

Exam Questions SCS-C02

AWS Certified Security - Specialty

<https://www.2passeasy.com/dumps/SCS-C02/>



NEW QUESTION 1

An AWS account that is used for development projects has a VPC that contains two subnets. The first subnet is named public-subnet-1 and has the CIDR block 192.168.1.0/24 assigned. The other subnet is named private-subnet-2 and has the CIDR block 192.168.2.0/24 assigned. Each subnet contains Amazon EC2 instances.

Each subnet is currently using the VPC's default network ACL. The security groups that the EC2 instances in these subnets use have rules that allow traffic between each instance where required. Currently, all network traffic flow is working as expected between the EC2 instances that are using these subnets. A security engineer creates a new network ACL that is named subnet-2-NACL with default entries. The security engineer immediately configures private-subnet-2 to use the new network ACL and makes no other changes to the infrastructure. The security engineer starts to receive reports that the EC2 instances in public-subnet-1 and public-subnet-2 cannot communicate with each other.

Which combination of steps should the security engineer take to allow the EC2 instances that are running in these two subnets to communicate again? (Select TWO.)

- A. Add an outbound allow rule for 192.168.2.0/24 in the VPC's default network ACL.
- B. Add an inbound allow rule for 192.168.2.0/24 in the VPC's default network ACL.
- C. Add an outbound allow rule for 192.168.2.0/24 in subnet-2-NACL.
- D. Add an inbound allow rule for 192.168.1.0/24 in subnet-2-NACL.
- E. Add an outbound allow rule for 192.168.1.0/24 in subnet-2-NACL.

Answer: CE

Explanation:

The AWS documentation states that you can add an outbound allow rule for 192.168.2.0/24 in subnet-2-NACL and add an outbound allow rule for 192.168.1.0/24 in subnet-2-NACL. This will allow the EC2 instances that are running in these two subnets to communicate again.

References: : Amazon VPC User Guide

NEW QUESTION 2

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

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

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

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

Answer: ACE

NEW QUESTION 3

A company is using Amazon Macie, AWS Firewall Manager, Amazon Inspector, and AWS Shield Advanced in its AWS account. The company wants to receive alerts if a DDoS attack occurs against the account.

Which solution will meet this requirement?

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

Answer: D

Explanation:

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

For more information, see Monitoring AWS Shield Advanced with Amazon CloudWatch and AWS Shield Advanced metrics and alarms.

NEW QUESTION 4

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

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

Answer: A

Explanation:

To rotate database credentials every 30 days, the most secure and efficient solution is to store the database credentials in AWS Secrets Manager and configure automatic credential rotation for every 30 days. Secrets Manager can handle the rotation of the credentials in both the secret and the database, and it can use AWS KMS to encrypt the credentials. Option B is incorrect because it requires creating a custom Lambda function to rotate the credentials, which is more effort than using Secrets Manager. Option C is incorrect because it stores the database credentials in an environment file or a configuration file, which is less secure than using Secrets Manager. Option D is incorrect because it combines the drawbacks of option B and option C. Verified References:

- > <https://docs.aws.amazon.com/secretsmanager/latest/userguide/rotating-secrets.html>
- > https://docs.aws.amazon.com/secretsmanager/latest/userguide/rotate-secrets_turn-on-for-other.html

NEW QUESTION 5

Your company has just set up a new central server in a VPC. There is a requirement for other teams who have their servers located in different VPC's in the same region to connect to the central server. Which of the below options is best suited to achieve this requirement. Please select:

- A. Set up VPC peering between the central server VPC and each of the teams VPCs.
- B. Set up IAM DirectConnect between the central server VPC and each of the teams VPCs.
- C. Set up an IPSec Tunnel between the central server VPC and each of the teams VPCs.
- D. None of the above options will work.

Answer: A

Explanation:

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. Instances in either VPC can communicate with each other as if they are within the same network. You can create a VPC peering connection between your own VPCs, or with a VPC in another IAM account within a single region.

Options B and C are invalid because you need to use VPC Peering Option D is invalid because VPC Peering is available

For more information on VPC Peering please see the below Link:

<http://docs.IAM.amazon.com/AmazonVPC/latest/UserGuide/vpc-peering.html>

The correct answer is: Set up VPC peering between the central server VPC and each of the teams VPCs. Submit your Feedback/Queries to our Experts

NEW QUESTION 6

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 7

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 8

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 instanc
- B. Send the custom logs to CloudTrail instead of CloudWatch.
- C. Add Amazon S3 to the trust policy of the EC2 instanc
- 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 instanc
- 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 documentation¹, 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 CloudWatch². 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 logs³. 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 instances⁴. 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 logs⁵. 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 9

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

Which statement should the security specialist include in the policy?

- A.

