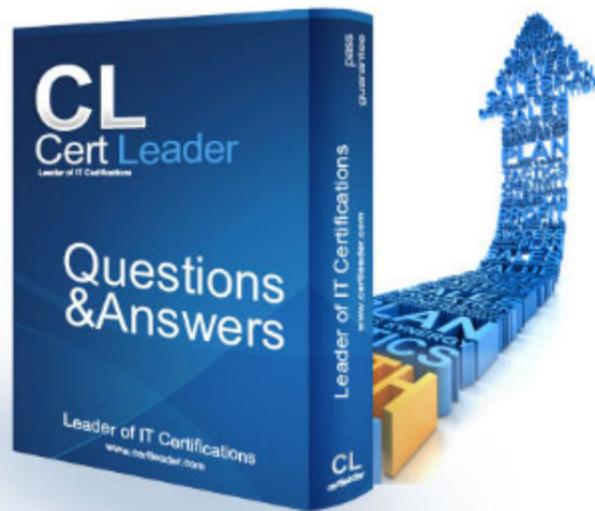


## SCS-C02 Dumps

### AWS Certified Security - Specialty

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



**NEW QUESTION 1**

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 2**

A company developed an application by using AWS Lambda, Amazon S3, Amazon Simple Notification Service (Amazon SNS), and Amazon DynamoDB. An external application puts objects into the company's S3 bucket and tags the objects with date and time. A Lambda function periodically pulls data from the company's S3 bucket based on date and time tags and inserts specific values into a DynamoDB table for further processing.

The data includes personally identifiable information (PII). The company must remove data that is older than 30 days from the S3 bucket and the DynamoDB table. Which solution will meet this requirement with the MOST operational efficiency?

- A. Update the Lambda function to add a TTL S3 flag to S3 object
- B. Create an S3 Lifecycle policy to expire objects that are older than 30 days by using the TTL S3 flag.
- C. Create an S3 Lifecycle policy to expire objects that are older than 30 day
- D. Update the Lambda function to add the TTL attribute in the DynamoDB tabl
- E. Enable TTL on the DynamoDB table to expire entires that are older than 30 days based on the TTL attribute.
- F. Create an S3 Lifecycle policy to expire objects that are older than 30 days and to add all prefixes to the S3 bucke
- G. Update the Lambda function to delete entries that are older than 30 days.
- H. Create an S3 Lifecycle policy to expire objects that are older than 30 days by using object tag
- I. Update the Lambda function to delete entries that are older than 30 days.

**Answer:** B

**NEW QUESTION 3**

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 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 4**

A company is implementing new compliance requirements to meet customer needs. According to the new requirements the company must not use any Amazon RDS DB instances or DB clusters that lack encryption of the underlying storage. The company needs a solution that will generate an email alert when an unencrypted DB instance or DB cluster is created. The solution also must terminate the unencrypted DB instance or DB cluster.

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

- A. Create an AWS Config managed rule to detect unencrypted ROS storag
- B. Configure an automatic remediation action to publish messages to an Amazon Simple Notification Service (Amazon SNS) topic that includes an AWS Lambda function and an email delivery target as subscriber
- C. Configure the Lambda function to delete the unencrypted resource.
- D. Create an AWS Config managed rule to detect unencrypted RDS storag
- E. Configure a manual remediation action to invoke an AWS Lambda functio
- F. Configure the Lambda function to publish messages to an Amazon Simple Notification Service (Amazon SNS) topic and to delete the unencrypted resource.
- G. Create an Amazon EventBridge rule that evaluates RDS event patterns and is initiated by the creation of DB instances or DB clusters Configure the rule to publish messages to an Amazon Simple Notification Service (Amazon SNS) topic that includes an AWS Lambda function and an email delivery target as subscriber
- H. Configure the Lambda function to delete the unencrypted resource.

- I. Create an Amazon EventBridge rule that evaluates RDS event patterns and is initiated by the creation of DB instances or DB cluster
- J. Configure the rule to invoke an AWS Lambda function
- K. Configure the Lambda function to publish messages to an Amazon Simple Notification Service (Amazon SNS) topic and to delete the unencrypted resource.

**Answer:** A

**Explanation:**

<https://docs.aws.amazon.com/config/latest/developerguide/rds-storage-encrypted.html>

**NEW QUESTION 5**

A security engineer wants to forward custom application-security logs from an Amazon EC2 instance to Amazon CloudWatch. The security engineer installs the CloudWatch agent on the EC2 instance and adds the path of the logs to the CloudWatch configuration file. However, CloudWatch does not receive the logs. The security engineer verifies that the awslogs service is running on the EC2 instance.

What should the security engineer do next to resolve the issue?

- A. Add AWS CloudTrail to the trust policy of the EC2 instance
- B. Send the custom logs to CloudTrail instead of CloudWatch.
- C. Add Amazon S3 to the trust policy of the EC2 instance
- D. Configure the application to write the custom logs to an S3 bucket that CloudWatch can use to ingest the logs.
- E. Add Amazon Inspector to the trust policy of the EC2 instance
- F. Use Amazon Inspector instead of the CloudWatch agent to collect the custom logs.
- G. Attach the CloudWatchAgentServerPolicy AWS managed policy to the EC2 instance role.

**Answer:** D

**Explanation:**

The correct answer is D. Attach the CloudWatchAgentServerPolicy AWS managed policy to the EC2 instance role.

According to the AWS documentation<sup>1</sup>, the CloudWatch agent is a software agent that you can install on your EC2 instances to collect system-level metrics and logs. To use the CloudWatch agent, you need to attach an IAM role or user to the EC2 instance that grants permissions for the agent to perform actions on your behalf. The CloudWatchAgentServerPolicy is an AWS managed policy that provides the necessary permissions for the agent to write metrics and logs to CloudWatch<sup>2</sup>. By attaching this policy to the EC2 instance role, the security engineer can resolve the issue of CloudWatch not receiving the custom application-security logs.

The other options are incorrect for the following reasons:

- > A. Adding AWS CloudTrail to the trust policy of the EC2 instance is not relevant, because CloudTrail is a service that records API activity in your AWS account, not custom application logs<sup>3</sup>. Sending the custom logs to CloudTrail instead of CloudWatch would not meet the requirement of forwarding them to CloudWatch.
- > B. Adding Amazon S3 to the trust policy of the EC2 instance is not necessary, because S3 is a storage service that does not require any trust relationship with EC2 instances<sup>4</sup>. Configuring the application to write the custom logs to an S3 bucket that CloudWatch can use to ingest the logs would be an alternative solution, but it would be more complex and costly than using the CloudWatch agent directly.
- > C. Adding Amazon Inspector to the trust policy of the EC2 instance is not helpful, because Inspector is a service that scans EC2 instances for software vulnerabilities and unintended network exposure, not custom application logs<sup>5</sup>. Using Amazon Inspector instead of the CloudWatch agent would not meet the requirement of forwarding them to CloudWatch.

