

SCS-C01 Dumps

AWS Certified Security- Specialty

<https://www.certleader.com/SCS-C01-dumps.html>



NEW QUESTION 1

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 2

A company manages three separate IAM accounts for its production, development, and test environments, Each Developer is assigned a unique IAM user under the development account. A new application hosted on an Amazon EC2 instance in the developer account requires read access to the archived documents stored in an Amazon S3 bucket in the production account. How should access be granted?

- A. Create an IAM role in the production account and allow EC2 instances in the development account to assume that role using the trust polic
- B. Provide read access for the required S3 bucket to this role.
- C. Use a custom identity broker to allow Developer IAM users to temporarily access the S3 bucket.
- D. Create a temporary IAM user for the application to use in the production account.
- E. Create a temporary IAM user in the production account and provide read access to Amazon S3. Generate the temporary IAM user's access key and secret key and store these on the EC2 instance used by the application in the development account.

Answer: A

Explanation:

<https://IAM.amazon.com/premiumsupport/knowledge-center/cross-account-access-s3/>

NEW QUESTION 3

A company deploys a set of standard IAM roles in AWS accounts. The IAM roles are based on job functions within the company. To balance operational efficiency and security, a security engineer implemented AWS Organizations SCPs to restrict access to critical security services in all company accounts. All of the company's accounts and OUs within AWS Organizations have a default FullAWSAccess SCP that is attached. The security engineer needs to ensure that no one can disable Amazon GuardDuty and AWS Security Hub. The security engineer also must not override other permissions that are granted by IAM policies that are defined in the accounts. Which SCP should the security engineer attach to the root of the organization to meet these requirements?

- A.

```
"Version": "2012-10-17",
"Statement": [
  {
    "Effect": "Deny",
    "Action": [
      "guardduty:DeleteDetector",
      "guardduty:UpdateDetector",
      "securityhub:DisableSecurityHub"
    ],
    "Resource": [
      "*"
    ]
  }
]
```
- B. A screenshot of a computer code Description automatically generated

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Deny",
      "Action": "*",
      "Resource": "*"
    },
    {
      "Effect": "Allow",
      "NotAction": [
        "guardduty:DeleteDetector",
        "guardduty:UpdateDetector",
        "securityhub:DisableSecurityHub"
      ],
      "Resource": [
        "*"
      ]
    }
  ]
}
```

C. A screenshot of a computer code Description automatically generated

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": "*",
      "Resource": "*"
    },
    {
      "Effect": "Deny",
      "NotAction": [
        "guardduty:DeleteDetector",
        "guardduty:UpdateDetector",
        "securityhub:DisableSecurityHub"
      ],
      "Resource": [
        "*"
      ]
    }
  ]
}
```