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

B.

- ```

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

```
- C. {
- ```

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

```
- D. {
- ```

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

```
- E. Option A  
 F. Option B  
 G. Option C  
 H. Option D

**Answer: D**

#### NEW QUESTION 10

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 material<sup>3</sup>.
- D. Using keyrings with the AWS Encryption SDK does not address the problem of throttling or cost optimization, as explained in option A. Moreover, using any of the wrapping keys in the multi-keyring to decrypt the data is not a valid option, because only one of the wrapping keys can decrypt a given data key. The wrapping key that encrypts a data key is stored in the encrypted data key structure, and only that wrapping key can decrypt it<sup>4</sup>.

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 10

A Security Engineer creates an Amazon S3 bucket policy that denies access to all users. A few days later, the Security Engineer adds an additional statement to the bucket policy to allow read-only access to one other employee. Even after updating the policy, the employee still receives an access denied message. What is the likely cause of this access denial?

- A. The ACL in the bucket needs to be updated
- B. The IAM policy does not allow the user to access the bucket
- C. It takes a few minutes for a bucket policy to take effect
- D. The allow permission is being overridden by the deny

**Answer: D**

#### NEW QUESTION 12

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 16

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

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

**Answer: C**

#### Explanation:

because this is a solution that can monitor each S3 bucket for unrestricted public write access and use IAM managed services. S3 is a service that provides object storage in the cloud. Systems Manager is a service that helps you automate and manage your AWS resources. You can use Systems Manager to monitor S3

bucket policies for public write access by using a State Manager association that runs a predefined document called AWS-FindS3BucketWithPublicWriteAccess. This document checks each S3 bucket in an account and reports any bucket that has public write access enabled. The other options are either not suitable or not feasible for meeting the requirements.

#### NEW QUESTION 20

A company is running workloads in a single IAM account on Amazon EC2 instances and Amazon EMR clusters a recent security audit revealed that multiple Amazon Elastic Block Store (Amazon EBS) volumes and snapshots are not encrypted. The company's security engineer is working on a solution that will allow users to deploy EC2 Instances and EMR clusters while ensuring that all new EBS volumes and EBS snapshots are encrypted at rest. The solution must also minimize operational overhead. Which steps should the security engineer take to meet these requirements?

- A. Create an Amazon Event Bridge (Amazon Cloud watch Events) event with an EC2 instance as the source and create volume as the event trigger.
- B. When the event is triggered invoke an IAM Lambda function to evaluate and notify the security engineer if the EBS volume that was created is not encrypted.
- C. Use a customer managed IAM policy that will verify that the encryption ag of the Createvolume context is set to true.
- D. Apply this rule to all users.
- E. Create an IAM Config rule to evaluate the configuration of each EC2 instance on creation or modification. Have the IAM Config rule trigger an IAM Lambda function to alert the security team and terminate the instance if the EBS volume is not encrypted.
- F. 5
- G. Use the IAM Management Console or IAM CLI to enable encryption by default for EBS volumes in each IAM Region where the company operates.

**Answer:** D

#### Explanation:

To ensure that all new EBS volumes and EBS snapshots are encrypted at rest and minimize operational overhead, the security engineer should do the following:

- Use the AWS Management Console or AWS CLI to enable encryption by default for EBS volumes in each AWS Region where the company operates. This allows the security engineer to automatically encrypt any new EBS volumes and snapshots created from those volumes, without requiring any additional actions from users.

#### NEW QUESTION 25

A development team is attempting to encrypt and decode a secure string parameter from the IAM Systems Manager Parameter Store using an IAM Key Management Service (IAM KMS) CMK. However, each attempt results in an error message being sent to the development team. Which CMK-related problems possibly account for the error? (Select two.)

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

**Answer:** AD

#### Explanation:

<https://docs.IAM.amazon.com/kms/latest/developerguide/services-parameter-store.html#parameter-store-cmk-fa>

#### NEW QUESTION 30

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 32

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 35

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

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

**Answer:** C

#### NEW QUESTION 39

A company hosts business-critical applications on Amazon EC2 instances in a VPC. The VPC uses default DHCP options sets. A security engineer needs to log all DNS queries that internal resources make in the VPC. The security engineer also must create a list of the most common DNS queries over time. Which solution will meet these requirements?

- A. Install the Amazon CloudWatch agent on each EC2 instance in the VPC
- B. Use the CloudWatch agent to stream the DNS query logs to an Amazon CloudWatch Logs log group
- C. Use CloudWatch metric filters to automatically generate metrics that list the most common DNS queries.
- D. Install a BIND DNS server in the VPC
- E. Create a bash script to list the DNS request number of common DNS queries from the BIND logs.
- F. Create VPC flow logs for all subnets in the VPC
- G. Stream the flow logs to an Amazon CloudWatch Logs log group
- H. Use CloudWatch Logs Insights to list the most common DNS queries for the log group in a custom dashboard.
- I. Configure Amazon Route 53 Resolver query logging
- J. Add an Amazon CloudWatch Logs log group as the destination
- K. Use Amazon CloudWatch Contributor Insights to analyze the data and create time series that display the most common DNS queries.

