/userguide/orgs_manage_accounts_access.html
- https://docs.aws.amazon.com/IAM/latest/UserGuide/tutorial_cross-account-with-roles.html

NEW QUESTION 61

A company needs to follow security best practices to deploy resources from an AWS CloudFormation template. The CloudFormation template must be able to configure sensitive database credentials.
The company already uses AWS Key Management Service (AWS KMS) and AWS Secrets Manager. Which solution will meet the requirements?

A. Use a dynamic reference in the CloudFormation template to reference the database credentials in Secrets Manager.
B. Use a parameter in the CloudFormation template to reference the database credential
C. Encrypt the CloudFormation template by using AWS KMS.
D. Use a SecureString parameter in the CloudFormation template to reference the database credentials in Secrets Manager.
E. Use a SecureString parameter in the CloudFormation template to reference an encrypted value in AWS KMS

Answer: A

Explanation:

- Option A: This option meets the requirements of following security best practices and configuring sensitive database credentials in the CloudFormation template. A dynamic reference is a way to specify external values that are stored and managed in other services, such as Secrets Manager, in the stack templates1. When using a dynamic reference, CloudFormation retrieves the value of the specified reference when necessary during stack and change set operations1. Dynamic references can be used for certain resources that support them, such as AWS::RDS::DBInstance1. By using a dynamic reference to reference the database credentials in Secrets Manager, the company can leverage the existing integration between these services and avoid hardcoding the secret information in the template. Secrets Manager is a service that helps you protect secrets needed to access your applications, services, and IT resources2. Secrets Manager enables you to rotate, manage, and retrieve database credentials, API keys, and other secrets throughout their lifecycle2.

NEW QUESTION 62

A company has thousands of AWS Lambda functions. While reviewing the Lambda functions, a security engineer discovers that sensitive information is being stored in environment variables and is viewable as plaintext in the Lambda console. The values of the sensitive information are only a few characters long.
What is the MOST cost-effective way to address this security issue?

A. Set up IAM policies from the Lambda console to hide access to the environment variables.
B. Use AWS Step Functions to store the environment variable
C. Access the environment variables at runtim
D. Use IAM permissions to restrict access to the environment variables to only the Lambda functions that require access.
E. Store the environment variables in AWS Secrets Manager, and access them at runtim
F. Use IAM permissions to restrict access to the secrets to only the Lambda functions that require access.
G. Store the environment variables in AWS Systems Manager Parameter Store as secure string parameters, and access them at runtim
H. Use IAM permissions to restrict access to the parameters to only the Lambda functions that require access.

Answer: D

Explanation:

Storing sensitive information in environment variables is not a secure practice, as anyone who has access to the Lambda console or the Lambda function code can view them as plaintext. To address this security issue, the security engineer needs to use a service that can store and encrypt the environment variables, and access them at runtime using IAM permissions. The most cost-effective way to do this is to use AWS Systems Manager Parameter Store, which is a service that provides secure, hierarchical storage for configuration data management and secrets management. Parameter Store allows you to store values as standard parameters (plaintext) or secure string parameters (encrypted). Secure string parameters use a AWS Key Management Service (AWS KMS) customer master key (CMK) to encrypt the parameter value. To access the parameter value at runtime, the Lambda function needs to have IAM permissions to decrypt the parameter using the KMS CMK.
The other options are incorrect because:

- Option A is incorrect because setting up IAM policies from the Lambda console to hide access to the environment variables will not prevent someone who has access to the Lambda function code from viewing them as plaintext. IAM policies can only control who can perform actions on AWS resources, not what they can see in the code or the console.
- Option B is incorrect because using AWS Step Functions to store the environment variables is not a secure or cost-effective solution. AWS Step Functions is a service that lets you coordinate multiple AWS services into serverless workflows. Step Functions does not provide any encryption or secrets management capabilities, and it will incur additional charges for each state transition in the workflow. Moreover, storing environment variables in Step Functions will make them visible in the execution history of the workflow, which can be accessed by anyone who has permission to view the Step Functions console or API.
- Option C is incorrect because storing the environment variables in AWS Secrets Manager and accessing them at runtime is not a cost-effective solution. AWS Secrets Manager is a service that helps you protect secrets needed to access your applications, services, and IT resources. Secrets Manager enables you to rotate, manage, and retrieve secrets throughout their lifecycle. While Secrets Manager can securely store and encrypt environment variables using KMS CMKs, it will incur higher charges than Parameter Store for storing and retrieving secrets. Unless the security engineer needs the advanced features of Secrets Manager, such as automatic rotation of secrets or integration with other AWS services, Parameter Store is a cheaper and simpler option.

NEW QUESTION 65

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 67
A company needs to store multiple years of financial records. The company wants to use Amazon S3 to store copies of these documents. The company must implement a solution to prevent the documents from being edited, replaced, or deleted for 7 years after the documents are stored in Amazon S3. The solution must also encrypt the documents at rest.
A security engineer creates a new S3 bucket to store the documents. What should the security engineer do next to meet these requirements?

A. Configure S3 server-side encryptio
B. Create an S3 bucket policy that has an explicit deny rule for all users for s3:DeleteObject and s3:PutObject API call
C. Configure S3 Object Lock to use governance mode with a retention period of 7 years.
D. Configure S3 server-side encryptio
E. Configure S3 Versioning on the S3 bucke
F. Configure S3 ObjectLock to use compliance mode with a retention period of 7 years.
G. Configure S3 Versionin
H. Configure S3 Intelligent-Tiering on the S3 bucket to move the documents to S3 Glacier Deep Archive storag
I. Use S3 server-side encryption immediatel
J. Expire the objects after 7 years.
K. Set up S3 Event Notifications and use S3 server-side encryptio
L. Configure S3 Event Notifications to target an AWS Lambda function that will review any S3 API call to the S3 bucket and deny the s3:DeleteObject and s3:PutObject API call
M. Remove the S3 event notification after 7 years.