References:

1: Collect metrics, logs, and traces with the CloudWatch agent - Amazon CloudWatch 2: CloudWatchAgentServerPolicy - AWS Managed Policy 3: What Is AWS CloudTrail? - AWS CloudTrail 4: Amazon S3 FAQs - Amazon Web Services 5: Automated Software Vulnerability Management - Amazon Inspector - AWS

**NEW QUESTION 6**

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 SDK
- 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 SDK
- 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<sup>1</sup>.

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<sup>2</sup>.
- > 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 material3.

➤ 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 it4.

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 7

A web application gives users the ability to log in verify their membership's validity and browse artifacts that are stored in an Amazon S3 bucket. When a user attempts to download an object, the application must verify the permission to access the object and allow the user to download the object from a custom domain name such as example.com.

What is the MOST secure way for a security engineer to implement this functionality?

- A. Configure read-only access to the object by using a bucket AC
- B. Remove the access after a set time has elapsed.
- C. Implement an IAM policy to give the user read access to the S3 bucket.
- D. Create an S3 presigned URL Provide the S3 presigned URL to the user through the application.
- E. Create an Amazon CloudFront signed UR
- F. Provide the CloudFront signed URL to the user through the application.

**Answer: D**

#### Explanation:

For this scenario you would need to set up static website hosting because a custom domain name is listed as a requirement. "Amazon S3 website endpoints do not support HTTPS or access points. If you want to use HTTPS, you can use Amazon CloudFront to serve a static website hosted on Amazon S3." This is not secure. <https://docs.aws.amazon.com/AmazonS3/latest/userguide/website-hosting-custom-domain-walkthrough.html> CloudFront signed URLs allow much more fine-grained control as well as HTTPS access with custom domain names:

<https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/private-content-signed-urls.html>

### NEW QUESTION 8

A company has AWS accounts in an organization in AWS Organizations. The organization includes a dedicated security account.

All AWS account activity across all member accounts must be logged and reported to the dedicated security account. The company must retain all the activity logs in a secure storage location within the dedicated security account for 2 years. No changes or deletions of the logs are allowed.

Which combination of steps will meet these requirements with the LEAST operational overhead? (Select TWO.)

- A. In the dedicated security account, create an Amazon S3 bucket
- B. Configure S3 Object Lock in compliance mode and a retention period of 2 years on the S3 bucket
- C. Set the bucket policy to allow the organization's management account to write to the S3 bucket.
- D. In the dedicated security account, create an Amazon S3 bucket
- E. Configure S3 Object Lock in compliance mode and a retention period of 2 years on the S3 bucket
- F. Set the bucket policy to allow the organization's member accounts to write to the S3 bucket.
- G. In the dedicated security account, create an Amazon S3 bucket that has an S3 Lifecycle configuration that expires objects after 2 year
- H. Set the bucket policy to allow the organization's member accounts to write to the S3 bucket.
- I. Create an AWS Cloud Trail trail for the organization
- J. Configure logs to be delivered to the logging Amazon S3 bucket in the dedicated security account.
- K. Turn on AWS CloudTrail in each account
- L. Configure logs to be delivered to an Amazon S3 bucket that is created in the organization's management account
- M. Forward the logs to the S3 bucket in the dedicated security account by using AWS Lambda and Amazon Kinesis Data Firehose.

**Answer: BD**

#### Explanation:

The correct answer is B and D. In the dedicated security account, create an Amazon S3 bucket. Configure S3 Object Lock in compliance mode and a retention period of 2 years on the S3 bucket. Set the bucket policy to allow the organization's member accounts to write to the S3 bucket. Create an AWS CloudTrail trail for the organization. Configure logs to be delivered to the logging Amazon S3 bucket in the dedicated security account.

According to the AWS documentation, AWS CloudTrail is a service that enables governance, compliance, operational auditing, and risk auditing of your AWS account. With CloudTrail, you can log, continuously monitor, and retain account activity related to actions across your AWS infrastructure. CloudTrail provides event history of your AWS account activity, including actions taken through the AWS Management Console, AWS SDKs, command line tools, and other AWS services.

To use CloudTrail 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 CloudTrail as a service principal for AWS Organizations, which lets you create an organization trail that applies to all accounts in your organization. An organization trail logs events for all AWS Regions and delivers the log files to an S3 bucket that you specify.

To create an organization trail, you need to use an administrator account, such as the organization's management account or a delegated administrator account. You can then configure the trail to deliver logs to an S3 bucket in the dedicated security account. This will ensure that all account activity across all member accounts and regions is logged and reported to the security account.

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 CloudTrail logs, you need to create an S3 bucket in the dedicated security account that will store the logs from the organization trail. You can then configure S3 Object Lock on the bucket to prevent objects from being deleted or overwritten for a fixed amount of time or indefinitely. You can also enable compliance mode on the bucket, which prevents any user, including the root user in your account, from deleting or modifying a locked object until it reaches its retention date.

To set a retention period of 2 years on the S3 bucket, you need to create a default retention configuration for the bucket that specifies a retention mode (either governance or compliance) and a retention period (either a number of days or a date). You can then set the bucket policy to allow the organization's member accounts to write to the S3 bucket. This will ensure that all logs are retained in a secure storage location within the security account for 2 years and no changes or deletions are allowed.

Option A is incorrect because setting the bucket policy to allow the organization's management account to write to the S3 bucket is not sufficient, as it will not grant access to the other member accounts in the organization.

Option C is incorrect because using an S3 Lifecycle configuration that expires objects after 2 years is not secure, as it will allow users to delete or modify objects before they expire.

Option E is incorrect because using Lambda and Kinesis Data Firehose to forward logs from one S3 bucket to another is not necessary, as CloudTrail can directly deliver logs to an S3 bucket in another account. It also introduces additional operational overhead and complexity.

#### NEW QUESTION 9

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 in 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 10

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

#### NEW QUESTION 11

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 13

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

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

**Answer:** D

#### Explanation:

The correct answer is D. The AWS Config rules have a MaximumExecutionFrequency value of 24 hours. According to the AWS documentation<sup>1</sup>, the MaximumExecutionFrequency parameter specifies the maximum frequency with which AWS Config runs evaluations for a rule. For AWS Config managed rules, this value can be one of the following:

- One\_Hour
- Three\_Hours
- Six\_Hours
- Twelve\_Hours
- TwentyFour\_Hours

If the rule is triggered by configuration changes, it will still run evaluations when AWS Config delivers the configuration snapshot. However, if the rule is triggered periodically, it will not run evaluations more often than the specified frequency.

In this case, the security engineer enabled four AWS Config managed rules that are triggered periodically. Therefore, these rules will only run evaluations every 24 hours, regardless of when the IAM credential report is generated. This means that the resources will display as noncompliant until the next evaluation cycle, which could take up to 24 hours after the IAM access keys are rotated.

The other options are incorrect because:

- A. The IAM credential report can be generated at any time, but it will not affect the compliance status of the resources until the next evaluation cycle of the AWS Config rules.
- B. The security engineer was able to invoke the IAM GenerateCredentialReport API operation, which means they have the GenerateCredentialReport permission. This permission is required to generate a credential report that lists all IAM users in an AWS account and their credential status<sup>2</sup>.
- C. The security engineer does not need the GetCredentialReport permission to enable or evaluate AWS Config rules. This permission is required to retrieve a credential report that was previously generated by using the GenerateCredentialReport operation<sup>2</sup>.

