

# Amazon-Web-Services

## Exam Questions SCS-C02

AWS Certified Security - Specialty



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

Which solution will meet this requirement?

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

**Answer:** D

#### Explanation:

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

For more information, see [Monitoring AWS Shield Advanced with Amazon CloudWatch and AWS Shield Advanced metrics and alarms](#).

### NEW QUESTION 3

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 4

A company's security team is building a solution for logging and visualization. The solution will assist the company with the large variety and velocity of data that it receives from IAM across multiple accounts. The security team has enabled IAM CloudTrail and VPC Flow Logs in all of its accounts. In addition, the company has an organization in IAM Organizations and has an IAM Security Hub master account.

The security team wants to use Amazon Detective. However, the security team cannot enable Detective and is unsure why. What must the security team do to enable Detective?

- A. Enable Amazon Macie so that Security Hub will allow Detective to process findings from Macie.
- B. Disable IAM Key Management Service (IAM KMS) encryption on CloudTrail logs in every member account of the organization.
- C. Enable Amazon GuardDuty on all member accounts. Try to enable Detective in 48 hours.
- D. Ensure that the principal that launches Detective has the organizations:ListAccounts permission.

**Answer:** D

#### NEW QUESTION 5

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

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

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

**Answer:** B

#### NEW QUESTION 6

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

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

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

**Answer:** D

#### Explanation:

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

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

The other options are incorrect for the following reasons:

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

References:

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

#### NEW QUESTION 7

A company is using an AWS Key Management Service (AWS KMS) AWS owned key in its application to encrypt files in an AWS account. The company's security team wants the ability to change to new key material for new files whenever a potential key breach occurs. A security engineer must implement a solution that gives

the security team the ability to change the key whenever the team wants to do so  
 Which solution will meet these requirements?

- A. Create a new customer managed key Add a key rotation schedule to the key Invoke the key rotation schedule every time the security team requests a key change
- B. Create a new AWS managed key Add a key rotation schedule to the key Invoke the key rotation schedule every time the security team requests a key change
- C. Create a key alias Create a new customer managed key every time the security team requests a key change Associate the alias with the new key
- D. Create a key alias Create a new AWS managed key every time the security team requests a key change Associate the alias with the new key

**Answer: A**

**Explanation:**

To meet the requirement of changing the key material for new files whenever a potential key breach occurs, the most appropriate solution would be to create a new customer managed key, add a key rotation schedule to the key, and invoke the key rotation schedule every time the security team requests a key change.  
 References: : Rotating AWS KMS keys - AWS Key Management Service

**NEW QUESTION 8**

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

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

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

**Answer: B**

**Explanation:**

The correct answer is B. Use data key caching. Use the local cache that the AWS Encryption SDK provides with a caching cryptographic materials manager. This answer is correct because data key caching can improve performance, reduce cost, and help the company stay within the service limits of AWS KMS. Data key caching stores data keys and related cryptographic material in a cache, and reuses them for encryption and decryption operations. This reduces the number of requests to AWS KMS endpoints and avoids throttling. The AWS Encryption SDK provides a local cache and a caching cryptographic materials manager (caching CMM) that interacts with the cache and enforces security thresholds that the company can 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 9**

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

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

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

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

**Answer: BD**

**Explanation:**

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

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

To use CloudTrail with multiple AWS accounts and regions, you need to enable AWS Organizations with all features enabled. This allows you to centrally manage your accounts and apply policies across your organization. You can also use CloudTrail as a service principal for AWS Organizations, which lets you create an organization trail that applies to all accounts in your organization. An organization trail logs events for all AWS Regions and delivers the log files to an S3 bucket that you specify.

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

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

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

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

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

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

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

#### NEW QUESTION 10

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

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

**Answer: C**

#### Explanation:

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

#### NEW QUESTION 10

A company 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 15

A company uses SAML federation to grant users access to AWS accounts. A company workload that is in an isolated AWS account runs on immutable infrastructure with no human access to Amazon EC2. The company requires a specialized user known as a break glass user to have access to the workload AWS account and instances in the case of SAML errors. A recent audit discovered that the company did not create the break glass user for the AWS account that contains the workload.

The company must create the break glass user. The company must log any activities of the break glass user and send the logs to a security team.

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

- A. Create a local individual break glass IAM user for the security team.

