



**Amazon**

## **Exam Questions AWS-Certified-Security-Specialty**

Amazon AWS Certified Security - Specialty

#### NEW QUESTION 1

An IT department currently has a Java web application deployed on Apache Tomcat running on Amazon EC2 instances. All traffic to the EC2 instances is sent through an internet-facing Application Load Balancer (ALB)

The Security team has noticed during the past two days thousands of unusual read requests coming from hundreds of IP addresses. This is causing the Tomcat server to run out of threads and reject new connections

Which the SIMPLEST change that would address this server issue?

- A. Create an Amazon CloudFront distribution and configure the ALB as the origin
- B. Block the malicious IPs with a network access list (NACL).
- C. Create an IAM Web Application Firewall (WAF). and attach it to the ALB
- D. Map the application domain name to use Route 53

**Answer:** A

#### Explanation:

this is the simplest change that can address the server issue. CloudFront is a service that provides a global network of edge locations that cache and deliver web content. Creating a CloudFront distribution and configuring the ALB as the origin can help reduce the load on the Tomcat server by serving cached content to the end users. CloudFront can also provide protection against distributed denial-of-service (DDoS) attacks by filtering malicious traffic at the edge locations. The other options are either ineffective or complex for solving the server issue.

#### NEW QUESTION 2

A company discovers a billing anomaly in its AWS account. A security consultant investigates the anomaly and discovers that an employee who left the company 30 days ago still has access to the account.

The company has not monitored account activity in the past.

The security consultant needs to determine which resources have been deployed or reconfigured by the employee as quickly as possible.

Which solution will meet these requirements?

- A. In AWS Cost Explorer, filter chart data to display results from the past 30 day
- B. Export the results to a data tabl
- C. Group the data table by re-source.
- D. Use AWS Cost Anomaly Detection to create a cost monito
- E. Access the detec-tion histor
- F. Set the time frame to Last 30 day
- G. In the search area, choose the service category.
- H. In AWS CloudTrail, filter the event history to display results from the past 30 day
- I. Create an Amazon Athena table that contains the dat
- J. Parti-tion the table by event source.
- K. Use AWS Audit Manager to create an assessment for the past 30 day
- L. Apply a usage-based framework to the assessmen
- M. Configure the assessment to as-sess by resource.

**Answer:** C

#### NEW QUESTION 3

A company has retail stores The company is designing a solution to store scanned copies of customer receipts on Amazon S3 Files will be between 100 KB and 5 MB in PDF format Each retail store must have a unique encryption key Each object must be encrypted with a unique key

Which solution will meet these requirements?

- A. Create a dedicated AWS Key Management Service (AWS KMS) customer managed key for each retail store Use the S3 Put operation to upload the objects to Amazon S3 Specify server-side encryption with AWS KMS keys (SSE-KMS) and the key ID of the store's key
- B. Create a new AWS Key Management Service (AWS KMS) customer managed key every day for each retail store Use the KMS Encrypt operation to encrypt objects Then upload the objects to Amazon S3
- C. Run the AWS Key Management Service (AWS KMS) GenerateDataKey operation every day for each retail store Use the data key and client-side encryption to encrypt the objects Then upload the objects to Amazon S3
- D. Use the AWS Key Management Service (AWS KMS) ImportKeyMaterial operation to import new key material to AWS KMS every day for each retail store Use a customer managed key and the KMS Encrypt operation to encrypt the objects Then upload the objects to Amazon S3

**Answer:** A

#### Explanation:

To meet the requirements of storing scanned copies of customer receipts on Amazon S3, where files will be between 100 KB and 5 MB in PDF format, each retail store must have a unique encryption key, and each object must be encrypted with a unique key, the most appropriate solution would be to create a dedicated AWS Key Management Service (AWS KMS) customer managed key for each retail store. Then, use the S3 Put operation to upload the objects to Amazon S3, specifying server-side encryption with AWS KMS keys (SSE-KMS) and the key ID of the store's key.

References: : Amazon S3 - Amazon Web Services : AWS Key Management Service - Amazon Web Services : Amazon S3 - Amazon Web Services : AWS Key Management Service - Amazon Web Service

#### NEW QUESTION 4

A company has a set of EC2 Instances hosted in IAM. The EC2 Instances have EBS volumes which is used to store critical information. There is a business continuity requirement to ensure high availability for the EBS volumes. How can you achieve this?

- A. Use lifecycle policies for the EBS volumes
- B. Use EBS Snapshots
- C. Use EBS volume replication
- D. Use EBS volume encryption

**Answer:** B

**Explanation:**

Data stored in Amazon EBS volumes is redundantly stored in multiple physical locations as part of normal operation of those services and at no additional charge. However, Amazon EBS replication is stored within the same availability zone, not across multiple zones; therefore, it is highly recommended that you conduct regular snapshots to Amazon S3 for long-term data durability Option A is invalid because there is no lifecycle policy for EBS volumes Option C is invalid because there is no EBS volume replication Option D is invalid because EBS volume encryption will not ensure business continuity For information on security for Compute Resources, please visit the below URL: [https://d1.awsstatic.com/whitepapers/Security/Security\\_Compute\\_Services\\_Whitepaper.pdf](https://d1.awsstatic.com/whitepapers/Security/Security_Compute_Services_Whitepaper.pdf)

**NEW QUESTION 5**

A security engineer is creating an AWS Lambda function. The Lambda function needs to use a role that is named LambdaAuditRole to assume a role that is named AcmeAuditFactoryRole in a different AWS account.

When the code is processed, the following error message appears: "An error occurred (AccessDenied) when calling the AssumeRole operation."

Which combination of steps should the security engineer take to resolve this error? (Select TWO.)

- A. Ensure that LambdaAuditRole has the sts:AssumeRole permission for AcmeAuditFactoryRole.
- B. Ensure that LambdaAuditRole has the AWSLambdaBasicExecutionRole managed policy attached.
- C. Ensure that the trust policy for AcmeAuditFactoryRole allows the sts:AssumeRole action from LambdaAuditRole.
- D. Ensure that the trust policy for LambdaAuditRole allows the sts:AssumeRole action from the lambda.amazonaws.com service.
- E. Ensure that the sts:AssumeRole API call is being issued to the us-east-1 Region endpoint.

