

SCS-C02 Dumps

AWS Certified Security - Specialty

<https://www.certleader.com/SCS-C02-dumps.html>



NEW QUESTION 1

A security engineer needs to develop a process to investigate and respond to potential security events on a company's Amazon EC2 instances. All the EC2 instances are backed by Amazon Elastic Block Store (Amazon EBS). The company uses AWS Systems Manager to manage all the EC2 instances and has installed Systems Manager Agent (SSM Agent) on all the EC2 instances.

The process that the security engineer is developing must comply with AWS security best practices and must meet the following requirements:

- A compromised EC2 instance's volatile memory and non-volatile memory must be preserved for forensic purposes.
- A compromised EC2 instance's metadata must be updated with corresponding incident ticket information.
- A compromised EC2 instance must remain online during the investigation but must be isolated to prevent the spread of malware.
- Any investigative activity during the collection of volatile data must be captured as part of the process. Which combination of steps should the security engineer take to meet these requirements with the LEAST operational overhead? (Select THREE.)

- A. Gather any relevant metadata for the compromised EC2 instance
- B. Enable termination protection
- C. Isolate the instance by updating the instance's security groups to restrict access
- D. Detach the instance from any Auto Scaling groups that the instance is a member of
- E. Deregister the instance from any Elastic Load Balancing (ELB) resources.
- F. Gather any relevant metadata for the compromised EC2 instance
- G. Enable termination protection
- H. Move the instance to an isolation subnet that denies all source and destination traffic
- I. Associate the instance with the subnet to restrict access
- J. Detach the instance from any Auto Scaling groups that the instance is a member of
- K. Deregister the instance from any Elastic Load Balancing (ELB) resources.
- L. Use Systems Manager Run Command to invoke scripts that collect volatile data.
- M. Establish a Linux SSH or Windows Remote Desktop Protocol (RDP) session to the compromised EC2 instance to invoke scripts that collect volatile data.
- N. Create a snapshot of the compromised EC2 instance's EBS volume for follow-up investigation
- O. Tag the instance with any relevant metadata and incident ticket information.
- P. Create a Systems Manager State Manager association to generate an EBS volume snapshot of the compromised EC2 instance
- Q. Tag the instance with any relevant metadata and incident ticket information.

Answer: ACE

NEW QUESTION 2

A security engineer needs to create an Amazon S3 bucket policy to grant least privilege read access to IAM user accounts that are named User1, User2, and User3. These IAM user accounts are members of the AuthorizedPeople IAM group. The security engineer drafts the following S3 bucket policy:

```
{
  "Version": "2012-10-17",
  "Id": "AuthorizedPeoplePolicy",
  "Statement": [
    {
      "Sid": "Actions-Authorized-People",
      "Effect": "Allow",
      "Action": [
        "s3:GetObject"
      ],
      "Resource": "arn:aws:s3:::authorized-people-bucket/*"
    }
  ]
}
```

When the security engineer tries to add the policy to the S3 bucket, the following error message appears: "Missing required field Principal." The security engineer is adding a Principal element to the policy. The addition must provide read access to only User1, User2, and User3. Which solution meets these requirements?

A)

```
"Principal": {
  "AWS": [
    "arn:aws:iam::1234567890:user/User1",
    "arn:aws:iam::1234567890:user/User2",
    "arn:aws:iam::1234567890:user/User3"
  ]
}
```

B)

```
"Principal": {
  "AWS": [
    "arn:aws:iam::1234567890:root"
  ]
}
```

C)

```
"Principal": {
  "AWS": [
    "*"
  ]
}
```

D)

```
"Principal": {
  "AWS": "arn:aws:iam::1234567890:group/AuthorizedPeople"
}
```

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

Answer: A

NEW QUESTION 3

A company hosts a public website on an Amazon EC2 instance. HTTPS traffic must be able to access the website. The company uses SSH for management of the web server.

The website is on the subnet 10.0.1.0/24. The management subnet is 192.168.100.0/24. A security engineer must create a security group for the EC2 instance. Which combination of steps should the security engineer take to meet these requirements in the MOST secure manner? (Select TWO.)

- A. Allow port 22 from source 0.0.0.0/0.
- B. Allow port 443 from source 0.0.0.0/0.
- C. Allow port 22 from 192.168.100.0/24.
- D. Allow port 22 from 10.0.1.0/24.
- E. Allow port 443 from 10.0.1.0/24.

Answer: BC

Explanation:

The correct answer is B and C.

* B. Allow port 443 from source 0.0.0.0/0.

This is correct because port 443 is used for HTTPS traffic, which must be able to access the website from any source IP address.

* C. Allow port 22 from 192.168.100.0/24.

This is correct because port 22 is used for SSH, which is the management protocol for the web server. The management subnet is 192.168.100.0/24, so only this subnet should be allowed to access port 22.

* A. Allow port 22 from source 0.0.0.0/0.

This is incorrect because it would allow anyone to access port 22, which is a security risk. SSH should be restricted to the management subnet only.

* D. Allow port 22 from 10.0.1.0/24.

This is incorrect because it would allow the website subnet to access port 22, which is unnecessary and a security risk. SSH should be restricted to the management subnet only.

* E. Allow port 443 from 10.0.1.0/24.

This is incorrect because it would limit the HTTPS traffic to the website subnet only, which defeats the purpose of having a public website.

NEW QUESTION 4

A company has developed a new Amazon RDS database application. The company must secure the RDS database credentials for encryption in transit and encryption at rest. The company also must rotate the credentials automatically on a regular basis.

Which solution meets these requirements?

- A. Use IAM Systems Manager Parameter Store to store the database credentials
- B. Configure automatic rotation of the credentials.
- C. Use IAM Secrets Manager to store the database credential
- D. Configure automatic rotation of the credentials
- E. Store the database credentials in an Amazon S3 bucket that is configured with server-side encryption with S3 managed encryption keys (SSE-S3) Rotate the credentials with IAM database authentication.
- F. Store the database credentials in Amazon S3 Glacier, and use S3 Glacier Vault Lock. Configure an IAM Lambda function to rotate the credentials on a scheduled basis

Answer: A

NEW QUESTION 5

A company is hosting a web application on Amazon EC2 instances behind an Application Load Balancer (ALB). The application has become the target of a DoS attack. Application logging shows that requests are coming from a small number of client IP addresses, but the addresses change regularly.

The company needs to block the malicious traffic with a solution that requires the least amount of ongoing effort.

Which solution meets these requirements?

- A. Create an AWS WAF rate-based rule, and attach it to the ALB.
- B. Update the security group that is attached to the ALB to block the attacking IP addresses.
- C. Update the ALB subnet's network ACL to block the attacking client IP addresses.
- D. Create a AWS WAF rate-based rule, and attach it to the security group of the EC2 instances.

Answer: A

NEW QUESTION 6

A Security Architect has been asked to review an existing security architecture and identify why the application servers cannot successfully initiate a connection to the database servers. The following summary describes the architecture:

* 1 An Application Load Balancer, an internet gateway, and a NAT gateway are configured in the public subnet

* 2 Database, application, and web servers are configured on three different private subnets.

* 3 The VPC has two route tables: one for the public subnet and one for all other subnets. The route table for the public subnet has a 0.0.0.0/0 route to the internet gateway. The route table for all other subnets has a 0.0.0.0/0 route to the NAT gateway. All private subnets can route to each other.

* 4 Each subnet has a network ACL implemented that limits all inbound and outbound connectivity to only the required ports and protocols.

* 5 There are 3 Security Groups (SGs) database application and web. Each group limits all inbound and outbound connectivity to the minimum required.

Which of the following accurately reflects the access control mechanisms the Architect should verify?

- A. Outbound SG configuration on database servers. Inbound SG configuration on application servers. Inbound and outbound network ACL configuration on the database subnet. Inbound and outbound network ACL configuration on the application server subnet.
- B. Inbound SG configuration on database servers. Outbound SG configuration on application servers. Inbound and outbound network ACL configuration on the

database subnet
Inbound and outbound network ACL configuration on the application server subnet
C. Inbound and outbound SG configuration on database servers
Inbound and outbound SG configuration on application servers
Inbound network ACL configuration on the database subnet
Outbound network ACL configuration on the application server subnet
D. Inbound SG configuration on database servers
Outbound SG configuration on application servers
Inbound network ACL configuration on the database subnet
Outbound network ACL configuration on the application server subnet.