- B. Create a trail in AWS CloudTrail that has Amazon CloudWatch Logs turned on.
- C. Use Amazon EventBridge to monitor local user activities.
- D. Create a break glass EC2 key pair for the AWS account.
- E. Provide the key pair to the security team.
- F. Use AWS CloudTrail to monitor key pair activities.
- G. Send notifications to the security team by using Amazon Simple Notification Service (Amazon SNS).
- H. Create a break glass IAM role for the account.
- I. Allow security team members to perform the AssumeRoleWithSAML operation.
- J. Create an AWS Cloud Trail trail that has Amazon CloudWatch Logs turned on.
- K. Use Amazon EventBridge to monitor security team activities.
- L. Create a local individual break glass IAM user on the operating system level of each workload instance. Configure unrestricted security groups on the instances to grant access to the break glass IAM users.
- M. Configure AWS Systems Manager Session Manager for Amazon EC2. Configure an AWS Cloud Trail filter based on Session Manager.
- N. Send the results to an Amazon Simple Notification Service (Amazon SNS) topic.

**Answer:** AE

**Explanation:**

The combination of solutions that will meet the requirements are:

- A. Create a local individual break glass IAM user for the security team. Create a trail in AWS CloudTrail that has Amazon CloudWatch Logs turned on. Use Amazon EventBridge to monitor local user activities. This is a valid solution because it allows the security team to access the workload AWS account and instances using a local IAM user that does not depend on SAML federation. It also enables logging and monitoring of the break glass user activities using AWS CloudTrail, Amazon CloudWatch Logs, and Amazon EventBridge.
- E. Configure AWS Systems Manager Session Manager for Amazon EC2. Configure an AWS CloudTrail filter based on Session Manager. Send the results to an Amazon Simple Notification Service (Amazon SNS) topic. This is a valid solution because it allows the security team to access the workload instances without opening any inbound ports or managing SSH keys or bastion hosts. It also enables logging and notification of the break glass user activities using AWS CloudTrail, Session Manager, and Amazon SNS.

The other options are incorrect because:

- B. Creating a break glass EC2 key pair for the AWS account and providing it to the security team is not a valid solution, because it requires opening inbound ports on the instances and managing SSH keys, which increases the security risk and complexity.
- C. Creating a break glass IAM role for the account and allowing security team members to perform the AssumeRoleWithSAML operation is not a valid solution, because it still depends on SAML federation, which might not work in case of SAML errors.
- D. Creating a local individual break glass IAM user on the operating system level of each workload instance and configuring unrestricted security groups on the instances to grant access to the break glass IAM users is not a valid solution, because it requires opening inbound ports on the instances and managing multiple local users, which increases the security risk and complexity.

References:

1: Creating an IAM User in Your AWS Account 2: Creating a Trail - AWS CloudTrail 3: Using Amazon EventBridge with AWS CloudTrail 4: Setting up Session Manager - AWS Systems Manager 5: Logging Session Manager sessions - AWS Systems Manager 6: Amazon Simple Notification Service 7: Connecting to your Linux instance using SSH - Amazon Elastic Compute Cloud 8: AssumeRoleWithSAML - AWS Security Token Service 9: IAM Users - AWS Identity and Access Management

**NEW QUESTION 17**

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 20**

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

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

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

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

- A. Configure Amazon Macie to identify any objects in DOC-EXAMPLE-BUCKET1 that contain PII and that were available to the access key.
- B. Use Amazon CloudWatch Logs Insights to identify any objects in DOC-EXAMPLE-BUCKET1 that contain PII and that were available to the access key.
- C. Use Amazon OpenSearch Service (Amazon Elasticsearch Service) to query the CloudTrail logs in DOC-EXAMPLE-BUCKET2 for API calls that used the access key.

key to access an object that contained PII.

D. Use Amazon Athena to query the CloudTrail logs in DOC-EXAMPLE-BUCKET2 for any API calls that used the access key to access an object that contained PII.

E. Use AWS Identity and Access Management Access Analyzer to identify any API calls that used the access key to access objects that contained PII in DOC-EXAMPLE-BUCKET1.

**Answer:** AD

#### NEW QUESTION 21

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

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

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

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

**Answer:** D

#### Explanation:

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

➤ Create a service control policy (SCP) that defines the maximum permissions that can be granted to any IAM role in the organization. An SCP is a type of policy that you can use with AWS Organizations to manage permissions for all accounts in your organization. SCPs restrict permissions for entities in member accounts, including each AWS account root user, IAM users, and roles. For more information, see [Service control policies overview](#).

➤ Create a permissions boundary for IAM roles that matches the SCP. A permissions boundary is an advanced feature for using a managed policy to set the maximum permissions that an identity-based policy can grant to an IAM entity. A permissions boundary allows an entity to perform only the actions that are allowed by both its identity-based policies and its permissions boundaries. For more information, see [Permissions boundaries for IAM entities](#).

➤ Add the SCP to the root organizational unit (OU) so that it applies to all accounts in the organization.