**Answer:** AC

**NEW QUESTION 6**

A company wants to monitor the deletion of customer managed CMKs A security engineer must create an alarm that will notify the company before a CMK is deleted The security engineer has configured the integration of IAM CloudTrail with Amazon CloudWatch

What should the security engineer do next to meet this requirement?

- A. Use inbound rule 100 to allow traffic on TCP port 443 Use inbound rule 200 to deny traffic on TCP port 3306 Use outbound rule 100 to allow traffic on TCP port 443
- B. Use inbound rule 100 to deny traffic on TCP port 3306. Use inbound rule 200 to allow traffic on TCP port range 1024-65535. Use outbound rule 100 to allow traffic on TCP port 443
- C. Use inbound rule 100 to allow traffic on TCP port range 1024-65535 Use inbound rule 200 to deny traffic on TCP port 3306 Use outbound rule 100 to allow traffic on TCP port 443
- D. Use inbound rule 100 to deny traffic on TCP port 3306 Use inbound rule 200 to allow traffic on TCP port 443 Use outbound rule 100 to allow traffic on TCP port 443

**Answer:** A

**NEW QUESTION 7**

A company is hosting a static website on Amazon S3 The company has configured an Amazon CloudFront distribution to serve the website contents The company has associated an IAM WAF web ACL with the CloudFront distribution. The web ACL ensures that requests originate from the United States to address compliance restrictions.

THE company is worried that the S3 URL might still be accessible directly and that requests can bypass the CloudFront distribution

Which combination of steps should the company take to remove direct access to the S3 URL? (Select TWO. )

- A. Select "Restrict Bucket Access" in the origin settings of the CloudFront distribution
- B. Create an origin access identity (OAI) for the S3 origin
- C. Update the S3 bucket policy to allow s3 GetObject with a condition that the IAM Referer key matches the secret value Deny all other requests
- D. Configure the S3 bucket policy so that only the origin access identity (OAI) has read permission for objects in the bucket
- E. Add an origin custom header that has the name Referer to the CloudFront distribution Give the header a secret value.

**Answer:** AD

**NEW QUESTION 8**

A company wants to establish separate IAM Key Management Service (IAM KMS) keys to use for different IAM services. The company's security engineer created the following key policy to allow the infrastructure deployment team to create encrypted Amazon Elastic Block Store (Amazon EBS) volumes by assuming the InfrastructureDeployment IAM role:

```
{
  "Version": "2012-10-17",
  "Id": "key-policy-eps",
  "Statement": [
    {
      "Sid": "Enable IAM User Permissions",
      "Effect": "Allow",
      "Principal": {
        "AWS": "arn:aws:iam::123456789012:root"
      },
      "Action": "kms:*",
      "Resource": "*"
    },
    {
      "Sid": "Allow use of the key",
      "Effect": "Allow",
      "Principal": {
        "AWS": "arn:aws:iam::123456789012:role/aws-reserved/sso.amazonaws.com/InfrastructureDeployment"
      },
      "Action": [
        "kms:Encrypt",
        "kms:Decrypt",
        "kms:ReEncrypt*",
        "kms:GenerateDataKey*",
        "kms:DescribeKey",
        "kms:CreateGrant",
        "kms:ListGrants",
        "kms:RevokeGrant"
      ],
      "Resource": "*",
      "Condition": {
        "StringEquals": {
          "kms:ViaService": "ec2.us-west-2.amazonaws.com"
        }
      }
    }
  ]
}
```

The security engineer recently discovered that IAM roles other than the InfrastructureDeployment role used this key (or other services). Which change to the policy should the security engineer make to resolve these issues?

- A. In the statement block that contains the Sid "Allow use of the key", under the "Condition" block, change StringEquals to StringLike.
- B. In the policy document, remove the statement block that contains the Sid "Enable IAM User Permissions". Add key management policies to the KMS policy.
- C. In the statement block that contains the Sid "Allow use of the Key", under the "Condition" block, change the Kms:ViaService value to ec2.us-east-1.amazonaws.com.
- D. In the policy document, add a new statement block that grants the kms:Disable permission to the security engineer's IAM role.

**Answer: C**

#### Explanation:

To resolve the issues, the security engineer should make the following change to the policy:

➤ In the statement block that contains the Sid "Allow use of the key", under the "Condition" block, change the Kms:ViaService value to ec2.us-east-1.amazonaws.com. This allows the security engineer to restrict the use of the key to only EC2 service in the us-east-1 region, and prevent other services from using the key.

#### NEW QUESTION 9

An Application team has requested a new IAM KMS master key for use with Amazon S3, but the organizational security policy requires separate master keys for different IAM services to limit blast radius.

How can an IAM KMS customer master key (CMK) be constrained to work with only Amazon S3?

- A. Configure the CMK key policy to allow only the Amazon S3 service to use the kms Encrypt action
- B. Configure the CMK key policy to allow IAM KMS actions only when the kms ViaService condition matches the Amazon S3 service name.
- C. Configure the IAM user's policy to allow KMS to pass a role to Amazon S3
- D. Configure the IAM user's policy to allow only Amazon S3 operations when they are combined with the CMK

**Answer: B**

#### Explanation:

the kms:ViaService condition key can be used to restrict a CMK to work with only a specific AWS

service. By configuring the CMK key policy to allow KMS actions only when the kms:ViaService condition matches the Amazon S3 service name, you can ensure that only Amazon S3 can use the CMK. The other options are either incorrect or insufficient for constraining a CMK to work with only Amazon S3.

#### NEW QUESTION 10

A company is expanding its group of stores. On the day that each new store opens, the company wants to launch a customized web application for that store. Each store's application will have a non-production environment and a production environment. Each environment will be deployed in a separate AWS account. The company uses AWS Organizations and has an OU that is used only for these accounts.

The company distributes most of the development work to third-party development teams. A security engineer needs to ensure that each team follows the company's

deployment plan for AWS resources. The security engineer also must limit access to the deployment plan to only the developers who need access. The security engineer already has created an AWS CloudFormation template that implements the deployment plan.

What should the security engineer do next to meet the requirements in the MOST secure way?