Answer: A

Explanation:

this is the accurate reflection of the access control mechanisms that the Architect should verify. Access control mechanisms are methods that regulate who can access what resources and how. Security groups and network ACLs are two types of access control mechanisms that can be applied to EC2 instances and subnets. Security groups are stateful, meaning they remember and return traffic that was previously allowed. Network ACLs are stateless, meaning they do not remember or return traffic that was previously allowed. Security groups and network ACLs can have inbound and outbound rules that specify the source, destination, protocol, and port of the traffic. By verifying the outbound security group configuration on database servers, the inbound security group configuration on application servers, and the inbound and outbound network ACL configuration on both the database and application server subnets, the Architect can check if there are any misconfigurations or conflicts that prevent the application servers from initiating a connection to the database servers. The other options are either inaccurate or incomplete for verifying the access control mechanisms.

NEW QUESTION 7

A company uses an Amazon S3 bucket to store reports Management has mandated that all new objects stored in this bucket must be encrypted at rest using server-side encryption with a client-specified IAM Key Management Service (IAM KMS) CMK owned by the same account as the S3 bucket. The IAM account number is 111122223333, and the bucket name is report bucket. The company's security specialist must write the S3 bucket policy to ensure the mandate can be implemented

Which statement should the security specialist include in the policy?

- A.

```
{
  "Effect": "Deny",
  "Principal": "*",
  "Action": "s3:PutObject",
  "Resource": "arn:aws:s3:::reportbucket/*",
  "Condition": {
    "StringEquals": {
      "s3:x-amz-server-side-encryption": "AES256"
    }
  }
}
```
- B.

```
{
  "Effect": "Deny",
  "Principal": "*",
  "Action": "s3:PutObject",
  "Resource": "arn:aws:s3:::reportbucket/*",
  "Condition": {
    "StringNotLike": {
      "s3:x-amz-server-side-encryption-aws-kms-key-id": "arn:aws:kms:*:111122223333:key/*"
    }
  }
}
```
- C.

```
{
  "Effect": "Deny",
  "Principal": "*",
  "Action": "s3:PutObject",
  "Resource": "arn:aws:s3:::reportbucket/*",
  "Condition": {
    "StringNotLike": {
      "s3:x-amz-server-side-encryption": "aws:kms"
    }
  }
}
```
- D.

```
{
  "Effect": "Deny",
  "Principal": "*",
  "Action": "s3:PutObject",
  "Resource": "arn:aws:s3:::reportbucket/*",
  "Condition": {
    "StringNotLikeIfExists": {
      "s3:x-amz-server-side-encryption-aws-kms-key-id": "arn:aws:kms:*:111122223333:key/*"
    }
  }
}
```
- E. Option A
- F. Option B
- G. Option C
- H. Option D

Answer: D

NEW QUESTION 8

A company plans to use AWS Key Management Service (AWS KMS) to implement an encryption strategy to protect data at rest. The company requires client-side encryption for company projects. The company is currently conducting multiple projects to test the company's use of AWS KMS. These tests have led to a sudden increase in the company's AWS resource consumption. The test projects include applications that issue multiple requests each second to KMS endpoints for encryption activities.

The company needs to develop a solution that does not throttle the company's ability to use AWS KMS. The solution must improve key usage for client-side encryption and must be cost optimized. Which solution will meet these requirements?

- A. Use keyrings with the AWS Encryption SD
- B. Use each keyring individually or combine keyrings into a multi-keyring
- C. Decrypt the data by using a keyring that has the primary key in the multi-keyring.
- D. Use data key caching
- E. Use the local cache that the AWS Encryption SDK provides with a caching cryptographic materials manager.
- F. Use KMS key rotation
- G. Use a local cache in the AWS Encryption SDK with a caching cryptographic materials manager.
- H. Use keyrings with the AWS Encryption SD
- I. Use each keyring individually or combine keyrings into a multi-keyring
- J. Use any of the wrapping keys in the multi-keyring to decrypt the data.

Answer: B

Explanation:

The correct answer is B. Use data key caching. Use the local cache that the AWS Encryption SDK provides with a caching cryptographic materials manager. This answer is correct because data key caching can improve performance, reduce cost, and help the company stay within the service limits of AWS KMS. Data key caching stores data keys and related cryptographic material in a cache, and reuses them for encryption and decryption operations. This reduces the number of requests to AWS KMS endpoints and avoids throttling. The AWS Encryption SDK provides a local cache and a caching cryptographic materials manager (caching CMM) that interacts with the cache and enforces security thresholds that the company can set.

The other options are incorrect because:

- > A. Using keyrings with the AWS Encryption SDK does not address the problem of throttling or cost optimization. Keyrings are used to generate, encrypt, and decrypt data keys, but they do not cache or reuse them. Using each keyring individually or combining them into a multi-keyring does not reduce the number of requests to AWS KMS endpoints.
- > C. Using KMS key rotation does not address the problem of throttling or cost optimization. Key rotation is a security practice that creates new cryptographic material for a KMS key every year, but it does not affect the data that the KMS key protects. Key rotation does not reduce the number of requests to AWS KMS endpoints, and it might incur additional costs for storing multiple versions of key material.
- > D. Using keyrings with the AWS Encryption SDK does not address the problem of throttling or cost optimization, as explained in option A. Moreover, using any of the wrapping keys in the multi-keyring to decrypt the data is not a valid option, because only one of the wrapping keys can decrypt a given data key. The wrapping key that encrypts a data key is stored in the encrypted data key structure, and only that wrapping key can decrypt it.

References:

- 1: Data key caching - AWS Encryption SDK
- 2: Using keyrings - AWS Encryption SDK
- 3: Rotating AWS KMS keys - AWS Key Management Service
- 4: How keyrings work - AWS Encryption SDK

NEW QUESTION 9

A corporation is preparing to acquire several companies. A Security Engineer must design a solution to ensure that newly acquired IAM accounts follow the corporation's security best practices. The solution should monitor each Amazon S3 bucket for unrestricted public write access and use IAM managed services. What should the Security Engineer do to meet these requirements?

- A. Configure Amazon Macie to continuously check the configuration of all S3 buckets.
- B. Enable IAM Config to check the configuration of each S3 bucket.
- C. Set up IAM Systems Manager to monitor S3 bucket policies for public write access.
- D. Configure an Amazon EC2 instance to have an IAM role and a cron job that checks the status of all S3 buckets.

Answer: C

Explanation:

because this is a solution that can monitor each S3 bucket for unrestricted public write access and use IAM managed services. S3 is a service that provides object storage in the cloud. Systems Manager is a service that helps you automate and manage your AWS resources. You can use Systems Manager to monitor S3 bucket policies for public write access by using a State Manager association that runs a predefined document called AWS-FindS3BucketWithPublicWriteAccess. This document checks each S3 bucket in an account and reports any bucket that has public write access enabled. The other options are either not suitable or not feasible for meeting the requirements.

NEW QUESTION 10

A company's Chief Security Officer has requested that a Security Analyst review and improve the security posture of each company IAM account. The Security Analyst decides to do this by improving IAM account root user security.

Which actions should the Security Analyst take to meet these requirements? (Select THREE.)

- A. Delete the access keys for the account root user in every account.
- B. Create an admin IAM user with administrative privileges and delete the account root user in every account.
- C. Implement a strong password to help protect account-level access to the IAM Management Console by the account root user.
- D. Enable multi-factor authentication (MFA) on every account root user in all accounts.
- E. Create a custom IAM policy to limit permissions to required actions for the account root user and attach the policy to the account root user.
- F. Attach an IAM role to the account root user to make use of the automated credential rotation in IAM STS.

Answer: ADE

Explanation:

because these are the actions that can improve IAM account root user security. IAM account root user is a user that has complete access to all AWS resources and services in an account. IAM account root user security is a set of best practices that help protect the account root user from unauthorized or accidental use. Deleting the access keys for the account root user in every account can help prevent programmatic access by the account root user, which reduces the risk of compromise or misuse. Enabling MFA on every account root user in all accounts can help add an extra layer of security for console access by requiring a verification code in addition to a password. Creating a custom IAM policy to limit permissions to required actions for the account root user and attaching the policy to the account root user can help enforce the principle of least privilege and restrict the account root user from performing unnecessary or dangerous actions. The other options are either invalid or ineffective for improving IAM account root user security.

NEW QUESTION 10

An Incident Response team is investigating an IAM access key leak that resulted in Amazon EC2 instances being launched. The company did not discover the incident until many months later. The Director of Information Security wants to implement new controls that will alert when similar incidents happen in the future. Which controls should the company implement to achieve this? (Select TWO.)