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

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

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

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

Verified References:

➤ [https://docs.aws.amazon.com/IAM/latest/UserGuide/access\\_policies\\_boundaries.html](https://docs.aws.amazon.com/IAM/latest/UserGuide/access_policies_boundaries.html)

#### NEW QUESTION 25

A security engineer is using AWS Organizations and wants to optimize SCPs. The security engineer needs to ensure that the SCPs conform to best practices. Which approach should the security engineer take to meet this requirement?

- A. Use AWS IAM Access Analyzer to analyze the policie
- B. View the findings from policy validation checks.
- C. Review AWS Trusted Advisor checks for all accounts in the organization.
- D. Set up AWS Audit Manage
- E. Run an assessment for all AWS Regions for all accounts.
- F. Ensure that Amazon Inspector agents are installed on all Amazon EC2 in-stances in all accounts.

**Answer:** A

#### NEW QUESTION 27

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

Which combination of steps should the company take to see findings from accounts that are outside the organization that includes the Security Hub administrator account? (Select TWO.)

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

**Answer:** CE

**Explanation:**

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

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

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

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

Verified References:

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

**NEW QUESTION 29**

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<sup>12</sup>.

The other options are incorrect because:

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

References:

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

**NEW QUESTION 34**

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

Please select:

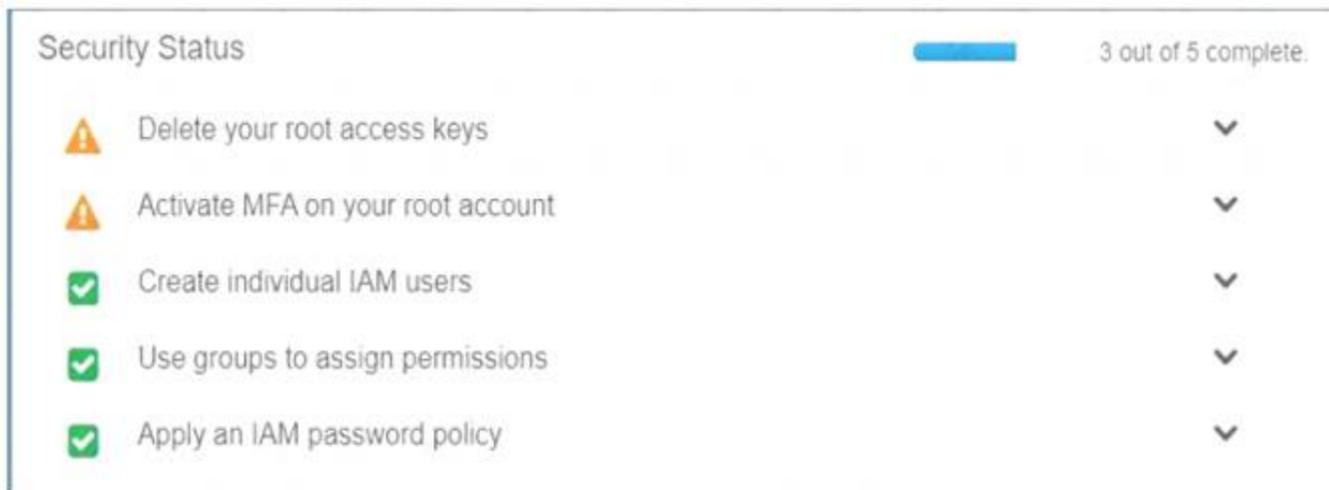
- A. Use short but complex password on the root account and any administrators.
- B. Use IAM Geo-Lock and disallow anyone from logging in except for in your city.
- C. Use MFA on all users and accounts, especially on the root account.
- D. Don't write down or remember the root account password after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

**Answer: C**

**Explanation:**

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

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



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

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

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

### NEW QUESTION 35

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

Which solution will meet these requirements?

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

**Answer: A**

### NEW QUESTION 40

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

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

The security team has enabled CloudTrail for AWS management events with a basic configuration in the AWS Region in which EventBridge is being tested. Verification of the EventBridge event pattern indicates that the pattern is set up correctly. The security team must implement a solution so that the s3:PutObjectAcl API call will invoke an EventBridge event. The solution must not generate false notifications.

Which solution will meet these requirements?

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

**Answer: D**

### Explanation:

The correct answer is D. Enable CloudTrail to monitor data events for read and write operations to S3 buckets. According to the AWS documentation<sup>1</sup>, CloudTrail data events are the resource operations performed on or within a resource. These are also known as data plane operations. Data events are often high-volume activities. For example, Amazon S3 object-level API activity (such as GetObject, DeleteObject, and PutObject) is a data event.