D. A screenshot of a computer code Description automatically generated

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "NotAction": [
        "guardduty:DeleteDetector",
        "guardduty:UpdateDetector",
        "securityhub:DisableSecurityHub"
      ],
      "Resource": [
        "*"
      ]
    }
  ]
}
```

Answer: A

NEW QUESTION 4

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 5

A company has two IAM accounts within IAM Organizations. In Account-1, Amazon EC2 Auto Scaling is launched using a service-linked role. In Account-2, Amazon EBS volumes are encrypted with an IAM KMS key. A Security Engineer needs to ensure that the service-linked role can launch instances with these encrypted volumes.

Which combination of steps should the Security Engineer take in both accounts? (Select TWO.)

- A. Allow Account-1 to access the KMS key in Account-2 using a key policy
- B. Attach an IAM policy to the service-linked role in Account-1 that allows these actions CreateGrant, DescribeKey, Encrypt, GenerateDataKey, Decrypt, and ReEncrypt
- C. Create a KMS grant for the service-linked role with these actions CreateGrant, DescribeKey, Encrypt, GenerateDataKey, Decrypt, and ReEncrypt
- D. Attach an IAM policy to the role attached to the EC2 instances with KMS actions and then allow Account-1 in the KMS key policy.
- E. Attach an IAM policy to the user who is launching EC2 instances and allow the user to access the KMS key policy of Account-2.

Answer: CD

Explanation:

because these are the steps that can ensure that the service-linked role can launch instances with encrypted volumes. A service-linked role is a type of IAM role that is linked to an AWS service and allows the service to perform actions on your behalf. A KMS grant is a mechanism that allows you to delegate permissions to use a customer master key (CMK) to a principal such as a service-linked role. A KMS grant specifies the actions that the principal can perform, such as encrypting and decrypting data. By creating a KMS grant for the service-linked role with the specified actions, you can allow the service-linked role to use the CMK in Account-2 to launch instances with encrypted volumes. By attaching an IAM policy to the role attached to the EC2 instances with KMS actions and then allowing Account-1 in the KMS key policy, you can also enable cross-account access to the CMK and allow the EC2 instances to use the encrypted volumes. The other options are either incorrect or unnecessary for meeting the requirement.

NEW QUESTION 6

A company has an encrypted Amazon Aurora DB cluster in the us-east-1 Region. The DB cluster is encrypted with an AWS Key Management Service (AWS KMS) customer managed key. To meet compliance requirements, the company needs to copy a DB snapshot to the us-west-1 Region. However, when the company tries to copy the snapshot to us-west-1 the company cannot access the key that was used to encrypt the original database.

What should the company do to set up the snapshot in us-west-1 with proper encryption?

- A. Use AWS Secrets Manager to store the customer managed key in us-west-1 as a secret. Use this secret to encrypt the snapshot in us-west-1.
- B. Create a new customer managed key in us-west-1. Use this new key to encrypt the snapshot in us-west-1.
- C. Create an IAM policy that allows access to the customer managed key in us-east-1. Specify `arn:aws:kms:us-east-1:*` as the principal.
- D. Create an IAM policy that allows access to the customer managed key in us-east-1. Specify `arn:aws:kms:us-west-1:*` as the principal.

Answer: B

Explanation:

"If you copy an encrypted snapshot across Regions, you must specify a KMS key valid in the destination AWS Region. It can be a Region-specific KMS key, or a multi-Region key." <https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/aurora-copy-snapshot.html#aurora-copy-sna>

NEW QUESTION 7

A Security Engineer is asked to update an AWS CloudTrail log file prefix for an existing trail. When attempting to save the change in the CloudTrail console, the Security Engineer receives the following error message: "There is a problem with the bucket policy." What will enable the Security Engineer to save the change?

- A. Create a new trail with the updated log file prefix, and then delete the original trail
- B. Update the existing bucket policy in the Amazon S3 console with the new log file prefix, and then update the log file prefix in the CloudTrail console.
- C. Update the existing bucket policy in the Amazon S3 console to allow the Security Engineer's Principal to perform PutBucketPolicy, and then update the log file prefix in the CloudTrail console.
- D. Update the existing bucket policy in the Amazon S3 console with the new log file prefix, and then update the log file prefix in the CloudTrail console.
- E. Update the existing bucket policy in the Amazon S3 console to allow the Security Engineer's Principal to perform GetBucketPolicy, and then update the log file prefix in the CloudTrail console.

Answer: C

Explanation:

The correct answer is C. Update the existing bucket policy in the Amazon S3 console with the new log file prefix, and then update the log file prefix in the CloudTrail console.

According to the AWS documentation¹, a bucket policy is a resource-based policy that you can use to grant access permissions to your Amazon S3 bucket and the objects in it. Only the bucket owner can associate a policy with a bucket. The permissions attached to the bucket apply to all of the objects in the bucket that are owned by the bucket owner.

When you create a trail in CloudTrail, you can specify an existing S3 bucket or create a new one to store your log files. CloudTrail automatically creates a bucket policy for your S3 bucket that grants CloudTrail write-only access to deliver log files to your bucket. The bucket policy also grants read-only access to AWS services that you can use to view and analyze your log data, such as Amazon Athena, Amazon CloudWatch Logs, and Amazon QuickSight.

If you want to update the log file prefix for an existing trail, you must also update the existing bucket policy in the S3 console with the new log file prefix. The log file prefix is part of the resource ARN that identifies the objects in your bucket that CloudTrail can access. If you don't update the bucket policy with the new log file prefix, CloudTrail will not be able to deliver log files to your bucket, and you will receive an error message when you try to save the change in the CloudTrail console.

The other options are incorrect because:

- A. Creating a new trail with the updated log file prefix, and then deleting the original trail is not necessary and may cause data loss or inconsistency. You can simply update the existing trail and its associated bucket policy with the new log file prefix.
- B. Updating the existing bucket policy in the S3 console to allow the Security Engineer's Principal to perform PutBucketPolicy is not relevant to this issue. The PutBucketPolicy action allows you to create or replace a policy on a bucket, but it does not affect CloudTrail's ability to deliver log files to your bucket. You still need to update the existing bucket policy with the new log file prefix.