- A. Enable VPC Flow Logs in all VPCs Create a scheduled IAM Lambda function that downloads and parses the logs, and sends an Amazon SNS notification for violations.
- B. Use IAM CloudTrail to make a trail, and apply it to all Regions Specify an Amazon S3 bucket to receive all the CloudTrail log files
- C. Add the following bucket policy to the company's IAM CloudTrail bucket to prevent log tampering{"Version": "2012-10-17-","Statement": { "Effect": "Deny","Action": "s3:PutObject", "Principal": "-", "Resource": "arn:IAM:s3:::cloudtrail/IAMLogs/111122223333/*"}}Create an Amazon S3 data event for an PutObject attempts, which sends notifications to an Amazon SNS topic.
- D. Create a Security Auditor role with permissions to access Amazon CloudWatch Logs m all Regions Ship the logs to an Amazon S3 bucket and make a lifecycle policy to ship the logs to Amazon S3 Glacier.
- E. Verify that Amazon GuardDuty is enabled in all Regions, and create an Amazon CloudWatch Events rule for Amazon GuardDuty findings Add an Amazon SNS topic as the rule's target

Answer: AE

NEW QUESTION 15

A company is designing a multi-account structure for its development teams. The company is using AWS Organizations and AWS Single Sign-On (AWS SSO). The company must implement a solution so that the development teams can use only specific AWS Regions and so that each AWS account allows access to only specific AWS services.

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

- A. Use AWS SSO to set up service-linked roles with IAM policy statements that include the Condition, Resource, and NotAction elements to allow access to only the Regions and services that are needed.
- B. Deactivate AWS Security Token Service (AWS STS) in Regions that the developers are not allowed to use.
- C. Create SCPs that include the Condition, Resource, and NotAction elements to allow access to only the Regions and services that are needed.
- D. For each AWS account, create tailored identity-based policies for AWS SS
- E. Use statements that include the Condition, Resource, and NotAction elements to allow access to only the Regions and services that are needed.

Answer: C

Explanation:

https://docs.aws.amazon.com/organizations/latest/userguide/orgs_manage_policies_scps_syntax.html#scp-eleme

NEW QUESTION 17

A company has a legacy application that runs on a single Amazon EC2 instance. A security audit shows that the application has been using an IAM access key within its code to access an Amazon S3 bucket that is named DOC-EXAMPLE-BUCKET1 in the same AWS account. This access key pair has the s3:GetObject permission to all objects in only this S3 bucket. The company takes the application offline because the application is not compliant with the company's security policies for accessing other AWS resources from Amazon EC2.

A security engineer validates that AWS CloudTrail is turned on in all AWS Regions. CloudTrail is sending logs to an S3 bucket that is named DOC-EXAMPLE-BUCKET2. This S3 bucket is in the same AWS account as DOC-EXAMPLE-BUCKET1. However, CloudTrail has not been configured to send logs to Amazon CloudWatch Logs.

The company wants to know if any objects in DOC-EXAMPLE-BUCKET1 were accessed with the IAM access key in the past 60 days. If any objects were accessed, the company wants to know if any of the objects that are text files (.txt extension) contained personally identifiable information (PII).

Which combination of steps should the security engineer take to gather this information? (Choose two.)

- A. Configure Amazon Macie to identify any objects in DOC-EXAMPLE-BUCKET1 that contain PII and that were available to the access key.
- B. Use Amazon CloudWatch Logs Insights to identify any objects in DOC-EXAMPLE-BUCKET1 that contain PII and that were available to the access key.
- C. Use Amazon OpenSearch Service (Amazon Elasticsearch Service) to query the CloudTrail logs in DOC-EXAMPLE-BUCKET2 for API calls that used the access key to access an object that contained PII.
- D. Use Amazon Athena to query the CloudTrail logs in DOC-EXAMPLE-BUCKET2 for any API calls that used the access key to access an object that contained PII.
- E. Use AWS Identity and Access Management Access Analyzer to identify any API calls that used the access key to access objects that contained PII in DOC-EXAMPLE-BUCKET1.

Answer: AD

NEW QUESTION 22

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 27

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 IAM 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 32

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 34

A security engineer must troubleshoot an administrator's inability to make an existing Amazon S3 bucket public in an account that is part of an organization n IAM Organizations. The administrator switched the role from the master account to a member account and then attempted to make one S3 bucket public. This action was immediately denied

Which actions should the security engineer take to troubleshoot the permissions issue? (Select TWO.)

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

Answer: DE

Explanation:

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

NEW QUESTION 37

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 41

A company used a lift-and-shift approach to migrate from its on-premises data centers to the AWS Cloud. The company migrated on-premises VMS to Amazon EC2 in-stances. Now the company wants to replace some of components that are running on the EC2 instances with managed AWS services that provide similar functionality.

Initially, the company will transition from load balancer software that runs on EC2 instances to AWS Elastic Load Balancers. A security engineer must ensure that after this transition, all the load balancer logs are centralized and searchable for auditing. The security engineer must also ensure that metrics are generated to show which ciphers are in use.

Which solution will meet these requirements?

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

Answer: C

Explanation:

- Amazon S3 is a service that provides scalable, durable, and secure object storage. You can use Amazon S3 to store and retrieve any amount of data from anywhere on the web1

- AWS Elastic Load Balancing is a service that distributes incoming application or network traffic across multiple targets, such as EC2 instances, containers, or IP addresses. You can use Elastic Load Balancing to increase the availability and fault tolerance of your applications2

- Elastic Load Balancing supports access logging, which captures detailed information about requests sent to your load balancer. Each log contains information such as the time the request was received, the client's IP address, latencies, request paths, and server responses. You can use access logs to analyze traffic patterns and troubleshoot issues3

- You can configure your load balancer to store access logs in an Amazon S3 bucket that you specify.

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

- Amazon Athena is a service that allows you to analyze data in Amazon S3 using standard SQL. You can use Athena to run ad-hoc queries and get results in seconds. Athena is serverless, so there is no infrastructure to manage and you pay only for the queries that you run.

- You can use Athena to search the access logs that are stored in your S3 bucket. You can create a table in Athena that maps to your S3 bucket and then run SQL queries on the table. You can also use the Athena console or API to view and download the query results.

- You can also use Athena to create queries for the required metrics, such as the number of requests per cipher or protocol. You can then publish the metrics to Amazon CloudWatch, which is a service that monitors and manages your AWS resources and applications. You can use CloudWatch to collect and track metrics, create alarms, and automate actions based on the state of your resources.

- By using this solution, you can meet the requirements of ensuring that all the load balancer logs are centralized and searchable for auditing and that metrics are generated to show which ciphers are in use.

NEW QUESTION 42

A company that uses AWS Organizations wants to see AWS Security Hub findings for many AWS accounts and AWS Regions. Some of the accounts are in the company's organization, and some accounts are in organizations that the company manages for customers. Although the company can see findings in the Security Hub administrator account for accounts in the company's organization, there are no findings from accounts in other organizations.

Which combination of steps should the company take to see findings from accounts that are outside the organization that includes the Security Hub administrator

account? (Select TWO.)

- A. Use a designated administration account to automatically set up member accounts.
- B. Create the AWS Service Role ForSecurity Hub service-linked role for Security Hub.
- C. Send an administration request from the member accounts.
- D. Enable Security Hub for all member accounts.
- E. Send invitations to accounts that are outside the company's organization from the Security Hub administrator account.

Answer: CE

Explanation:

To see Security Hub findings for accounts that are outside the organization that includes the Security Hub administrator account, the following steps are required:

➤ Send invitations to accounts that are outside the company's organization from the Security Hub administrator account. This will allow the administrator account to view and manage findings from those accounts. The administrator account can send invitations by using the Security Hub console, API, or CLI. For more information, see [Sending invitations to member accounts](#).

➤ Send an administration request from the member accounts. This will allow the member accounts to accept the invitation from the administrator account and establish a relationship with it. The member accounts can send administration requests by using the Security Hub console, API, or CLI. For more information, see [Sending administration requests](#).

This solution will enable the company to see Security Hub findings for many AWS accounts and AWS Regions, including accounts that are outside its own organization.

The other options are incorrect because they either do not establish a relationship between the administrator and member accounts (A, B), do not enable Security Hub for all member accounts (D), or do not use a valid service for Security Hub (F).

Verified References:

➤ <https://docs.aws.amazon.com/securityhub/latest/userguide/securityhub-member-accounts.html>

NEW QUESTION 45

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 ensure 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 mode
- 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 process
- E. Place objects in the S3 buckets.
- F. Create new S3 buckets with S3 Object Lock enabled in governance mode
- G. Place objects in the S3 buckets.
- H. Create new S3 buckets with S3 Object Lock enabled in governance mode
- I. Add a legal hold to the S3 bucket
- J. Place objects in the S3 buckets.