By default, trails do not log data events. To record CloudTrail data events, you must explicitly add the supported resources or resource types for which you want to collect activity. For more information, see Logging data events in the Amazon S3 User Guide<sup>2</sup>. In this case, the security team wants EventBridge to watch for the s3:PutObjectAcl API invocation logs from CloudTrail. This API uses the acl subresource to set the access control list (ACL) permissions for a new or existing object in an S3 bucket<sup>3</sup>. This is a data event that affects the S3 object resource type. Therefore, the security team must enable CloudTrail to monitor data events for read and write operations to S3 buckets in order to invoke an EventBridge event for this API call. The other options are incorrect because:

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

References:

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

**NEW QUESTION 44**

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 48**

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

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

**Answer:** C

**Explanation:**

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

**NEW QUESTION 51**

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

- A. Use IPv6 addresses that are configured for hostnames.
- B. Configure external DNS resolvers as internal resolvers that are visible only to IAM.
- C. Use IAM DNS resolvers for all EC2 instances.
- D. Configure a third-party DNS resolver with logging for all EC2 instances.

**Answer:** C

**Explanation:**

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

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

**NEW QUESTION 55**

A company has multiple departments. Each department has its own IAM account. All these accounts belong to the same organization in IAM Organizations. A large .csv file is stored in an Amazon S3 bucket in the sales department's IAM account. The company wants to allow users from the other accounts to access the .csv file's content through the combination of IAM Glue and Amazon Athena. However, the company does not want to allow users from the other accounts to access other files in the same folder. Which solution will meet these requirements?

- A. Apply a user policy in the other accounts to allow IAM Glue and Athena to access the .csv file.
- B. Use S3 Select to restrict access to the .csv file.
- C. In IAM Glue Data Catalog, use S3 Select as the source of the IAM Glue database.
- D. Define an IAM Glue Data Catalog resource policy in IAM Glue to grant cross-account S3 object access to the .csv file.
- E. Grant IAM Glue access to Amazon S3 in a resource-based policy that specifies the organization as the principal.

**Answer:** A

**NEW QUESTION 60**

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

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

**Answer:** A

**Explanation:**

The IAM Documentation mentions the following

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

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

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

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

**NEW QUESTION 63**

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 67**

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 69**

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

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

A)

Apply the following SCP:

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

B)

Apply the following SCP:

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

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

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

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

**Answer: ACE**

#### NEW QUESTION 70

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 74

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 77**

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

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

**Answer: CE**

**NEW QUESTION 82**

A company's public Application Load Balancer (ALB) recently experienced a DDoS attack. To mitigate this issue, the company deployed Amazon CloudFront in front of the ALB so that users would not directly access the Amazon EC2 instances behind the ALB.

The company discovers that some traffic is still coming directly into the ALB and is still being handled by the EC2 instances.

Which combination of steps should the company take to ensure that the EC2 instances will receive traffic only from CloudFront? (Choose two.)

- A. Configure CloudFront to add a cache key policy to allow a custom HTTP header that CloudFront sends to the ALB.
- B. Configure CloudFront to add a custom: HTTP header to requests that CloudFront sends to the ALB.
- C. Configure the ALB to forward only requests that contain the custom HTTP header.
- D. Configure the ALB and CloudFront to use the X-Forwarded-For header to check client IP addresses.
- E. Configure the ALB and CloudFront to use the same X.509 certificate that is generated by AWS Certificate Manager (ACM).

**Answer: BC**

**Explanation:**

To prevent users from directly accessing an Application Load Balancer and allow access only through CloudFront, complete these high-level steps: Configure CloudFront to add a custom HTTP header to requests that it sends to the Application Load Balancer. Configure the Application Load Balancer to only forward requests that contain the custom HTTP header. (Optional) Require HTTPS to improve the security of this solution.

<https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/restrict-access-to-load-balancer.html>

**NEW QUESTION 83**

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

Which solution will meet these requirements?

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

**Answer: B**

**NEW QUESTION 85**

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

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

- A. Use a dynamic reference in the CloudFormation template to reference the database credentials in Secrets Manager.

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

**Answer:** A

**Explanation:**

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

**NEW QUESTION 87**

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

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

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

**Answer:** C

**Explanation:**

The correct answer is C. The GuardDuty integration with Security Hub was never activated in the AWS account where the finding was generated. According to the AWS documentation<sup>1</sup>, GuardDuty findings are automatically sent to Security Hub only if the GuardDuty integration with Security Hub is enabled in the same account and Region. This means that the security tooling account, which is the delegated administrator for both GuardDuty and Security Hub, must enable the GuardDuty integration with Security Hub in each member account and Region where GuardDuty is enabled. Otherwise, the findings from GuardDuty will not be visible in Security Hub.