- A. Create an AWS Service Catalog portfolio in the organization's management account
- B. Upload the CloudFormation template
- C. Add the template to the portfolio's product list
- D. Share the portfolio with the OU.
- E. Use the CloudFormation CLI to create a module from the CloudFormation template
- F. Register the module as a private extension in the CloudFormation registry
- G. Publish the extension
- H. In the OU, create an SCP that allows access to the extension.
- I. Create an AWS Service Catalog portfolio in the organization's management account
- J. Upload the CloudFormation template
- K. Add the template to the portfolio's product list
- L. Create an IAM role that has a trust policy that allows cross-account access to the portfolio for users in the OU account
- M. Attach the AWSServiceCatalogEndUserFullAccess managed policy to the role.
- N. Use the CloudFormation CLI to create a module from the CloudFormation template
- O. Register the module as a private extension in the CloudFormation registry
- P. Publish the extension
- Q. Share the extension with the OU

**Answer:** A

**Explanation:**

The correct answer is A. Create an AWS Service Catalog portfolio in the organization's management account. Upload the CloudFormation template. Add the template to the portfolio's product list. Share the portfolio with the OU.

According to the AWS documentation, AWS Service Catalog is a service that allows you to create and manage catalogs of IT services that are approved for use on AWS. You can use Service Catalog to centrally manage commonly deployed IT services and help achieve consistent governance and compliance requirements, while enabling users to quickly deploy only the approved IT services they need.

To use Service Catalog with multiple AWS accounts, 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 Service Catalog as a service principal for AWS Organizations, which lets you share your portfolios with organizational units (OUs) or accounts in your organization.

To create a Service Catalog portfolio, you need to use an administrator account, such as the organization's management account. You can upload your CloudFormation template as a product in your portfolio, and define constraints and tags for it. You can then share your portfolio with the OU that contains the accounts for the web applications. This will allow the developers in those accounts to launch products from the shared portfolio using the Service Catalog end user console.

Option B is incorrect because CloudFormation modules are reusable components that encapsulate one or more resources and their configurations. They are not meant to be used as templates for deploying entire stacks of resources. Moreover, sharing a module with an OU does not grant access to launch stacks from it.

Option C is incorrect because creating an IAM role that has a trust policy that allows cross-account access to the portfolio is not secure. It would allow any user in the OU accounts to assume the role and access the portfolio, regardless of their job function or access requirements.

Option D is incorrect because sharing a module with an OU does not grant access to launch stacks from it. It also does not limit access to the deployment plan to only the developers who need access.

**NEW QUESTION 10**

Auditors for a health care company have mandated that all data volumes be encrypted at rest. Infrastructure is deployed mainly via IAM CloudFormation, however third-party frameworks and manual deployment are required on some legacy systems.

What is the BEST way to monitor, on a recurring basis, whether all EBS volumes are encrypted?

- A. On a recurring basis, update an IAM user policy to require that EC2 instances are created with an encrypted volume
- B. Configure an IAM Config rule to run on a recurring basis for volume encryption
- C. Set up Amazon Inspector rules for volume encryption to run on a recurring schedule
- D. Use CloudWatch Logs to determine whether instances were created with an encrypted volume

**Answer:** B

**Explanation:**

To support answer B, use the reference <https://d1.IAMstatic.com/whitepapers/IAM-security-whitepaper.pdf>. "For example, IAM Config provides a managed IAM Config Rules to ensure that encryption is turned on for all EBS volumes in your account."

**NEW QUESTION 12**

A developer has created an AWS Lambda function in a company's development account. The Lambda function requires the use of an AWS Key Management Service (AWS KMS) customer managed key that exists in a security account that the company's security team controls. The developer obtains the ARN of the KMS key from a previous Lambda function in the development account. The previous Lambda function had been working properly with the KMS key.

When the developer uses the ARN and tests the new Lambda function, an error message states that access is denied to the KMS key in the security account. The developer tests the previous Lambda function that uses the same KMS key and discovers that the previous Lambda function still can encrypt data as expected.

A security engineer must resolve the problem so that the new Lambda function in the development account can use the KMS key from the security account.

Which combination of steps should the security engineer take to meet these requirements? (Select TWO.)

- A. In the security account, configure an IAM role for the new Lambda function
- B. Attach an IAM policy that allows access to the KMS key in the security account.
- C. In the development account, configure an IAM role for the new Lambda function
- D. Attach a key policy that allows access to the KMS key in the security account.
- E. In the development account, configure an IAM role for the new Lambda function
- F. Attach an IAM policy that allows access to the KMS key in the security account.
- G. Configure a key policy for the KMS key in the security account to allow access to the IAM role of the new Lambda function in the security account.
- H. Configure a key policy for the KMS key in the security account to allow access to the IAM role of the new Lambda function in the development account.

**Answer:** CE

**Explanation:**

To allow cross-account access to a KMS key, the key policy of the KMS key must grant permission to the external account or principal, and the IAM policy of the external account or principal must delegate the key policy permission. In this case, the new Lambda function in the development account needs to use the KMS key in the security account, so the key policy of the KMS key must allow access to the IAM role of the new Lambda function in the development account (option E),

and the IAM role of the new Lambda function in the development account must have an IAM policy that allows access to the KMS key in the security account (option C). Option A is incorrect because it creates an IAM role for the new Lambda function in the security account, not in the development account. Option B is incorrect because it attaches a key policy to an IAM role, which is not valid. Option D is incorrect because it allows access to the IAM role of the new Lambda function in the security account, not in the development account. Verified References:

➤ <https://docs.aws.amazon.com/autoscaling/ec2/userguide/key-policy-requirements-EBS-encryption.html>

#### NEW QUESTION 15

A security engineer receives an IAM abuse email message. According to the message, an Amazon EC2 instance that is running in the security engineer's IAM account is sending phishing email messages.

The EC2 instance is part of an application that is deployed in production. The application runs on many EC2 instances behind an Application Load Balancer. The instances run in an Amazon EC2 Auto Scaling group across multiple subnets and multiple Availability Zones.

The instances normally communicate only over the HTTP, HTTPS, and MySQL protocols. Upon investigation, the security engineer discovers that email messages are being sent over port 587. All other traffic is normal.

The security engineer must create a solution that contains the compromised EC2 instance, preserves forensic evidence for analysis, and minimizes application downtime. Which combination of steps must the security engineer take to meet these requirements? (Select THREE.)