**Answer:** D

#### Explanation:

<https://aws.amazon.com/blogs/aws/log-your-vpc-dns-queries-with-route-53-resolver-query-logs/>

#### NEW QUESTION 42

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 status2.
- C. The security engineer does not need the GetCredentialReport permission to enable or evaluate AWS Config rules. This permission is required to retrieve a credential report that was previously generated by using the GenerateCredentialReport operation2.

References:

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

#### NEW QUESTION 46

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 51

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

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

**Answer: D**

#### Explanation:

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

#### NEW QUESTION 52

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

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

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

**Answer: DE**

#### Explanation:

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

**NEW QUESTION 54**

A company that uses AWS Organizations wants to see AWS Security Hub findings for many AWS accounts and AWS Regions. Some of the accounts are in the company's organization, and some accounts are in organizations that the company manages for customers. Although the company can see findings in the Security Hub administrator account for accounts in the company's organization, there are no findings from accounts in other organizations. Which combination of steps should the company take to see findings from accounts that are outside the organization that includes the Security Hub administrator account? (Select TWO.)

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

**Answer:** CE

**Explanation:**

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

- > Send invitations to accounts that are outside the company's organization from the Security Hub administrator account. This will allow the administrator account to view and manage findings from those accounts. The administrator account can send invitations by using the Security Hub console, API, or CLI. For more information, see Sending invitations to member accounts.
- > Send an administration request from the member accounts. This will allow the member accounts to accept the invitation from the administrator account and establish a relationship with it. The member accounts can send administration requests by using the Security Hub console, API, or CLI. For more information, see Sending administration requests.

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

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

Verified References:

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

**NEW QUESTION 59**

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

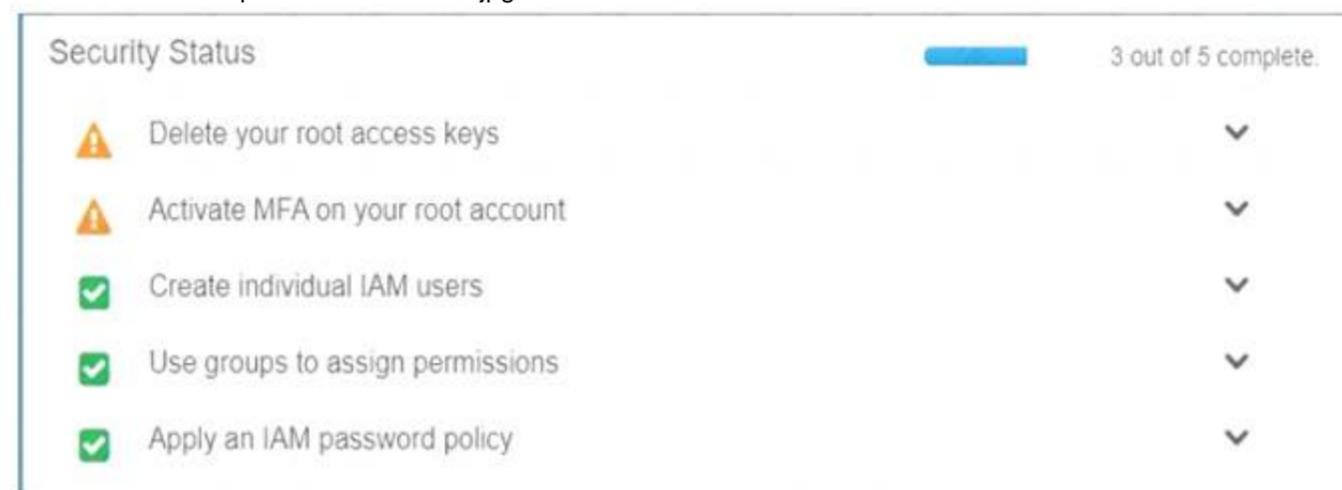
- A. Use short but complex password on the root account and any administrators.
- B. Use IAM Geo-Lock and disallow anyone from logging in except for in your city.
- C. Use MFA on all users and accounts, especially on the root account.
- D. Don't write down or remember the root account password after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

**Answer:** C

**Explanation:**

Multi-factor authentication can add one more layer of security to your IAM account Even when you go to your Security Credentials dashboard one of the items is to enable MFA on your root account

C:\Users\wk\Desktop\mudassar\Untitled.jpg



Option A is invalid because you need to have a good password policy Option B is invalid because there is no IAM Geo-Lock Option D is invalid because this is not a recommended practices For more information on MFA, please visit the below URL