The other options are incorrect because:

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

References:

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

**NEW QUESTION 92**

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

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

The security engineer deploys a custom Lambda function in the VPC. The custom Lambda function will be responsible for rotating the secret in Secrets Manager. The security engineer edits the DB instance's security group to allow connections from this function. When the function is invoked, the function cannot communicate with Secrets Manager to rotate the secret properly.

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

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

**Answer:** D

**Explanation:**

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

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

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

Lambda function's private subnet during the configuration process to allow the function to use the endpoint2.

The other options are incorrect for the following reasons:

- A. An egress-only internet gateway is a VPC component that allows outbound communication over IPv6 from instances in the VPC to the internet, and prevents the internet from initiating an IPv6 connection with the instances3. However, this option does not meet the requirement that the VPC must not send or receive network traffic through the internet. Moreover, an egress-only internet gateway is for use with IPv6 traffic only, and Secrets Manager does not support IPv6 addresses2.
- B. A NAT gateway is a VPC component that enables instances in a private subnet to connect to the internet or other AWS services, but prevents the internet from initiating connections with those instances4. However, this option does not meet the requirement that the VPC must not send or receive network traffic through the internet. Additionally, a NAT gateway requires an elastic IP address, which is a public IPv4 address4.
- C. A VPC peering connection is a networking connection between two VPCs that enables you to route traffic between them using private IPv4 addresses or IPv6 addresses5. However, this option does not work because Secrets Manager does not have a default VPC that can be peered with. Furthermore, a VPC peering connection does not provide a private connection to Secrets Manager APIs without an internet gateway or other devices2.

#### NEW QUESTION 97

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 101

A security engineer is designing an IAM policy for a script that will use the AWS CLI. The script currently assumes an IAM role that is attached to three AWS managed IAM policies: AmazonEC2FullAccess, AmazonDynamoDBFullAccess, and AmazonVPCFullAccess.

The security engineer needs to construct a least privilege IAM policy that will replace the AWS managed IAM policies that are attached to this role.

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

- A. In AWS CloudTrail, create a trail for management event
- B. Run the script with the existing AWS managed IAM policies
- C. Use IAM Access Analyzer to generate a new IAM policy that is based on access activity in the trail
- D. Replace the existing AWS managed IAM policies with the generated IAM policy for the role.
- E. Remove the existing AWS managed IAM policies from the role
- F. Attach the IAM Access Analyzer Role Policy Generator to the role
- G. Run the script
- H. Return to IAM Access Analyzer and generate a least privilege IAM policy
- I. Attach the new IAM policy to the role.
- J. Create an account analyzer in IAM Access Analyzer
- K. Create an archive rule that has a filter that checks whether the PrincipalArn value matches the ARN of the role
- L. Run the script
- M. Remove the existing AWS managed IAM policies from the role.
- N. In AWS CloudTrail, create a trail for management event
- O. Remove the existing AWS managed IAM policies from the role
- P. Run the script
- Q. Find the authorization failure in the trail event that is associated with the script
- R. Create a new IAM policy that includes the action and resource that caused the authorization failure
- S. Repeat the process until the script succeeds
- T. Attach the new IAM policy to the role.

**Answer: A**

#### NEW QUESTION 105

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

Which bucket policy statement meets these requirements?

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

```

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

```

B. A computer code with black text Description automatically generated

```

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

```

C. A computer code with black text Description automatically generated

```

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

```

D. A computer code with black text Description automatically generated

```

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

```

**Answer: B**

**Explanation:**

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

**NEW QUESTION 106**

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

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

- A. Configure access logging for the required API stage.
- B. Configure an AWS CloudTrail trail destination for API Gateway event
- C. Configure filters on the userIdentity, userAgent, and sourceIPAddress fields.
- D. Configure an Amazon S3 destination for API Gateway log
- E. Run Amazon Athena queries to analyze API access information.

- F. Use Amazon CloudWatch Logs Insights to analyze API access information.
- G. Select the Enable Detailed CloudWatch Metrics option on the required API stage.

**Answer:** CD

#### NEW QUESTION 107

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

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

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

**Answer:** B

#### NEW QUESTION 108

A company has an AWS account that includes an Amazon S3 bucket. The S3 bucket uses server-side encryption with AWS KMS keys (SSE-KMS) to encrypt all the objects at rest by using a customer managed key. The S3 bucket does not have a bucket policy.

An IAM role in the same account has an IAM policy that allows s3 List\* and s3 Get\* permissions for the S3 bucket. When the IAM role attempts to access an object in the S3 bucket the role receives an access denied message.

Why does the IAM role not have access to the objects that are in the S3 bucket?