- D. Updating the existing bucket policy in the S3 console to allow the Security Engineer's Principal to perform GetBucketPolicy is not relevant to this issue. The GetBucketPolicy action allows you to retrieve a policy on a bucket, but it does not affect CloudTrail's ability to deliver log files to your bucket. You still need to update the existing bucket policy with the new log file prefix.

References:

1: Using bucket policies - Amazon Simple Storage Service

NEW QUESTION 8

A company wants to migrate its static primary domain website to AWS. The company hosts the website and DNS servers internally. The company wants the website to enforce SSL/TLS encryption block IP addresses from outside the United States (US), and take advantage of managed services whenever possible. Which solution will meet these requirements?

- A. Migrate the website to Amazon S3 Import a public SSL certificate to an Application Load Balancer
- B. Balancer with rules to block traffic from outside the US Migrate DNS to Amazon Route 53.
- C. Migrate the website to Amazon EC2 Import a public SSL certificate that is created by AWS Certificate Manager (ACM) to an Application Load Balancer with rules to block traffic from outside the US Update DNS accordingly.
- D. Migrate the website to Amazon S3. Import a public SSL certificate to Amazon CloudFront Use AWS WAF rules to block traffic from outside the US Update DNS accordingly
- E. Migrate the website to Amazon S3 Import a public SSL certificate that is created by AWS Certificate Manager (ACM) to Amazon CloudFront
- F. CloudFront Configure CloudFront to block traffic from outside the US
- G. Migrate DNS to Amazon Route 53.

Answer: D

Explanation:

To migrate the static website to AWS and meet the requirements, the following steps are required:

- Migrate the website to Amazon S3, which is a highly scalable and durable object storage service that can host static websites. To do this, create an S3 bucket with the same name as the domain name of the website, enable static website hosting for the bucket, upload the website files to the bucket, and configure the bucket policy to allow public read access to the objects. For more information, see [Hosting a static website on Amazon S3](#).
- Import a public SSL certificate that is created by AWS Certificate Manager (ACM) to Amazon CloudFront, which is a global content delivery network (CDN) service that can improve the performance and security of web applications. To do this, request or import a public SSL certificate for the domain name of the website using ACM, create a CloudFront distribution with the S3 bucket as the origin, and associate the SSL certificate with the distribution. For more information, see [Using alternate domain names and HTTPS](#).
- Configure CloudFront to block traffic from outside the US, which is one of the requirements. To do this, create a CloudFront web ACL using AWS WAF, which is a web application firewall service that lets you control access to your web applications. In the web ACL, create a rule that uses a geo match condition to block requests that originate from countries other than the US. Associate the web ACL with the CloudFront distribution. For more information, see [How AWS WAF works with Amazon CloudFront features](#).

➤ Migrate DNS to Amazon Route 53, which is a highly available and scalable cloud DNS service that can route traffic to various AWS services. To do this, register or transfer your domain name to Route 53, create a hosted zone for your domain name, and create an alias record that points your domain name to your CloudFront distribution. For more information, see [Routing traffic to an Amazon CloudFront web distribution by using your domain name](#).

The other options are incorrect because they either do not implement SSL/TLS encryption for the website (A), do not use managed services whenever possible (B), or do not block IP addresses from outside the US (C). Verified References:

- <https://docs.aws.amazon.com/AmazonS3/latest/userguide/HostingWebsiteOnS3Setup.html>
- <https://docs.aws.amazon.com/waf/latest/developerguide/waf-cloudfront.html>
- <https://docs.aws.amazon.com/Route53/latest/DeveloperGuide/routing-to-cloudfront-distribution.html>

NEW QUESTION 9

A company is designing a multi-account structure for its development teams. The company is using AWS Organizations and AWS Single Sign-On (AWS SSO). The company must implement a solution so that the development teams can use only specific AWS Regions and so that each AWS account allows access to only specific AWS services.

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

- A. Use AWS SSO to set up service-linked roles with IAM policy statements that include the Condition, Resource, and NotAction elements to allow access to only the Regions and services that are needed.
- B. Deactivate AWS Security Token Service (AWS STS) in Regions that the developers are not allowed to use.
- C. Create SCPs that include the Condition, Resource, and NotAction elements to allow access to only the Regions and services that are needed.
- D. For each AWS account, create tailored identity-based policies for AWS SSO
- E. Use statements that include the Condition, Resource, and NotAction elements to allow access to only the Regions and services that are needed.

Answer: C

Explanation:

https://docs.aws.amazon.com/organizations/latest/userguide/orgs_manage_policies_scps_syntax.html#scp-element

NEW QUESTION 10

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

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

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

- A. Update the EC2 instance security group to add a rule that allows outbound traffic on port 443 for 0.0.0.0/0.

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

Answer: BCF

NEW QUESTION 10

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

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

Answer: D

Explanation:

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

NEW QUESTION 12

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

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

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

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

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

Answer: BE

NEW QUESTION 16

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

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

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

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

Answer: D

Explanation:

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

- Create a service control policy (SCP) that defines the maximum permissions that can be granted to any IAM role in the organization. An SCP is a type of policy that you can use with AWS Organizations to manage permissions for all accounts in your organization. SCPs restrict permissions for entities in member accounts, including each AWS account root user, IAM users, and roles. For more information, see [Service control policies overview](#).