[http://docs.IAM.amazon.com/IAM/latest/UserGuide/id\\_credentials\\_mfa.html](http://docs.IAM.amazon.com/IAM/latest/UserGuide/id_credentials_mfa.html)

The correct answer is: Use MFA on all users and accounts, especially on the root account. Submit your Feedback/Queries to our Experts

**NEW QUESTION 60**

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

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

- A. Designate an AWS account as a delegated administrator for Security Hu
- B. Publish events to Amazon CloudWatch from the delegated administrator account, all member accounts, and required Regions that are enabled for Security Hub findings.
- C. Designate an AWS account in an organization in AWS Organizations as a delegated administrator for Security Hu
- D. Publish events to Amazon EventBridge from the delegated administrator account, all member accounts, and required Regions that are enabled for Security Hub

findings.

- E. In each Region, create an Amazon EventBridge rule to deliver findings to an Amazon Kinesis data stream.
- F. Configure the Kinesis data streams to output the logs to a single Amazon S3 bucket.
- G. In each Region, create an Amazon EventBridge rule to deliver findings to an Amazon Kinesis Data Firehose delivery stream.
- H. Configure the Kinesis Data Firehose delivery streams to deliver the logs to a single Amazon S3 bucket.
- I. Use AWS Glue DataBrew to crawl the Amazon S3 bucket and build the schema.
- J. Use AWS Glue Data Catalog to query the data and create views to flatten nested attributes.
- K. Build Amazon QuickSight dashboards by using Amazon Athena.
- L. Partition the Amazon S3 data.
- M. Use AWS Glue to crawl the S3 bucket and build the schema.
- N. Use Amazon Athena to query the data and create views to flatten nested attributes.
- O. Build Amazon QuickSight dashboards that use the Athena views.

**Answer:** BDF

**Explanation:**

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

**NEW QUESTION 62**

A security team is working on a solution that will use Amazon EventBridge (Amazon CloudWatch Events) to monitor new Amazon S3 objects. The solution will monitor for public access and for changes to any S3 bucket policy or setting that result in public access. The security team configures EventBridge to watch for specific API calls that are logged from AWS CloudTrail. EventBridge has an action to send an email notification through Amazon Simple Notification Service (Amazon SNS) to the security team immediately with details of the API call.

Specifically, the security team wants EventBridge to watch for the s3:PutObjectAcl, s3:DeleteBucketPolicy, and s3:PutBucketPolicy API invocation logs from CloudTrail. While developing the solution in a single account, the security team discovers that the s3:PutObjectAcl API call does not invoke an EventBridge event. However, the s3:DeleteBucketPolicy API call and the s3:PutBucketPolicy API call do invoke an event.

The security team has enabled CloudTrail for AWS management events with a basic configuration in the AWS Region in which EventBridge is being tested.

Verification of the EventBridge event pattern indicates that the pattern is set up correctly. The security team must implement a solution so that the s3:PutObjectAcl API call will invoke an EventBridge event. The solution must not generate false notifications.

Which solution will meet these requirements?

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

**Answer:** D

**Explanation:**

The correct answer is D. Enable CloudTrail to monitor data events for read and write operations to S3 buckets. According to the AWS documentation<sup>1</sup>, CloudTrail data events are the resource operations performed on or within a resource. These are also known as data plane operations. Data events are often high-volume activities. For example, Amazon S3 object-level API activity (such as GetObject, DeleteObject, and PutObject) is a data event. By default, trails do not log data events. To record CloudTrail data events, you must explicitly add the supported resources or resource types for which you want to collect activity. For more information, see Logging data events in the Amazon S3 User Guide<sup>2</sup>. In this case, the security team wants EventBridge to watch for the s3:PutObjectAcl API invocation logs from CloudTrail. This API uses the acl subresource to set the access control list (ACL) permissions for a new or existing object in an S3 bucket<sup>3</sup>. This is a data event that affects the S3 object resource type. Therefore, the security team must enable CloudTrail to monitor data events for read and write operations to S3 buckets in order to invoke an EventBridge event for this API call. The other options are incorrect because:

- A. Modifying the EventBridge event pattern by selecting Amazon S3 and All Events as the event type will not capture the s3:PutObjectAcl API call, because this is a data event and not a management event. Management events provide information about management operations that are performed on resources in your AWS account. These are also known as control plane operations<sup>4</sup>.
- B. Modifying the EventBridge event pattern by selecting Amazon S3 and Bucket Level Operations as the event type will not capture the s3:PutObjectAcl API call, because this is a data event that affects the S3 object resource type and not the S3 bucket resource type. Bucket level operations are management events that affect the configuration or metadata of an S3 bucket<sup>5</sup>.
- C. Enabling CloudTrail Insights to identify unusual API activity will not help the security team monitor new S3 objects or changes to any S3 bucket policy or setting that result in public access. CloudTrail Insights helps AWS users identify and respond to unusual activity associated with API calls and API error rates by continuously analyzing CloudTrail management events<sup>6</sup>. It does not analyze data events or generate EventBridge events.