References:

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

### NEW QUESTION 15

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 DynamoDB
- 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 key
- D. Create a Lambda function that 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 EKS
- 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 EKS
- I. Create an AWS Config rule to send alerts for resources that are not encrypted with the key.

**Answer:** B

### NEW QUESTION 19

A company uses AWS Organizations to manage a multi-account AWS environment in a single AWS Region. The organization's management account is named management-01. The company has turned on AWS Config in all accounts in the organization. The company has designated an account named security-01 as the delegated administrator for AWS Config.

All accounts report the compliance status of each account's rules to the AWS Config delegated administrator account by using an AWS Config aggregator. Each account administrator can configure and manage the account's own AWS Config rules to handle each account's unique compliance requirements.

A security engineer needs to implement a solution to automatically deploy a set of 10 AWS Config rules to all existing and future AWS accounts in the organization. The solution must turn on AWS Config automatically during account creation.

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

- A. Create an AWS CloudFormation template that contains the 10 required AWS Config rules
- B. Deploy the template by using CloudFormation StackSets in the security-01 account.

- C. Create a conformance pack that contains the 10 required AWS Config rule
- D. Deploy the conformance pack from the security-01 account.
- E. Create a conformance pack that contains the 10 required AWS Config rule
- F. Deploy the conformance pack from the management-01 account.
- G. Create an AWS CloudFormation template that will activate AWS Confi
- H. De-ploy the template by using CloudFormation StackSets in the security-01 ac-count.
- I. Create an AWS CloudFormation template that will activate AWS Confi
- J. De-ploy the template by using CloudFormation StackSets in the management-01 account.

**Answer:** BE

**NEW QUESTION 21**

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 22**

A security engineer is configuring a mechanism to send an alert when three or more failed sign-in attempts to the AWS Management Console occur during a 5-minute period. The security engineer creates a trail in AWS CloudTrail to assist in this work.

Which solution will meet these requirements?

- A. In CloudTrail, turn on Insights events on the trai
- B. Configure an alarm on the insight with eventName matching ConsoleLogin and errorMessage matching "Failed authentication". Configure a threshold of 3 and a period of 5 minutes.
- C. Configure CloudTrail to send events to Amazon CloudWatch Log
- D. Create a metric filter for the relevant log grou
- E. Create a filter pattern with eventName matching ConsoleLogin and errorMessage matching "Failed authentication". Create a CloudWatch alarm with a threshold of 3 and a period of 5 minutes.
- F. Create an Amazon Athena table from the CloudTrail event
- G. Run a query for eventName matching ConsoleLogin and for errorMessage matching "Failed authentication". Create a notification action from the query to send an Amazon Simple Notification Service (Amazon SNS) notification when the count equals 3 within a period of 5 minutes.
- H. In AWS Identity and Access Management Access Analyzer, create a new analyze

I. Configure the analyzer to send an Amazon Simple Notification Service (Amazon SNS) notification when a failed sign-in event occurs 3 times for any IAM user within a period of 5 minutes.

**Answer:** B

**Explanation:**

The correct answer is B. Configure CloudTrail to send events to Amazon CloudWatch Logs. Create a metric filter for the relevant log group. Create a filter pattern with eventName matching ConsoleLogin and errorMessage matching "Failed authentication". Create a CloudWatch alarm with a threshold of 3 and a period of 5 minutes.

This answer is correct because it meets the requirements of sending an alert when three or more failed sign-in attempts to the AWS Management Console occur during a 5-minute period. By configuring CloudTrail to send events to CloudWatch Logs, the security engineer can create a metric filter that matches the desired pattern of failed sign-in events. Then, by creating a CloudWatch alarm based on the metric filter, the security engineer can set a threshold of 3 and a period of 5 minutes, and choose an action such as sending an email or an Amazon Simple Notification Service (Amazon SNS) message when the alarm is triggered<sup>12</sup>.

The other options are incorrect because:

- A. Turning on Insights events on the trail and configuring an alarm on the insight is not a solution, because Insights events are used to analyze unusual activity in management events, such as spikes in API call volume or error rates. Insights events do not capture failed sign-in attempts to the AWS Management Console<sup>3</sup>.
- C. Creating an Amazon Athena table from the CloudTrail events and running a query for failed sign-in events is not a solution, because it does not provide a mechanism to send an alert based on the query results. Amazon Athena is an interactive query service that allows analyzing data in Amazon S3 using standard SQL, but it does not support creating notifications or alarms from queries<sup>4</sup>.
- D. Creating an analyzer in AWS Identity and Access Management Access Analyzer and configuring it to send an Amazon SNS notification when a failed sign-in event occurs 3 times for any IAM user within a period of 5 minutes is not a solution, because IAM Access Analyzer is not a service that monitors sign-in events, but a service that helps identify resources that are shared with external entities. IAM Access Analyzer does not generate findings for failed sign-in attempts to the AWS Management Console<sup>5</sup>.

References:

1: Sending CloudTrail Events to CloudWatch Logs - AWS CloudTrail 2: Creating Alarms Based on Metric Filters - Amazon CloudWatch 3: Analyzing unusual activity in management events - AWS CloudTrail 4: What is Amazon Athena? - Amazon Athena 5: Using AWS Identity and Access Management Access Analyzer - AWS Identity and Access Management

**NEW QUESTION 25**

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 30**

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

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

**Answer:** B

**Explanation:**

The correct answer is B. Create an Amazon Simple Notification Service (Amazon SNS) topic. Subscribe the third-party ticketing email system to the SNS topic. Create an Amazon EventBridge rule that includes an event pattern that matches High severity GuardDuty findings. Specify the SNS topic as the target for the Event-Bridge rule.

According to the AWS documentation<sup>1</sup>, you can use Amazon EventBridge to create rules that match events from GuardDuty and route them to targets such as Amazon SNS topics. You can use event patterns to filter events based on criteria such as severity, type, or resource. For example, you can create a rule that matches only High severity findings and sends them to an SNS topic that is subscribed by a third-party ticketing email system. This way, you can automate the creation of tickets for High severity findings and notify the security team.

**NEW QUESTION 31**

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

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

**Answer:** BCE

#### NEW QUESTION 32

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 IAM 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 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. Use Systems Manager Inventory to select the EC2 instance and connect.
- B. Assign an IAM policy to the IAM 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 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. Use Systems Manager Session Manager to select the EC2 instance and connect.
- D. Assign an IAM policy to the IAM 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 35

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 are 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 39

A company uses AWS Organizations and has production workloads across multiple AWS accounts. A security engineer needs to design a solution that will proactively monitor for suspicious behavior across all the accounts that contain production workloads.

The solution must automate remediation of incidents across the production accounts. The solution also must publish a notification to an Amazon Simple Notification Service (Amazon SNS) topic when a critical security finding is detected. In addition, the solution must send all security incident logs to a dedicated account.