- A. Add an outbound rule to the security group that is attached to the compromised EC2 instance to deny traffic to 0.0.0.0/0 and port 587.
- B. Add an outbound rule to the network ACL for the subnet that contains the compromised EC2 instance to deny traffic to 0.0.0.0/0 and port 587.
- C. Gather volatile memory from the compromised EC2 instance
- D. Suspend the compromised EC2 instance from the Auto Scaling group
- E. Then take a snapshot of the compromised EC2 instance
- F. v
- G. Take a snapshot of the compromised EC2 instance
- H. Suspend the compromised EC2 instance from the Auto Scaling group
- I. Then gather volatile memory from the compromised EC2 instance.
- J. Move the compromised EC2 instance to an isolated subnet that has a network ACL that has no inbound rules or outbound rules.
- K. Replace the existing security group that is attached to the compromised EC2 instance with a new security group that has no inbound rules or outbound rules.

**Answer:** ACE

#### NEW QUESTION 17

A company finds that one of its Amazon EC2 instances suddenly has a high CPU usage. The company does not know whether the EC2 instance is compromised or whether the operating system is performing background cleanup.

Which combination of steps should a security engineer take before investigating the issue? (Select THREE.)

- A. Disable termination protection for the EC2 instance if termination protection has not been disabled.
- B. Enable termination protection for the EC2 instance if termination protection has not been enabled.
- C. Take snapshots of the Amazon Elastic Block Store (Amazon EBS) data volumes that are attached to the EC2 instance.
- D. Remove all snapshots of the Amazon Elastic Block Store (Amazon EBS) data volumes that are attached to the EC2 instance.
- E. Capture the EC2 instance metadata, and then tag the EC2 instance as under quarantine.
- F. Immediately remove any entries in the EC2 instance metadata that contain sensitive information.

**Answer:** BCE

#### Explanation:

[https://d1.awsstatic.com/WWPS/pdf/aws\\_security\\_incident\\_response.pdf](https://d1.awsstatic.com/WWPS/pdf/aws_security_incident_response.pdf)

#### NEW QUESTION 20

A company uses Amazon RDS for MySQL as a database engine for its applications. A recent security audit revealed an RDS instance that is not compliant with company policy for encrypting data at rest. A security engineer at the company needs to ensure that all existing RDS databases are encrypted using server-side encryption and that any future deviations from the policy are detected.

Which combination of steps should the security engineer take to accomplish this? (Select TWO.)

- A. Create an IAM Config rule to detect the creation of unencrypted RDS database
- B. Create an Amazon EventBridge (Amazon CloudWatch Events) rule to trigger on the IAM Config rules compliance state change and use Amazon Simple Notification Service (Amazon SNS) to notify the security operations team.
- C. Use IAM System Manager State Manager to detect RDS database encryption configuration drift
- D. Create an Amazon EventBridge (Amazon CloudWatch Events) rule to track state changes and use Amazon Simple Notification Service (Amazon SNS) to notify the security operations team.
- E. Create a read replica for the existing unencrypted RDS database and enable replica encryption in the process
- F. Once the replica becomes active, promote it into a standalone database instance and terminate the unencrypted database instance.
- G. Take a snapshot of the unencrypted RDS database
- H. Copy the snapshot and enable snapshot encryption in the process
- I. Restore the database instance from the newly created encrypted snapshot
- J. Terminate the unencrypted database instance.
- K. Enable encryption for the identified unencrypted RDS instance by changing the configurations of the existing database

**Answer:** AD

#### NEW QUESTION 25

A company's cloud operations team is responsible for building effective security for IAM cross-account access. The team asks a security engineer to help troubleshoot why some developers in the developer account (123456789012) in the developers group are not able to assume a cross-account role (ReadS3) into a production account (999999999999) to read the contents of an Amazon S3 bucket (productionapp). The two account policies are as follows:

Developer account 123456789012:

Developer group permissions:

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": "sts:AssumeRole",
      "Resource": "arn:aws:iam::999999999999:role/ReadS3"
    }
  ]
}
```

Production account 999999999999:

Production account ReadS3 role policy:

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": "s3:ListAllMyBuckets",
      "Resource": "*"
    },
    {
      "Effect": "Allow",
      "Action": [
        "s3:ListBucket",
        "s3:GetBucketLocation"
      ]
    }
  ]
}
```

Production account ReadS3 role policy - trust relationship:

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Principal": {
        "AWS": "arn:aws:iam::888888888888:root"
      },
      "Action": "sts:AssumeRole",
      "Condition": {
        "BoolIfExists": {
          "aws:MultiFactorAuthPresent": "true"
        }
      }
    }
  ]
}
```

Which recommendations should the security engineer make to resolve this issue? (Select TWO.)

- A. Ask the developers to change their password and use a different web browser.
- B. Ensure that developers are using multi-factor authentication (MFA) when they log in to their developer account as the developer role.
- C. Modify the production account ReadS3 role policy to allow the PutBucketPolicy action on the productionapp S3 bucket.
- D. Update the trust relationship policy on the production account S3 role to allow the account number of the developer account.
- E. Update the developer group permissions in the developer account to allow access to the productionapp S3 bucket.

**Answer:** AD

#### NEW QUESTION 28

A company is using Amazon Elastic Container Service (Amazon ECS) to run its container-based application on AWS. The company needs to ensure that the container images contain no severe vulnerabilities. The company also must ensure that only specific IAM roles and specific AWS accounts can access the container images.

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

- A. Pull images from the public container registr
- B. Publish the images to Amazon Elastic Container Registry (Amazon ECR) repositories with scan on push configured in a centralized AWS account
- C. Use a CI/CD pipeline to deploy the images to different AWS account
- D. Use identity-based policies to restrict access to which IAM principals can access the images.
- E. Pull images from the public container registr
- F. Publish the images to a private container registry that is hosted on Amazon EC2 instances in a centralized AWS account
- G. Deploy host-based container scanning tools to EC2 instances that run Amazon EC



- H. Restrict access to the container images by using basic authentication over HTTPS.
- I. Pull images from the public container registr
- J. Publish the images to Amazon Elastic Container Registry (Amazon ECR) repositories with scan on push configured in a centralized AWS account
- K. Use a CI/CD pipeline to deploy the images to different AWS account
- L. Use repository policies and identity-based policies to restrict access to which IAM principals and accounts can access the images.
- M. Pull images from the public container registr
- N. Publish the images to AWS CodeArtifact repositories in a centralized AWS account
- O. Use a CI/CD pipeline to deploy the images to different AWS account
- P. Use repository policies and identity-based policies to restrict access to which IAM principals and accounts can access the images.

**Answer: C**

**Explanation:**

The correct answer is C. Pull images from the public container registry. Publish the images to Amazon Elastic Container Registry (Amazon ECR) repositories with scan on push configured in a centralized AWS account.

Use a CI/CD pipeline to deploy the images to different AWS accounts. Use repository policies and identity-based policies to restrict access to which IAM principals and accounts can access the images.

This solution meets the requirements because:

- Amazon ECR is a fully managed container registry service that supports Docker and OCI images and artifacts<sup>1</sup>. It integrates with Amazon ECS and other AWS services to simplify the development and deployment of container-based applications.
- Amazon ECR provides image scanning on push, which uses the Common Vulnerabilities and Exposures (CVEs) database from the open-source Clair project to detect software vulnerabilities in container images<sup>2</sup>. The scan results are available in the AWS Management Console, AWS CLI, or AWS SDKs<sup>2</sup>.
- Amazon ECR supports cross-account access to repositories, which allows sharing images across multiple AWS accounts<sup>3</sup>. This can be achieved by using repository policies, which are resource-based policies that specify which IAM principals and accounts can access the repositories and what actions they can perform<sup>4</sup>. Additionally, identity-based policies can be used to control which IAM roles in each account can access the repositories<sup>5</sup>.

The other options are incorrect because:

- A. This option does not use repository policies to restrict cross-account access to the images, which is a requirement. Identity-based policies alone are not sufficient to control access to Amazon ECR repositories<sup>5</sup>.
- B. This option does not use Amazon ECR, which is a fully managed service that provides image scanning and cross-account access features. Hosting a private container registry on EC2 instances would require more management overhead and additional security measures.
- D. This option uses AWS CodeArtifact, which is a fully managed artifact repository service that supports Maven, npm, NuGet, PyPI, and generic package formats<sup>6</sup>. However, AWS CodeArtifact does not support Docker or OCI container images, which are required for Amazon ECS applications.

**NEW QUESTION 32**

There are currently multiple applications hosted in a VPC. During monitoring it has been noticed that multiple port scans are coming in from a specific IP Address block. The internal security team has requested that all offending IP Addresses be denied for the next 24 hours. Which of the following is the best method to quickly and temporarily deny access from the specified IP Address's.

Please select:

- A. Create an AD policy to modify the Windows Firewall settings on all hosts in the VPC to deny access from the IP Address block.
- B. Modify the Network ACLs associated with all public subnets in the VPC to deny access from the IP Address block.
- C. Add a rule to all of the VPC Security Groups to deny access from the IP Address block.
- D. Modify the Windows Firewall settings on all AMI'S that your organization uses in that VPC to deny access from the IP address block.

**Answer: B**

**Explanation:**

NACL acts as a firewall at the subnet level of the VPC and we can deny the offending IP address block at the subnet level using NACL rules to block the incoming traffic to the VPC instances. Since NACL rules are applied as per the Rule numbers make sure that this rule number should take precedence over other rule numbers if there are any such rules that will allow traffic from these IP ranges. The lowest rule number has more precedence over a rule that has a higher number. The IAM Documentation mentions the following as a best practices for IAM users

For extra security, enable multi-factor authentication (MFA) for privileged IAM users (users who are allowed access to sensitive resources or APIs). With MFA, users have a device that generates a unique authentication code (a one-time password, or OTP). Users must provide both their normal credentials (like their user name and password) and the OTP. The MFA device can either be a special piece of hardware, or it can be a virtual device (for example, it can run in an app on a smartphone).

Options C is invalid because these options are not available Option D is invalid because there is not root access for users

For more information on IAM best practices, please visit the below URL: <https://docs.IAM.amazon.com/IAM/latest/UserGuide/best-practices.html>

The correct answer is: Modify the Network ACLs associated with all public subnets in the VPC to deny access from the IP Address block.

omit your Feedback/Queries to our Experts

**NEW QUESTION 34**

A company has deployed servers on Amazon EC2 instances in a VPC. External vendors access these servers over the internet. Recently, the company deployed a new application on EC2 instances in a new CIDR range. The company needs to make the application available to the vendors.

A security engineer verified that the associated security groups and network ACLs are allowing the required ports in the inbound diction. However, the vendors cannot connect to the application.

Which solution will provide the vendors access to the application?

- A. Modify the security group that is associated with the EC2 instances to have the same outbound rules as inbound rules.
- B. Modify the network ACL that is associated with the CIDR range to allow outbound traffic to ephemeral ports.
- C. Modify the inbound rules on the internet gateway to allow the required ports.
- D. Modify the network ACL that is associated with the CIDR range to have the same outbound rules as inbound rules.

**Answer: B**

**Explanation:**

The correct answer is B. Modify the network ACL that is associated with the CIDR range to allow outbound traffic to ephemeral ports.

This answer is correct because network ACLs are stateless, which means that they do not automatically allow return traffic for inbound connections. Therefore, the network ACL that is associated with the CIDR range of the new application must have outbound rules that allow traffic to ephemeral ports, which are the temporary ports used by the vendors' machines to communicate with the application servers. Ephemeral ports are typically in the range of 1024-65535. If the network ACL



does not have such rules, the vendors will not be able to connect to the application.

The other options are incorrect because:

- A. Modifying the security group that is associated with the EC2 instances to have the same outbound rules as inbound rules is not a solution, because security groups are stateful, which means that they automatically allow return traffic for inbound connections. Therefore, there is no need to add outbound rules to the security group for the vendors to access the application2.