Answer: B

NEW QUESTION 69
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 74
A company has launched an Amazon EC2 instance with an Amazon Elastic Block Store (Amazon EBS) volume in the us-east-1 Region The volume is encrypted with an AWS Key Management Service (AWS KMS) customer managed key that the company's security team created The security team has created an 1AM key policy and has assigned the policy to the key The security team has also created an 1AM instance profile and has assigned the profile to the instance
The EC2 instance will not start and transitions from the pending state to the shutting-down state to the terminated state
Which combination of steps should a security engineer take to troubleshoot this issue? (Select TWO)

A. Verify that the KMS key policy specifies a deny statement that prevents access to the key by using the aws SourcelP condition key Check that the range includes the EC2 instance IP address that is associated with the EBS volume
B. Verify that the KMS key that is associated with the EBS volume is set to the Symmetric key type
C. Verify that the KMS key that is associated with the EBS volume is in the Enabled state
D. Verify that the EC2 role that is associated with the instance profile has the correct 1AM instance policy to launch an EC2 instance with the EBS volume
E. Verify that the key that is associated with the EBS volume has not expired and needs to be rotated

Answer: CD

Explanation:
To troubleshoot the issue of an EC2 instance failing to start and transitioning to a terminated state when it has an EBS volume encrypted with an AWS KMS

customer managed key, a security engineer should take the following steps:
* C. Verify that the KMS key that is associated with the EBS volume is in the Enabled state. If the key is not enabled, it will not function properly and could cause the EC2 instance to fail.
* D. Verify that the EC2 role that is associated with the instance profile has the correct IAM instance policy to launch an EC2 instance with the EBS volume. If the instance does not have the necessary permissions, it may not be able to mount the volume and could cause the instance to fail.
Therefore, options C and D are the correct answers.

NEW QUESTION 76
A company usesAWS Organizations to run workloads in multiple AWS accounts Currently the individual team members at the company access all Amazon EC2 instances remotely by using SSH or Remote Desktop Protocol (RDP) The company does not have any audit trails and security groups are occasionally open The company must secure access management and implement a centralized togging solution
Which solution will meet these requirements MOST securely?

A. Configure trusted access for AWS System Manager in Organizations Configure a bastion host from the management account Replace SSH and RDP by using Systems Manager Session Manager from the management account Configure Session Manager logging to Amazon CloudWatch Logs
B. Replace SSH and RDP with AWS Systems Manager Session Manager Install Systems Manager Agent (SSM Agent) on the instances Attach the
C. AmazonSSMManagedlnstanceCore role to the instances Configure session data streaming to Amazon CloudWatch Logs Create a separate logging account that has appropriate cross-account permissions to audit the log data
D. Install a bastion host in the management account Reconfigure all SSH and RDP to allow access only from the bastion host Install AWS Systems Manager Agent (SSM Agent) on the bastion host Attach the AmazonSSMManagedlnstanceCore role to the bastion host Configure session data streaming to Amazon CloudWatch Logs in a separate logging account to audit log data
E. Replace SSH and RDP with AWS Systems Manager State Manager Install Systems Manager Agent (SSM Agent) on the instances Attach theAmazonSSMManagedlnstanceCore role to the instances Configure session data streaming to AmazonCloudTrail Use CloudTrail Insights to analyze the trail data

Answer: C

Explanation:
To meet the requirements of securing access management and implementing a centralized logging solution, the most secure solution would be to:
- Install a bastion host in the management account.
- Reconfigure all SSH and RDP to allow access only from the bastion host.
- Install AWS Systems Manager Agent (SSM Agent) on the bastion host.
- Attach the AmazonSSMManagedlnstanceCore role to the bastion host.
- Configure session data streaming to Amazon CloudWatch Logs in a separate logging account to audit log data

This solution provides the following security benefits:
- It uses AWS Systems Manager Session Manager instead of traditional SSH and RDP protocols, which provides a secure method for accessing EC2 instances without requiring inbound firewall rules or open ports.
- It provides audit trails by configuring Session Manager logging to Amazon CloudWatch Logs and creating a separate logging account to audit the log data.
- It uses the AWS Systems Manager Agent to automate common administrative tasks and improve the security posture of the instances.
- The separate logging account with cross-account permissions provides better data separation and improves security posture.

https://aws.amazon.com/solutions/implementations/centralized-logging/

NEW QUESTION 78
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 81
A company is using AWS WAF to protect a customized public API service that is based on Amazon EC2 instances. The API uses an Application Load Balancer. The AWS WAF web ACL is configured with an AWS Managed Rules rule group. After a software upgrade to the API and the client application, some types of requests are no longer working and are causing application stability issues. A security engineer discovers that AWS WAF logging is not turned on for the web ACL.
The security engineer needs to immediately return the application to service, resolve the issue, and ensure that logging is not turned off in the future. The security engineer turns on logging for the web ACL and specifies Amazon Cloud-Watch Logs as the destination.
Which additional set of steps should the security engineer take to meet the re-quirements?

A. Edit the rules in the web ACL to include rules with Count action
B. Review the logs to determine which rule is blocking the reques
C. Modify the IAM policy of all AWS WAF administrators so that they cannot remove the log-ging configuration for any AWS WAF web ACLs.
D. Edit the rules in the web ACL to include rules with Count action
E. Review the logs to determine which rule is blocking the reques
F. Modify the AWS WAF resource policy so that AWS WAF administrators cannot remove the log-ging configuration for any AWS WAF web ACLs.
G. Edit the rules in the web ACL to include rules with Count and Challenge action

H. Review the logs to determine which rule is blocking the reques
I. Modify the AWS WAF resource policy so that AWS WAF administrators cannot remove the logging configuration for any AWS WAF web ACLs.
J. Edit the rules in the web ACL to include rules with Count and Challenge action
K. Review the logs to determine which rule is blocking the reques
L. Modify the IAM policy of all AWS WAF administrators so that they cannot remove the logging configuration for any AWS WAF web ACLs.

