

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 company in France uses Amazon Cognito with the Cognito Hosted UI as an identity broker for sign-in and sign-up processes. The company is marketing an application and expects that all the application's users will come from France.

When the company launches the application the company's security team observes fraudulent sign-ups for the application. Most of the fraudulent registrations are from users outside of France.

The security team needs a solution to perform custom validation at sign-up. Based on the results of the validation the solution must accept or deny the registration request.

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

- A. Create a pre sign-up AWS Lambda trigger.
- B. Associate the Amazon Cognito function with the Amazon Cognito user pool.
- C. Use a geographic match rule statement to configure an AWS WAF web ACL.
- D. Associate the web ACL with the Amazon Cognito user pool.
- E. Configure an app client for the application's Amazon Cognito user pool.
- F. Use the app client ID to validate the requests in the hosted UI.
- G. Update the application's Amazon Cognito user pool to configure a geographic restriction setting.
- H. Use Amazon Cognito to configure a social identity provider (IdP) to validate the requests on the hosted UI.

Answer: B

Explanation:

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

NEW QUESTION 3

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 4

A company stores sensitive documents in Amazon S3 by using server-side encryption with an IAM Key Management Service (IAM KMS) CMK. A new requirement mandates that the CMK that is used for these documents can be used only for S3 actions.

Which statement should the company add to the key policy to meet this requirement?

A)

```
{
  "Effect": "Deny",
  "Principal": "*",
  "Action": "kms:*",
  "Resource": "*",
  "Condition": {
    "StringNotEquals": {
      "kms:CallerAccount": "s3.amazonaws.com"
    }
  }
}
```

B)

```
{
  "Effect": "Deny",
  "Principal": "*",
  "Action": "s3:*",
  "Resource": "*",
  "Condition": {
    "StringNotEquals": {
      "kms:ViaService": "kms.*amazonaws.com"
    }
  }
}
```

A. Option A

B. Option B

Answer: A

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

An AWS account administrator created an IAM group and applied the following managed policy to require that each individual user authenticate using multi-factor authentication:

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": "ec2:*",
      "Resource": "*"
    },
    {
      "Sid": "BlockAnyAccessUnlessSignedInWithMFA",
      "Effect": "Deny",
      "Action": "ec2:*",
      "Resource": "*",
      "Condition": {
        "BoolIfExists": {
          "aws:MultiFactorAuthPresent": false
        }
      }
    }
  ]
}
```

After implementing the policy, the administrator receives reports that users are unable to perform Amazon EC2 commands using the AWS CLI. What should the administrator do to resolve this problem while still enforcing multi-factor authentication?

- A. Change the value of `aws:MultiFactorAuthPresent` to true.
- B. Instruct users to run the `aws sts get-session-token` CLI command and pass the multi-factor authentication `--serial-number` and `--token-code` parameter
- C. Use these resulting values to make API/CLI calls.
- D. Implement federated API/CLI access using SAML 2.0, then configure the identity provider to enforce multi-factor authentication.
- E. Create a role and enforce multi-factor authentication in the role trust policy
- F. Instruct users to run the `sts assume-role` CLI command and pass `--serial-number` and `--token-code` parameter
- G. Store the resulting values in environment variable
- H. Add `sts:AssumeRole` to `NotAction` in the policy.

Answer: B

Explanation:

The correct answer is B. Instruct users to run the `aws sts get-session-token` CLI command and pass the multi-factor authentication `--serial-number` and `--token-code` parameters. Use these resulting values to make API/CLI calls.

According to the AWS documentation¹, the `aws sts get-session-token` CLI command returns a set of temporary credentials for an AWS account or IAM user. The credentials consist of an access key ID, a secret access key, and a security token. These credentials are valid for the specified duration only. The session duration for IAM users can be between 15 minutes and 36 hours, with a default of 12 hours.

You can use the `--serial-number` and `--token-code` parameters to provide the MFA device serial number and the MFA code from the device. The MFA device must be associated with the user who is making the `get-session-token` call. If you do not provide these parameters when your IAM user or role has a policy that requires MFA, you will receive an Access Denied error. The temporary security credentials that are returned by the `get-session-token` command can then be used to make subsequent API or CLI calls that require MFA authentication. You can use environment variables or a profile in your AWS CLI configuration file to specify the temporary credentials.

Therefore, this solution will resolve the problem of users being unable to perform EC2 commands using the AWS CLI, while still enforcing MFA.

The other options are incorrect because:

- > A. Changing the value of `aws:MultiFactorAuthPresent` to true will not work, because this is a condition key that is evaluated by AWS when a request is made. You cannot set this value manually in your policy or request. You must provide valid MFA information to AWS for this condition key to be true.
- > C. Implementing federated API/CLI access using SAML 2.0 may work, but it requires more operational effort than using the `get-session-token` command. You would need to configure a SAML identity provider and trust relationship with AWS, and use a custom SAML client to request temporary credentials from AWS STS. This solution may also introduce additional security risks if the identity provider is compromised.
- > D. Creating a role and enforcing MFA in the role trust policy may work, but it also requires more operational effort than using the `get-session-token` command. You would need to create a role for each user or group that needs to perform EC2 commands, and specify a trust policy that requires MFA. You would also need to grant the users permission to assume the role, and instruct them to use the `sts assume-role` command instead of the `get-session-token` command.

References:

1: `get-session-token` — AWS CLI Command Reference

NEW QUESTION 7

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

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

Which solution meets these requirements?

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

Answer: A

NEW QUESTION 8

A company uses a third-party identity provider and SAML-based SSO for its AWS accounts. After the third-party identity provider renewed an expired signing certificate, users saw the following message when trying to log in:

Error: Response Signature Invalid (Service: AWSSecurityTokenService; Status Code: 400; Error Code: InvalidIdentityToken)

A security engineer needs to provide a solution that corrects the error and minimizes operational overhead. Which solution meets these requirements?

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

Answer: C

Explanation:

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

NEW QUESTION 9

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

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

Answer: A

Explanation:

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

NEW QUESTION 10

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

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

Answer: D

Explanation:

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

According to the AWS 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 10

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 11

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

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

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

Answer: B

Explanation:

The correct answer is B. Use data key caching. Use the local cache that the AWS Encryption SDK provides with a caching cryptographic materials manager. This answer is correct because data key caching can improve performance, reduce cost, and help the company stay within the service limits of AWS KMS. Data key caching stores data keys and related cryptographic material in a cache, and reuses them for encryption and decryption operations. This reduces the number of requests to AWS KMS endpoints and avoids throttling. The AWS Encryption SDK provides a local cache and a caching cryptographic materials manager (caching CMM) that interacts with the cache and enforces security thresholds that the company can set1.

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 endpoints2.
- C. Using KMS key rotation does not address the problem of throttling or cost optimization. Key rotation is a security practice that creates new cryptographic material for a KMS key every year, but it does not affect the data that the KMS key protects. Key rotation does not reduce the number of requests to AWS KMS endpoints, and it might incur additional costs for storing multiple versions of key material3.
- D. Using keyrings with the AWS Encryption SDK does not address the problem of throttling or cost optimization, as explained in option A. Moreover, using any of the wrapping keys in the multi-keyring to decrypt the data is not a valid option, because only one of the wrapping keys can decrypt a given data key. The wrapping key that encrypts a data key is stored in the encrypted data key structure, and only that wrapping key can decrypt it4.

References:

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

NEW QUESTION 13

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 15

A company has several workloads running on AWS. Employees are required to authenticate using on-premises ADFS and SSO to access the AWS Management Console. Developers migrated an existing legacy web application to an Amazon EC2 instance. Employees need to access this application from anywhere on the internet, but currently, there is no authentication system built into the application. How should the Security Engineer implement employee-only access to this system without changing the application?

- A. Place the application behind an Application Load Balancer (ALB). Use Amazon Cognito as authentication for the ALB.
- B. Define a SAML-based Amazon Cognito user pool and connect it to ADFS.
- C. Implement AWS SSO in the master account and link it to ADFS as an identity provider.
- D. Define the EC2 instance as a managed resource, then apply an IAM policy on the resource.
- E. Define an Amazon Cognito identity pool, then install the connector on the Active Directory server.
- F. Use the Amazon Cognito SDK on the application instance to authenticate the employees using their Active Directory user names and passwords.
- G. Create an AWS Lambda custom authorizer as the authenticator for a reverse proxy on Amazon EC2. Ensure the security group on Amazon EC2 only allows access from the Lambda function.

Answer: A

Explanation:

<https://docs.aws.amazon.com/elasticloadbalancing/latest/application/listener-authenticate-users.html>

NEW QUESTION 18

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 CloudWatch 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 flag 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 19

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 23

A security team is developing an application on an Amazon EC2 instance to get objects from an Amazon S3 bucket. All objects in the S3 bucket are encrypted with an AWS Key Management Service (AWS KMS) customer managed key. All network traffic for requests that are made within the VPC is restricted to the AWS infrastructure. This traffic does not traverse the public internet.

The security team is unable to get objects from the S3 bucket. Which factors could cause this issue? (Select THREE.)

- A. The IAM instance profile that is attached to the EC2 instance does not allow the s3 ListBucket action to the S3 bucket in the AWS accounts.
- B. The IAM instance profile that is attached to the EC2 instance does not allow the s3 ListParts action to the S3 bucket in the AWS accounts.
- C. The KMS key policy that encrypts the object in the S3 bucket does not allow the kms ListKeys action to the EC2 instance profile ARN.
- D. The KMS key policy that encrypts the object in the S3 bucket does not allow the kms Decrypt action to the EC2 instance profile ARN.
- E. The security group that is attached to the EC2 instance is missing an outbound rule to the S3 managed prefix list over port 443.
- F. The security group that is attached to the EC2 instance is missing an inbound rule from the S3 managed prefix list over port 443.

Answer: ADE

Explanation:

<https://docs.aws.amazon.com/vpc/latest/userguide/security-group-rules.html>

To get objects from an S3 bucket that are encrypted with a KMS customer managed key, the security team needs to have the following factors in place:

- The IAM instance profile that is attached to the EC2 instance must allow the s3:GetObject action to the S3 bucket or object in the AWS account. This permission is required to read the object from S3. Option A is incorrect because it specifies the s3:ListBucket action, which is only required to list the objects in the bucket, not to get them.
- The KMS key policy that encrypts the object in the S3 bucket must allow the kms:Decrypt action to the EC2 instance profile ARN. This permission is required to decrypt the object using the KMS key. Option D is correct.
- The security group that is attached to the EC2 instance must have an outbound rule to the S3 managed prefix list over port 443. This rule is required to allow HTTPS traffic from the EC2 instance to S3 within the AWS infrastructure. Option E is correct. Option B is incorrect because it specifies the s3:ListParts action, which is only required for multipart uploads, not for getting objects. Option C is incorrect because it specifies the kms:ListKeys action, which is not required for getting objects. Option F is incorrect because it specifies an inbound rule from the S3 managed prefix list, which is not required for getting objects. Verified References:
- <https://docs.aws.amazon.com/kms/latest/developerguide/control-access.html>
- <https://docs.aws.amazon.com/vpc/latest/userguide/vpc-endpoints-s3.html>

NEW QUESTION 25

A company is attempting to conduct forensic analysis on an Amazon EC2 instance, but the company is unable to connect to the instance by using AWS Systems Manager Session Manager. The company has installed AWS Systems Manager Agent (SSM Agent) on the EC2 instance.

The EC2 instance is in a subnet in a VPC that does not have an internet gateway attached. The company has associated a security group with the EC2 instance. The security group does not have inbound or outbound rules. The subnet's network ACL allows all inbound and outbound traffic.

Which combination of actions will allow the company to conduct forensic analysis on the EC2 instance without compromising forensic data? (Select THREE.)

- A. Update the EC2 instance security group to add a rule that allows outbound traffic on port 443 for 0.0.0.0/0.
- B. Update the EC2 instance security group to add a rule that allows inbound traffic on port 443 to the VPC's CIDR range.
- C. Create an EC2 key pair
- D. Associate the key pair with the EC2 instance.
- E. Create a VPC interface endpoint for Systems Manager in the VPC where the EC2 instance is located.
- F. Attach a security group to the VPC interface endpoint
- G. Allow inbound traffic on port 443 to the VPC's CIDR range.
- H. Create a VPC interface endpoint for the EC2 instance in the VPC where the EC2 instance is located.

Answer: BCF

NEW QUESTION 30

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

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

Answer: D

Explanation:

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

NEW QUESTION 33

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

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

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

Answer: DE

Explanation:

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

NEW QUESTION 37

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

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

Which solution will meet these requirements?

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

Answer: C

Explanation:

- > Amazon S3 is a service that provides scalable, durable, and secure object storage. You can use Amazon S3 to store and retrieve any amount of data from anywhere on the web1
- > AWS Elastic Load Balancing is a service that distributes incoming application or network traffic across multiple targets, such as EC2 instances, containers, or IP addresses. You can use Elastic Load Balancing to increase the availability and fault tolerance of your applications2

- > Elastic Load Balancing supports access logging, which captures detailed information about requests sent to your load balancer. Each log contains information such as the time the request was received, the client's IP address, latencies, request paths, and server responses. You can use access logs to analyze traffic patterns and troubleshoot issues.
- > You can configure your load balancer to store access logs in an Amazon S3 bucket that you specify. You can also specify the interval for publishing the logs, which can be 5 or 60 minutes. The logs are stored in a hierarchical folder structure by load balancer name, IP address, year, month, day, and time.
- > Amazon Athena is a service that allows you to analyze data in Amazon S3 using standard SQL. You can use Athena to run ad-hoc queries and get results in seconds. Athena is serverless, so there is no infrastructure to manage and you pay only for the queries that you run.
- > You can use Athena to search the access logs that are stored in your S3 bucket. You can create a table in Athena that maps to your S3 bucket and then run SQL queries on the table. You can also use the Athena console or API to view and download the query results.
- > You can also use Athena to create queries for the required metrics, such as the number of requests per cipher or protocol. You can then publish the metrics to Amazon CloudWatch, which is a service that monitors and manages your AWS resources and applications. You can use CloudWatch to collect and track metrics, create alarms, and automate actions based on the state of your resources.
- > By using this solution, you can meet the requirements of ensuring that all the load balancer logs are centralized and searchable for auditing and that metrics are generated to show which ciphers are in use.

NEW QUESTION 41

A company created an IAM account for its developers to use for testing and learning purposes. Because the IAM account will be shared among multiple teams of developers, the company wants to restrict the ability to stop and terminate Amazon EC2 instances so that a team can perform these actions only on the instances it owns.

Developers were instructed to tag all their instances with a Team tag key and use the team name in the tag value. One of the first teams to use this account is Business Intelligence. A security engineer needs to develop a highly scalable solution for providing developers with access to the appropriate resources within the account. The security engineer has already created individual IAM roles for each team.

Which additional configuration steps should the security engineer take to complete the task?

A. For each team, create an IAM policy similar to the one that follows. Populate the ec2: ResourceTag/Team condition key with a proper team name. Attach resulting policies to the corresponding IAM roles.