Answer: A

NEW QUESTION 49

An international company wants to combine AWS Security Hub findings across all the company's AWS Regions and from multiple accounts. In addition, the company

wants to create a centralized custom dashboard to correlate these findings with operational data for deeper analysis and insights. The company needs an analytics tool to search and visualize Security Hub findings. Which combination of steps will meet these requirements? (Select THREE.)

- A. Designate an AWS account as a delegated administrator for Security Hub
- B. Publish events to Amazon CloudWatch from the delegated administrator account, all member accounts, and required Regions that are enabled for Security Hub findings.
- C. Designate an AWS account in an organization in AWS Organizations as a delegated administrator for Security Hub
- D. Publish events to Amazon EventBridge from the delegated administrator account, all member accounts, and required Regions that are enabled for Security Hub findings.
- E. In each Region, create an Amazon EventBridge rule to deliver findings to an Amazon Kinesis data stream
- F. Configure the Kinesis data streams to output the logs to a single Amazon S3 bucket.
- G. In each Region, create an Amazon EventBridge rule to deliver findings to an Amazon Kinesis Data Firehose delivery stream
- H. Configure the Kinesis Data Firehose delivery streams to deliver the logs to a single Amazon S3 bucket.
- I. Use AWS Glue DataBrew to crawl the Amazon S3 bucket and build the schema
- J. Use AWS Glue Data Catalog to query the data and create views to flatten nested attributes
- K. Build Amazon QuickSight dashboards by using Amazon Athena.
- L. Partition the Amazon S3 data
- M. Use AWS Glue to crawl the S3 bucket and build the schema
- N. Use Amazon Athena to query the data and create views to flatten nested attributes
- O. Build Amazon QuickSight dashboards that use the Athena views.

Answer: BDF

Explanation:

The correct answer is B, D, and F. Designate an AWS account in an organization in AWS Organizations as a delegated administrator for Security Hub. Publish events to Amazon EventBridge from the delegated administrator account, all member accounts, and required Regions that are enabled for Security Hub findings. In each Region, create an Amazon EventBridge rule to deliver findings to an Amazon Kinesis Data Firehose delivery stream. Configure the Kinesis Data Firehose delivery streams to deliver the logs to a single Amazon S3 bucket. Partition the Amazon S3 data. Use AWS Glue to crawl the S3 bucket and build the schema. Use Amazon Athena to query the data and create views to flatten nested attributes. Build Amazon QuickSight dashboards that use the Athena views.

According to the AWS documentation, AWS Security Hub is a service that provides you with a comprehensive view of your security state across your AWS accounts, and helps you check your environment against security standards and best practices. You can use Security Hub to aggregate security findings from various sources, such as AWS services, partner products, or your own applications.

To use Security Hub with multiple AWS accounts and Regions, you need to enable AWS Organizations with all features enabled. This allows you to centrally manage your accounts and apply policies across your organization. You can also use Security Hub as a service principal for AWS Organizations, which lets you designate a delegated administrator account for Security Hub. The delegated administrator account can enable Security Hub automatically in all existing and future accounts in your organization, and can view and manage findings from all accounts.

According to the AWS documentation, Amazon EventBridge is a serverless event bus that makes it easy to connect applications using data from your own applications, integrated software as a service (SaaS) applications, and AWS services. You can use EventBridge to create rules that match events from various sources and route them to targets for processing.

To use EventBridge with Security Hub findings, you need to enable Security Hub as an event source in EventBridge. This will allow you to publish events from Security Hub to EventBridge in the same Region. You can then create EventBridge rules that match Security Hub findings based on criteria such as severity, type, or resource. You can also specify targets for your rules, such as Lambda functions, SNS topics, or Kinesis Data Firehose delivery streams.

According to the AWS documentation, Amazon Kinesis Data Firehose is a fully managed service that delivers real-time streaming data to destinations such as Amazon S3, Amazon Redshift, Amazon Elasticsearch Service (Amazon ES), and Splunk. You can use Kinesis Data Firehose to transform and enrich your data before delivering it to your destination.

To use Kinesis Data Firehose with Security Hub findings, you need to create a Kinesis Data Firehose delivery stream in each Region where you have enabled Security Hub. You can then configure the delivery stream to receive events from EventBridge as a source, and deliver the logs to a single S3 bucket as a destination. You can also enable data transformation or compression on the delivery stream if needed.

According to the AWS documentation, Amazon S3 is an object storage service that offers scalability, data availability, security, and performance. You can use S3 to store and retrieve any amount of data from anywhere on the web. You can also use S3 features such as lifecycle management, encryption, versioning, and replication to optimize your storage.

To use S3 with Security Hub findings, you need to create an S3 bucket that will store the logs from Kinesis Data Firehose delivery streams. You can then partition the data in the bucket by using prefixes such as account ID or Region. This will improve the performance and cost-effectiveness of querying the data.

According to the AWS documentation, AWS Glue is a fully managed extract, transform, and load (ETL) service that makes it easy to prepare and load your data for analytics. You can use Glue to crawl your data sources, identify data formats, and suggest schemas and transformations. You can also use Glue Data Catalog as a central metadata repository for your data assets.

To use Glue with Security Hub findings, you need to create a Glue crawler that will crawl the S3 bucket and build the schema for the data. The crawler will create tables in the Glue Data Catalog that you can query using standard SQL.

According to the AWS documentation, Amazon Athena is an interactive query service that makes it easy to analyze data in Amazon S3 using standard SQL. Athena is serverless, so there is no infrastructure to manage, and you pay only for the queries that you run. You can use Athena with Glue Data Catalog as a metadata store for your tables.

To use Athena with Security Hub findings, you need to create views in Athena that will flatten nested attributes in the data. For example, you can create views that extract fields such as account ID, Region, resource type, resource ID, finding type, finding title, and finding description from the JSON data. You can then query the views using SQL and join them with other tables if needed.

According to the AWS documentation, Amazon QuickSight is a fast, cloud-powered business intelligence service that makes it easy to deliver insights to everyone in your organization. You can use QuickSight to create and publish interactive dashboards that include machine learning insights. You can also use QuickSight to connect to various data sources, such as Athena, S3, or RDS.

To use QuickSight with Security Hub findings, you need to create QuickSight dashboards that use the Athena views as data sources. You can then visualize and analyze the findings using charts, graphs, maps, or tables. You can also apply filters, calculations, or aggregations to the data. You can then share the dashboards with your users or embed them in your applications.

NEW QUESTION 51

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

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

Answer: ACD

NEW QUESTION 56

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.amazonaws.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.amazonaws.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 58

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

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

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

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

Answer: C

Explanation:

The correct answer is C. Change the Resource element to "arn:aws:s3:::DOC-EXAMPLE-BUCKET/*".

The reason is that the Resource element in the bucket policy specifies which objects in the bucket are affected by the policy. In this case, the policy only applies to the bucket itself, not the objects inside it. Therefore, the Lambda function cannot access the objects with the s3:GetObject permission. To fix this, the Resource element should include a wildcard (*) to match all objects in the bucket. This way, the policy grants the Lambda function permission to read any object in the bucket.

The other options are incorrect for the following reasons:

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

NEW QUESTION 62

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 Manager AWS Management Console : AWS Identity and Access Management - AWS Management Console : Amazon Elastic Compute Cloud - Amazon Web Services : Amazon Linux 2 - Amazon Web Services : AWS Systems Manager - AWS Management Console : AWS Systems Manager - AWS Management Console : AWS Systems Manager - AWS Management Console

NEW QUESTION 66

A company deployed Amazon GuardDuty in the us-east-1 Region. The company wants all DNS logs that relate to the company's Amazon EC2 instances to be inspected. What should a security engineer do to ensure that the EC2 instances are logged?

- A. Use IPv6 addresses that are configured for hostnames.
- B. Configure external DNS resolvers as internal resolvers that are visible only to IAM.
- C. Use IAM DNS resolvers for all EC2 instances.

D. Configure a third-party DNS resolver with logging for all EC2 instances.

Answer: C

Explanation:

To ensure that the EC2 instances are logged, the security engineer should do the following:

- Use AWS DNS resolvers for all EC2 instances. This allows the security engineer to use Amazon-provided DNS servers that resolve public DNS hostnames to private IP addresses within their VPC, and that log DNS queries in Amazon CloudWatch Logs.

NEW QUESTION 71

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 73

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 security 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 trail
- 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 75

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 Store
- B. In the template, replace all references to the value with `{{resolve:ssm:MySSMParameterName:}}`.
- C. Store the API key value in AWS Secrets Manager
- D. In the template, replace all references to the value with `{{resolve:secretsmanager:MySecretId:SecretString}}`.
- E. Store the API key value in Amazon DynamoDB
- 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 bucket
- H. In the template, replace all references to the value with `{{resolve:s3:MyBucketName:MyObjectName}}`.