Answer: A

Explanation:
This answer is correct because it meets the requirements of returning the application to service, resolving the issue, and ensuring that logging is not turned off in the future. By editing the rules in the web ACL to include rules with Count actions, the security engineer can test the effect of each rule without blocking or allowing requests. By reviewing the logs, the security engineer can identify which rule is causing the problem and modify or delete it accordingly. By modifying the IAM policy of all AWS WAF administrators, the security engineer can restrict their permissions to prevent them from removing the logging configuration for any AWS WAF web ACLs.

NEW QUESTION 84
A company's security team needs to receive a notification whenever an AWS access key has not been rotated in 90 or more days. A security engineer must develop a solution that provides these notifications automatically.
Which solution will meet these requirements with the LEAST amount of effort?

A. Deploy an AWS Config managed rule to run on a periodic basis of 24 hour
B. Select theaccess-keys-rotated managed rule, and set the maxAccessKeyAge parameter to 90 day
C. Create an Amazon EventBridge (Amazon CloudWatch Events) rule with an event pattern that matches the compliance type of NON_COMPLIANT from AWS Config for the managed rul
D. Configure EventBridge (CloudWatch Events) to send an Amazon Simple Notification Service (Amazon SNS) notification to the security team.
E. Create a script to export a .csv file from the AWS Trusted Advisor check for IAM access key rotation.Load the script into an AWS Lambda function that will upload the .csv file to an Amazon S3 bucke
F. Create an Amazon Athena table query that runs when the .csv file is uploaded to the S3 bucke
G. Publish the results for any keys older than 90 days by using an invocation of an Amazon Simple Notification Service (Amazon SNS) notification to the security team.
H. Create a script to download the IAM credentials report on a periodic basi
I. Load the script into an AWS Lambda function that will run on a schedule through Amazon EventBridge (Amazon CloudWatch Events). Configure the Lambda script to load the report into memory and to filter the report for recordsin which the key was last rotated at least 90 days ag
J. If any records are detected, send an Amazon Simple Notification Service (Amazon SNS) notification to the security team.
K. Create an AWS Lambda function that queries the IAM API to list all the user
L. Iterate through the users by using the ListAccessKeys operatio
M. Verify that the value in the CreateDate field is not at least 90 days ol
N. Send an Amazon Simple Notification Service (Amazon SNS) notification to the security team if the value is at least 90 days ol
O. Create an Amazon EventBridge (Amazon CloudWatch Events) rule to schedule the Lambda function to run each day.

Answer: A

NEW QUESTION 86
A developer 15 building a serverless application hosted on IAM that uses Amazon Redshift in a data store. The application has separate modules for read/write and read-only functionality. The modules need their own database users tor 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 read/write.
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 IAM poky for each module Specify the ARN of an Amazon Redshift database user that allows the GetClusterCredentials API call
D. Create focal database users for each module
E. Configure an IAM policy for each module Specify the ARN of an IAM user that allows the GetClusterCredentials API call

Answer: CD

Explanation:
To grant appropriate access to the application modules, the security engineer should do the following:
- Configure an IAM policy for each module. Specify the ARN of an Amazon Redshift database user that allows the GetClusterCredentials API call. This allows the application modules to use temporary credentials to access the database with the permissions of the specified user.
- Create local database users for each module. This allows the security engineer to create separate users for read/write and read-only functionality, and to assign them different privileges on the database tables.

NEW QUESTION 89
You need to create a policy and apply it for just an individual user. How could you accomplish this in the right way?
Please select:

A. Add an IAM managed policy for the user
B. Add a service policy for the user
C. Add an IAM role for the user
D. Add an inline policy for the user

Answer: D

Explanation:
Options A and B are incorrect since you need to add an inline policy just for the user Option C is invalid because you don't assign an IAM role to a user
The IAM Documentation mentions the following
An inline policy is a policy that's embedded in a principal entity (a user, group, or role)—that is, the policy is an inherent part of the principal entity. You can create a policy and embed it in a principal entity, either when you create the principal entity or later.

For more information on IAM Access and Inline policies, just browse to the below URL: https://docs.IAM.amazon.com/IAM/latest/UserGuide/access
The correct answer is: Add an inline policy for the user Submit your Feedback/Queries to our Experts

NEW QUESTION 94
A company is running its workloads in a single AWS Region and uses AWS Organizations. A security engineer must implement a solution to prevent users from launching resources in other Regions.
Which solution will meet these requirements with the LEAST operational overhead?

A. Create an IAM policy that has an aws RequestedRegion condition that allows actions only in the designated Region Attach the policy to all users.
B. Create an I AM policy that has an aws RequestedRegion condition that denies actions that are not in the designated Region Attach the policy to the AWS account in AWS Organizations.
C. Create an IAM policy that has an aws RequestedRegion condition that allows the desired actions Attach the policy only to the users who are in the designated Region.
D. Create an SCP that has an aws RequestedRegion condition that denies actions that are not in the designated Regio
E. Attach the SCP to the AWS account in AWS Organizations.

Answer: D

Explanation:
Although you can use a IAM policy to prevent users launching resources in other regions. The best practice is to use SCP when using AWS organizations.
https://docs.aws.amazon.com/organizations/latest/userguide/orgs_manage_policies_scps_examples_general.htm

NEW QUESTION 96
You have an S3 bucket defined in IAM. You want to ensure that you encrypt the data before sending it across the wire. What is the best way to achieve this.
Please select:

A. Enable server side encryption for the S3 bucke
B. This request will ensure that the data is encrypted first.
C. Use the IAM Encryption CLI to encrypt the data first
D. Use a Lambda function to encrypt the data before sending it to the S3 bucket.
E. Enable client encryption for the bucket

Answer: B

Explanation:
One can use the IAM Encryption CLI to encrypt the data before sending it across to the S3 bucket. Options A and C are invalid because this would still mean that data is transferred in plain text Option D is invalid because you cannot just enable client side encryption for the S3 bucket For more information on Encrypting and Decrypting data, please visit the below URL:
https://IAM.amazonxom/blogs/securirv/how4o-encrvpt-and-decrypt-your-data-with-the-IAM-encryption-cl The correct answer is: Use the IAM Encryption CLI to encrypt the data first Submit your Feedback/Queries to our Experts