```

{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "NotAction": [
        "ec2:StopInstances",
        "ec2:TerminateInstances"
      ],
      "Resource": "*"
    },
    {
      "Effect": "Allow",
      "Action": [
        "ec2:StopInstances",
        "ec2:TerminateInstances"
      ],
      "Resource": "*",
      "Condition": {
        "StringEquals": {
          "ec2:ResourceTag/Team": "BusinessIntelligence"
        }
      }
    }
  ]
}

```

B. For each team create an IAM policy similar to the one that follows. Populate the IAM TagKeys/Team condition key with a proper team name.

C. Attach the resulting policies to the corresponding IAM roles.

```

{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "NotAction": [
        "ec2:StopInstances",
        "ec2:TerminateInstances"
      ],
      "Resource": "*"
    },
    {
      "Effect": "Allow",
      "Action": [
        "ec2:StopInstances",
        "ec2:TerminateInstances"
      ],
      "Resource": "*",
      "Condition": {
        "ForAnyValue:StringEquals": {
          "aws:TagKeys/Team": "BusinessIntelligence"
        }
      }
    }
  ]
}

```

D. Tag each IAM role with a Team tag key.

E. and use the team name in the tag value.

F. Create an IAM policy similar to the one that follows, and attach it to all the IAM roles used by developers.

```

{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "NotAction": [
        "ec2:StopInstances",
        "ec2:TerminateInstances"
      ],
      "Resource": "*"
    },
    {
      "Effect": "Allow",
      "Action": [
        "ec2:StopInstances",
        "ec2:TerminateInstances"
      ],
      "Resource": "*",
      "Condition": {
        "StringEquals": {
          "ec2:ResourceTag/Team": "${aws:PrincipalTag/Team}"
        }
      }
    }
  ]
}

```

- G. Tag each IAM role with the Team key, and use the team name in the tag value
- H. Create an IAM policy similar to the one that follows, and apply it to all the IAM roles used by developers.