References:

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

**NEW QUESTION 66**

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

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

**Answer: B**

**NEW QUESTION 69**

A company is implementing a new application in a new IAM account. A VPC and subnets have been created for the application. The application has been peered to an existing VPC in another account in the same IAM Region for database access. Amazon EC2 instances will regularly be created and terminated in the application VPC, but only some of them will need access to the databases in the peered VPC over TCP port 1521. A security engineer must ensure that only the EC2 instances that need access to the databases can access them through the network. How can the security engineer implement this solution?

- A. Create a new security group in the database VPC and create an inbound rule that allows all traffic from the IP address range of the application VPC
- B. Add a new network ACL rule on the database subnet
- C. Configure the rule to TCP port 1521 from the IP address range of the application VPC
- D. Attach the new security group to the database instances that the application instances need to access.
- E. Create a new security group in the application VPC with an inbound rule that allows the IP address range of the database VPC over TCP port 1521. Create a new security group in the database VPC with an inbound rule that allows the IP address range of the application VPC over port 1521. Attach the new security group to the database instances and the application instances that need database access.
- F. Create a new security group in the application VPC with no inbound rule
- G. Create a new security group in the database VPC with an inbound rule that allows TCP port 1521 from the new application security group in the application VPC
- H. Attach the application security group to the application instances that need database access, and attach the database security group to the database instances.
- I. Create a new security group in the application VPC with an inbound rule that allows the IP address range of the database VPC over TCP port 1521. Add a new network ACL rule on the database subnet
- J. Configure the rule to allow all traffic from the IP address range of the application VPC
- K. Attach the new security group to the application instances that need database access.

**Answer: C**

**NEW QUESTION 72**

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 75**

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 79**

A company is planning to use Amazon Elastic File System (Amazon EFS) with its on-premises servers. The company has an existing IAM Direct Connect connection established between its on-premises data center and an IAM Region. Security policy states that the company's on-premises firewall should only have specific IP addresses added to the allow list and not a CIDR range. The company also wants to restrict access so that only certain data center-based servers have access to Amazon EFS.

How should a security engineer implement this solution?

- A. Add the file-system-id efs IAM-region amazonIAM.com URL to the allow list for the data center firewall. Install the IAM CLI on the data center-based servers to mount the EFS file system. In the EFS security group, add the data center IP range to the allow list. Mount the EFS using the EFS file system name.
- B. Assign an Elastic IP address to Amazon EFS and add the Elastic IP address to the allow list for the data center firewall. Install the IAM CLI on the data center-based servers to mount the EFS file system. In the EFS security group, add the IP addresses of the data center servers to the allow list. Mount the EFS using the Elastic IP address.
- C. Add the EFS file system mount target IP addresses to the allow list for the data center firewall. In the EFS security group, add the data center server IP addresses to the allow list. Use the Linux terminal to mount the EFS file system using the IP address of one of the mount targets.
- D. Assign a static range of IP addresses for the EFS file system by contacting IAM Support. In the EFS security group, add the data center server IP addresses to the allow list. Use the Linux terminal to mount the EFS file system using one of the static IP addresses.

**Answer:** B

**Explanation:**

To implement the solution, the security engineer should do the following:

- Assign an Elastic IP address to Amazon EFS and add the Elastic IP address to the allow list for the data center firewall. This allows the security engineer to use a specific IP address for the EFS file system that can be added to the firewall rules, instead of a CIDR range or a URL.
- Install the AWS CLI on the data center-based servers to mount the EFS file system. This allows the security engineer to use the mount helper provided by AWS CLI to mount the EFS file system with encryption in transit.
- In the EFS security group, add the IP addresses of the data center servers to the allow list. This allows the security engineer to restrict access to the EFS file system to only certain data center-based servers.
- Mount the EFS using the Elastic IP address. This allows the security engineer to use the Elastic IP address as the DNS name for mounting the EFS file system.

**NEW QUESTION 83**

A company has recently recovered from a security incident that required the restoration of Amazon EC2 instances from snapshots.

After performing a gap analysis of its disaster recovery procedures and backup strategies, the company is concerned that, next time, it will not be able to recover the EC2 instances if the AWS account was compromised and Amazon EBS snapshots were deleted.

All EBS snapshots are encrypted using an AWS KMS CMK. Which solution would solve this problem?