- C. Modifying the inbound rules on the internet gateway to allow the required ports is not a solution, because internet gateways do not have inbound or outbound rules. Internet gateways are VPC components that enable communication between instances in a VPC and the internet. They do not filter traffic based on ports or protocols3.
- D. Modifying the network ACL that is associated with the CIDR range to have the same outbound rules as inbound rules is not a solution, because it does not address the issue of ephemeral ports. The outbound rules of the network ACL must match the ephemeral port range of the vendors' machines, not necessarily the inbound rules of the network ACL4.

References:

1: Ephemeral port - Wikipedia 2: Security groups for your VPC - Amazon Virtual Private Cloud 3: Internet gateways - Amazon Virtual Private Cloud 4: Network ACLs - Amazon Virtual Private Cloud

#### NEW QUESTION 36

A company has an application that uses dozens of Amazon DynamoDB tables to store data. Auditors find that the tables do not comply with the company's data protection policy.

The company's retention policy states that all data must be backed up twice each month: once at midnight on the 15th day of the month and again at midnight on the 25th day of the month. The company must retain the backups for 3 months.

Which combination of steps should a security engineer take to meet these re-quirements? (Select TWO.)

- A. Use the DynamoDB on-demand backup capability to create a backup pla
- B. Con-figure a lifecycle policy to expire backups after 3 months.
- C. Use AWS DataSync to create a backup pla
- D. Add a backup rule that includes a retention period of 3 months.
- E. Use AVVS Backup to create a backup pla
- F. Add a backup rule that includes a retention period of 3 months.
- G. Set the backup frequency by using a cron schedule expressio
- H. Assign each DynamoDB table to the backup plan.
- I. Set the backup frequency by using a rate schedule expressio
- J. Assign each DynamoDB table to the backup plan.

**Answer:** AD

#### NEW QUESTION 41

A company has a web-based application using Amazon CloudFront and running on Amazon Elastic Container Service (Amazon ECS) behind an Application Load Balancer (ALB). The ALB is terminating TLS and balancing load across ECS service tasks A security engineer needs to design a solution to ensure that application content is accessible only through CloudFront and that I is never accessible directly.

How should the security engineer build the MOST secure solution?

- A. Add an origin custom header Set the viewer protocol policy to HTTP and HTTPS Set the origin protocol pokey to HTTPS only Update the application to validate the CloudFront custom header
- B. Add an origin custom header Set the viewer protocol policy to HTTPS only Set the origin protocol policy to match viewer Update the application to validate the CloudFront custom header.
- C. Add an origin custom header Set the viewer protocol policy to redirect HTTP to HTTPS Set the origin protocol policy to HTTP only Update the application to validate the CloudFront custom header.
- D. Add an origin custom header Set the viewer protocol policy to redirect HTTP to HTTP
- E. Set the origin protocol policy to HTTPS only Update the application to validate the CloudFront custom header

**Answer:** D

#### Explanation:

To ensure that application content is accessible only through CloudFront and not directly, the security engineer should do the following:

- Add an origin custom header. This is a header that CloudFront adds to the requests that it sends to the origin, but viewers cannot see or modify.
- Set the viewer protocol policy to redirect HTTP to HTTPS. This ensures that the viewers always use HTTPS when they access the website through CloudFront.
- Set the origin protocol policy to HTTPS only. This ensures that CloudFront always uses HTTPS when it connects to the origin.
- Update the application to validate the CloudFront custom header. This means that the application checks if the request has the custom header and only responds if it does. Otherwise, it denies or ignores the request. This prevents users from bypassing CloudFront and accessing the content directly on the origin.

#### NEW QUESTION 42

A company has two teams, and each team needs to access its respective Amazon S3 buckets. The company anticipates adding more teams that also will have their own S3 buckets. When the company adds these teams, team members will need the ability to be assigned to multiple teams. Team members also will need the ability to change teams. Additional S3 buckets can be created or deleted.

An IAM administrator must design a solution to accomplish these goals. The solution also must be scalable and must require the least possible operational overhead.

Which solution meets these requirements?

- A. Add users to groups that represent the team
- B. Create a policy for each team that allows the team to access its respective S3 buckets onl
- C. Attach the policy to the corresponding group.
- D. Create an IAM role for each tea
- E. Create a policy for each team that allows the team to access its respective S3 buckets onl
- F. Attach the policy to the corresponding role.
- G. Create IAM roles that are labeled with an access tag value of a tea
- H. Create one policy that allows dynamic access to S3 buckets with the same ta
- I. Attach the policy to the IAM role

- J. Tag the S3 buckets accordingly.
- K. Implement a role-based access control (RBAC) authorization mode
- L. Create the corresponding policies, and attach them to the IAM users.

**Answer:** A

#### NEW QUESTION 43

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 46

A security engineer needs to create an Amazon S3 bucket policy to grant least privilege read access to IAM user accounts that are named User=1, User2. and User3. These IAM user accounts are members of the AuthorizedPeople IAM group. The security engineer drafts the following S3 bucket policy:

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

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

A)

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

B)

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

C)

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

D)

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

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

**Answer:** A

#### NEW QUESTION 51

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

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

**Answer:** A

#### Explanation:

A VPC peering connection is a networking connection between two VPCs that enables you to route traffic between them using private IPv4 addresses or IPv6 addresses. Instances in either VPC can communicate with each other as if they are within the same network. You can create a VPC peering connection between your own VPCs, or with a VPC in another IAM account within a single region.

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

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

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

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

#### NEW QUESTION 55

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

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

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

**Answer:** BC

#### Explanation:

The correct answer is B and C.

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

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

\* C. Allow port 22 from 192.168.100.0/24.

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

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

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

\* D. Allow port 22 from 10.0.1.0/24.

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

\* E. Allow port 443 from 10.0.1.0/24.

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

#### NEW QUESTION 58

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

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

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

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

**Answer: B**

#### Explanation:

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

According to the AWS documentation<sup>1</sup>, the aws sts get-session-token CLI command returns a set of temporary credentials for an AWS account or IAM user. The credentials consist of an access key ID, a secret access key, and a security token. These credentials are valid for the specified duration only. The session duration for IAM users can be between 15 minutes and 36 hours, with a default of 12 hours.

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

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

The other options are incorrect because:

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

References:

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

#### NEW QUESTION 62

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

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

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



J. Configure the rule to invoke an AWS Lambda function

K. Configure the Lambda function to publish messages to an Amazon Simple Notification Service (Amazon SNS) topic and to delete the unencrypted resource.

**Answer:** A

**Explanation:**

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

**NEW QUESTION 63**

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

Which statement should the security specialist include in the policy?

- A. 

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

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