 - Create a permissions boundary for IAM roles that matches the SCP. A permissions boundary is an advanced feature for using a managed policy to set the maximum permissions that an identity-based policy can grant to an IAM entity. A permissions boundary allows an entity to perform only the actions that are allowed by both its identity-based policies and its permissions boundaries. For more information, see [Permissions boundaries for IAM entities](#).
 - Add the SCP to the root organizational unit (OU) so that it applies to all accounts in the organization. This will ensure that no IAM role can exceed the permissions defined by the SCP, regardless of how it is created or modified.
 - Instruct the application teams to attach the permissions boundary to any IAM role they create. This will prevent them from creating IAM roles that can escalate their own privileges or access resources they are not authorized to access. This solution will meet the requirements with the least operational overhead, as it leverages AWS Organizations and IAM features to delegate and limit IAM role creation without requiring manual reviews or approvals. The other options are incorrect because they either do not allow application teams to provision their own IAM roles (A), do not limit the scope of IAM roles or prevent privilege escalation (B), or do not take advantage of managed services whenever possible ©.
- Verified References:
- https://docs.aws.amazon.com/IAM/latest/UserGuide/access_policies_boundaries.html

NEW QUESTION 20

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

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

Answer: A

NEW QUESTION 24

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 en-sure that objects cannot be overwritten or deleted by any user, including the AWS account root user.

Which solution will meet these requirements?

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

Answer: A

NEW QUESTION 28

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

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

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

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

- D. "NumericGreaterThan" : { " aws : MultiFactorAuthAge " : "7200"
E. "NumericLessThan" : { "MaxSessionDuration " : "7200"}

Answer: AC

Explanation:

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

- Option A: "Bool" : { "aws:MultiFactorAuthPresent" : "true" } is a valid condition that checks if the principal (the IAM user) has authenticated with MFA before making the request. This condition will enforce MFA for the IAM users to access the specified services. This condition key is supported by all AWS services that support IAM policies1.
- Option B: "Bool" : { "aws:MultiFactorAuthPresent" : "false" } is the opposite of option A. This condition will allow access only if the principal has not authenticated with MFA, which is not the desired requirement. This condition key is supported by all AWS services that support IAM policies1.
- Option C: "NumericLessThan" : { "aws:MultiFactorAuthAge" : "7200" } is a valid condition that checks if the time since the principal authenticated with MFA is less than 7200 seconds (2 hours). This condition will enforce the session duration limit for the IAM users. This condition key is supported by all AWS services that support IAM policies1.
- Option D: "NumericGreaterThan" : { "aws:MultiFactorAuthAge" : "7200" } is the opposite of option C. This condition will allow access only if the time since the principal authenticated with MFA is more than 7200 seconds (2 hours), which is not the desired requirement. This condition key is supported by all AWS services that support IAM policies1.
- Option E: "NumericLessThan" : { "MaxSessionDuration" : "7200" } is not a valid condition key. MaxSessionDuration is a property of an IAM role, not a condition key. It specifies the maximum session duration (in seconds) for the role, which can be between 3600 and 43200 seconds (1 to 12 hours). This property can be set when creating or modifying a role, but it cannot be used as a condition in a policy2.

NEW QUESTION 29

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

Please select:

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

Answer: C

Explanation:

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

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

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

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

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

NEW QUESTION 32

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

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

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

Answer: A

Explanation:

To grant appropriate access to separate modules for read-write and read-only functionality in a serverless application hosted on AWS that uses Amazon Redshift as a data store, a security engineer should configure cluster security groups for each application module to control access to database users that are required for read-only and readwrite, and configure an IAM policy for each module specifying the ARN of an IAM user that allows the GetClusterCredentials API call.

References: : Amazon Redshift - Amazon Web Services : Amazon Redshift - Amazon Web Services : Identity and Access Management - AWS Management Console : AWS Identity and Access Management - AWS Management Console

NEW QUESTION 34

A company wants to remove all SSH keys permanently from a specific subset of its Amazon Linux 2 Amazon EC2 instances that are using the same IAM instance profile However three individuals who have IAM user accounts will need to access these instances by using an SSH session to perform critical duties

How can a security engineer provide the access to meet these requirements'?

- A. Assign an IAM policy to the instance profile to allow the EC2 instances to be managed by AWS Systems Manager Provide the IAM user accounts with permission to use Systems Manager Remove the SSH keys from the EC2 instances Use Systems Manager Inventory to select the EC2 instance and connect
- B. Assign an IAM policy to the IAM user accounts to provide permission to use AWS Systems Manager Run Command Remove the SSH keys from the EC2 instances Use Run Command to open an SSH connection to the EC2 instance

- C. Assign an 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 Manage AWS Management Console : AWS Identity and Access Management - AWS Management Console : Am Elastic Compute Cloud - Amazon Web Services : Amazon Linux 2 - Amazon Web Services : AWS Syst Manager - AWS Management Console : AWS Systems Manager - AWS Management Console : AWS Systems Manager - AWS Management Console

NEW QUESTION 36

A company has recently recovered from a security incident that required the restoration of Amazon EC2 instances from snapshots. After performing a gap analysis of its disaster recovery procedures and backup strategies, the company is concerned that, next time, it will not be able to recover the EC2 instances if the AWS account was compromised and Amazon EBS snapshots were deleted. All EBS snapshots are encrypted using an AWS KMS CMK. Which solution would solve this problem?