Which solution will meet these requirements?

- A. Activate Amazon GuardDuty in each production account.
- B. In a dedicated logging account.
- C. Aggregate all GuardDuty logs from each production account.
- D. Remediate incidents by configuring GuardDuty to directly invoke an AWS Lambda function.
- E. Configure the Lambda function to also publish notifications to the SNS topic.
- F. Activate AWS Security Hub in each production account.
- G. In a dedicated logging account.
- H. Aggregate all Security Hub findings from each production account.
- I. Remediate incidents by using AWS Config and AWS Systems Manager.

- J. Configure Systems Manager to also pub11Sh notifications to the SNS topic.
- K. Activate Amazon GuardDuty in each production account
- L. In a dedicated logging account
- M. aggregate all GuardDuty logs from each production account Remediate incidents by using Amazon EventBridge to invoke a custom AWS Lambda function from the GuardDuty finding
- N. Configure the Lambda function to also publish notifications to the SNS topic.
- O. Activate AWS Security Hub in each production account
- P. In a dedicated logging account
- Q. aggregate all Security Hub findings from each production account
- R. Remediate incidents by using Amazon EventBridge to invoke a custom AWS Lambda function from the Security Hub finding
- S. Configure the Lambda function to also publish notifications to the SNS topic.

**Answer: D**

**Explanation:**

The correct answer is D.

To design a solution that will proactively monitor for suspicious behavior across all the accounts that contain production workloads, the security engineer needs to use a service that can aggregate and analyze security findings from multiple sources. AWS Security Hub is a service that provides a comprehensive view of your security posture across your AWS accounts and enables you to check your environment against security standards and best practices. Security Hub also integrates with other AWS services, such as Amazon GuardDuty, AWS Config, and AWS Systems Manager, to collect and correlate security findings.

To automate remediation of incidents across the production accounts, the security engineer needs to use a service that can trigger actions based on events. Amazon EventBridge is a serverless event bus service that allows you to connect your applications with data from a variety of sources. EventBridge can use rules to match events and route them to targets for processing. You can use EventBridge to invoke a custom AWS Lambda function from the Security Hub findings. Lambda is a serverless compute service that lets you run code without provisioning or managing servers.

To publish a notification to an Amazon SNS topic when a critical security finding is detected, the security engineer needs to use a service that can send messages to subscribers. Amazon SNS is a fully managed messaging service that enables you to decouple and scale microservices, distributed systems, and serverless applications. SNS can deliver messages to a variety of endpoints, such as email, SMS, or HTTP. You can configure the Lambda function to also publish notifications to the SNS topic.

To send all security incident logs to a dedicated account, the security engineer needs to use a service that can aggregate and store log data from multiple sources. AWS Security Hub allows you to aggregate security findings from multiple accounts into a single account using the delegated administrator feature. This feature enables you to designate an AWS account as the administrator for Security Hub in an organization. The administrator account can then view and manage Security Hub findings from all member accounts.

Therefore, option D is correct because it meets all the requirements of the solution. Option A is incorrect because GuardDuty does not provide a comprehensive view of your security posture across your AWS accounts. GuardDuty is primarily a threat detection service that monitors for malicious or unauthorized behavior. Option B is incorrect because Config and Systems Manager are not designed to automate remediation of incidents based on Security Hub findings. Config is a service that enables you to assess, audit, and evaluate the configurations of your AWS resources, while Systems Manager is a service that allows you to manage your infrastructure on AWS at scale. Option C is incorrect because GuardDuty does not provide a comprehensive view of your security posture across your AWS accounts.

References:

- > AWS Security Hub
- > Amazon EventBridge
- > AWS Lambda
- > Amazon SNS
- > Aggregating Security Hub findings across accounts

**NEW QUESTION 41**

Within a VPC, a corporation runs an Amazon RDS Multi-AZ DB instance. The database instance is connected to the internet through a NAT gateway via two subnets.

Additionally, the organization has application servers that are hosted on Amazon EC2 instances and use the RDS database. These EC2 instances have been deployed onto two more private subnets inside the same VPC. These EC2 instances connect to the internet through a default route via the same NAT gateway. Each VPC subnet has its own route table.

The organization implemented a new security requirement after a recent security examination. Never allow the database instance to connect to the internet. A security engineer must perform this update promptly without interfering with the network traffic of the application servers. How will the security engineer be able to comply with these requirements?

- A. Remove the existing NAT gateway
- B. Create a new NAT gateway that only the application server subnets can use.
- C. Configure the DB instance's inbound network ACL to deny traffic from the security group ID of the NAT gateway.
- D. Modify the route tables of the DB instance subnets to remove the default route to the NAT gateway.
- E. Configure the route table of the NAT gateway to deny connections to the DB instance subnets.

**Answer: C**

**Explanation:**

Each subnet has a route table, so modify the routing associated with DB instance subnets to prevent internet access.

**NEW QUESTION 43**

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

What is the FASTEST way for the security engineer to identify the federated user?

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

H. Identify the corresponding role and run another query to filter the AssumeRoleWithWebIdentity event for the user name.

**Answer:** B

**Explanation:**

The fastest way to identify the federated user who terminated a production Amazon EC2 instance is to filter the IAM CloudTrail event history for the TerminateInstances event and identify the assumed IAM role. Then, review the AssumeRoleWithSAML event call in CloudTrail to identify the corresponding username. This method does not require any additional tools or queries, and it directly links the IAM role with the federated user.

Option A is incorrect because the role session name may not be the same as the federated user name, and it may not be unique or descriptive enough to identify the user.

Option C is incorrect because the IAM Access Advisor tab only shows when a role was last accessed, not by whom or for what purpose. It also does not show the specific time of access, only the date.

Option D is incorrect because using Amazon Athena to run SQL queries on the IAM CloudTrail logs is not the fastest way to identify the federated user, as it requires creating a table schema and running multiple queries. It also assumes that the federation is done using web identity providers, not SAML providers, as indicated by the AssumeRoleWithWebIdentity event.

References:

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

**NEW QUESTION 45**

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 documentation<sup>1</sup>, 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 command<sup>2</sup>. 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 46**

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 {

**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 lifecycle<sup>1</sup>. Secrets Manager integrates with AWS CloudFormation, which allows you to reference secrets from your templates using the {{resolve:secretsmanager:...}} syntax<sup>2</sup>. 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 Store<sup>3</sup>. 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 tables<sup>4</sup>. 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 buckets<sup>5</sup>. 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 49

A security engineer must use AWS Key Management Service (AWS KMS) to design a key management solution for a set of Amazon Elastic Block Store (Amazon EBS) volumes that contain sensitive data. The solution needs to ensure that the key material automatically expires in 90 days.

Which solution meets these criteria?