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

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

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

**Answer:** D

**NEW QUESTION 68**

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 69

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 74

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

Which solution will meet these requirements?

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

**Answer:** A

#### NEW QUESTION 76

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

- A. Designate an AWS account as a delegated administrator for Security Hub
- B. Publish events to Amazon CloudWatch from the delegated administrator account, all member accounts, and required Regions that are enabled for Security Hub findings.
- C. Designate an AWS account in an organization in AWS Organizations as a delegated administrator for Security Hub
- D. Publish events to Amazon EventBridge from the delegated administrator account, all member accounts, and required Regions that are enabled for Security Hub findings.
- E. In each Region, create an Amazon EventBridge rule to deliver findings to an Amazon Kinesis data stream

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

**Answer:** BDF

**Explanation:**

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

**NEW QUESTION 79**

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 84

A company wants to remove all SSH keys permanently from a specific subset of its Amazon Linux 2 Amazon EC2 instances that are using the same 1AM instance profile. However, three individuals who have IAM user accounts will need to access these instances by using an SSH session to perform critical duties. How can a security engineer provide the access to meet these requirements?

- A. Assign an 1AM policy to the instance profile to allow the EC2 instances to be managed by AWS Systems Manager. Provide the 1AM user accounts with permission to use Systems Manager. Remove the SSH keys from the EC2 instances. Use Systems Manager Inventory to select the EC2 instance and connect.
- B. Assign an 1AM policy to the 1AM user accounts to provide permission to use AWS Systems Manager. Run Command. Remove the SSH keys from the EC2 instances. Use Run Command to open an SSH connection to the EC2 instance.
- C. Assign an 1AM policy to the instance profile to allow the EC2 instances to be managed by AWS Systems Manager. Provide the 1AM user accounts with permission to use Systems Manager. Remove the SSH keys from the EC2 instances. Use Systems Manager Session Manager to select the EC2 instance and connect.
- D. Assign an 1AM policy to the 1AM user accounts to provide permission to use the EC2 service in the AWS Management Console. Remove the SSH keys from the EC2 instances. Connect to the EC2 instance as the ec2-user through the AWS Management Console's EC2 SSH client method.

**Answer:** C

#### Explanation:

To provide access to the three individuals who have IAM user accounts to access the Amazon Linux 2 Amazon EC2 instances that are using the same IAM instance profile, the most appropriate solution would be to assign an IAM policy to the instance profile to allow the EC2 instances to be managed by AWS Systems Manager, provide the IAM user accounts with permission to use Systems Manager, remove the SSH keys from the EC2 instances, and use Systems Manager Session Manager to select the EC2 instance and connect.

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

#### NEW QUESTION 87

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

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

**Answer:** B

#### Explanation:

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

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

#### NEW QUESTION 90

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 93

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 97

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 101

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

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

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

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

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

**Answer: C**

#### Explanation:

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

#### NEW QUESTION 103

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

Which solution meets these criteria?

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

**Answer: A**

#### Explanation:

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

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

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

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

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

The other options are incorrect for the following reasons:

- \* B. A customer managed CMK that uses AWS provided key material does not expire automatically. You can enable automatic rotation of the key material every year, but this does not prevent access to the data encrypted with the previous key material. You would need to manually delete the CMK and its backing key material to make the data inaccessible<sup>3</sup>.
- \* C. An AWS managed CMK is a KMS key that is created, owned, and managed by an AWS service on your behalf. You have limited control over the key configuration, permissions, rotation, and deletion. You cannot use an AWS managed CMK to encrypt data in other AWS services or applications. You also cannot set an expiration date for the key material of an AWS managed CMK<sup>4</sup>.
- \* D. Operation system-native encryption that uses GnuPG is not a solution that uses AWS KMS. GnuPG is a command line tool that implements the OpenPGP standard for encrypting and signing data. It does not integrate with Amazon EBS or other AWS services. It also does not provide a way to automatically expire the key material used for encryption<sup>5</sup>.

References:

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

#### NEW QUESTION 104

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

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

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

- F. Configure the Java application to catch a connection failure and make a call to AWS Secrets Manager to retrieve updated credentials when the password is rotate
- G. Grant permission to the instance role associated with the EC2 instance to access Secrets Manager.

**Answer:** CE

**Explanation:**

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

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

By configuring the Java application to catch a connection failure and make a call to AWS Secrets Manager to retrieve updated credentials, you can avoid hard-coding the credentials in your application code or configuration files. This way, your application can dynamically obtain the latest credentials from Secrets Manager whenever the password is rotated, without needing to restart or redeploy the application. To enable this, you need to grant permission to the instance role associated with the EC2 instance to access Secrets Manager using IAM policies<sup>2</sup>. You can also use the AWS SDK for Java to integrate your application with Secrets Manager<sup>3</sup>.

**NEW QUESTION 107**

A company needs a security engineer to implement a scalable solution for multi-account authentication and authorization. The solution should not introduce additional user-managed architectural components. Native IAM features should be used as much as possible The security engineer has set up IAM Organizations w1th all features activated and IAM SSO enabled.

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

- A. Use AD Connector to create users and groups for all employees that require access to IAM accounts. Assign AD Connector groups to IAM accounts and link to the IAM roles in accordance with the employees' job functions and access requirements Instruct employees to access IAM accounts by using the IAM Directory Service user portal.
- B. Use an IAM SSO default directory to create users and groups for all employees that require access to IAM account
- C. Assign groups to IAM accounts and link to permission sets in accordance with the employees' job functions and access requirement
- D. Instruct employees to access IAM accounts by using the IAM SSO user portal.
- E. Use an IAM SSO default directory to create users and groups for all employees that require access to IAM account
- F. Link IAM SSO groups to the IAM users present in all accounts to inherit existing permission
- G. Instruct employees to access IAM accounts by using the IAM SSO user portal.
- H. Use IAM Directory Service for Microsoft Active Directory to create users and groups for all employees that require access to IAM accounts Enable IAM Management Console access in the created directory and specify IAM SSO as a source of information for integrated accounts and permission set
- I. Instruct employees to access IAM accounts by using the IAM Directory Service user portal.

**Answer:** B

**NEW QUESTION 110**

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

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

**Answer:** C

**Explanation:**

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

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

This answer is correct because disabling the Network Source/Destination check allows the virtual security appliance to route traffic that is not addressed to or from itself. By default, this check is enabled on all EC2 instances, and it prevents them from forwarding traffic that does not match their own IP or MAC addresses. However, for a virtual security appliance that acts as a router or a firewall, this check needs to be disabled, otherwise it will drop the traffic that it is supposed to route<sup>12</sup>.