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 76

A security engineer is defining the controls required to protect the IAM account root user credentials in an IAM Organizations hierarchy. The controls should also limit the impact in case these credentials have been compromised.

Which combination of controls should the security engineer propose? (Select THREE.)

A)

Apply the following SCP:

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "GRRESTRICTROOTUSER",
      "Effect": "Deny",
      "Action": "*",
      "Resource": [
        "*"
      ],
      "Condition": {
        "StringLike": {
          "aws:PrincipalArn": [
            "arn:aws:iam::*:root"
          ]
        }
      }
    }
  ]
}
```

B)

Apply the following SCP:

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "GRRESTRICTROOTUSER",
      "Effect": "Deny",
      "Principal": "arn:aws:iam::*:root",
      "Action": "*",
      "Resource": [
        "*"
      ]
    }
  ]
}
```

- C) Enable multi-factor authentication (MFA) for the root user.
- D) Set a strong randomized password and store it in a secure location.
- E) Create an access key ID and secret access key, and store them in a secure location.
- F) Apply the following permissions boundary to the root user:

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "GRRESTRICTROOTUSER",
      "Effect": "Deny",
      "Action": "*",
      "Resource": [
        "*"
      ],
      "Condition": {
        "StringLike": {
          "aws:PrincipalArn": [
            "arn:aws:iam::*:root"
          ]
        }
      }
    }
  ]
}
```

- A. Option A
- B. Option B
- C. Option C
- D. Option D
- E. Option E
- F. Option F

Answer: ACE

NEW QUESTION 78

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 to 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 logs 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": "sns.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 80

A security engineer receives a notice from the AWS Abuse team about suspicious activity from a Linux-based Amazon EC2 instance that uses Amazon Elastic Block Store (Amazon EBS)-based storage. The instance is making connections to known malicious addresses. The instance is in a development account within a VPC that is in the us-east-1 Region. The VPC contains an internet gateway and has a subnet in us-east-1a and us-east-1b. Each subnet is associated with a route table that uses the internet gateway as a default route. Each subnet also uses the default network ACL. The suspicious EC2 instance runs within the us-east-1b subnet. During an initial investigation, a security engineer discovers that the suspicious instance is the only instance that runs in the subnet.

Which response will immediately mitigate the attack and help investigate the root cause?

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

Answer: B

Explanation:

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

NEW QUESTION 84

A company purchased a subscription to a third-party cloud security scanning solution that integrates with AWS Security Hub. A security engineer needs to implement a solution that will remediate the findings from the third-party scanning solution automatically. Which solution will meet this requirement?

- A. Set up an Amazon EventBridge rule that reacts to new Security Hub findings.
- B. Configure an AWS Lambda function as the target for the rule to remediate the findings.
- C. Set up a custom action in Security Hub.
- D. Configure the custom action to call AWS Systems Manager Automation runbooks to remediate the findings.
- E. Set up a custom action in Security Hub.
- F. Configure an AWS Lambda function as the target for the custom action to remediate the findings.
- G. Set up AWS Config rules to use AWS Systems Manager Automation runbooks to remediate the findings.

Answer: A

NEW QUESTION 86

A company's security engineer is developing an incident response plan to detect suspicious activity in an AWS account for VPC hosted resources. The security engineer needs to provide visibility for as many AWS Regions as possible.

Which combination of steps will meet these requirements MOST cost-effectively? (Select TWO.)

- A. Turn on VPC Flow Logs for all VPCs in the account.
- B. Activate Amazon GuardDuty across all AWS Regions.
- C. Activate Amazon Detective across all AWS Regions.
- D. Create an Amazon Simple Notification Service (Amazon SNS) topic
- E. Create an Amazon EventBridge rule that responds to findings and publishes the findings to the SNS topic.
- F. Create an AWS Lambda function
- G. Create an Amazon EventBridge rule that invokes the Lambda function to publish findings to Amazon Simple Email Service (Amazon SES).

Answer: BD

Explanation:

To detect suspicious activity in an AWS account for VPC hosted resources, the security engineer needs to use a service that can monitor network traffic and API calls across all AWS Regions. Amazon GuardDuty is a threat detection service that can do this by analyzing VPC Flow Logs, AWS CloudTrail event logs, and DNS logs. By activating GuardDuty across all AWS Regions, the security engineer can provide visibility for as many regions as possible. GuardDuty generates findings that contain details about the potential threats detected in the account. To respond to these findings, the security engineer needs to create a mechanism that can notify the relevant stakeholders or take remedial actions. One way to do this is to use Amazon EventBridge, which is a serverless event bus service that can connect AWS services and third-party applications. By creating an EventBridge rule that responds to GuardDuty findings and publishes them to an Amazon Simple Notification Service (Amazon SNS) topic, the security engineer can enable subscribers of the topic to receive notifications via email, SMS, or other methods. This is a cost-effective solution that does not require any additional infrastructure or code.

NEW QUESTION 88

A business requires a forensic logging solution for hundreds of Docker-based apps running on Amazon EC2. The solution must analyze logs in real time, provide message replay, and persist logs.

Which Amazon Web Offerings (IAM) services should be employed to satisfy these requirements? (Select two.)

- A. Amazon Athena
- B. Amazon Kinesis
- C. Amazon SQS
- D. Amazon Elasticsearch
- E. Amazon EMR

Answer: BD

NEW QUESTION 92

What are the MOST secure ways to protect the AWS account root user of a recently opened AWS account? (Select TWO.)

- A. Use the AWS account root user access keys instead of the AWS Management Console.
- B. Enable multi-factor authentication for the AWS IAM users with the AdministratorAccess managed policy attached to them.
- C. Enable multi-factor authentication for the AWS account root user.
- D. Use AWS KMS to encrypt all AWS account root user and AWS IAM access keys and set automatic rotation to 30 days.
- E. Do not create access keys for the AWS account root user; instead, create AWS IAM users.

Answer: CE

NEW QUESTION 95

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 secret
- B. Update the IAM principals in the role trust policy as required.
- C. Deploy a VPC endpoint for Secrets Manager
- D. Create and attach an endpoint policy that specifies the IAM principals that are allowed to access the secret
- E. Update the list of IAM principals as required.
- F. Use a tag-based approach by attaching a resource policy to the secret
- 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 explicitly
- J. Attach the policies to an IAM group
- K. Add all IAM principals to the IAM group
- L. Remove principals from the group when they need access
- M. Add the principals to the group again when access is no longer allowed.

Answer: C

NEW QUESTION 97

A company is building an application on IAM that will store sensitive information. The company has a support team with access to the IT infrastructure, including databases. The company's security engineer must introduce measures to protect the sensitive data against any data breach while minimizing management overhead. The credentials must be regularly rotated.

What should the security engineer recommend?

- A. Enable Amazon RDS encryption to encrypt the database and snapshot
- B. Enable Amazon Elastic Block Store (Amazon EBS) encryption on Amazon EC2 instance
- C. Include the database credential in the EC2 user data field

- D. Use an IAM Lambda function to rotate database credential
- E. Set up TLS for the connection to the database.
- F. Install a database on an Amazon EC2 Instance
- G. Enable third-party disk encryption to encrypt the Amazon Elastic Block Store (Amazon EBS) volume
- H. Store the database credentials in IAM CloudHSM with automatic rotation
- I. Set up TLS for the connection to the database.
- J. Enable Amazon RDS encryption to encrypt the database and snapshot
- K. Enable Amazon Elastic Block Store (Amazon EBS) encryption on Amazon EC2 instance
- L. Store the database credentials in IAM Secrets Manager with automatic rotation
- M. Set up TLS for the connection to the RDS hosted database.
- N. Set up an IAM CloudHSM cluster with IAM Key Management Service (IAM KMS) to store KMS keys. Set up Amazon RDS encryption using IAM KMS to encrypt the database
- O. Store database credentials in the IAM Systems Manager Parameter Store with automatic rotation
- P. Set up TLS for the connection to the RDS hosted database.

Answer: C

Explanation:

To protect the sensitive data against any data breach and minimize management overhead, the security engineer should recommend the following solution:

- Enable Amazon RDS encryption to encrypt the database and snapshots. This allows the security engineer to use AWS Key Management Service (AWS KMS) to encrypt data at rest for the database and any backups or replicas.
- Enable Amazon Elastic Block Store (Amazon EBS) encryption on Amazon EC2 instances. This allows the security engineer to use AWS KMS to encrypt data at rest for the EC2 instances and any snapshots or volumes.
- Store the database credentials in AWS Secrets Manager with automatic rotation. This allows the security engineer to encrypt and manage secrets centrally, and to configure automatic rotation schedules for them.
- Set up TLS for the connection to the RDS hosted database. This allows the security engineer to encrypt data in transit between the EC2 instances and the database.

NEW QUESTION 101

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 role
- 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 account
- I. Create an identity policy that allows the sts: AssumeRole action
- 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 104