- A. The IAM role does not have permission to use the KMS CreateKey operation.
- B. The S3 bucket lacks a policy that allows access to the customer managed key that encrypts the objects.
- C. The IAM role does not have permission to use the customer managed key that encrypts the objects that are in the S3 bucket.
- D. The ACL of the S3 objects does not allow read access for the objects when the objects are encrypted at rest.

**Answer:** C

#### Explanation:

When using server-side encryption with AWS KMS keys (SSE-KMS), the requester must have both Amazon S3 permissions and AWS KMS permissions to access the objects. The Amazon S3 permissions are for the bucket and object operations, such as s3:ListBucket and s3:GetObject. The AWS KMS permissions are for the key operations, such as kms:GenerateDataKey and kms:Decrypt. In this case, the IAM role has the necessary Amazon S3 permissions, but not the AWS KMS permissions to use the customer managed key that encrypts the objects. Therefore, the IAM role receives an access denied message when trying to access the objects. Verified References:

- > <https://docs.aws.amazon.com/AmazonS3/latest/userguide/troubleshoot-403-errors.html>
- > <https://repost.aws/knowledge-center/s3-access-denied-error-kms>
- > <https://repost.aws/knowledge-center/cross-account-access-denied-error-s3>

#### NEW QUESTION 110

A company is using AWS Organizations to manage multiple AWS accounts for its human resources, finance, software development, and production departments. All the company's developers are part of the software development AWS account.

The company discovers that developers have launched Amazon EC2 instances that were preconfigured with software that the company has not approved for use. The company wants to implement a solution to ensure that developers can launch EC2 instances with only approved software applications and only in the software development AWS account.

Which solution will meet these requirements?

- A. In the software development account, create AMIs of preconfigured instances that include only approved software
- B. Include the AMI IDs in the condition section of an AWS CloudFormation template to launch the appropriate AMI based on the AWS Region
- C. Provide the developers with the CloudFormation template to launch EC2 instances in the software development account.
- D. Create an Amazon EventBridge rule that runs when any EC2 RunInstances API event occurs in the software development account
- E. Specify AWS Systems Manager Run Command as a target of the rule
- F. Configure Run Command to run a script that will install all approved software onto the instances that the developers launch.
- G. Use an AWS Service Catalog portfolio that contains EC2 products with appropriate AMIs that include only approved software
- H. Grant the developers permission to portfolio access only the Service Catalog to launch a product in the software development account.
- I. In the management account, create AMIs of preconfigured instances that include only approved software
- J. Use AWS CloudFormation StackSets to launch the AMIs across any AWS account in the organization
- K. Grant the developers permission to launch the stack sets within the management account.

**Answer:** C

#### NEW QUESTION 112

An IAM user receives an Access Denied message when the user attempts to access objects in an Amazon S3 bucket. The user and the S3 bucket are in the same AWS account. The S3 bucket is configured to use server-side encryption with AWS KMS keys (SSE-KMS) to encrypt all of its objects at rest by using a customer managed key from the same AWS account. The S3 bucket has no bucket policy defined. The IAM user has been granted permissions through an IAM policy that

allows the kms:Decrypt permission to the customer managed key. The IAM policy also allows the s3:List\* and s3:Get\* permissions for the S3 bucket and its objects.

Which of the following is a possible reason that the IAM user cannot access the objects in the S3 bucket?

- A. The IAM policy needs to allow the kms:DescribeKey permission.
- B. The S3 bucket has been changed to use the AWS managed key to encrypt objects at rest.
- C. An S3 bucket policy needs to be added to allow the IAM user to access the objects.
- D. The KMS key policy has been edited to remove the ability for the AWS account to have full access to the key.

**Answer: D**

**Explanation:**

The possible reason that the IAM user cannot access the objects in the S3 bucket is D. The KMS key policy has been edited to remove the ability for the AWS account to have full access to the key.

This answer is correct because the KMS key policy is the primary way to control access to the KMS key, and it must explicitly allow the AWS account to have full access to the key. If the KMS key policy has been edited to remove this permission, then the IAM policy that grants kms:Decrypt permission to the IAM user has no effect, and the IAM user cannot decrypt the objects in the S3 bucket<sup>12</sup>.

The other options are incorrect because:

- A. The IAM policy does not need to allow the kms:DescribeKey permission, because this permission is not required for decrypting objects in S3 using SSE-KMS. The kms:DescribeKey permission allows getting information about a KMS key, such as its creation date, description, and key state<sup>3</sup>.
- B. The S3 bucket has not been changed to use the AWS managed key to encrypt objects at rest, because this would not cause an Access Denied message for the IAM user. The AWS managed key is a default KMS key that is created and managed by AWS for each AWS account and Region. The IAM user does not need any permissions on this key to use it for SSE-KMS<sup>4</sup>.
- C. An S3 bucket policy does not need to be added to allow the IAM user to access the objects, because the IAM user already has s3:List\* and s3:Get\* permissions for the S3 bucket and its objects through an IAM policy. An S3 bucket policy is an optional way to grant cross-account access or public access to an S3 bucket<sup>5</sup>.