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

Answer: C

Explanation:

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

NEW QUESTION 37

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 accoun
- B. In a dedicated logging accoun
- C. aggregate all GuardDuty logs from each production accoun
- D. Remediate incidents by configuring GuardDuty to directly invoke an AWS Lambda functio
- E. Configure the Lambda function to also publish notifications to the SNS topic.
- F. Activate AWS security Hub in each production accoun
- G. In a dedicated logging accoun
- H. aggregate all security Hub findings from each production accoun
- I. Remediate incidents by ustng AWS Config and AWS Systems Manage
- J. Configure Systems Manager to also pub11Sh notifications to the SNS topic.
- K. Activate Amazon GuardDuty in each production accoun
- L. In a dedicated logging accoun
- 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 accoun
- P. In a dedicated logging accoun
- Q. aggregate all Security Hub findings from each production accoun
- 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 42

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

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

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

Answer: B

Explanation:

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

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

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

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

References:

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

NEW QUESTION 45

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.

A customer managed CMK can use either AWS provided key material or customer provided key material. AWS provided key material is generated by AWS KMS and never leaves the service unencrypted. Customer provided key material is generated outside of AWS KMS and imported into a customer managed CMK. You

can specify an expiration date for the imported key material, after which the CMK becomes unusable until you reimport new key material². To meet the criteria of automatically expiring the key material in 90 days, you need to use customer provided key material and set the expiration date accordingly. This way, you can ensure that the data encrypted with the CMK will not be accessible after 90 days unless you reimport new key material and re-encrypt the data. The other options are incorrect for the following reasons:

- * B. A customer managed CMK that uses AWS provided key material does not expire automatically. You can enable automatic rotation of the key material every year, but this does not prevent access to the data encrypted with the previous key material. You would need to manually delete the CMK and its backing key material to make the data inaccessible³.
- * C. An AWS managed CMK is a KMS key that is created, owned, and managed by an AWS service on your behalf. You have limited control over the key configuration, permissions, rotation, and deletion. You cannot use an AWS managed CMK to encrypt data in other AWS services or applications. You also cannot set an expiration date for the key material of an AWS managed CMK⁴.
- * D. Operation system-native encryption that uses GnuPG is not a solution that uses AWS KMS. GnuPG is a command line tool that implements the OpenPGP standard for encrypting and signing data. It does not integrate with Amazon EBS or other AWS services. It also does not provide a way to automatically expire the key material used for encryption⁵.

References:

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

NEW QUESTION 47

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

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

What should the Security Engineer do to troubleshoot this issue?

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

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

B)

Add the following statement to the CMK key policy:

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

C)

Add the following statement to the CMK key policy:

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

D)

Add the following statement to the CMK key policy:

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

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

Answer: D

NEW QUESTION 50

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

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

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

Answer: CE

Explanation:

AWS Secrets Manager is a service that helps you manage, retrieve, and rotate secrets such as database credentials, API keys, and other sensitive information. By configuring automatic rotation of credentials in AWS Secrets Manager, you can ensure that your secrets are changed regularly and securely, without requiring manual intervention or application downtime. You can also specify the rotation frequency and the rotation function that performs the logic of changing the credentials on the database and updating the secret in Secrets Manager¹.

* 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 policies². You can also use the AWS SDK for Java to integrate your application with Secrets Manager³.

NEW QUESTION 55

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 with 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 requirements.
- 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 sets.
- I. Instruct employees to access IAM accounts by using the IAM Directory Service user portal.

Answer: B

NEW QUESTION 60

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

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

Answer: C

Explanation:

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

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

This answer is correct because disabling the Network Source/Destination check allows the virtual security appliance to route traffic that is not addressed to or from itself. By default, this check is enabled on all EC2 instances, and it prevents them from forwarding traffic that does not match their own IP or MAC addresses. However, for a virtual security appliance that acts as a router or a firewall, this check needs to be disabled, otherwise it will drop the traffic that it is supposed to route¹².

The other options are incorrect because:

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

References:

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

NEW QUESTION 64

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

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

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

Answer: BD

NEW QUESTION 66

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 70

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

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

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

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

Answer: BC

Explanation:

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

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

NEW QUESTION 74

The Security Engineer is managing a traditional three-tier web application that is running on Amazon EC2 instances. The application has become the target of increasing numbers of malicious attacks from the Internet.

What steps should the Security Engineer take to check for known vulnerabilities and limit the attack surface? (Choose two.)

- A. Use AWS Certificate Manager to encrypt all traffic between the client and application servers.
- B. Review the application security groups to ensure that only the necessary ports are open.
- C. Use Elastic Load Balancing to offload Secure Sockets Layer encryption.
- D. Use Amazon Inspector to periodically scan the backend instances.
- E. Use AWS Key Management Services to encrypt all the traffic between the client and application servers.

Answer: BD

Explanation:

The steps that the Security Engineer should take to check for known vulnerabilities and limit the attack surface are:

➤ B. Review the application security groups to ensure that only the necessary ports are open. This is a good practice to reduce the exposure of the EC2 instances to potential attacks from the Internet. Application security groups are a feature of Azure that allow you to group virtual machines and define network security policies based on those groups¹.

➤ D. Use Amazon Inspector to periodically scan the backend instances. This is a service that helps you to identify vulnerabilities and exposures in your EC2 instances and applications. Amazon Inspector can perform automated security assessments based on predefined or custom rules packages².

NEW QUESTION 78

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

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

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

Answer: DE

Explanation:

these are the recommended actions to take when you receive an abuse notice from AWS⁸. You should review the abuse notice to see what content or activity was reported and detach the internet gateway from the VPC to isolate the affected instances from the internet. You should also remove any rules that allow inbound traffic from 0.0.0.0/0 from the security groups and create a network access control list (NACL) rule to deny all traffic inbound from the internet. You should then delete the compromised instances and any associated resources

that you did not create. The other options are either inappropriate or unnecessary for responding to the abuse notice.

NEW QUESTION 79

During a manual review of system logs from an Amazon Linux EC2 instance, a Security Engineer noticed that there are sudo commands that were never properly alerted or reported on the Amazon CloudWatch Logs agent

Why were there no alerts on the sudo commands?

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

Answer: B

Explanation:

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

to push the logs to CloudWatch, then there would be no alerts on the sudo commands. The other options are either irrelevant or invalid for explaining why there were no alerts on the sudo commands.

NEW QUESTION 80

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

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

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

Answer: C

Explanation:

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

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

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

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

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

The other options are incorrect because:

- A. Using VPC Traffic Mirroring would not capture the DNS queries that go to the on-premises DNS servers, because Traffic Mirroring only copies network traffic from an elastic network interface of an EC2 instance to a target for analysis³. Traffic Mirroring does not include traffic that goes through a Route 53 Resolver outbound endpoint, which is used to forward queries to on-premises DNS servers⁴. Therefore, this solution would not meet the requirements.
- B. Configuring VPC flow logs on all relevant VPCs would not capture the DNS name that was requested in Route 53 Resolver, because flow logs only record information about the IP traffic going to and from network interfaces in a VPC⁵. 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 network⁶. 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 85

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

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

Which solution will meet these requirements?

- A. Disable the FTP rule in GuardDuty in the Region where the FTP server is deployed.
- B. Add the FTP server to a trusted IP list
- C. Deploy the list to GuardDuty to stop receiving the notifications.
- D. Create a suppression rule in GuardDuty to filter findings by automatically archiving new findings that match the specified criteria.
- E. Create an AWS Lambda function that has the appropriate permissions to delete the finding whenever a new occurrence is reported.

Answer: C

Explanation:

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

NEW QUESTION 87

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

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

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

Answer: CD

NEW QUESTION 89

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

Which solution meets these requirements in the MOST secure way?

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

Answer: C

Explanation:

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

References: : AWS Lambda Developer Guide

NEW QUESTION 90

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

Which solution will meet these requirements MOST securely?

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

Answer: C

Explanation:

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

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

This solution provides the following security benefits:

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

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

NEW QUESTION 93

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

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

Answer: C

Explanation:

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

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

NEW QUESTION 98

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

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

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

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

Answer: ACE

NEW QUESTION 99

A security engineer needs to create an IAM Key Management Service (IAM KMS) key that will be used to encrypt all data stored in a company's Amazon S3 Buckets in the us-west-1 Region. The key will use server-side encryption. Usage of the key must be limited to requests coming from Amazon S3 within the company's account.

Which statement in the KMS key policy will meet these requirements?

A)