A security engineer needs to set up an Amazon CloudFront distribution for an Amazon S3 bucket that hosts a static website. The security engineer must allow only specified IP addresses to access the website. The security engineer also must prevent users from accessing the website directly by using S3 URLs.

Which solution will meet these requirements?

- A. Generate an S3 bucket policy
- B. Specify cloudfront.amazonaws.com as the principal
- C. Use the aws:SourceIp condition key to allow access only if the request comes from the specified IP addresses.
- D. Create a CloudFront origin access identity (OAI). Create the S3 bucket policy so that only the OAI has access
- E. Create an AWS WAF web ACL and add an IP set rule
- F. Associate the web ACL with the CloudFront distribution.
- G. Implement security groups to allow only the specified IP addresses access and to restrict S3 bucket access by using the CloudFront distribution.
- H. Create an S3 bucket access point to allow access from only the CloudFront distribution
- I. Create an AWS WAF web ACL and add an IP set rule
- J. Associate the web ACL with the CloudFront distribution.

Answer: B

NEW QUESTION 106

A company has hundreds of AWS accounts in an organization in AWS Organizations. The company operates out of a single AWS Region. The company has a dedicated security tooling AWS account in the organization. The security tooling account is configured as the organization's delegated administrator for Amazon GuardDuty and AWS Security Hub. The company has configured the environment to automatically enable GuardDuty and Security Hub for existing AWS accounts and new AWS accounts.

The company is performing control tests on specific GuardDuty findings to make sure that the company's security team can detect and respond to security events. The security team launched an Amazon EC2 instance and attempted to run DNS requests against a test domain, example.com, to generate a DNS finding. However, the GuardDuty finding was never created in the Security Hub delegated administrator account. Why was the finding was not created in the Security Hub delegated administrator account?

- A. VPC flow logs were not turned on for the VPC where the EC2 instance was launched.
- B. The VPC where the EC2 instance was launched had the DHCP option configured for a custom OpenDNS resolver.
- C. The GuardDuty integration with Security Hub was never activated in the AWS account where the finding was generated.
- D. Cross-Region aggregation in Security Hub was not configured.

Answer: C

Explanation:

The correct answer is C. The GuardDuty integration with Security Hub was never activated in the AWS account where the finding was generated. According to the AWS documentation¹, GuardDuty findings are automatically sent to Security Hub only if the GuardDuty integration with Security Hub is enabled in the same account and Region. This means that the security tooling account, which is the delegated administrator for both GuardDuty and Security Hub, must enable the GuardDuty integration with Security Hub in each member account and Region where GuardDuty is enabled. Otherwise, the findings from GuardDuty will not be visible in Security Hub.

The other options are incorrect because:

- VPC flow logs are not required for GuardDuty to generate DNS findings. GuardDuty uses VPC DNS logs, which are automatically enabled for all VPCs, to detect malicious or unauthorized DNS activity.
- The DHCP option configured for a custom OpenDNS resolver does not affect GuardDuty's ability to generate DNS findings. GuardDuty uses its own threat intelligence sources to identify malicious domains, regardless of the DNS resolver used by the EC2 instance.
- Cross-Region aggregation in Security Hub is not relevant for this scenario, because the company operates out of a single AWS Region. Cross-Region aggregation allows Security Hub to aggregate findings from multiple Regions into a single Region.

References:

1: Managing GuardDuty accounts with AWS Organizations : Amazon GuardDuty Findings : How Amazon GuardDuty Works : Cross-Region aggregation in AWS Security Hub

NEW QUESTION 108

A company is using Amazon Route 53 Resolver for its hybrid DNS infrastructure. The company has set up Route 53 Resolver forwarding rules for authoritative domains that are hosted on on-premises DNS servers.

A new security mandate requires the company to implement a solution to log and query DNS traffic that goes to the on-premises DNS servers. The logs must show details of the source IP address of the instance from which the query originated. The logs also must show the DNS name that was requested in Route 53 Resolver. Which solution will meet these requirements?

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

Answer: C

Explanation:

The correct answer is C. Configure Route 53 Resolver query logging on all relevant VPCs. Send the logs to Amazon CloudWatch Logs. Use CloudWatch Insights to run queries on the source IP address and DNS name.

According to the AWS documentation¹, Route 53 Resolver query logging lets you log the DNS queries that Route 53 Resolver handles for your VPCs. You can send the logs to CloudWatch Logs, Amazon S3, or Kinesis Data Firehose. The logs include information such as the following:

- The AWS Region where the VPC was created
- The ID of the VPC that the query originated from
- The IP address of the instance that the query originated from
- The instance ID of the resource that the query originated from
- The date and time that the query was first made
- The DNS name requested (such as prod.example.com)
- The DNS record type (such as A or AAAA)
- The DNS response code, such as NoError or ServFail
- The DNS response data, such as the IP address that is returned in response to the DNS query

You can use CloudWatch Insights to run queries on your log data and analyze the results using graphs and statistics². You can filter and aggregate the log data based on any field, and use operators and functions to perform calculations and transformations. For example, you can use CloudWatch Insights to find out how many queries were made for a specific domain name, or which instances made the most queries.

Therefore, this solution meets the requirements of logging and querying DNS traffic that goes to the on-premises DNS servers, showing details of the source IP address of the instance from which the query originated, and the DNS name that was requested in Route 53 Resolver.

The other options are incorrect because:

- > A. Using VPC Traffic Mirroring would not capture the DNS queries that go to the on-premises DNS servers, because Traffic Mirroring only copies network traffic from an elastic network interface of an EC2 instance to a target for analysis³. Traffic Mirroring does not include traffic that goes through a Route 53 Resolver outbound endpoint, which is used to forward queries to on-premises DNS servers⁴. Therefore, this solution would not meet the requirements.
- > B. Configuring VPC flow logs on all relevant VPCs would not capture the DNS name that was requested in Route 53 Resolver, because flow logs only record information about the IP traffic going to and from network interfaces in a VPC⁵. Flow logs do not include any information about the content or payload of a packet, such as a DNS query or response. Therefore, this solution would not meet the requirements.
- > D. Modifying the Route 53 Resolver rules on the authoritative domains that forward to the on-premises DNS servers would not enable logging of DNS queries, because Resolver rules only specify how to forward queries for specified domain names to your network⁶. Resolver rules do not have any logging functionality by themselves. Therefore, this solution would not meet the requirements. References:
1: Resolver query logging - Amazon Route 53 2: Analyzing log data with CloudWatch Logs Insights - Amazon CloudWatch 3: What is Traffic Mirroring? - Amazon Virtual Private Cloud 4: Outbound Resolver endpoints - Amazon Route 53 5: Logging IP traffic using VPC Flow Logs - Amazon Virtual Private Cloud 6: Managing forwarding rules - Amazon Route 53

NEW QUESTION 110

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

Please select:

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

Answer: A

Explanation:

IAM Direct Connect makes it easy to establish a dedicated network connection from your premises to IAM. Using IAM Direct Connect you can establish private connectivity between IAM and your datacenter, office, or colocation environment which in many cases can reduce your network costs, increase bandwidth throughput and provide a more consistent network experience than Internet-based connections.

Options B and C are invalid because these options will not reduce network latency Options D is invalid because this is only used to connect 2 VPC's

For more information on IAM direct connect, just browse to the below URL: <https://IAM.amazon.com/directconnect>

The correct answer is: Provision a Direct Connect connection to an IAM region using a Direct Connect partner. omit your Feedback/Queries to our Experts

NEW QUESTION 112

A company has multiple Amazon S3 buckets encrypted with customer-managed CMKs Due to regulatory requirements the keys must be rotated every year. The company's Security Engineer has enabled automatic key rotation for the CMKs; however the company wants to verify that the rotation has occurred. What should the Security Engineer do to accomplish this?

- A. Filter IAM CloudTrail logs for KeyRotaton events
- B. Monitor Amazon CloudWatchn Events for any IAM KMS CMK rotation events
- C. Using the IAM CL
- D. run the IAM kms get-key-rotation-status operation with the --key-id parameter to check the CMK rotation date
- E. Use Amazon Athena to query IAM CloudTrail logs saved in an S3 bucket to filter Generate New Key events

Answer: C

Explanation:

the aws kms get-key-rotation-status command returns a boolean value that indicates whether automatic rotation of the customer master key (CMK) is enabled¹. This command also shows the date and time when the CMK was last rotated². The other options are not valid ways to check the CMK rotation status.

NEW QUESTION 116

A security engineer needs to configure an Amazon S3 bucket policy to restrict access to an S3 bucket that is named DOC-EXAMPLE-BUCKET. The policy must allow access to only DOC-EXAMPLE-BUCKET from only the following endpoint: vpce-1a2b3c4d. The policy must deny all access to DOC-EXAMPLE-BUCKET if the specified endpoint is not used.

Which bucket policy statement meets these requirements?

- A. A computer code with black text Description automatically generated