References:

1: Key policies in AWS KMS 2: Using server-side encryption with AWS KMS keys (SSE-KMS) 3: AWS KMS API Permissions Reference 4: Using server-side encryption with Amazon S3 managed keys (SSE-S3) 5: Bucket policy examples

**NEW QUESTION 114**

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

How should the security engineer prevent unauthorized access to the EC2 instances?

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

**Answer: C**

**Explanation:**

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

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

**NEW QUESTION 118**

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

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

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

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

**Answer: ACE**

**NEW QUESTION 123**

Which of the following bucket policies will ensure that objects being uploaded to a bucket called 'demo' are encrypted.

Please select:

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

```
"Version":"2012-10-17",
  "Id":"PutObj",
  "Statement":[{"Sid":"DenyUploads",
    "Effect":"Deny",
    "Principal":"*",
    "Action":"s3:PutObject",
    "Resource":"arn:aws:s3:::demo/*",
    "Condition":{"StringNotEquals":{"s3:x-amz-server-side-encryption":"aws:kms"}
  }
}
]
```

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

```
"Version":"2012-10-17",
  "Id":"PutObj",
  "Statement":[{"Sid":"DenyUploads",
    "Effect":"Deny",
    "Principal":"*",
    "Action":"s3:PutObject",
    "Resource":"arn:aws:s3:::demo/*",
    "Condition":{"StringEquals":{"s3:x-amz-server-side-encryption":"aws:kms"}
  }
}
]
```

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

```
"Version":"2012-10-17",
  "Id":"PutObj",
  "Statement":[{"Sid":"DenyUploads",
    "Effect":"Deny",
    "Principal":"*",
    "Action":"s3:PutObject",
    "Resource":"arn:aws:s3:::demo/*"
  }
}
]
```

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

```
"Version":"2012-10-17",
  "Id":"PutObj",
  "Statement":[{"Sid":"DenyUploads",
    "Effect":"Deny",
    "Principal":"*",
    "Action":"s3:PutObjectEncrypted",
    "Resource":"arn:aws:s3:::demo/*"
  }
}
]
```

**Answer: A**

**Explanation:**

The condition of "s3:x-amz-server-side-encryption":"IAM:kms" ensures that objects uploaded need to be encrypted. Options B,C and D are invalid because you have to ensure the condition of ns3:x-amz-server-side-encryption":"IAM:kms" is present

For more information on IAM KMS best practices, just browse to the below URL: <https://dl.IAMstatic.com/whitepapers/IAM-kms-best-practices.pdf>  
Submit your Feedback/Queries to our Expert

#### NEW QUESTION 125

A developer is building a serverless application hosted on IAM that uses Amazon Redshift in a data store. The application has separate modules for read/write and read-only functionality. The modules need their own database users for compliance reasons.

Which combination of steps should a security engineer implement to grant appropriate access? (Select TWO )

- A. Configure cluster security groups for each application module to control access to database users that are required for read-only and read/write.
- B. Configure a VPC endpoint for Amazon Redshift. Configure an endpoint policy that maps database users to each application module, and allow access to the tables that are required for read-only and read/write.
- C. Configure an IAM policy for each module. Specify the ARN of an Amazon Redshift database user that allows the GetClusterCredentials API call.
- D. Create local database users for each module.
- E. Configure an IAM policy for each module. Specify the ARN of an IAM user that allows the GetClusterCredentials API call.

**Answer:** CD

#### Explanation:

To grant appropriate access to the application modules, the security engineer should do the following:

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

#### NEW QUESTION 126

A company uses Amazon EC2 Linux instances in the AWS Cloud. A member of the company's security team recently received a report about common vulnerability identifiers on the instances.

A security engineer needs to verify patching and perform remediation if the instances do not have the correct patches installed. The security engineer must determine which EC2 instances are at risk and must implement a solution to automatically update those instances with the applicable patches.

What should the security engineer do to meet these requirements?

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

**Answer:** A

#### Explanation:

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

#### NEW QUESTION 130

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

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

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

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

**Answer:** B

#### Explanation:

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

The `chmod` command takes a numeric argument that represents the permissions for the owner, group, and others in octal notation. Each digit corresponds to a set of permissions: read (4), write (2), and execute (1). The digits are added together to get the final permissions for each category. For example, `0400` means that the owner has read permission (4) and no other permissions (0), and the group and others have no permissions at all (0).