```
{
  "Effect": "Allow",
  "Principal": {
    "AWS": "*"
  },
  "Action": [
    "kms:Encrypt",
    "kms:Decrypt",
    "kms:ReEncrypt*",
    "kms:GenerateDataKey*",
    "kms:DescribeKey"
  ],
  "Resource": "*",
  "Condition": {
    "StringEquals": {
      "kms:ViaService": "s3.us-west-1.amazonaws.com",
      "kms:CallerAccount": "<CustomerAccountID>"
    }
  }
}
```

B)

```
{
  "Effect": "Allow",
  "Principal": {
    "AWS": "s3.us-west-1.amazonaws.com"
  },
  "Action": [
    "kms:Encrypt",
    "kms:Decrypt",
    "kms:ReEncrypt*",
    "kms:GenerateDataKey*",
    "kms:DescribeKey"
  ],
  "Resource": "*",
  "Condition": {
    "StringEquals": {
      "kms:CallerAccount": "<CustomerAccountID>"
    }
  }
}
```

C)

```
{
  "Effect": "Allow",
  "Principal": {
    "AWS": "*"
  },
  "Action": [
    "kms:Encrypt",
    "kms:Decrypt",
    "kms:ReEncrypt*",
    "kms:GenerateDataKey*",
    "kms:DescribeKey"
  ],
  "Resource": "*",
  "Condition": {
    "StringEquals": {
      "kms:EncryptionContext:aws:s3:arn": [
        "arn:aws:s3:::"
      ]
    }
  }
}
```

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

Answer: A

NEW QUESTION 101

A company's security team needs to receive a notification whenever an AWS access key has not been rotated in 90 or more days. A security engineer must develop a solution that provides these notifications automatically.

Which solution will meet these requirements with the LEAST amount of effort?

- A. Deploy an AWS Config managed rule to run on a periodic basis of 24 hour
- B. Select the access-keys-rotated managed rule, and set the maxAccessKeyAge parameter to 90 day
- C. Create an Amazon EventBridge (Amazon CloudWatch Events) rule with an event pattern that matches the compliance type of NON_COMPLIANT from AWS Config for the managed rule
- D. Configure EventBridge (CloudWatch Events) to send an Amazon Simple Notification Service (Amazon SNS) notification to the security team.
- E. Create a script to export a .csv file from the AWS Trusted Advisor check for IAM access key rotation. Load the script into an AWS Lambda function that will upload the .csv file to an Amazon S3 bucket
- F. Create an Amazon Athena table query that runs when the .csv file is uploaded to the S3 bucket
- G. Publish the results for any keys older than 90 days by using an invocation of an Amazon Simple Notification Service (Amazon SNS) notification to the security team.
- H. Create a script to download the IAM credentials report on a periodic basis
- I. Load the script into an AWS Lambda function that will run on a schedule through Amazon EventBridge (Amazon CloudWatch Events). Configure the Lambda script to load the report into memory and to filter the report for records in which the key was last rotated at least 90 days ago
- J. If any records are detected, send an Amazon Simple Notification Service (Amazon SNS) notification to the security team.
- K. Create an AWS Lambda function that queries the IAM API to list all the users
- L. Iterate through the users by using the ListAccessKeys operation
- M. Verify that the value in the CreateDate field is not at least 90 days old
- N. Send an Amazon Simple Notification Service (Amazon SNS) notification to the security team if the value is at least 90 days old
- O. Create an Amazon EventBridge (Amazon CloudWatch Events) rule to schedule the Lambda function to run each day.

Answer: A

NEW QUESTION 103

A security engineer needs to run an AWS CloudFormation script. The CloudFormation script builds AWS infrastructure to support a stack that includes web servers

and a MySQL database. The stack has been deployed in pre-production environments and is ready for production. The production script must comply with the principle of least privilege. Additionally, separation of duties must exist between the security engineer's IAM account and CloudFormation.

Which solution will meet these requirements?

- A. Use IAM Access Analyzer policy generation to generate a policy that allows the CloudFormation script to run and manage the stack
- B. Attach the policy to a new IAM role
- C. Modify the security engineer's IAM permissions to be able to pass the new role to CloudFormation.
- D. Create an IAM policy that allows ec2:* and rds:* permission
- E. Attach the policy to a new IAM role. Modify the security engineer's IAM permissions to be able to assume the new role.
- F. Use IAM Access Analyzer policy generation to generate a policy that allows the CloudFormation script to run and manage the stack
- G. Modify the security engineer's IAM permissions to be able to run the CloudFormation script.
- H. Create an IAM policy that allows ec2:* and rds:* permission
- I. Attach the policy to a new IAM role
- J. Use the IAM policy simulator to confirm that the policy allows the AWS API calls that are necessary to build the stack
- K. Modify the security engineer's IAM permissions to be able to pass the new role to CloudFormation.

Answer: A

Explanation:

The correct answer is A. Use IAM Access Analyzer policy generation to generate a policy that allows the CloudFormation script to run and manage the stack. Attach the policy to a new IAM role. Modify the security engineer's IAM permissions to be able to pass the new role to CloudFormation.

According to the AWS documentation, IAM Access Analyzer is a service that helps you identify the resources in your organization and accounts, such as Amazon S3 buckets or IAM roles, that are shared with an external entity. You can also use IAM Access Analyzer to generate fine-grained policies that grant least privilege access based on access activity and access attempts.

To use IAM Access Analyzer policy generation, you need to enable IAM Access Analyzer in your account or organization. You can then use the IAM console or the AWS CLI to generate a policy for a resource based on its access activity or access attempts. You can review and edit the generated policy before applying it to the resource.

To use IAM Access Analyzer policy generation with CloudFormation, you can follow these steps:

- > Run the CloudFormation script in a pre-production environment and monitor its access activity or access attempts using IAM Access Analyzer.
- > Use IAM Access Analyzer policy generation to generate a policy that allows the CloudFormation script to run and manage the stack. The policy will include only the permissions that are necessary for the script to function.
- > Attach the policy to a new IAM role that has a trust relationship with CloudFormation. This will allow CloudFormation to assume the role and execute the script.
- > Modify the security engineer's IAM permissions to be able to pass the new role to CloudFormation.

This will allow the security engineer to launch the stack using the role.

- > Run the CloudFormation script in the production environment using the new role.

This solution will meet the requirements of least privilege and separation of duties, as it will limit the permissions of both CloudFormation and the security engineer to only what is needed for running and managing the stack.

Option B is incorrect because creating an IAM policy that allows ec2:* and rds:* permissions is not following the principle of least privilege, as it will grant more permissions than necessary for running and managing the stack. Moreover, modifying the security engineer's IAM permissions to be able to assume the new role is not ensuring separation of duties, as it will allow the security engineer to bypass CloudFormation and directly access the resources.