- A. A customer managed CMK that uses customer provided key material
- B. A customer managed CMK that uses AWS provided key material
- C. An AWS managed CMK
- D. Operation system-native encryption that uses GnuPG

**Answer:** A

#### Explanation:

```
https://awscli.amazonaws.com/v2/documentation/api/latest/reference/kms/import-key-material.html aws kms import-key-material \
--key-id 1234abcd-12ab-34cd-56ef-1234567890ab \
--encrypted-key-material fileb://EncryptedKeyMaterial.bin \
--import-token fileb://ImportToken.bin \
--expiration-model KEY_MATERIAL_EXPIRES \
--valid-to 2021-09-21T19:00:00Z
```

The correct answer is A. A customer managed CMK that uses customer provided key material.

A customer managed CMK is a KMS key that you create, own, and manage in your AWS account. You have full control over the key configuration, permissions, rotation, and deletion. You can use a customer managed CMK to encrypt and decrypt data in AWS services that are integrated with AWS KMS, such as Amazon EBS<sup>1</sup>.

A customer managed CMK can use either AWS provided key material or customer provided key material. AWS provided key material is generated by AWS KMS and never leaves the service unencrypted. Customer provided key material is generated outside of AWS KMS and imported into a customer managed CMK. You can specify an expiration date for the imported key material, after which the CMK becomes unusable until you reimport new key material<sup>2</sup>.

To meet the criteria of automatically expiring the key material in 90 days, you need to use customer provided key material and set the expiration date accordingly. This way, you can ensure that the data encrypted with the CMK will not be accessible after 90 days unless you reimport new key material and re-encrypt the data. The other options are incorrect for the following reasons:

\* B. A customer managed CMK that uses AWS provided key material does not expire automatically. You can enable automatic rotation of the key material every year, but this does not prevent access to the data encrypted with the previous key material. You would need to manually delete the CMK and its backing key material to make the data inaccessible<sup>3</sup>.

\* C. An AWS managed CMK is a KMS key that is created, owned, and managed by an AWS service on your behalf. You have limited control over the key configuration, permissions, rotation, and deletion. You cannot use an AWS managed CMK to encrypt data in other AWS services or applications. You also cannot set an expiration date for the key material of an AWS managed CMK<sup>4</sup>.

\* D. Operation system-native encryption that uses GnuPG is not a solution that uses AWS KMS. GnuPG is a command line tool that implements the OpenPGP standard for encrypting and signing data. It does not integrate with Amazon EBS or other AWS services. It also does not provide a way to automatically expire the key material used for encryption<sup>5</sup>.

References:

1: Customer Managed Keys - AWS Key Management Service 2: [Importing Key Material in AWS Key Management Service (AWS KMS) - AWS Key Management Service] 3: [Rotating Customer Master Keys - AWS Key Management Service] 4: [AWS Managed Keys - AWS Key Management Service] 5: The GNU Privacy Guard

#### NEW QUESTION 53

A Security Engineer is building a Java application that is running on Amazon EC2. The application communicates with an Amazon RDS instance and authenticates with a user name and password.

Which combination of steps can the Engineer take to protect the credentials and minimize downtime when the credentials are rotated? (Choose two.)

- A. Have a Database Administrator encrypt the credentials and store the ciphertext in Amazon S3. Grant permission to the instance role associated with the EC2 instance to read the object and decrypt the ciphertext.
- B. Configure a scheduled job that updates the credential in AWS Systems Manager Parameter Store and notifies the Engineer that the application needs to be restarted.
- C. Configure automatic rotation of credentials in AWS Secrets Manager.
- D. Store the credential in an encrypted string parameter in AWS Systems Manager Parameter Store
- E. Grant permission to the instance role associated with the EC2 instance to access the parameter and the AWS KMS key that is used to encrypt it.
- F. Configure the Java application to catch a connection failure and make a call to AWS Secrets Manager to retrieve updated credentials when the password is rotate
- G. Grant permission to the instance role associated with the EC2 instance to access Secrets Manager.

**Answer:** CE

#### Explanation:

AWS Secrets Manager is a service that helps you manage, retrieve, and rotate secrets such as database credentials, API keys, and other sensitive information. By configuring automatic rotation of credentials in AWS Secrets Manager, you can ensure that your secrets are changed regularly and securely, without requiring manual intervention or application downtime. You can also specify the rotation frequency and the rotation function that performs the logic of changing the credentials on the database and updating the secret in Secrets Manager<sup>1</sup>.

\* E. Configure the Java application to catch a connection failure and make a call to AWS Secrets Manager to retrieve updated credentials when the password is rotated. Grant permission to the instance role associated with the EC2 instance to access Secrets Manager.

By configuring the Java application to catch a connection failure and make a call to AWS Secrets Manager to retrieve updated credentials, you can avoid hard-coding the credentials in your application code or configuration files. This way, your application can dynamically obtain the latest credentials from Secrets Manager whenever the password is rotated, without needing to restart or redeploy the application. To enable this, you need to grant permission to the instance role associated with the EC2 instance to access Secrets Manager using IAM policies<sup>2</sup>. You can also use the AWS SDK for Java to integrate your application with

Secrets Manager3.

**NEW QUESTION 56**

A Systems Engineer is troubleshooting the connectivity of a test environment that includes a virtual security appliance deployed inline. In addition to using the virtual security appliance, the Development team wants to use security groups and network ACLs to accomplish various security requirements in the environment. What configuration is necessary to allow the virtual security appliance to route the traffic?

- A. Disable network ACLs.
- B. Configure the security appliance's elastic network interface for promiscuous mode.
- C. Disable the Network Source/Destination check on the security appliance's elastic network interface
- D. Place the security appliance in the public subnet with the internet gateway

**Answer: C**

**Explanation:**

<https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/using-eni.html#eni-basics> Source/destination checking "You must disable source/destination checks if the instance runs services such as network address translation, routing, or firewalls."

The correct answer is C. Disable the Network Source/Destination check on the security appliance's elastic network interface.

This answer is correct because disabling the Network Source/Destination check allows the virtual security appliance to route traffic that is not addressed to or from itself. By default, this check is enabled on all EC2 instances, and it prevents them from forwarding traffic that does not match their own IP or MAC addresses.

However, for a virtual security appliance that acts as a router or a firewall, this check needs to be disabled, otherwise it will drop the traffic that it is supposed to route12.

The other options are incorrect because:

- A. Disabling network ACLs is not a solution, because network ACLs are optional layers of security for the subnets in a VPC. They can be used to allow or deny traffic based on IP addresses and ports, but they do not affect the routing behavior of the virtual security appliance3.
- B. Configuring the security appliance's elastic network interface for promiscuous mode is not a solution, because promiscuous mode is a mode for a network interface that causes it to pass all traffic it receives to the CPU, rather than passing only the frames that it is programmed to receive. Promiscuous mode is normally used for packet sniffing or monitoring, but it does not enable the network interface to route traffic4.
- D. Placing the security appliance in the public subnet with the internet gateway is not a solution, because it does not address the routing issue of the virtual security appliance. The security appliance can be placed in either a public or a private subnet, depending on the network design and security requirements, but it still needs to have the Network Source/Destination check disabled to route traffic properly5.