```

{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "NotAction": [
        "ec2:StopInstances",
        "ec2:TerminateInstances"
      ],
      "Resource": "*"
    },
    {
      "Effect": "Allow",
      "Action": [
        "ec2:StopInstances",
        "ec2:TerminateInstances"
      ],
      "Resource": "*",
      "Condition": {
        "ForAnyValue:StringEquals": {
          "aws:TagKeys/Team": "${aws:PrincipalTag/Team}"
        }
      }
    }
  ]
}

```

Answer: A

NEW QUESTION 44

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

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

Answer: B

Explanation:

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

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

The other options are incorrect because:

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

References:

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

NEW QUESTION 47

An international company wants to combine AWS Security Hub findings across all the company's AWS Regions and from multiple accounts. In addition, the company wants to create a centralized custom dashboard to correlate these findings with operational data for deeper analysis and insights. The company needs an analytics tool to search and visualize Security Hub findings. Which combination of steps will meet these requirements? (Select THREE.)

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

Answer: BDF

Explanation:

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

NEW QUESTION 49

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

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

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

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

Answer: AC

Explanation:

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

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

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

NEW QUESTION 50

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

Please select:

- A. A Bastion host should be on a private subnet and never a public subnet due to security concerns
- B. A Bastion host sits on the outside of an internal network and is used as a gateway into the private network and is considered the critical strong point of the network
- C. Bastion hosts allow users to log in using RDP or SSH and use that session to SSH into internal network to access private subnet resources.
- D. A Bastion host should maintain extremely tight security and monitoring as it is available to the public

Answer: C

Explanation:

A bastion host is a special purpose computer on a network specifically designed and configured to withstand attacks. The computer generally hosts a single application, for example a proxy server, and all other services are removed or limited to reduce the threat to the computer.

In IAM, A bastion host is kept on a public subnet. Users log on to the bastion host via SSH or RDP and then use that session to manage other hosts in the private subnets.

Options A and B are invalid because the bastion host needs to sit on the public network. Option D is invalid because bastion hosts are not used for monitoring For more information on bastion hosts, just browse to the below URL:

<https://docs.IAM.amazon.com/quickstart/latest/linux-bastion/architecture.html>

The correct answer is: Bastion hosts allow users to log in using RDP or SSH and use that session to SSH into internal network to access private subnet resources. Submit your Feedback/Queries to our Experts

NEW QUESTION 55

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

Specifically, the security team wants EventBridge to watch for the s3:PutObjectAcl, s3:DeleteBucketPolicy, and s3:PutBucketPolicy API invocation logs from

CloudTrail. While developing the solution in a single account, the security team discovers that the s3:PutObjectAcl API call does not invoke an EventBridge event. However, the s3:DeleteBucketPolicy API call and the s3:PutBucketPolicy API call do invoke an event. The security team has enabled CloudTrail for AWS management events with a basic configuration in the AWS Region in which EventBridge is being tested. Verification of the EventBridge event pattern indicates that the pattern is set up correctly. The security team must implement a solution so that the s3:PutObjectAcl API call will invoke an EventBridge event. The solution must not generate false notifications. Which solution will meet these requirements?

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

Answer: D

Explanation:

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

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

References:

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

NEW QUESTION 60

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

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

Answer: ACD

NEW QUESTION 61

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 65

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

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

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

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

Answer: C

Explanation:

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

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

The other options are incorrect for the following reasons:

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

NEW QUESTION 67

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 72

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

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

Which solution will meet these requirements?

- A. Activate Amazon GuardDuty in each production account.
- B. In a dedicated logging account.
- C. Aggregate all GuardDuty logs from each production account.
- D. Remediate incidents by configuring GuardDuty to directly invoke an AWS Lambda function.
- E. Configure the Lambda function to also publish notifications to the SNS topic.

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

Answer: D

Explanation:

The correct answer is D.

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

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

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

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

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

References:

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

NEW QUESTION 77

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

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

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

How will the security engineer be able to comply with these requirements?

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

Answer: C

Explanation:

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

NEW QUESTION 82

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 role
- 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 time
- E. Review the IAM Access Advisor tab for all federated roles
- 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 event
- 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 84

A startup company is using a single AWS account that has resources in a single AWS Region. A security engineer configures an AWS Cloud Trail trail in the same Region to deliver log files to an Amazon S3 bucket by using the AWS CLI.