```

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

```

- B. A computer code with black text Description automatically generated

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

C. A computer code with black text Description automatically generated

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

D. A computer code with black text Description automatically generated

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

Answer: B

Explanation:

<https://docs.aws.amazon.com/AmazonS3/latest/userguide/example-bucket-policies-vcpe-endpoint.html>

NEW QUESTION 119

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 encryption
- 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 encryption
- E. Configure S3 Versioning on the S3 bucket
- F. Configure S3 ObjectLock to use compliance mode with a retention period of 7 years.
- G. Configure S3 Versioning
- H. Configure S3 Intelligent-Tiering on the S3 bucket to move the documents to S3 Glacier Deep Archive storage
- I. Use S3 server-side encryption immediately
- J. Expire the objects after 7 years.
- K. Set up S3 Event Notifications and use S3 server-side encryption
- 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 124

A company has a web server in the AWS Cloud. The company will store the content for the web server in an Amazon S3 bucket. A security engineer must use an Amazon CloudFront distribution to speed up delivery of the content. None of the files can be publicly accessible from the S3 bucket directly.

Which solution will meet these requirements?

- A. Configure the permissions on the individual files in the S3 bucket so that only the CloudFront distribution has access to them.
- B. Create an origin access identity (OAI). Associate the OAI with the CloudFront distributio
- C. Configure the S3 bucket permissions so that only the OAI can access the files in the S3 bucket.
- D. Create an S3 role in AWS Identity and Access Management (IAM). Allow only the CloudFront distribution to assume the role to access the files in the S3 bucket.
- E. Create an S3 bucket policy that uses only the CloudFront distribution ID as the principal and the Amazon Resource Name (ARN) as the target.

Answer: B

NEW QUESTION 129

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

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

Answer: AD

NEW QUESTION 134

A security engineer is working with a company to design an ecommerce application. The application will run on Amazon EC2 instances that run in an Auto Scaling group behind an Application Load Balancer (ALB). The application will use an Amazon RDS DB instance for its database.

The only required connectivity from the internet is for HTTP and HTTPS traffic to the application. The application must communicate with an external payment provider that allows traffic only from a preconfigured allow list of IP addresses. The company must ensure that communications with the external payment provider are not interrupted as the environment scales.

Which combination of actions should the security engineer recommend to meet these requirements? (Select THREE.)

- A. Deploy a NAT gateway in each private subnet for every Availability Zone that is in use.
- B. Place the DB instance in a public subnet.
- C. Place the DB instance in a private subnet.
- D. Configure the Auto Scaling group to place the EC2 instances in a public subnet.
- E. Configure the Auto Scaling group to place the EC2 instances in a private subnet.
- F. Deploy the ALB in a private subnet.

Answer: ACE

NEW QUESTION 137

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

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

Answer: B

Explanation:

Option A is incorrect since you don't add a user to the IAM Role Option C is incorrect since you don't assign multiple users to a policy Option D is incorrect since this is not an ideal approach

An IAM group is used to collectively manage users who need the same set of permissions. By having groups, it becomes easier to manage permissions. So if you change the permissions on the group scale, it will affect all the users in that group

For more information on IAM Groups, just browse to the below URL: https://docs.IAM.amazon.com/IAM/latest/UserGuide/id_eroups.html

The correct answer is: Use the IAM groups and add users, based upon their role, to different groups and apply the policy to group

Submit your Feedback/Queries to our Experts

NEW QUESTION 138

A security engineer needs to create an IAM Key Management Service (IAM KMS) key that will be used to encrypt all data stored in a company's Amazon S3 Buckets in the us-west-1 Region. The key will use server-side encryption. Usage of the key must be limited to requests coming from Amazon S3 within the company's account.

Which statement in the KMS key policy will meet these requirements?

A)

```
{
  "Effect": "Allow",
  "Principal": {
    "AWS": "*"
  },
  "Action": [
    "kms:Encrypt",
    "kms:Decrypt",
    "kms:ReEncrypt*",
    "kms:GenerateDataKey*",
    "kms:DescribeKey"
  ],
  "Resource": "*",
  "Condition": {
    "StringEquals": {
      "kms:ViaService": "s3.us-west-1.amazonaws.com",
      "kms:CallerAccount": "<CustomerAccountID>"
    }
  }
}
```

B)

```
{
  "Effect": "Allow",
  "Principal": {
    "AWS": "s3.us-west-1.amazonaws.com"
  },
  "Action": [
    "kms:Encrypt",
    "kms:Decrypt",
    "kms:ReEncrypt*",
    "kms:GenerateDataKey*",
    "kms:DescribeKey"
  ],
  "Resource": "*",
  "Condition": {
    "StringEquals": {
      "kms:CallerAccount": "<CustomerAccountID>"
    }
  }
}
```

C)

```
{
  "Effect": "Allow",
  "Principal": {
    "AWS": "*"
  },
  "Action": [
    "kms:Encrypt",
    "kms:Decrypt",
    "kms:ReEncrypt*",
    "kms:GenerateDataKey*",
    "kms:DescribeKey"
  ],
  "Resource": "*",
  "Condition": {
    "StringEquals": {
      "kms:EncryptionContext:aws:s3:arn": [
        "arn:aws:s3:::"
      ]
    }
  }
}
```

- A. Option A
- B. Option B
- C. Option C

Answer: A

NEW QUESTION 141

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 143

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 IAM 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 Region
- 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 144

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

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

Answer: A

Explanation:

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

NEW QUESTION 146

A company has a guideline that mandates the encryption of all Amazon S3 bucket data in transit. A security engineer must implement an S3 bucket policy that denies any S3 operations if data is not encrypted.
Which S3 bucket policy will meet this requirement?

```
A. {
    "Version": "2012-10-17",
    "Statement": [{
        "Sid": "AllowSSLRequestOnly",
        "Action": "s3:*",
        "Effect": "Deny",
        "Resource": [
            "arn:aws:s3:::DOC-EXAMPLE-BUCKET",
            "arn:aws:s3:::DOC-EXAMPLE-BUCKET/*"
        ],
        "Condition": {
            "Bool": {
                "aws:SecureTransport": "true"
            }
        }
    }],
    "Principal": "*"
}
```

B.

```
{
  "Version": "2012-10-17",
  "Statement": [{
    "Sid": "AllowSSLRequestOnly",
    "Action": "s3:*",
    "Effect": "Deny",
    "Resource": [
      "arn:aws:s3:::DOC-EXAMPLE-BUCKET",
      "arn:aws:s3:::DOC-EXAMPLE-BUCKET/*"
    ],
    "Condition": {
      "Bool": {
        "aws:SecureTransport": "false"
      }
    },
    "Principal": "*"
  }]
}

C. {
  "Version": "2012-10-17",
  "Statement": [{
    "Sid": "AllowSSLRequestOnly",
    "Action": "s3:*",
    "Effect": "Deny",
    "Resource": [
      "arn:aws:s3:::DOC-EXAMPLE-BUCKET",
      "arn:aws:s3:::DOC-EXAMPLE-BUCKET/*"
    ],
    "Condition": {
      "StringNotEquals": {
        "s3:x-amz-server-side-encryption": "AES256"
      }
    },
    "Principal": "*"
  }]
}
```

D. A screenshot of a computer code Description automatically generated {