The other options are incorrect because:

- A. Disabling network ACLs is not a solution, because network ACLs are optional layers of security for the subnets in a VPC. They can be used to allow or deny traffic based on IP addresses and ports, but they do not affect the routing behavior of the virtual security appliance<sup>3</sup>.
- B. Configuring the security appliance's elastic network interface for promiscuous mode is not a solution, because promiscuous mode is a mode for a network interface that causes it to pass all traffic it receives to the CPU, rather than passing only the frames that it is programmed to receive. Promiscuous mode is normally used for packet sniffing or monitoring, but it does not enable the network interface to route traffic<sup>4</sup>.
- D. Placing the security appliance in the public subnet with the internet gateway is not a solution, because it does not address the routing issue of the virtual security appliance. The security appliance can be placed in either a public or a private subnet, depending on the network design and security requirements, but it still needs to have the Network Source/Destination check disabled to route traffic properly<sup>5</sup>.

References:

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

**NEW QUESTION 111**

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

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

- A. Turn on VPC Flow Logs for all VPCs in the account.



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

**Answer:** BD

**Explanation:**

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

**NEW QUESTION 114**

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 119**

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

What should the security engineer recommend?

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

**Answer:** C

**Explanation:**

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

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

**NEW QUESTION 120**

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

- A. Place the RDS instance in a public subnet and an IAM Lambda function outside the VPC
- B. Schedule the Lambda function to run every 3 months to rotate the secrets.
- C. Place the RDS instance in a private subnet and an IAM Lambda function inside the VPC in the private subnet
- D. Configure the private subnet to use a NAT gateway
- E. Schedule the Lambda function to run every 3 months to rotate the secrets.
- F. Place the RDS instance in a private subnet and an IAM Lambda function outside the VPC
- G. Configure the private subnet to use an internet gateway



- H. Schedule the Lambda function to run every 3 months to rotate the secrets.
- I. Place the RDS instance in a private subnet and an IAM Lambda function inside the VPC in the private subne
- J. Schedule the Lambda function to run quarterly to rotate the secrets.
- K. Place the RDS instance in a private subnet and an IAM Lambda function inside the VPC in the private subne
- L. Configure a Secrets Manager interface endpoint
- M. Schedule the Lambda function to run every 3 months to rotate the secrets.

**Answer:** BE

**Explanation:**

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

**NEW QUESTION 123**

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 124**

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

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

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

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

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

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

**Answer:** D

**Explanation:**

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

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

The correct answer is D. Configure a Secrets Manager interface VPC endpoint. Include the Lambda function's private subnet during the configuration process.

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

The other options are incorrect for the following reasons:

- A. An egress-only internet gateway is a VPC component that allows outbound communication over IPv6 from instances in the VPC to the internet, and prevents the internet from initiating an IPv6 connection with the instances<sup>3</sup>. However, this option does not meet the requirement that the VPC must not send or receive network traffic through the internet. Moreover, an egress-only internet gateway is for use with IPv6 traffic only, and Secrets Manager does not support IPv6 addresses<sup>2</sup>.
- B. A NAT gateway is a VPC component that enables instances in a private subnet to connect to the internet or other AWS services, but prevents the internet from initiating connections with those instances<sup>4</sup>. However, this option does not meet the requirement that the VPC must not send or receive network traffic through the internet. Additionally, a NAT gateway requires an elastic IP address, which is a public IPv4 address<sup>4</sup>.
- C. A VPC peering connection is a networking connection between two VPCs that enables you to route traffic between them using private IPv4 addresses or

IPv6 addresses<sup>5</sup>. However, this option does not work because Secrets Manager does not have a default VPC that can be peered with. Furthermore, a VPC peering connection does not provide a private connection to Secrets Manager APIs without an internet gateway or other devices<sup>2</sup>.

#### NEW QUESTION 126

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

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

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

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

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

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

**Answer:** C

#### Explanation:

The correct answer is C. Use AWS Resource Access Manager (AWS RAM) to share Account-A's VPC

subnets with the remaining member accounts. Configure the member accounts to use the shared subnets to launch workloads.

This answer is correct because AWS RAM is a service that helps you securely share your AWS resources across AWS accounts, within your organization or organizational units (OUs), and with IAM roles and users for supported resource types<sup>1</sup>. One of the supported resource types is VPC subnets<sup>2</sup>, which means you can share the subnets in Account-A's VPC with the other member accounts using AWS RAM. This way, you can meet the requirements of using a centrally managed VPC, avoiding duplicate VPCs in each account, and launching workloads in shared subnets. You can also control the access to the shared subnets by using IAM policies and resource-based policies<sup>3</sup>, which can prevent one account from modifying another account's resources.

The other options are incorrect because:

- A. Using a CloudFormation template in the member accounts to launch workloads and using the Fn::ImportValue function to obtain the subnet ID values is not a solution, because Fn::ImportValue can only import values that have been exported by another stack within the same region<sup>4</sup>. This means that you cannot use Fn::ImportValue to reference the subnet IDs that are exported by Account-A's CloudFormation template, unless all the member accounts are in the same region as Account-A. This option also does not avoid creating duplicate VPCs in each account, which is one of the requirements.
- B. Using a transit gateway in the VPC within Account-A and configuring the member accounts to use the transit gateway to access the subnets in Account-A to launch workloads is not a solution, because a transit gateway does not allow you to launch workloads in another account's subnets. A transit gateway is a network transit hub that enables you to route traffic between your VPCs and on-premises networks<sup>5</sup>, but it does not enable you to share subnets across accounts.
- D. Creating a peering connection between Account-A and the remaining member accounts and configuring the member accounts to use the subnets in Account-A through the VPC peering connection to launch workloads is not a solution, because a VPC peering connection does not allow you to launch workloads in another account's subnets. A VPC peering connection is a networking connection between two VPCs that enables you to route traffic between them privately<sup>6</sup>, but it does not enable you to share subnets across accounts.

References:

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

#### NEW QUESTION 130

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 135

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

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

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

**Answer: C**

#### Explanation:

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

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

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

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

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

The other options are incorrect because:

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

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

### NEW QUESTION 136

.....

## 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

### AWS-Certified-Security-Specialty Practice Exam Features:

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

**100% Actual & Verified — Instant Download, Please Click**  
**[Order The AWS-Certified-Security-Specialty Practice Test Here](#)**