Because of expansion, the company adds resources in multiple Regions. The security engineer notices that the logs from the new Regions are not reaching the S3 bucket.

What should the security engineer do to fix this issue with the LEAST amount of operational overhead?

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

Answer: D

Explanation:

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

NEW QUESTION 87

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

Which solution meets these criteria?

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

Answer: A

Explanation:

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

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

A customer managed CMK is a KMS key that you create, own, and manage in your AWS account. You have full control over the key configuration, permissions, rotation, and deletion. You can use a customer managed CMK to encrypt and decrypt data in AWS services that are integrated with AWS KMS, such as Amazon EBS¹.

A customer managed CMK can use either AWS provided key material or customer provided key material. AWS provided key material is generated by AWS KMS and never leaves the service unencrypted. Customer provided key material is generated outside of AWS KMS and imported into a customer managed CMK. You can specify an expiration date for the imported key material, after which the CMK becomes unusable until you reimport new key material².

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

* B. A customer managed CMK that uses AWS provided key material does not expire automatically. You can enable automatic rotation of the key material every year, but this does not prevent access to the data encrypted with the previous key material. You would need to manually delete the CMK and its backing key material to make the data inaccessible³.

* C. An AWS managed CMK is a KMS key that is created, owned, and managed by an AWS service on your behalf. You have limited control over the key configuration, permissions, rotation, and deletion. You cannot use an AWS managed CMK to encrypt data in other AWS services or applications. You also cannot set an expiration date for the key material of an AWS managed CMK4.

* D. Operation system-native encryption that uses GnuPG is not a solution that uses AWS KMS. GnuPG is a command line tool that implements the OpenPGP standard for encrypting and signing data. It does not integrate with Amazon EBS or other AWS services. It also does not provide a way to automatically expire the key material used for encryption5.

References:

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

NEW QUESTION 89

A Security Engineer is troubleshooting an issue with a company's custom logging application. The application logs are written to an Amazon S3 bucket with event notifications enabled to send events to an Amazon SNS topic. All logs are encrypted at rest using an IAM KMS CMK. The SNS topic is subscribed to an encrypted Amazon SQS queue. The logging application polls the queue for new messages that contain metadata about the S3 object. The application then reads the content of the object from the S3 bucket for indexing.

The Logging team reported that Amazon CloudWatch metrics for the number of messages sent or received is showing zero. No logs are being received.

What should the Security Engineer do to troubleshoot this issue?

A) Add the following statement to the IAM managed CMKs:

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

B)

Add the following statement to the CMK key policy:

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

C)

Add the following statement to the CMK key policy:

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

D)

Add the following statement to the CMK key policy:

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

A. Option A

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

Answer: D

NEW QUESTION 90

A company is running an application in The eu-west-1 Region. The application uses an IAM Key Management Service (IAM KMS) CMK to encrypt sensitive data. The company plans to deploy the application in the eu-north-1 Region.

A security engineer needs to implement a key management solution for the application deployment in the new Region. The security engineer must minimize changes to the application code.

Which change should the security engineer make to the IAM KMS configuration to meet these requirements?

- A. Update the key policies in eu-west-1. Point the application in eu-north-1 to use the same CMK as the application in eu-west-1.
- B. Allocate a new CMK to eu-north-1 to be used by the application that is deployed in that Region.
- C. Allocate a new CMK to eu-north-1. Create the same alias name for both key
- D. Configure the application deployment to use the key alias.
- E. Allocate a new CMK to eu-north-1. Create an alias for eu-'-1. Change the application code to point to the alias for eu-'-1.

Answer: B

NEW QUESTION 95

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-1 b subnet. During an initial investigation, a security engineer discovers that the suspicious instance is the only instance that runs in the subnet.

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

- A. Log in to the suspicious instance and use the netstat command to identify remote connections. Use the IP addresses from these remote connections to create deny rules in the security group of the instance. Install diagnostic tools on the instance for investigation. Update the outbound network ACL for the subnet 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 98

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.

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.
- 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.
- 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.

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 99

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 find-ing
- B. Configure an AWS Lambda function as the target for the rule to reme-diate the findings.
- C. Set up a custom action in Security Hu
- D. Configure the custom action to call AWS Systems Manager Automation runbooks to remediate the findings.
- E. Set up a custom action in Security Hu
- 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 103

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

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

Answer: A

NEW QUESTION 106

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

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

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

Answer: BD

NEW QUESTION 111

A company accidentally deleted the private key for an Amazon Elastic Block Store (Amazon EBS)-backed Amazon EC2 instance. A security engineer needs to regain access to the instance.

Which combination of steps will meet this requirement? (Choose two.)

- A. Stop the instanc
- B. Detach the root volum
- C. Generate a new key pair.
- D. Keep the instance runnin
- E. Detach the root volum
- F. Generate a new key pair.
- G. When the volume is detached from the original instance, attach the volume to another instance as a data volum
- H. Modify the authorized_keys file with a new public ke
- I. Move the volume back to the original instanc
- J. Start the instance.
- K. When the volume is detached from the original instance, attach the volume to another instance as a data volum