References:

1: Enabling or disabling source/destination checks - Amazon Elastic Compute Cloud 2: Virtual security appliance - Wikipedia 3: Network ACLs - Amazon Virtual Private Cloud 4: Promiscuous mode - Wikipedia 5: NAT instances - Amazon Virtual Private Cloud

**NEW QUESTION 60**

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 65**

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

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

**Answer: C**

**NEW QUESTION 66**

A company is using AWS Organizations to manage multiple accounts. The company needs to allow an IAM user to use a role to access resources that are in another organization's AWS account.

Which combination of steps must the company perform to meet this requirement? (Select TWO.)

- A. Create an identity policy that allows the sts: AssumeRole action in the AWS account that contains the resource
- B. Attach the identity policy to the IAM user.
- C. Ensure that the sts: AssumeRole action is allowed by the SCPs of the organization that owns the resources that the IAM user needs to access.
- D. Create a role in the AWS account that contains the resource

- E. Create an entry in the role's trust policy that allows the IAM user to assume the rol
- F. Attach the trust policy to the role.
- G. Establish a trust relationship between the IAM user and the AWS account that contains the resources.
- H. Create a role in the IAM user's AWS account
- I. Create an identity policy that allows the sts: AssumeRole actio
- J. Attach the identity policy to the role.

**Answer:** BC

**Explanation:**

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

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

Verified References:

- <https://repost.aws/knowledge-center/cross-account-access-iam>
- [https://docs.aws.amazon.com/organizations/latest/userguide/orgs\\_manage\\_accounts\\_access.html](https://docs.aws.amazon.com/organizations/latest/userguide/orgs_manage_accounts_access.html)
- [https://docs.aws.amazon.com/IAM/latest/UserGuide/tutorial\\_cross-account-with-roles.html](https://docs.aws.amazon.com/IAM/latest/UserGuide/tutorial_cross-account-with-roles.html)

**NEW QUESTION 67**

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

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

**Answer:** BE

**Explanation:**

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

**NEW QUESTION 72**

A security engineer needs to see 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 polic
- B. Specify cloudfront amazonaws com as the principa
- 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 acces
- E. Create an AWS WAF web ACL and add an IP set rul
- 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 distributio
- I. Create an AWS WAF web ACL and add an IP set rul
- J. Associate the web ACL with the CloudFront distribution.

**Answer:** B

**NEW QUESTION 73**

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

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

**Answer: D**

**Explanation:**

You can establish a private connection between your VPC and Secrets Manager by creating an interface VPC endpoint. Interface endpoints are powered by AWS PrivateLink, a technology that enables you to privately access Secrets Manager APIs without an internet gateway, NAT device, VPN connection, or AWS Direct Connect connection. Reference:

<https://docs.aws.amazon.com/secretsmanager/latest/userguide/vpc-endpoint-overview.html>

The correct answer is D. Configure a Secrets Manager interface VPC endpoint. Include the Lambda function's private subnet during the configuration process. A Secrets Manager interface VPC endpoint is a private connection between the VPC and Secrets Manager that does not require an internet gateway, NAT device, VPN connection, or AWS Direct Connect connection<sup>1</sup>. By configuring a Secrets Manager interface VPC endpoint, the security engineer can enable the custom Lambda function to communicate with Secrets Manager without sending or receiving network traffic through the internet. The security engineer must include the Lambda function's private subnet during the configuration process to allow the function to use the endpoint<sup>2</sup>.

The other options are incorrect for the following reasons:

- A. An egress-only internet gateway is a VPC component that allows outbound communication over IPv6 from instances in the VPC to the internet, and prevents the internet from initiating an IPv6 connection with the instances<sup>3</sup>. However, this option does not meet the requirement that the VPC must not send or receive network traffic through the internet. Moreover, an egress-only internet gateway is for use with IPv6 traffic only, and Secrets Manager does not support IPv6 addresses<sup>2</sup>.
- B. A NAT gateway is a VPC component that enables instances in a private subnet to connect to the internet or other AWS services, but prevents the internet from initiating connections with those instances<sup>4</sup>. However, this option does not meet the requirement that the VPC must not send or receive network traffic through the internet. Additionally, a NAT gateway requires an elastic IP address, which is a public IPv4 address<sup>4</sup>.
- C. A VPC peering connection is a networking connection between two VPCs that enables you to route traffic between them using private IPv4 addresses or IPv6 addresses<sup>5</sup>. However, this option does not work because Secrets Manager does not have a default VPC that can be peered with. Furthermore, a VPC peering connection does not provide a private connection to Secrets Manager APIs without an internet gateway or other devices<sup>2</sup>.

**NEW QUESTION 74**

During a manual review of system logs from an Amazon Linux EC2 instance, a Security Engineer noticed that there are sudo commands that were never properly alerted or reported on the Amazon CloudWatch Logs agent. Why were there no alerts on the sudo commands?

- A. There is a security group blocking outbound port 80 traffic that is preventing the agent from sending the logs
- B. The IAM instance profile on the EC2 instance was not properly configured to allow the CloudWatchLogs agent to push the logs to CloudWatch
- C. CloudWatch Logs status is set to ON versus SECURE, which prevents it from pulling in OS security event logs
- D. The VPC requires that all traffic go through a proxy, and the CloudWatch Logs agent does not support a proxy configuration.

**Answer: B**

**Explanation:**

the reason why there were no alerts on the sudo commands. Sudo commands are commands that allow a user to execute commands as another user, usually the superuser or root. CloudWatch Logs agent is a software agent that can send log data from an EC2 instance to CloudWatch Logs, a service that monitors and stores log data. The CloudWatch Logs agent needs an IAM instance profile, which is a container for an IAM role that allows applications running on an EC2 instance to make API requests to AWS services. If the IAM instance profile on the EC2 instance was not properly configured to allow the CloudWatch Logs agent to push the logs to CloudWatch, then there would be no alerts on the sudo commands. The other options are either irrelevant or invalid for explaining why there were no alerts on the sudo commands.

**NEW QUESTION 77**

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 CloudWatch 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 enabled<sup>1</sup>. This command also shows the date and time when the CMK was last rotated<sup>2</sup>. The other options are not valid ways to check the CMK rotation status.

**NEW QUESTION 82**

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

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

```

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

```

B. A computer code with black text Description automatically generated

```

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

```

C. A computer code with black text Description automatically generated

```

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

```

D. A computer code with black text Description automatically generated

```

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

```

**Answer:** B

**Explanation:**

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

**NEW QUESTION 83**

A company has enabled Amazon GuardDuty in all AWS Regions as part of its security monitoring strategy. In one of its VPCs, the company hosts an Amazon EC2 instance that works as an FTP server. A high number of clients from multiple locations contact the FTP server. GuardDuty identifies this activity as a brute force attack because of the high number of connections that happen every hour.

The company has flagged the finding as a false positive, but GuardDuty continues to raise the issue. A security engineer must improve the signal-to-noise ratio without compromising the company's visibility of potential anomalous behavior.

Which solution will meet these requirements?