NEW QUESTION 98
An organization wants to log all IAM API calls made within all of its IAM accounts, and must have a central place to analyze these logs. What steps should be taken to meet these requirements in the MOST secure manner? (Select TWO)

A. Turn on IAM CloudTrail in each IAM account
B. Turn on CloudTrail in only the account that will be storing the logs
C. Update the bucket ACL of the bucket in the account that will be storing the logs so that other accounts can log to it
D. Create a service-based role for CloudTrail and associate it with CloudTrail in each account
E. Update the bucket policy of the bucket in the account that will be storing the logs so that other accounts can log to it

Answer: AE

Explanation:
these are the steps that can meet the requirements in the most secure manner. CloudTrail is a service that records AWS API calls and delivers log files to an S3 bucket. Turning on CloudTrail in each IAM account can help capture all IAM API calls made within those accounts. Updating the bucket policy of the bucket in the account that will be storing the logs can help grant other accounts permission to write log files to that bucket. The other options are either unnecessary or insecure for logging and analyzing IAM API calls.

NEW QUESTION 101
A company needs to retain tog data archives for several years to be compliant with regulations. The tog data is no longer used but It must be retained
What Is the MOST secure and cost-effective solution to meet these requirements?

A. Archive the data to Amazon S3 and apply a restrictive bucket policy to deny the s3 DeleteOotect API
B. Archive the data to Amazon S3 Glacier and apply a Vault Lock policy
C. Archive the data to Amazon S3 and replicate it to a second bucket in a second IAM Region Choose the S3 Standard-Infrequent Access (S3 Standard-1A) storage class and apply a restrictive bucket policy to deny the s3 DeleteObject API
D. Migrate the log data to a 16 T8 Amazon Elastic Block Store (Amazon EBS) volume Create a snapshot of the EBS volume

Answer: B

Explanation:
To securely and cost-effectively retain log data archives for several years, the company should do the following:

- Archive the data to Amazon S3 Glacier and apply a Vault Lock policy. This allows the company to use a low-cost storage class that is designed for long-term archival of data that is rarely accessed. It also allows the company to enforce compliance controls on their S3 Glacier vault by locking a vault access policy that cannot be changed.

NEW QUESTION 102
A company needs to encrypt all of its data stored in Amazon S3. The company wants to use IAM Key Management Service (IAM KMS) to create and manage its encryption keys. The company's security policies require the ability to Import the company's own key material for the keys, set an expiration date on the keys, and delete keys immediately, if needed.
How should a security engineer set up IAM KMS to meet these requirements?

A. Configure IAM KMS and use a custom key stor
B. Create a customer managed CMK with no key material Import the company's keys and key material into the CMK
C. Configure IAM KMS and use the default Key store Create an IAM managed CMK with no key material Import the company's key material into the CMK
D. Configure IAM KMS and use the default key store Create a customer managed CMK with no key material import the company's key material into the CMK
E. Configure IAM KMS and use a custom key stor
F. Create an IAM managed CMK with no key material.Import the company's key material into the CMK.

Answer: A

Explanation:
To meet the requirements of importing their own key material, setting an expiration date on the keys, and deleting keys immediately, the security engineer should do the following:

➢ Configure AWS KMS and use a custom key store. This allows the security engineer to use a key manager outside of AWS KMS that they own and manage, such as an AWS CloudHSM cluster or an external key manager.

➢ Create a customer managed CMK with no key material. Import the company’s keys and key material into the CMK. This allows the security engineer to use their own key material for encryption and decryption operations, and to specify an expiration date for it.

NEW QUESTION 107
A company is developing an ecommerce application. The application uses Amazon EC2 instances and an Amazon RDS MySQL database. For compliance reasons, data must be secured in transit and at rest. The company needs a solution that minimizes operational overhead and minimizes cost.
Which solution meets these requirements?

A. Use TLS certificates from AWS Certificate Manager (ACM) with an Application Load Balancer.Deploy self-signed certificates on the EC2 instance
B. Ensure that the database client software uses a TLS connection to Amazon RD
C. Enable encryption of the RDS DB instanc
D. Enable encryption on the Amazon Elastic Block Store (Amazon EBS) volumes that support the EC2 instances.
E. Use TLS certificates from a third-party vendor with an Application Load Balance
F. Install the same certificates on the EC2 instance
G. Ensure that the database client software uses a TLS connection to Amazon RD
H. Use AWS Secrets Manager for client-side encryption of application data.
I. Use AWS CloudHSM to generate TLS certificates for the EC2 instance
J. Install the TLS certificates on the EC2 instance
K. Ensure that the database client software uses a TLS connection to Amazon RD
L. Use the encryption keys form CloudHSM for client-side encryption of application data.
M. Use Amazon CloudFront with AWS WA
N. Send HTTP connections to the origin EC2 instance
O. Ensure that the database client software uses a TLS connection to Amazon RD
P. Use AWS Key Management Service (AWS KMS) for client-side encryption of application data before the data is stored in the RDS database.

Answer: A

NEW QUESTION 109
A company hosts an application on Amazon EC2 that is subject to specific rules for regulatory compliance. One rule states that traffic to and from the workload must be inspected for network-level attacks. This involves inspecting the whole packet.
To comply with this regulatory rule, a security engineer must install intrusion detection software on a c5n.4xlarge EC2 instance. The engineer must then configure the software to monitor traffic to and from the application instances.
What should the security engineer do next?

A. Place the network interface in promiscuous mode to capture the traffic.
B. Configure VPC Flow Logs to send traffic to the monitoring EC2 instance using a Network Load Balancer.
C. Configure VPC traffic mirroring to send traffic to the monitoring EC2 instance using a Network Load Balancer.
D. Use Amazon Inspector to detect network-level attacks and trigger an IAM Lambda function to send the suspicious packets to the EC2 instance.

Answer: D

NEW QUESTION 111
A company uses a third-party application to store encrypted data in Amazon S3. The company uses another third-party application trial decrypts the data from Amazon S3 to ensure separation of duties Between the applications A Security Engineer warns to separate the permissions using IAM roles attached to Amazon EC2 instances. The company prefers to use native IAM services.
Which encryption method will meet these requirements?