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

**Answer:** C

**Explanation:**

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

**NEW QUESTION 87**

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

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

**Answer:** A

**Explanation:**

The IAM Documentation mentions the following

To determine whether a log file was modified, deleted, or unchanged after CloudTrail delivered it you can use CloudTrail log file integrity validation. This feature is built using industry standard algorithms: SHA-256 for hashing and SHA-256 with RSA for digital signing. This makes it computationally infeasible to modify, delete or forge CloudTrail log files without detection. You can use the IAM CLI to validate the files in the location where CloudTrail delivered them

Validated log files are invaluable in security and forensic investigations. For example, a validated log file enables you to assert positively that the log file itself has not changed, or that particular user credentials performed specific API activity. The CloudTrail log file integrity validation process also lets you know if a log file has been deleted or changed, or assert positively that no log files were delivered to your account during a given period of time.

Options B.C and D is invalid because you need to check for log File Integrity Validation for cloudtrail logs For more information on Cloudtrail log file validation, please visit the below URL: <http://docs.IAM.amazon.com/IAMcloudtrail/latest/userguide/cloudtrail-log-file-validation-intro.html>

The correct answer is: Use CloudTrail Log File Integrity Validation. omit your Feedback/Queries to our Expert

**NEW QUESTION 91**

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 gatewa
- 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 96**

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

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

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

**Answer:** B

**Explanation:**

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

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

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

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

## References:

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

**NEW QUESTION 101**

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

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

A)

Apply the following SCP:

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

B)

Apply the following SCP:

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

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

```

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

```

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

**Answer:** ACE

#### NEW QUESTION 106

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 rotated.
- 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 Manager<sup>3</sup>.

#### NEW QUESTION 111

A security engineer receives a notice from the AWS Abuse team about suspicious activity from a Linux-based Amazon EC2 instance that uses Amazon Elastic Block Store (Amazon EBS)-based storage. The instance is making connections to known malicious addresses.

The instance is in a development account within a VPC that is in the us-east-1 Region. The VPC contains an internet gateway and has a subnet in us-east-1a and us-east-1b. Each subnet is associated with a route table that uses the internet gateway as a default route. Each subnet also uses the default network ACL. The suspicious EC2 instance runs within the us-east-1b subnet. During an initial investigation, a security engineer discovers that the suspicious instance is the only instance that runs in the subnet.

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

- A. Log in to the suspicious instance and use the netstat command to identify remote connections. Use the IP addresses from these remote connections to create deny rules in the security group of the instance. Install diagnostic tools on the instance for investigation. Update the outbound network ACL for the subnet in us-east-1b to explicitly deny all connections as the first rule during the investigation of the instance.
- B. Update the outbound network ACL for the subnet in us-east-1b to explicitly deny all connections as the first rule. Replace the security group with a new security group that allows connections only from a diagnostics security group. Update the outbound network ACL for the us-east-1b subnet to remove the deny all rule. Launch a new EC2 instance that has diagnostic tools. Assign the new security group to the new EC2 instance. Use the new EC2 instance to investigate the suspicious instance.

- C. Ensure that the Amazon Elastic Block Store (Amazon EBS) volumes that are attached to the suspicious EC2 instance will not delete upon termination
- Terminate the instance Launch a new EC2 instance in us-east-1a that has diagnostic tools Mount the EBS volumes from the terminated instance for investigation
- D. Create an AWS WAF web ACL that denies traffic to and from the suspicious instance Attach the AWS WAF web ACL to the instance to mitigate the attack Log in to the instance and install diagnostic tools to investigate the instance

**Answer: B**

**Explanation:**

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

**NEW QUESTION 116**

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 route<sup>12</sup>.

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 appliance<sup>3</sup>.
- > 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 traffic<sup>4</sup>.
- > 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 properly<sup>5</sup>.

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 118**

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

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

**Answer: A**

**NEW QUESTION 121**

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

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

**Answer: A**

**NEW QUESTION 122**

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

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

**Answer: C**

**Explanation:**

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

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

**NEW QUESTION 125**

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

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

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

**Answer: BC**

**Explanation:**

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

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

Verified References:

- <https://repost.aws/knowledge-center/cross-account-access-iam>
- [https://docs.aws.amazon.com/organizations/latest/userguide/orgs\\_manage\\_accounts\\_access.html](https://docs.aws.amazon.com/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 128**

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

Which solution will meet these requirements?

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

Answer: B

#### NEW QUESTION 129

A company needs to follow security best practices to deploy resources from an AWS CloudFormation template. The CloudFormation template must be able to configure sensitive database credentials.

The company already uses AWS Key Management Service (AWS KMS) and AWS Secrets Manager. Which solution will meet the requirements?