- A. Disable the FTP rule in GuardDuty in the Region where the FTP server is deployed.
- B. Add the FTP server to a trusted IP list

- C. Deploy the list to GuardDuty to stop receiving the notifications.
- D. Create a suppression rule in GuardDuty to filter findings by automatically archiving new findings that match the specified criteria.
- E. Create an AWS Lambda function that has the appropriate permissions to delete the finding whenever a new occurrence is reported.

**Answer:** C

**Explanation:**

"When you create an Amazon GuardDuty filter, you choose specific filter criteria, name the filter and can enable the auto-archiving of findings that the filter matches. This allows you to further tune GuardDuty to your unique environment, without degrading the ability to identify threats. With auto-archive set, all findings are still generated by GuardDuty, so you have a complete and immutable history of all suspicious activity."

**NEW QUESTION 84**

A company uses Amazon API Gateway to present REST APIs to users. An API developer wants to analyze API access patterns without the need to parse the log files.

Which combination of steps will meet these requirements with the LEAST effort? (Select TWO.)

- A. Configure access logging for the required API stage.
- B. Configure an AWS CloudTrail trail destination for API Gateway event
- C. Configure filters on the userIdentity, userAgent, and sourceIPAddress fields.
- D. Configure an Amazon S3 destination for API Gateway log
- E. Run Amazon Athena queries to analyze API access information.
- F. Use Amazon CloudWatch Logs Insights to analyze API access information.
- G. Select the Enable Detailed CloudWatch Metrics option on the required API stage.

**Answer:** CD

**NEW QUESTION 88**

A company is building a data processing application that uses AWS Lambda functions. The application's Lambda functions need to communicate with an Amazon RDS DB instance that is deployed within a VPC in the same AWS account.

Which solution meets these requirements in the MOST secure way?

- A. Configure the DB instance to allow public access. Update the DB instance security group to allow access from the Lambda public address space for the AWS Region.
- B. Deploy the Lambda functions inside the VPC. Attach a network ACL to the Lambda subnet. Provide outbound rule access to the VPC CIDR range only. Update the DB instance security group to allow traffic from 0.0.0.0/0.
- C. Deploy the Lambda functions inside the VPC. Attach a security group to the Lambda functions. Provide outbound rule access to the VPC CIDR range only. Update the DB instance security group to allow traffic from the Lambda security group.
- D. Peer the Lambda default VPC with the VPC that hosts the DB instance to allow direct network access without the need for security groups.

**Answer:** C

**Explanation:**

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

References: : AWS Lambda Developer Guide

**NEW QUESTION 90**

A company needs a forensic-logging solution for hundreds of applications running in Docker on Amazon EC2. The solution must perform real-time analytics on the logs. The solution must support the replay of messages and must persist the logs.

Which IAM services should be used to meet these requirements? (Select TWO.)

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

**Answer:** BD

**Explanation:**

Amazon Kinesis and Amazon Elasticsearch are both suitable for forensic-logging solutions. Amazon Kinesis can collect, process, and analyze streaming data in real time. Amazon Elasticsearch can store, search, and analyze log data using the popular open-source tool Elasticsearch. The other options are not designed for forensic-logging purposes. Amazon Athena is a query service that can analyze data in S3, Amazon SQS is a message queue service that can decouple and scale microservices, and Amazon EMR is a big data platform that can run Apache Spark and Hadoop clusters.

**NEW QUESTION 95**

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 99**

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

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

**Answer: C**

**Explanation:**

To meet the requirements of securing access management and implementing a centralized logging solution, the most secure solution would be to:

- Install a bastion host in the management account.
- Reconfigure all SSH and RDP to allow access only from the bastion host.
- Install AWS Systems Manager Agent (SSM Agent) on the bastion host.
- Attach the AmazonSSMManagedInstanceCore role to the bastion host.
- Configure session data streaming to Amazon CloudWatch Logs in a separate logging account to audit log data.

This solution provides the following security benefits:

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

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

**NEW QUESTION 101**

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

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

**Answer: C**

**Explanation:**

To prevent unauthorized access to the EC2 instances, the security engineer should do the following:

- Restrict SSH access in the security group to only known corporate IP addresses. This allows the security engineer to use a virtual firewall that controls inbound and outbound traffic for their EC2 instances, and limit SSH access to only trusted sources.

**NEW QUESTION 102**

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](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 106**

A company is using IAM Organizations. The company wants to restrict IAM usage to the eu-west-1 Region for all accounts under an OU that is named "development." The solution must persist restrictions to existing and new IAM accounts under the development OU.

- A. Include the following SCP on the development OU:

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "DenyNonDefaultRegions",
      "Effect": "Deny",
      "NotAction": [
        <Desired Global Services> ],
      "Resource": "*",
      "Condition": {
        "StringNotEquals": {
          "aws:RequestedRegion": [
            "eu-west-1"
          ]
        }
      },
      "ArnNotLike": {
        "aws:PrincipalARN": "arn:aws:iam::*:role/AWSExecution"
      }
    }
  ]
}
```

- B. Include the following SCP on the development account:

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "DenyNonDefaultRegions",
      "Effect": "Deny",
      "NotAction": [
        <Desired Global Services> ],
      "Resource": "*",
      "Condition": {
        "StringNotEquals": {
          "aws:RequestedRegion": [
            "eu-west-1"
          ]
        }
      },
      "ArnNotLike": {
        "aws:PrincipalARN": "arn:aws:iam::*:role/AWSExecution"
      }
    }
  ]
}
```

C. Include the following SCP on the development OU

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "DenyNonDefaultRegions",
      "Effect": "Deny",
      "NotAction": [
        <Desired Global Services> ],
      "Resource": "*",
      "Condition": {
        "StringEquals": {
          "aws:RequestedRegion": [
            "eu-west-1"
          ]
        }
      },
      "ArnNotLike": {
        "aws:PrincipalARN": "arn:aws:iam::*:role/AWSExecution"
      }
    }
  ]
}
```

D. Include the following SCP on the development OU

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "DenyNonDefaultRegions",
      "Effect": "Allow",
      "NotAction": [
        <Desired Global Services> ],
      "Resource": "*",
      "Condition": {
        "StringNotEquals": {
          "aws:RequestedRegion": [
            "us-east-1"
          ]
        }
      },
      "ArnNotLike": {
        "aws:PrincipalARN": "arn:aws:iam::*:role/AWSExecution"
      }
    }
  ]
}
```

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

**Answer:** A

**NEW QUESTION 109**

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

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

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

**Answer:** A

**Explanation:**

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

**NEW QUESTION 110**

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 112**

A business stores website images in an Amazon S3 bucket. The firm serves the photos to end users through Amazon CloudFront. The firm learned lately that the photographs are being accessible from nations in which it does not have a distribution license.

Which steps should the business take to safeguard the photographs and restrict their distribution? (Select two.)

- A. Update the S3 bucket policy to restrict access to a CloudFront origin access identity (OAI).
- B. Update the website DNS record to use an Amazon Route 53 geolocation record deny list of countries where the company lacks a license.
- C. Add a CloudFront geo restriction deny list of countries where the company lacks a license.
- D. Update the S3 bucket policy with a deny list of countries where the company lacks a license.
- E. Enable the Restrict Viewer Access option in CloudFront to create a deny list of countries where the company lacks a license.