A. Use encrypted Amazon EBS volumes with Amazon default keys (IAM EBS)
B. Use server-side encryption with customer-provided keys (SSE-C)
C. Use server-side encryption with IAM KMS managed keys (SSE-KMS)
D. Use server-side encryption with Amazon S3 managed keys (SSE-S3)

Answer: C

NEW QUESTION 115
A company's security engineer wants to receive an email alert whenever Amazon GuardDuty, AWS Identity and Access Management Access Analyzer, or Amazon Made generate a high-severity security finding. The company uses AWS Control Tower to govern all of its accounts. The company also uses AWS Security Hub with all of the AWS service integrations turned on.

Which solution will meet these requirements with the LEAST operational overhead?

A. Set up separate AWS Lambda functions for GuardDuty, 1AM Access Analyzer, and Macie to call each service's public API to retrieve high-severity finding
B. Use Amazon Simple Notification Service (Amazon SNS) to send the email alert
C. Create an Amazon EventBridge rule to invoke the functions on a schedule.
D. Create an Amazon EventBridge rule with a pattern that matches Security Hub findings events with high severit
E. Configure the rule to send the findings to a target Amazon Simple Notification Service (Amazon SNS) topi
F. Subscribe the desired email addresses to the SNS topic.
G. Create an Amazon EventBridge rule with a pattern that matches AWS Control Tower events with high severit
H. Configure the rule to send the findings to a target Amazon Simple Notification Service (Amazon SNS) topi
I. Subscribe the desired email addresses to the SNS topic.
J. Host an application on Amazon EC2 to call the GuardDuty, 1AM Access Analyzer, and Macie APIs.Within the application, use the Amazon Simple Notification Service (Amazon SNS) API to retrieve high-severity findings and to send the findings to an SNS topi
K. Subscribe the desired email addresses to the SNS topic.

Answer: B

Explanation:
The AWS documentation states that you can create an Amazon EventBridge rule with a pattern that matches Security Hub findings events with high severity. You can then configure the rule to send the findings to a target Amazon Simple Notification Service (Amazon SNS) topic. You can subscribe the desired email addresses to the SNS topic. This method is the least operational overhead way to meet the requirements.
References: : AWS Security Hub User Guide

NEW QUESTION 116
A company is developing a highly resilient application to be hosted on multiple Amazon EC2 instances . The application will store highly sensitive user data in Amazon RDS tables
The application must
• Include migration to a different IAM Region in the application disaster recovery plan.
• Provide a full audit trail of encryption key administration events
• Allow only company administrators to administer keys.
• Protect data at rest using application layer encryption
A Security Engineer is evaluating options for encryption key management
Why should the Security Engineer choose IAM CloudHSM over IAM KMS for encryption key management in this situation?

A. The key administration event logging generated by CloudHSM is significantly more extensive than IAM KMS.
B. CloudHSM ensures that only company support staff can administer encryption keys, whereas IAM KMS allows IAM staff to administer keys
C. The ciphertext produced by CloudHSM provides more robust protection against brute force decryption attacks than the ciphertext produced by IAM KMS
D. CloudHSM provides the ability to copy keys to a different Region, whereas IAM KMS does not

Answer: B

Explanation:
CloudHSM allows full control of your keys such including Symmetric (AES), Asymmetric (RSA), Sha-256, SHA 512, Hash Based, Digital Signatures (RSA). On the other hand, AWS Key Management Service is a
multi-tenant key storage that is owned and managed by AWS1.
References: 1: What are the differences between AWS Cloud HSM and KMS?

NEW QUESTION 117
A company needs to use HTTPS when connecting to its web applications to meet compliance requirements. These web applications run in Amazon VPC on Amazon EC2 instances behind an Application Load Balancer (ALB). A security engineer wants to ensure that the load balancer win only accept connections over port 443. even if the ALB is mistakenly configured with an HTTP listener
Which configuration steps should the security engineer take to accomplish this task?

A. Create a security group with a rule that denies Inbound connections from 0.0.0 0/0 on port 00. Attach this security group to the ALB to overwrite more permissive rules from the ALB's default securitygroup.
B. Create a network ACL that denies inbound connections from 0 0.0.0/0 on port 80 Associate the network ACL with the VPC s internet gateway
C. Create a network ACL that allows outbound connections to the VPC IP range on port 443 only.Associate the network ACL with the VPC's internet gateway.
D. Create a security group with a single inbound rule that allows connections from 0.0.0 0/0 on port 443.Ensure this security group is the only one associated with the ALB

Answer: D

Explanation:
To ensure that the load balancer only accepts connections over port 443, the security engineer should do the following:
➢ Create a security group with a single inbound rule that allows connections from 0.0.0.0/0 on port 443.
This means that the security group allows HTTPS traffic from any source IP address.
➢ Ensure this security group is the only one associated with the ALB. This means that the security group overrides any other rules that might allow HTTP traffic on port 80.

NEW QUESTION 121
A Development team has built an experimental environment to test a simple stale web application It has built an isolated VPC with a private and a public subnet. The public subnet holds only an Application Load Balancer a NAT gateway, and an internet gateway. The private subnet holds ail of the Amazon EC2 instances There are 3 different types of servers Each server type has its own Security Group that limits access lo only required connectivity. The Security Groups nave both inbound and outbound rules applied Each subnet has both inbound and outbound network ACls applied to limit access to only required connectivity
Which of the following should the team check if a server cannot establish an outbound connection to the
internet? (Select THREE.)

A. The route tables and the outbound rules on the appropriate private subnet security group
B. The outbound network ACL rules on the private subnet and the Inbound network ACL rules on the public subnet

C. The outbound network ACL rules on the private subnet and both the inbound and outbound rules on the public subnet
D. The rules on any host-based firewall that may be applied on the Amazon EC2 instances
E. The Security Group applied to the Application Load Balancer and NAT gateway
F. That the 0.0.0./0 route in the private subnet route table points to the internet gateway in the public subnet

Answer: CEF

Explanation:
because these are the factors that could affect the outbound connection to the internet from a server in a private subnet. The outbound network ACL rules on the private subnet and both the inbound and outbound rules on the public subnet must allow the traffic to pass through8. The security group applied to the application load balancer and NAT gateway must also allow the traffic from the private subnet9. The 0.0.0.0/0 route in the private subnet route table must point to the NAT gateway in the public subnet, not the internet gateway10. The other options are either irrelevant or incorrect for troubleshooting the outbound connection issue.