- L. Modify the authorized_keys file with a new private ke
- M. Move the volume back to the original instanc
- N. Start the instance.
- O. When the volume is detached from the original instance, attach the volume to another instance as a data volum
- P. Modify the authorized_keys file with a new public ke
- Q. Move the volume back to the original instance that is running.

Answer: AC

Explanation:

If you lose the private key for an EBS-backed instance, you can regain access to your instance. You must stop the instance, detach its root volume and attach it to another instance as a data volume, modify the authorized_keys file with a new public key, move the volume back to the original instance, and restart the instance. <https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/TroubleshootingInstancesConnecting.html#replacing>

NEW QUESTION 113

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

Which solution will meet these requirements?

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

Answer: B

NEW QUESTION 114

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

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

Answer: B

Explanation:

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

NEW QUESTION 117

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

The company issues a new security policy that contains the following requirements:

- No AWS account should use a VPC within the AWS account for workloads.
- The company should use a centrally managed VPC that all AWS accounts can access to launch workloads in subnets.
- No AWS account should be able to modify another AWS account's application resources within the centrally managed VPC.
- The centrally managed VPC should reside in an existing AWS account that is named Account-A within an organization.

The company uses an AWS CloudFormation template to create a VPC that contains multiple subnets in Account-A. This template exports the subnet IDs through the CloudFormation Outputs section.

Which solution will complete the security setup to meet these requirements?

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

Answer: C

Explanation:

The correct answer is C. Use AWS Resource Access Manager (AWS RAM) to share Account-A's VPC subnets with the remaining member accounts. Configure the member accounts to use the shared subnets to launch workloads.

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

The other options are incorrect because:

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

References:

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

NEW QUESTION 119

A company uses AWS Organizations. The company wants to implement short-term credentials for third-party AWS accounts to use to access accounts within the company's organization. Access is for the AWS Management Console and third-party software-as-a-service (SaaS) applications. Trust must be enhanced to prevent two external accounts from using the same credentials. The solution must require the least possible operational effort. Which solution will meet these requirements?

- A. Use a bearer token authentication with OAuth or SAML to manage and share a central Amazon Cognito user pool across multiple Amazon API Gateway APIs.
- B. Implement AWS IAM Identity Center (AWS Single Sign-On), and use an identity source of choice. Grant access to users and groups from other accounts by using permission sets that are assigned by account.
- C. Create a unique IAM role for each external account
- D. Create a trust policy
- E. Use AWS Secrets Manager to create a random external key.
- F. Create a unique IAM role for each external account
- G. Create a trust policy that includes a condition that uses the sts:ExternalId condition key.

Answer: D

Explanation:

The correct answer is D.

To implement short-term credentials for third-party AWS accounts, you can use IAM roles and trust policies. A trust policy is a JSON policy document that defines who can assume the role. You can specify the AWS account ID of the third-party account as a principal in the trust policy, and use the sts:ExternalId condition key to enhance the security of the role. The sts:ExternalId condition key is a unique identifier that is agreed upon by both parties and included in the AssumeRole request. This way, you can prevent the "confused deputy" problem, where an unauthorized party can use the same role as a legitimate party.

Option A is incorrect because bearer token authentication with OAuth or SAML is not suitable for granting access to AWS accounts and resources. Amazon Cognito and API Gateway are used for building web and mobile applications that require user authentication and authorization.

Option B is incorrect because AWS IAM Identity Center (AWS Single Sign-On) is a service that simplifies the management of access to multiple AWS accounts and cloud applications for your workforce users. It does not support granting access to third-party AWS accounts.

Option C is incorrect because using AWS Secrets Manager to create a random external key is not necessary and adds operational complexity. You can use the sts:ExternalId condition key instead to provide a unique identifier for each external account.

NEW QUESTION 123

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

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

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

Answer: C

Explanation:

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

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

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

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

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

The other options are incorrect because:

- A. Using VPC Traffic Mirroring would not capture the DNS queries that go to the on-premises DNS servers, because Traffic Mirroring only copies network traffic from an elastic network interface of an EC2 instance to a target for analysis³. Traffic Mirroring does not include traffic that goes through a Route 53 Resolver outbound endpoint, which is used to forward queries to on-premises DNS servers⁴. Therefore, this solution would not meet the requirements.
- B. Configuring VPC flow logs on all relevant VPCs would not capture the DNS name that was requested in Route 53 Resolver, because flow logs only record

information about the IP traffic going to and from network interfaces in a VPC. 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. 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 126

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 128

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

Please select:

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

Answer: A

Explanation:

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

Options B and C are invalid because these options will not reduce network latency. Option D is invalid because this is only used to connect 2 VPC's. For more information on IAM Direct Connect, just browse to the below URL: <https://IAM.amazon.com/directconnect>

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

NEW QUESTION 131

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 enabled¹. This command also shows the date and time when the CMK was last rotated². The other options are not valid ways to check the CMK rotation status.

NEW QUESTION 136

.....

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