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

Verified References:

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

#### NEW QUESTION 134

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

The team members access the account by assuming a role that has a specific set of permissions that are necessary for the job responsibilities of the team.

members. All team members have permissions to perform operations on the stacks.

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

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

**Answer:** BDF

#### NEW QUESTION 135

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

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

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

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

Which solution will meet these requirements?

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

**Answer:** D

#### Explanation:

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

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

#### NEW QUESTION 138

An organization must establish the ability to delete an IAM KMS Customer Master Key (CMK) within a

24-hour timeframe to keep it from being used for encrypt or decrypt operations Which of the following actions will address this requirement?

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

**Answer:** C

#### Explanation:

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

#### NEW QUESTION 139

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

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

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

**Answer:** ADF

#### NEW QUESTION 141

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

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

**Answer: C**

#### NEW QUESTION 143

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

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

**Answer: A**

#### Explanation:

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

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

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

Verified References:

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

#### NEW QUESTION 145

A security engineer is configuring account-based access control (ABAC) to allow only specific principals to put objects into an Amazon S3 bucket. The principals already have access to Amazon S3.

The security engineer needs to configure a bucket policy that allows principals to put objects into the S3 bucket only if the value of the Team tag on the object matches the value of the Team tag that is associated with the principal. During testing, the security engineer notices that a principal can still put objects into the S3 bucket when the tag values do not match.

Which combination of factors are causing the PutObject operation to succeed when the tag values are different? (Select TWO.)

- A. The principal's identity-based policy grants access to put objects into the S3 bucket with no conditions.
- B. The principal's identity-based policy overrides the condition because the identity-based policy contains an explicit allow.
- C. The S3 bucket's resource policy does not deny access to put objects.
- D. The S3 bucket's resource policy cannot allow actions to the principal.
- E. The bucket policy does not apply to principals in the same zone of trust.

**Answer: AC**

#### Explanation:

The correct answer is A and C.

When using ABAC, the principal's identity-based policy and the S3 bucket's resource policy are both evaluated to determine the effective permissions. If either policy grants access to the principal, the action is allowed. If either policy denies access to the principal, the action is denied. Therefore, to enforce the tag-based condition, both policies must deny access when the tag values do not match.

In this case, the principal's identity-based policy grants access to put objects into the S3 bucket with no conditions (A), which means that the policy does not check

for the tag values. This policy overrides the condition in the bucket policy because an explicit allow always takes precedence over an implicit deny. The bucket policy can only allow or deny actions to the principal based on the condition, but it cannot override the identity-based policy. The S3 bucket's resource policy does not deny access to put objects ©, which means that it also does not check for the tag values. The bucket policy can only allow or deny actions to the principal based on the condition, but it cannot override the identity-based policy. Therefore, the combination of factors A and C are causing the PutObject operation to succeed when the tag values are different.

References:

- > Using ABAC with Amazon S3
- > Bucket policy examples

**NEW QUESTION 150**

A developer signed in to a new account within an IAM Organization organizational unit (OU) containing multiple accounts. Access to the Amazon S3 service is restricted with the following SCP.

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Deny",
      "Action": "s3:*",
      "Resource": "*"
    }
  ]
}
```

How can the security engineer provide the developer with Amazon S3 access without affecting other account?

- A. Move the SCP to the root OU of organization to remove the restriction to access Amazon S3.
- B. Add an IAM policy for the developer, which grants S3 access.
- C. Create a new OU without applying the SCP restricting S3 access.
- D. Move the developer account to this new OU.
- E. Add an allow list for the developer account for the S3 service.

**Answer: C**

**NEW QUESTION 155**

A company's Security Engineer has been tasked with restricting a contractor's IAM account access to the company's Amazon EC2 console without providing access to any other AWS services. The contractor's IAM account must not be able to gain access to any other AWS service, even if the IAM account is assigned additional permissions based on IAM group membership.

What should the Security Engineer do to meet these requirements?

- A. Create an Inline IAM user policy that allows for Amazon EC2 access for the contractor's IAM user.
- B. Create an IAM permissions boundary policy that allows Amazon EC2 access.
- C. Associate the contractor's IAM account with the IAM permissions boundary policy.
- D. Create an IAM group with an attached policy that allows for Amazon EC2 access.
- E. Associate the contractor's IAM account with the IAM group.
- F. Create an IAM role that allows for EC2 and explicitly denies all other services.
- G. Instruct the contractor to always assume this role.

**Answer: B**

**NEW QUESTION 160**

.....

## **Thank You for Trying Our Product**

### **We offer two products:**

1st - We have Practice Tests Software with Actual Exam Questions

2nd - Questions and Answers in PDF Format

### **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

**100% Actual & Verified — Instant Download, Please Click**  
**[Order The SCS-C02 Practice Test Here](#)**