NEW QUESTION 126
A website currently runs on Amazon EC2, wan mostly statics content on the site. Recently the site was subjected to a DDoS attack a security engineer was (asked was redesigning the edge security to help
Mitigate this risk in the future.
What are some ways the engineer could achieve this (Select THREE)?

A. Use IAM X-Ray to inspect the trafc going to the EC2 instances.
B. Move the static content to Amazon S3, and front this with an Amazon Cloud Front distribution.
C. Change the security group conguration to block the source of the attack trafc
D. Use IAM WAF security rules to inspect the inbound trafc.
E. Use Amazon Inspector assessment templates to inspect the inbound traffic.
F. Use Amazon Route 53 to distribute trafc.

Answer: BDF

Explanation:
To redesign the edge security to help mitigate the DDoS attack risk in the future, the engineer could do the following:

- Move the static content to Amazon S3, and front this with an Amazon CloudFront distribution. This allows the engineer to use a global content delivery network that can cache static content at edge locations and reduce the load on the origin servers.
- Use AWS WAF security rules to inspect the inbound traffic. This allows the engineer to use web application firewall rules that can filter malicious requests based on IP addresses, headers, body, or URI strings, and block them before they reach the web servers.
- Use Amazon Route 53 to distribute traffic. This allows the engineer to use a scalable and highly available DNS service that can route traffic based on different policies, such as latency, geolocation, or health checks.

NEW QUESTION 129
A company uses AWS Signer with all of the company's AWS Lambda functions. A developer recently stopped working for the company. The company wants to ensure that all the code that the developer wrote can no longer be deployed to the Lambda functions.
Which solution will meet this requirement?

A. Revoke all versions of the signing profile assigned to the developer.
B. Examine the developer's IAM role
C. Remove all permissions that grant access to Signer.
D. Re-encrypt all source code with a new AWS Key Management Service (AWS KMS) key.
E. Use Amazon CodeGuru to profile all the code that the Lambda functions use.

Answer: A

Explanation:
The correct answer is A. Revoke all versions of the signing profile assigned to the developer.
According to the AWS documentation1, AWS Signer is a fully managed code-signing service that helps you ensure the trust and integrity of your code. You can use Signer to sign code artifacts, such as Lambda deployment packages, with code-signing certificates that you control and manage.
A signing profile is a collection of settings that Signer uses to sign your code artifacts. A signing profile includes information such as the following:

- The type of signature that you want to create (for example, a code-signing signature).
- The signing algorithm that you want Signer to use to sign your code.
- The code-signing certificate and its private key that you want Signer to use to sign your code.

You can create multiple versions of a signing profile, each with a different code-signing certificate. You can also revoke a version of a signing profile if you no longer want to use it for signing code artifacts.
In this case, the company wants to ensure that all the code that the developer wrote can no longer be deployed to the Lambda functions. One way to achieve this is to revoke all versions of the signing profile that was assigned to the developer. This will prevent Signer from using that signing profile to sign any new code artifacts, and also invalidate any existing signatures that were created with that signing profile. This way, the company can ensure that only trusted and authorized code can be deployed to the Lambda functions.
The other options are incorrect because:

- B. Examining the developer's IAM roles and removing all permissions that grant access to Signer may not be sufficient to prevent the deployment of the developer's code. The developer may have already signed some code artifacts with a valid signing profile before leaving the company, and those signatures may still be accepted by Lambda unless the signing profile is revoked.
- C. Re-encrypting all source code with a new AWS Key Management Service (AWS KMS) key may not be effective or practical. AWS KMS is a service that lets you create and manage encryption keys for your data. However, Lambda does not require encryption keys for deploying code artifacts, only valid signatures from Signer. Therefore, re-encrypting the source code may not prevent the deployment of the developer's code if it has already been signed with a valid signing profile. Moreover, re-encrypting all source code may be time-consuming and disruptive for other developers who are working on the same code base.
- D. Using Amazon CodeGuru to profile all the code that the Lambda functions use may not help with preventing the deployment of the developer's code. Amazon CodeGuru is a service that provides intelligent recommendations to improve your code quality and identify an application's most expensive lines of code. However, CodeGuru does not perform any security checks or validations on your code artifacts, nor does it interact with Signer or Lambda in any way. Therefore, using CodeGuru may not prevent unauthorized or untrusted code from being deployed to the Lambda functions.

References:
1: What is AWS Signer? - AWS Signer

NEW QUESTION 134
A company's application team needs to host a MySQL database on IAM. According to the company's security policy, all data that is stored on IAM must be encrypted at rest. In addition, all cryptographic material must be compliant with FIPS 140-2 Level 3 validation.
The application team needs a solution that satisfies the company's security requirements and minimizes
operational overhead.
Which solution will meet these requirements?

A. Host the database on Amazon RD
B. Use Amazon Elastic Block Store (Amazon EBS) for encryption.Use an IAM Key Management Service (IAM KMS) custom key store that is backed by IAM CloudHSM for key management.
C. Host the database on Amazon RD
D. Use Amazon Elastic Block Store (Amazon EBS) for encryption.Use an IAM managed CMK in IAM Key Management Service (IAM KMS) for key management.
E. Host the database on an Amazon EC2 instanc
F. Use Amazon Elastic Block Store (Amazon EBS) for encryptio
G. Use a customer managed CMK in IAM Key Management Service (IAM KMS) for key management.
H. Host the database on an Amazon EC2 instanc
I. Use Transparent Data Encryption (TDE) for encryption and key management.