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

Answer: A

#### Explanation:

➤ Option A: This option meets the requirements of following security best practices and configuring sensitive database credentials in the CloudFormation template. A dynamic reference is a way to specify external values that are stored and managed in other services, such as Secrets Manager, in the stack templates<sup>1</sup>. When using a dynamic reference, CloudFormation retrieves the value of the specified reference when necessary during stack and change set operations<sup>1</sup>. Dynamic references can be used for certain resources that support them, such as `AWS::RDS::DBInstance`<sup>1</sup>. By using a dynamic reference to reference the database credentials in Secrets Manager, the company can leverage the existing integration between these services and avoid hardcoding the secret information in the template. Secrets Manager is a service that helps you protect secrets needed to access your applications, services, and IT resources<sup>2</sup>. Secrets Manager enables you to rotate, manage, and retrieve database credentials, API keys, and other secrets throughout their lifecycle<sup>2</sup>.

#### NEW QUESTION 134

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

The company is performing control tests on specific GuardDuty findings to make sure that the company's security team can detect and respond to security events.

The security team launched an Amazon EC2 instance and attempted to run DNS requests against a test domain, `example.com`, to generate a DNS finding.

However, the GuardDuty finding was never created in the Security Hub delegated administrator account.

Why was the finding was not created in the Security Hub delegated administrator account?

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

Answer: C

#### Explanation:

The correct answer is C. The GuardDuty integration with Security Hub was never activated in the AWS account where the finding was generated.

According to the AWS documentation<sup>1</sup>, GuardDuty findings are automatically sent to Security Hub only if the GuardDuty integration with Security Hub is enabled in the same account and Region. This means that the security tooling account, which is the delegated administrator for both GuardDuty and Security Hub, must enable the GuardDuty integration with Security Hub in each member account and Region where GuardDuty is enabled. Otherwise, the findings from GuardDuty will not be visible in Security Hub.

The other options are incorrect because:

➤ VPC flow logs are not required for GuardDuty to generate DNS findings. GuardDuty uses VPC DNS logs, which are automatically enabled for all VPCs, to detect malicious or unauthorized DNS activity.

➤ The DHCP option configured for a custom OpenDNS resolver does not affect GuardDuty's ability to generate DNS findings. GuardDuty uses its own threat intelligence sources to identify malicious domains, regardless of the DNS resolver used by the EC2 instance.

➤ Cross-Region aggregation in Security Hub is not relevant for this scenario, because the company operates out of a single AWS Region. Cross-Region aggregation allows Security Hub to aggregate findings from multiple Regions into a single Region.

References:

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

#### NEW QUESTION 136

A company is running an Amazon RDS for MySQL DB instance in a VPC. The VPC must not send or receive network traffic through the internet.

A security engineer wants to use AWS Secrets Manager to rotate the DB instance credentials automatically. Because of a security policy, the security engineer cannot use the standard AWS Lambda function that Secrets Manager provides to rotate the credentials.

The security engineer deploys a custom Lambda function in the VPC. The custom Lambda function will be responsible for rotating the secret in Secrets Manager.

The security engineer edits the DB instance's security group to allow connections from this function. When the function is invoked, the function cannot communicate with Secrets Manager to rotate the secret properly.

What should the security engineer do so that the function can rotate the secret?

- A. Add an egress-only internet gateway to the VP
- B. Allow only the Lambda function's subnet to route traffic through the egress-only internet gateway.
- C. Add a NAT gateway to the VP
- D. Configure only the Lambda function's subnet with a default route through the NAT gateway.
- E. Configure a VPC peering connection to the default VPC for Secrets Manage
- F. Configure the Lambda function's subnet to use the peering connection for routes.
- G. Configure a Secrets Manager interface VPC 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 140**

A company's Security Team received an email notification from the Amazon EC2 Abuse team that one or more of the company's Amazon EC2 instances may have been compromised

Which combination of actions should the Security team take to respond to (be current modem? (Select TWO.)

- A. Open a support case with the IAM Security team and ask them to remove the malicious code from the affected instance
- B. Respond to the notification and list the actions that have been taken to address the incident
- C. Delete all IAM users and resources in the account
- D. Detach the internet gateway from the VPC remove aft rules that contain 0.0.0.0/0 from the security groups, and create a NACL rule to deny all traffic Inbound from the internet
- E. Delete the identified compromised instances and delete any associated resources that the Security team did not create.

**Answer:** DE

**Explanation:**

these are the recommended actions to take when you receive an abuse notice from AWS<sup>8</sup>. You should review the abuse notice to see what content or activity was reported and detach the internet gateway from the VPC to isolate the affected instances from the internet. You should also remove any rules that allow inbound traffic from 0.0.0.0/0 from the security groups and create a network access control list (NACL) rule to deny all traffic inbound from the internet. You should then delete the compromised instances and any associated resources that you did not create. The other options are either inappropriate or unnecessary for responding to the abuse notice.