```
{
  "Version": "2012-10-17",
  "Statement": [{
    "Sid": "AllowSSLRequestOnly",
    "Action": "s3:*",
    "Effect": "Deny",
    "Resource": [
      "arn:aws:s3:::DOC-EXAMPLE-BUCKET",
      "arn:aws:s3:::DOC-EXAMPLE-BUCKET/*"
    ],
    "Condition": {
      "StringNotEquals": {
        "s3:x-amz-server-side-encryption": true
      }
    },
    "Principal": "*"
  }]
}
```

Answer: B

Explanation:

<https://aws.amazon.com/blogs/security/how-to-use-bucket-policies-and-apply-defense-in-depth-to-help-secure-y>

NEW QUESTION 150

A company is evaluating the use of AWS Systems Manager Session Manager to gain access to the company's Amazon EC2 instances. However, until the company implements the change, the company must protect the key file for the EC2 instances from read and write operations by any other users.

When a security administrator tries to connect to a critical EC2 Linux instance during an emergency, the security administrator receives the following error. "Error Unprotected private key file - Permissions for 'ssh/my_private_key.pem' are too open".

Which command should the security administrator use to modify the private key file permissions to resolve this error?

- A. `chmod 0040 ssh/my_private_key.pem`
- B. `chmod 0400 ssh/my_private_key.pem`
- C. `chmod 0004 ssh/my_private_key.pem`
- D. `chmod 0777 ssh/my_private_key.pem`

Answer: B

Explanation:

The error message indicates that the private key file permissions are too open, meaning that other users can read or write to the file. This is a security risk, as the private key should be accessible only by the owner of the file. To fix this error, the security administrator should use the `chmod` command to change the permissions of the private key file to `0400`, which means that only the owner can read the file and no one else can read or write to it.

The `chmod` command takes a numeric argument that represents the permissions for the owner, group, and others in octal notation. Each digit corresponds to a set

of permissions: read (4), write (2), and execute (1). The digits are added together to get the final permissions for each category. For example, 0400 means that the owner has read permission (4) and no other permissions (0), and the group and others have no permissions at all (0).

The other options are incorrect because they either do not change the permissions at all (D), or they give too much or too little permissions to the owner, group, or others (A, C).

Verified References:

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

NEW QUESTION 154

A company uses several AWS CloudFormation stacks to handle the deployment of a suite of applications. The leader of the company's application development team notices that the stack deployments fail with permission errors when some team members try to deploy the stacks. However, other team members can deploy the stacks successfully.

The team members access the account by assuming a role that has a specific set of permissions that are necessary for the job responsibilities of the team members. All team members have permissions to perform operations on the stacks.

Which combination of steps will ensure consistent deployment of the stacks MOST securely? (Select THREE.)

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

Answer: BDF

NEW QUESTION 156

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

How should the security team securely store the API key?

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

Answer: C

Explanation:

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

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

NEW QUESTION 160

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 store.
- 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 store.
- 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 162

A Security Engineer has been tasked with enabling IAM Security Hub to monitor Amazon EC2 instances for CVE in a single IAM account. The Engineer has already enabled IAM Security Hub and Amazon Inspector in the IAM Management Console and has installed the Amazon Inspector agent on an EC2 instance that needs to be monitored.

Which additional steps should the Security Engineer take to meet this requirement?

- A. Configure the Amazon Inspector agent to use the CVE rule package
- B. Configure the Amazon Inspector agent to use the CVE rule package. Configure Security Hub to ingest from IAM Inspector by writing a custom resource policy
- C. Configure the Security Hub agent to use the CVE rule package. Configure IAM Inspector to ingest from Security Hub by writing a custom resource policy
- D. Configure the Amazon Inspector agent to use the CVE rule package. Install an additional Integration library. Allow the Amazon Inspector agent to communicate with Security Hub

Answer: D

Explanation:

You need to configure the Amazon Inspector agent to use the CVE rule package, which is a set of rules that check for vulnerabilities and exposures on your EC2 instances⁵. You also need to install an additional integration library that enables communication between the Amazon Inspector agent and Security Hub⁶. Security Hub is a service that provides you with a comprehensive view of your security state in AWS and helps you check your environment against security industry standards and best practices⁷. The other options are either incorrect or incomplete for meeting the requirement.

NEW QUESTION 163

A security engineer is designing a cloud architecture to support an application. The application runs on Amazon EC2 instances and processes sensitive information, including credit card numbers.

The application will send the credit card numbers to a component that is running in an isolated environment. The component will encrypt, store, and decrypt the numbers.

The component then will issue tokens to replace the numbers in other parts of the application.

The component of the application that manages the tokenization process will be deployed on a separate set of EC2 instances. Other components of the application must not be able to store or access the credit card numbers.

Which solution will meet these requirements?

- A. Use EC2 Dedicated Instances for the tokenization component of the application.
- B. Place the EC2 instances that manage the tokenization process into a partition placement group.
- C. Create a separate VPC
- D. Deploy new EC2 instances into the separate VPC to support the data tokenization.
- E. Deploy the tokenization code onto AWS Nitro Enclaves that are hosted on EC2 instances.

Answer: D

Explanation:

AWS Nitro Enclaves are isolated and hardened virtual machines that run on EC2 instances and provide a secure environment for processing sensitive data. Nitro Enclaves have no persistent storage, interactive access, or external networking, and they can only communicate with the parent instance through a secure local channel. Nitro Enclaves also support cryptographic attestation, which allows verifying the identity and integrity of the enclave and its code. Nitro Enclaves are ideal for implementing data protection solutions such as tokenization, encryption, and key management.

Using Nitro Enclaves for the tokenization component of the application meets the requirements of isolating the sensitive data from other parts of the application, encrypting and storing the credit card numbers securely, and issuing tokens to replace the numbers. Other components of the application will not be able to access or store the credit card numbers, as they are only available within the enclave.

NEW QUESTION 168

A security engineer wants to evaluate configuration changes to a specific AWS resource to ensure that the resource meets compliance standards. However, the security engineer is concerned about a situation in which several configuration changes are made to the resource in quick succession. The security engineer wants to record only the latest configuration of that resource to indicate the cumulative impact of the set of changes.

Which solution will meet this requirement in the MOST operationally efficient way?

- A. Use AWS CloudTrail to detect the configuration changes by filtering API calls to monitor the changes. Use the most recent API call to indicate the cumulative impact of multiple calls
- B. Use AWS Config to detect the configuration changes and to record the latest configuration in case of multiple configuration changes.
- C. Use Amazon CloudWatch to detect the configuration changes by filtering API calls to monitor the change
- D. Use the most recent API call to indicate the cumulative impact of multiple calls.
- E. Use AWS Cloud Map to detect the configuration change
- F. Generate a report of configuration changes from AWS Cloud Map to track the latest state by using a sliding time window.

Answer: B

Explanation:

AWS Config is a service that enables you to assess, audit, and evaluate the configurations of your AWS resources. AWS Config continuously monitors and records your AWS resource configurations and allows you to automate the evaluation of recorded configurations against desired configurations.

To evaluate configuration changes to a specific AWS resource and ensure that it meets compliance standards, the security engineer should use AWS Config to detect the configuration changes and to record the latest configuration in case of multiple configuration changes. This will allow the security engineer to view the current state of the resource and its compliance status, as well as its configuration history and timeline.

AWS Config records configuration changes as ConfigurationItems, which are point-in-time snapshots of the resource's attributes, relationships, and metadata. If multiple configuration changes occur within a short period of time, AWS Config records only the latest ConfigurationItem for that resource. This indicates the cumulative impact of the set of changes on the resource's configuration.

This solution will meet the requirement in the most operationally efficient way, as it leverages AWS Config's features to monitor, record, and evaluate resource configurations without requiring additional tools or services.

The other options are incorrect because they either do not record the latest configuration in case of multiple configuration changes (A, C), or do not use a valid service for evaluating resource configurations (D).

Verified References:

➤ <https://docs.aws.amazon.com/config/latest/developerguide/WhatIsConfig.html>

➤ <https://docs.aws.amazon.com/config/latest/developerguide/config-item-table.html>

NEW QUESTION 173

An organization must establish the ability to delete an IAM KMS Customer Master Key (CMK) within a 24- hour timeframe to keep it from being used for encrypt or decrypt operations Which of the following actions will address this requirement?

- A. Manually rotate a key within KMS to create a new CMK immediately
- B. Use the KMS import key functionality to execute a delete key operation
- C. Use the schedule key deletion function within KMS to specify the minimum wait period for deletion
- D. Change the KMS CMK alias to immediately prevent any services from using the CMK.

Answer: C

Explanation:

the schedule key deletion function within KMS allows you to specify a waiting period before deleting a customer master key (CMK)⁴. The minimum waiting period is 7 days and the maximum is 30 days⁵. This function prevents the CMK from being used for encryption or decryption operations during the waiting period⁴. The other options are either invalid or ineffective for deleting a CMK within a 24-hour timeframe.

NEW QUESTION 177

.....

Thank You for Trying Our Product

* 100% Pass or Money Back

All our products come with a 90-day Money Back Guarantee.

* One year free update

You can enjoy free update one year. 24x7 online support.

* Trusted by Millions

We currently serve more than 30,000,000 customers.

* Shop Securely

All transactions are protected by VeriSign!

100% Pass Your SCS-C02 Exam with Our Prep Materials Via below:

<https://www.certleader.com/SCS-C02-dumps.html>