Answer: B

NEW QUESTION 135
An ecommerce company has a web application architecture that runs primarily on containers. The application containers are deployed on Amazon Elastic Container Service (Amazon ECS). The container images for the application are stored in Amazon Elastic Container Registry (Amazon ECR).
The company's security team is performing an audit of components of the application architecture. The security team identifies issues with some container images that are stored in the container repositories.
The security team wants to address these issues by implementing continual scanning and on-push scanning of the container images. The security team needs to implement a solution that makes any findings from these scans visible in a centralized dashboard. The security team plans to use the dashboard to view these findings along with other security-related findings that they intend to generate in the future.
There are specific repositories that the security team needs to exclude from the scanning process. Which solution will meet these requirements?

A. Use Amazon Inspecto
B. Create inclusion rules in Amazon ECR to match repos-itories that need to be scanne
C. Push Amazon Inspector findings to AWS Se-curity Hub.
D. Use ECR basic scanning of container image
E. Create inclusion rules in Ama-zon ECR to match repositories that need to be scanne
F. Push findings to AWS Security Hub.
G. Use ECR basic scanning of container image
H. Create inclusion rules in Ama-zon ECR to match repositories that need to be scanne
I. Push findings to Amazon Inspector.
J. Use Amazon Inspecto
K. Create inclusion rules in Amazon Inspector to match repositories that need to be scanne
L. Push Amazon Inspector findings to AWS Config.

Answer: A

NEW QUESTION 138
A security engineer is trying to use Amazon EC2 Image Builder to create an image of an EC2 instance. The security engineer has configured the pipeline to send logs to an Amazon S3 bucket. When the security engineer runs the pipeline, the build fails with the following error: “AccessDenied: Access Denied status code: 403”.
The security engineer must resolve the error by implementing a solution that complies with best practices for least privilege access.
Which combination of steps will meet these requirements? (Choose two.)

A. Ensure that the following policies are attached to the IAM role that the security engineer is using: EC2InstanceProfileForImageBuilder, EC2InstanceProfileForImageBuilderECRContainerBuilds, and AmazonSSMManagedInstanceCore.
B. Ensure that the following policies are attached to the instance profile for the EC2 instance: EC2InstanceProfileForImageBuilder, EC2InstanceProfileForImageBuilderECRContainerBuilds, and AmazonSSMManagedInstanceCore.
C. Ensure that the AWSImageBuilderFullAccess policy is attached to the instance profile for the EC2 instance.
D. Ensure that the security engineer’s IAM role has the s3:PutObject permission for the S3 bucket.
E. Ensure that the instance profile for the EC2 instance has the s3:PutObject permission for the S3 bucket.

Answer: BE

Explanation:
The most likely cause of the error is that the instance profile for the EC2 instance does not have the s3:PutObject permission for the S3 bucket. This permission is needed to upload logs to the bucket. Therefore, the security engineer should ensure that the instance profile has this permission.
One possible solution is to attach the AWSImageBuilderFullAccess policy to the instance profile for the EC2 instance. This policy grants full access to Image Builder resources and related AWS services, including the s3:PutObject permission for any bucket with “imagebuilder” in its name. However, this policy may grant more permissions than necessary, which violates the principle of least privilege.
Another possible solution is to create a custom policy that only grants the s3:PutObject permission for the specific S3 bucket that is used for logging. This policy can be attached to the instance profile along with the other policies that are required for Image Builder functionality: EC2InstanceProfileForImageBuilder, EC2InstanceProfileForImageBuilderECRContainerBuilds, and AmazonSSMManagedInstanceCore. This solution follows the principle of least privilege more closely than the previous one.
➢ Ensure that the following policies are attached to the instance profile for the EC2 instance: EC2InstanceProfileForImageBuilder, EC2InstanceProfileForImageBuilderECRContainerBuilds, and AmazonSSMManagedInstanceCore.
➢ Ensure that the instance profile for the EC2 instance has the s3:PutObject permission for the S3 bucket.
This can be done by either attaching the AWSImageBuilderFullAccess policy or creating a custom policy with this permission.
1: Using managed policies for EC2 Image Builder - EC2 Image Builder 2: PutObject - Amazon Simple Storage Service 3: AWSImageBuilderFullAccess - AWS Managed Policy

NEW QUESTION 139
A company is using AWS to run a long-running analysis process on data that is stored in Amazon S3 buckets. The process runs on a fleet of Amazon EC2 instances that are in an Auto Scaling group. The EC2 instances are deployed in a private subnet Of a VPC that does not have internet access. The EC2 instances and the S3 buckets are in the same AWS account
The EC2 instances access the S3 buckets through an S3 gateway endpoint that has the default access policy. Each EC2 instance is associated With an instance profile role that has a policy that explicitly allows the s3:GetObject action and the s3:PutObject action for only the required S3 buckets.
The company learns that one or more of the EC2 instances are compromised and are exfiltrating data to an S3 bucket that is outside the companys organization in AWS Organizations. A security engtneer must implement a solution to stop this exfiltration of data and to keep the EC2 processing job functional.
Which solution will meet these requirements?

A. Update the policy on the S3 gateway endpoint to allow the S3 actions CY11y if the values of the aws:ResourceOrglD and aws:PrincipalOrglD condition keys match the companys values.
B. Update the policy on the instance profile role to allow the S3 actions only if the value of the aws:ResourceOrglD condition key matches the company's value.
C. Add a network ACL rule to the subnet of the EC2 instances to block outgoing connections on port 443.
D. Apply an SCP on the AWS account to allow the $3 actions only if the values of the aws:ResourceOrglD and aws:PrincipalOrglD condition keys match the company's values.