**NEW QUESTION 145**

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

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

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

**Answer:** C

**Explanation:**

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

According to the AWS documentation<sup>1</sup>, Route 53 Resolver query logging lets you log the DNS queries that Route 53 Resolver handles for your VPCs. You can send the logs to CloudWatch Logs, Amazon S3, or Kinesis Data Firehose. The logs include information such as the following:

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

The DNS record type (such as A or AAAA)

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

You can use CloudWatch Insights to run queries on your log data and analyze the results using graphs and statistics<sup>2</sup>. You can filter and aggregate the log data based on any field, and use operators and functions to perform calculations and transformations. For example, you can use CloudWatch Insights to find out how many queries were made for a specific domain name, or which instances made the most queries.

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

The other options are incorrect because:

- A. Using VPC Traffic Mirroring would not capture the DNS queries that go to the on-premises DNS servers, because Traffic Mirroring only copies network traffic from an elastic network interface of an EC2 instance to a target for analysis<sup>3</sup>. Traffic Mirroring does not include traffic that goes through a Route 53 Resolver outbound endpoint, which is used to forward queries to on-premises DNS servers<sup>4</sup>. Therefore, this solution would not meet the requirements.
- B. Configuring VPC flow logs on all relevant VPCs would not capture the DNS name that was requested in Route 53 Resolver, because flow logs only record information about the IP traffic going to and from network interfaces in a VPC<sup>5</sup>. Flow logs do not include any information about the content or payload of a packet, such as a DNS query or response. Therefore, this solution would not meet the requirements.
- D. Modifying the Route 53 Resolver rules on the authoritative domains that forward to the on-premises DNS servers would not enable logging of DNS queries, because Resolver rules only specify how to forward queries for specified domain names to your network<sup>6</sup>. Resolver rules do not have any logging functionality by themselves. Therefore, this solution would not meet the requirements. References:

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

#### NEW QUESTION 149

A company is using Amazon Elastic Container Service (Amazon ECS) to deploy an application that deals with sensitive data. During a recent security audit, the company identified a security issue in which Amazon RDS credentials were stored with the application code in the company's source code repository. A security engineer needs to develop a solution to ensure that database credentials are stored securely and rotated periodically. The credentials should be accessible to the application only. The engineer also needs to prevent database administrators from sharing database credentials as plaintext with other teammates. The solution must also minimize administrative overhead.

Which solution meets these requirements?

- A. Use the IAM Systems Manager Parameter Store to generate database credential
- B. Use an IAM profile for ECS tasks to restrict access to database credentials to specific containers only.
- C. Use IAM Secrets Manager to store database credential
- D. Use an IAM inline policy for ECS tasks to restrict access to database credentials to specific containers only.
- E. Use the IAM Systems Manager Parameter Store to store database credential
- F. Use IAM roles for ECS tasks to restrict access to database credentials to specific containers only.
- G. Use IAM Secrets Manager to store database credential
- H. Use IAM roles for ECS tasks to restrict access to database credentials to specific containers only.

**Answer:** D

#### Explanation:

To ensure that database credentials are stored securely and rotated periodically, the security engineer should do the following:

- Use AWS Secrets Manager to store database credentials. This allows the security engineer to encrypt and manage secrets centrally, and to configure automatic rotation schedules for them.
- Use IAM roles for ECS tasks to restrict access to database credentials to specific containers only. This allows the security engineer to grant fine-grained permissions to ECS tasks based on their roles, and to avoid sharing credentials as plaintext with other teammates.

#### NEW QUESTION 151

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 KeyRotation events
- B. Monitor Amazon CloudWatch Events for any IAM KMS CMK rotation events
- C. Using the IAM CLI
- 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 156

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 161**

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, 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 165**

A company needs to store multiple years of financial records. The company wants to use Amazon S3 to store copies of these documents. The company must implement a solution to prevent the documents from being edited, replaced, or deleted for 7 years after the documents are stored in Amazon S3. The solution must also encrypt the documents at rest.

A security engineer creates a new S3 bucket to store the documents. What should the security engineer do next to meet these requirements?

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

**Answer:** B

**NEW QUESTION 169**

.....

## THANKS FOR TRYING THE DEMO OF OUR PRODUCT

Visit Our Site to Purchase the Full Set of Actual SCS-C02 Exam Questions With Answers.

We Also Provide Practice Exam Software That Simulates Real Exam Environment And Has Many Self-Assessment Features. Order the SCS-C02 Product From:

<https://www.2passeasy.com/dumps/SCS-C02/>

### Money Back Guarantee

#### **SCS-C02 Practice Exam Features:**

- \* SCS-C02 Questions and Answers Updated Frequently
- \* SCS-C02 Practice Questions Verified by Expert Senior Certified Staff
- \* SCS-C02 Most Realistic Questions that Guarantee you a Pass on Your FirstTry
- \* SCS-C02 Practice Test Questions in Multiple Choice Formats and Updatesfor 1 Year