**Answer:** AC

**Explanation:**

For Enable Geo-Restriction, choose Yes. For Restriction Type, choose Whitelist to allow access to certain countries, or choose Blacklist to block access from certain countries. <https://IAM.amazon.com/premiumsupport/knowledge-center/cloudfront-geo-restriction/>

**NEW QUESTION 117**

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

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

**Answer:** B

**Explanation:**

One can use the IAM Encryption CLI to encrypt the data before sending it across to the S3 bucket. Options A and C are invalid because this would still mean that data is transferred in plain text Option D is invalid because you cannot just enable client side encryption for the S3 bucket For more information on Encrypting and Decrypting data, please visit the below URL:

<https://IAM.amazon.com/blogs/security/how-to-encrypt-and-decrypt-your-data-with-the-IAM-encryption-cli> The correct answer is: Use the IAM Encryption CLI to encrypt the data first Submit your Feedback/Queries to our Experts

**NEW QUESTION 121**

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

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

**Answer:** AE

**Explanation:**

these are the steps that can meet the requirements in the most secure manner. CloudTrail is a service that records AWS API calls and delivers log files to an S3 bucket. Turning on CloudTrail in each IAM account can help capture all IAM API calls made within those accounts. Updating the bucket policy of the bucket in the

account that will be storing the logs can help grant other accounts permission to write log files to that bucket. The other options are either unnecessary or insecure for logging and analyzing IAM API calls.

**NEW QUESTION 124**

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

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

**Answer: B**

**Explanation:**

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

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

**NEW QUESTION 125**

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 130**

A Security Engineer has been tasked with enabling IAM Security Hub to monitor Amazon EC2 instances for CVEs 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 the EC2 instances that need 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 instances. You also need to install an additional integration library that enables communication between the Amazon Inspector agent and Security Hub. 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 practices. The other options are either incorrect or incomplete for meeting the requirement.

**NEW QUESTION 135**

A company's application team wants to replace an internal application with a new IAM architecture that consists of Amazon EC2 instances, an IAM Lambda function, and an Amazon S3 bucket in a single IAM Region. After an architecture review, the security team mandates that no application network traffic can traverse the public internet at any point. The security team already has an SCP in place for the company's organization in IAM Organizations to restrict the creation of internet gateways, NAT gateways, and egress-only gateways.

Which combination of steps should the application team take to meet these requirements? (Select THREE.)

- A. Create an S3 endpoint that has a full-access policy for the application's VPC.
- B. Create an S3 access point for the S3 bucket
- C. Include a policy that restricts the network origin to VPCs.
- D. Launch the Lambda function
- E. Enable the block public access configuration.
- F. Create a security group that has an outbound rule over port 443 with a destination of the S3 endpoint. Associate the security group with the EC2 instances.
- G. Create a security group that has an outbound rule over port 443 with a destination of the S3 access point. Associate the security group with the EC2 instances.

H. Launch the Lambda function in a VPC.

**Answer:** ADF

**NEW QUESTION 138**

A company is hosting multiple applications within a single VPC in its IAM account. The applications are running behind an Application Load Balancer that is associated with an IAM WAF web ACL. The company's security team has identified that multiple port scans are originating from a specific range of IP addresses on the internet.

A security engineer needs to deny access from the offending IP addresses. Which solution will meet these requirements?

- A. Modify the IAM WAF web ACL with an IP set match rule statement to deny incoming requests from the IP address range.
- B. Add a rule to all security groups to deny the incoming requests from the IP address range.
- C. Modify the IAM WAF web ACL with a rate-based rule statement to deny the incoming requests from the IP address range.
- D. Configure the IAM WAF web ACL with regex match condition
- E. Specify a pattern set to deny the incoming requests based on the match condition

**Answer:** A

**Explanation:**

Note that the IP is known and the question wants us to deny access from that particular address and so we can use IP set match policy of WAF to block access.

**NEW QUESTION 141**

An ecommerce company is developing new architecture for an application release. The company needs to implement TLS for incoming traffic to the application. Traffic for the application will originate from the internet TLS does not have to be implemented in an end-to-end configuration because the company is concerned about impacts on performance. The incoming traffic types will be HTTP and HTTPS The application uses ports 80 and 443.

What should a security engineer do to meet these requirements?

- A. Create a public Application Load Balance
- B. Create two listeners one listener on port 80 and one listener on port 443. Create one target group
- C. Create a rule to forward traffic from port 80 to the listener on port 443 Provision a public TLS certificate in AWS Certificate Manager (ACM). Attach the certificate to the listener on port 443.
- D. Create a public Application Load Balance
- E. Create two listeners one listener on port 80 and one listener on port 443. Create one target group
- F. Create a rule to forward traffic from port 80 to the listener on port 443 Provision a public TLS certificate in AWS Certificate Manager (ACM). Attach the certificate to the listener on port 80.
- G. Create a public Network Load Balance
- H. Create two listeners one listener on port 80 and one listener on port 443. Create one target group
- I. Create a rule to forward traffic from port 80 to the listener on port 443. Set the protocol for the listener on port 443 to TLS.
- J. Create a public Network Load Balance
- K. Create a listener on port 443. Create one target group
- L. Create a rule to forward traffic from port 443 to the target group
- M. Set the protocol for the listener on port 443 to TLS.

**Answer:** A

**Explanation:**

An Application Load Balancer (ALB) is a type of load balancer that operates at the application layer (layer 7) of the OSI model. It can distribute incoming traffic based on the content of the request, such as the host header, path, or query parameters. An ALB can also terminate TLS connections and decrypt requests from clients before sending them to the targets.

To implement TLS for incoming traffic to the application, the following steps are required:

- Create a public ALB in a public subnet and register the EC2 instances as targets in a target group.
- Create two listeners for the ALB, one on port 80 for HTTP traffic and one on port 443 for HTTPS traffic.
- Create a rule for the listener on port 80 to redirect HTTP requests to HTTPS using the same host, path, and query parameters.
- Provision a public TLS certificate in AWS Certificate Manager (ACM) for the domain name of the application. ACM is a service that lets you easily provision, manage, and deploy public and private SSL/TLS certificates for use with AWS services and your internal connected resources.
- Attach the certificate to the listener on port 443 and configure the security policy to negotiate secure connections between clients and the ALB.
- Configure the security groups for the ALB and the EC2 instances to allow inbound traffic on ports 80 and 443 from the internet and outbound traffic on any port to the EC2 instances.

This solution will meet the requirements of implementing TLS for incoming traffic without impacting performance or requiring end-to-end encryption. The ALB will handle the TLS termination and decryption, while forwarding unencrypted requests to the EC2 instances.

Verified References:

- <https://docs.aws.amazon.com/elasticloadbalancing/latest/application/introduction.html>
- <https://docs.aws.amazon.com/elasticloadbalancing/latest/application/create-https-listener.html>
- <https://docs.aws.amazon.com/acm/latest/userguide/acm-overview.html>

**NEW QUESTION 146**

.....

## 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>