Answer: D

Explanation:
The correct answer is D.
To stop the data exfiltration from the compromised EC2 instances, the security engineer needs to implement a solution that can deny access to any S3 bucket that is outside the company’s organization. The solution should also allow the EC2 instances to access the required S3 buckets within the company’s organization for the analysis process.
Option A is incorrect because updating the policy on the S3 gateway endpoint will not affect the access to S3 buckets that are outside the company’s organization. The S3 gateway endpoint only applies to S3 buckets that are in the same AWS Region as the VPC. The compromised EC2 instances can still access S3 buckets in other Regions or other AWS accounts through the internet gateway or NAT device.
Option B is incorrect because updating the policy on the instance profile role will not prevent the compromised EC2 instances from using other credentials or methods to access S3 buckets outside the company’s organization. The instance profile role only applies to requests that are made using the credentials of that role. The compromised EC2 instances can still use other IAM users, roles, or access keys to access S3 buckets outside the company’s organization.
Option C is incorrect because adding a network ACL rule to block outgoing connections on port 443 will also block legitimate connections to S3 buckets within the company’s organization. The network ACL rule will prevent the EC2 instances from accessing any S3 bucket through HTTPS, regardless of whether it is inside or outside the company’s organization.
Option D is correct because applying an SCP on the AWS account will effectively deny access to any S3 bucket that is outside the company’s organization. The SCP will apply to all IAM users, roles, and resources in the AWS account, regardless of how they access S3. The SCP will use the aws:ResourceOrgID and aws:PrincipalOrgID condition keys to check whether the S3 bucket and the principal belong to the same organization as the AWS account. If they do not match, the SCP will deny the S3 actions.
References:

- Using service control policies
- AWS Organizations service control policy examples

NEW QUESTION 144
A company's on-premises networks are connected to VPCs using an IAM Direct Connect gateway. The company's on-premises application needs to stream data using an existing Amazon Kinesis Data Firehose delivery stream. The company's security policy requires that data be encrypted in transit using a private network.
How should the company meet these requirements?

A. Create a VPC endpoint tor Kinesis Data Firehos
B. Configure the application to connect to the VPC endpoint.
C. Configure an IAM policy to restrict access to Kinesis Data Firehose using a source IP condition.Configure the application to connect to the existing Firehose delivery stream.
D. Create a new TLS certificate in IAM Certificate Manager (ACM). Create a public-facing Network Load Balancer (NLB) and select the newly created TLS certificat
E. Configure the NLB to forward all traffic to Kinesis Data Firehos
F. Configure the application to connect to the NLB.
G. Peer the on-premises network with the Kinesis Data Firehose VPC using Direct Connec
H. Configure the application to connect to the existing Firehose delivery stream.

Answer: A

Explanation:
To stream data using an existing Amazon Kinesis Data Firehose delivery stream and encrypt it in transit using a private network, the company should do the following:

- Create a VPC endpoint for Kinesis Data Firehose. This allows the company to use a private connection between their VPC and Kinesis Data Firehose without requiring an internet gateway or NAT device.
- Configure the application to connect to the VPC endpoint. This allows the application to stream data using Kinesis Data Firehose over AWS PrivateLink, which encrypts all traffic with TLS.

NEW QUESTION 145
A company has a set of EC2 Instances hosted in IAM. The EC2 Instances have EBS volumes which is used to store critical information. There is a business continuity requirement to ensure high availability for the EBS volumes. How can you achieve this?

A. Use lifecycle policies for the EBS volumes
B. Use EBS Snapshots
C. Use EBS volume replication
D. Use EBS volume encryption

Answer: B

Explanation:
Data stored in Amazon EBS volumes is redundantly stored in multiple physical locations as part of normal operation of those services and at no additional charge.

However, Amazon EBS replication is stored within the same availability zone, not across multiple zones; therefore, it is highly recommended that you conduct regular snapshots to Amazon S3 for long-term data durability Option A is invalid because there is no lifecycle policy for EBS volumes Option C is invalid because there is no EBS volume replication Option D is invalid because EBS volume encryption will not ensure business continuity For information on security for Compute Resources, please visit the below URL: https://d1.awsstatic.com/whitepapers/Security/Security_Compute_Services_Whitepaper.pdf

NEW QUESTION 146
A company has a large fleet of Linux Amazon EC2 instances and Windows EC2 instances that run in private subnets. The company wants all remote administration to be performed as securely as possible in the AWS Cloud.
Which solution will meet these requirements?

A. Do not use SSH-RSA private keys during the launch of new instance
B. Implement AWS Systems Manager Session Manager.
C. Generate new SSH-RSA private keys for existing instance
D. Implement AWS Systems Manager Session Manager.
E. Do not use SSH-RSA private keys during the launch of new instance
F. Configure EC2 Instance Connect.
G. Generate new SSH-RSA private keys for existing instance
H. Configure EC2 Instance Connect.

Answer: A

Explanation:
AWS Systems Manager Session Manager is a fully managed service that allows you to securely and remotely administer your EC2 instances without the need to open inbound ports, maintain bastion hosts, or manage SSH keys. Session Manager provides an interactive browser-based shell or CLI access to your instances, as well as port forwarding and auditing capabilities. Session Manager works with both Linux and Windows instances, and supports hybrid environments and edge devices.
EC2 Instance Connect is a feature that allows you to use SSH to connect to your Linux instances using
short-lived keys that are generated on demand and delivered securely through the AWS metadata service. EC2 Instance Connect does not require any additional software installation or configuration on the instance, but it does require you to use SSH-RSA keys during the launch of new instances.
The correct answer is to use Session Manager, as it provides more security and flexibility than EC2 Instance Connect, and does not require SSH-RSA keys or inbound ports. Session Manager also works with Windows instances, while EC2 Instance Connect does not.
Verified References:

- https://docs.aws.amazon.com/systems-manager/latest/userguide/session-manager.html
- https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/Connect-using-EC2-Instance-Connect.html
- https://repost.aws/questions/QUnV4R9EoeSdW0GT3cKBUR7w/what-is-the-difference-between-ec-2-ins

NEW QUESTION 151
......

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

SCS-C01 Practice Exam Features:

* SCS-C01 Questions and Answers Updated Frequently

* SCS-C01 Practice Questions Verified by Expert Senior Certified Staff

* SCS-C01 Most Realistic Questions that Guarantee you a Pass on Your FirstTry

* SCS-C01 Practice Test Questions in Multiple Choice Formats and Updatesfor 1 Year

100% Actual & Verified — Instant Download, Please Click
Order The SCS-C01 Practice Test Here