Option C is incorrect because modifying the security engineer's IAM permissions to be able to run the CloudFormation script is not ensuring separation of duties, as it will allow the security engineer to execute the script without using CloudFormation.

Option D is incorrect because creating an IAM policy that allows ec2:* and rds:* permissions is not following the principle of least privilege, as it will grant more permissions than necessary for running and managing the stack. Using the IAM policy simulator to confirm that the policy allows the AWS API calls that are necessary to build the stack is not sufficient, as it will not generate a fine-grained policy based on access activity or access attempts.

NEW QUESTION 105

A company uses AWS Organizations. The company has teams that use an AWS CloudHSM hardware security module (HSM) that is hosted in a central AWS account. One of the teams creates its own new dedicated AWS account and wants to use the HSM that is hosted in the central account.

How should a security engineer share the HSM that is hosted in the central account with the new dedicated account?

- A. Use AWS Resource Access Manager (AWS RAM) to share the VPC subnet ID of the HSM that is hosted in the central account with the new dedicated account
- B. Configure the CloudHSM security group to accept inbound traffic from the private IP addresses of client instances in the new dedicated account.
- C. Use AWS Identity and Access Management (IAM) to create a cross-account role to access the CloudHSM cluster that is in the central account. Create a new IAM user in the new dedicated account. Assign the cross-account role to the new IAM user.
- D. Use AWS IAM Identity Center (AWS Single Sign-On) to create an AWS Security Token Service (AWS STS) token to authenticate from the new dedicated account to the central account
- E. Use the cross-account permissions that are assigned to the STS token to invoke an operation on the HSM in the central account.
- F. Use AWS Resource Access Manager (AWS RAM) to share the ID of the HSM that is hosted in the central account with the new dedicated account
- G. Configure the CloudHSM security group to accept inbound traffic from the private IP addresses of client instances in the new dedicated account.

Answer: A

Explanation:

<https://aws.amazon.com/premiumsupport/knowledge-center/cloudhsm-share-clusters/#:~:text=In%20the%20nav>

NEW QUESTION 108

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

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

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

Answer: CD

Explanation:

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

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

NEW QUESTION 113

A company hosts multiple externally facing applications, each isolated in its own IAM account. The company's Security team has enabled IAM WAF, IAM Config, and Amazon GuardDuty on all accounts. The company's Operations team has also joined all of the accounts to IAM Organizations and established centralized logging for CloudTrail, IAM Config, and GuardDuty. The company wants the Security team to take a reactive remediation in one account, and automate implementing this remediation as proactive prevention in all the other accounts.

How should the Security team accomplish this?

- A. Update the IAM WAF rules in the affected account and use IAM Firewall Manager to push updated IAM WAF rules across all other accounts.
- B. Use GuardDuty centralized logging and Amazon SNS to set up alerts to notify all application teams of security incidents.
- C. Use GuardDuty alerts to write an IAM Lambda function that updates all accounts by adding additional NACLs on the Amazon EC2 instances to block known malicious IP addresses.
- D. Use IAM Shield Advanced to identify threats in each individual account and then apply the account-based protections to all other accounts through Organizations.

Answer: C

NEW QUESTION 118

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

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

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

Answer: AC

Explanation:

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

NEW QUESTION 123

A company is deploying an Amazon EC2-based application. The application will include a custom health-checking component that produces health status data in JSON format. A Security Engineer must

implement a secure solution to monitor application availability in near-real time by analyzing the health status data.

Which approach should the Security Engineer use?

- A. Use Amazon CloudWatch monitoring to capture Amazon EC2 and networking metrics. Visualize metrics using Amazon CloudWatch dashboards.
- B. Run the Amazon Kinesis Agent to write the status data to Amazon Kinesis Data Firehose. Store the streaming data from Kinesis Data Firehose in Amazon Redshift.
- C. (Then run a script on the pool data and analyze the data in Amazon Redshift.)
- D. Write the status data directly to a public Amazon S3 bucket from the health-checking component. Configure S3 events to invoke an IAM Lambda function that analyzes the data.
- E. Generate events from the health-checking component and send them to Amazon CloudWatch Events. Include the status data as event payload.
- F. Use CloudWatch Events rules to invoke an IAM Lambda function that analyzes the data.

Answer: A

Explanation:

Amazon CloudWatch monitoring is a service that collects and tracks metrics from AWS resources and applications, and provides visualization tools and alarms to monitor performance and availability¹. The health status data in JSON format can be sent to CloudWatch as custom metrics², and then displayed in CloudWatch dashboards³. The other options are either inefficient or insecure for monitoring application availability in near-real time.

NEW QUESTION 127

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

The team members access the account by assuming a role that has a specific set of permissions that are necessary for the job responsibilities of the team members. All team members have permissions to perform operations on the stacks.

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

- A. Create a service role that has a composite principal that contains each service that needs the necessary permission.
- B. Configure the role to allow the sts:AssumeRole action.
- C. Create a service role that has cloudformation.amazonaws.com as the service principal.
- D. Configure the role to allow the sts:AssumeRole action.
- E. For each required set of permissions, add a separate policy to the role to allow those permissions.
- F. Add the ARN of each CloudFormation stack in the resource field of each policy.

- G. For each required set of permissions, add a separate policy to the role to allow those permission
- H. Add the ARN of each service that needs the per-missions in the resource field of the corresponding policy.
- I. Update each stack to use the service role.
- J. Add a policy to each member role to allow the iam:PassRole actio
- K. Set the policy's resource field to the ARN of the service role.

Answer: BDF

NEW QUESTION 128

A company is evaluating the use of AWS Systems Manager Session Manager to gam access to the company's Amazon EC2 instances. However, until the company implements the change, the company must protect the key file for the EC2 instances from read and write operations by any other users. When a security administrator tries to connect to a critical EC2 Linux instance during an emergency, the security administrator receives the following error. "Error Unprotected private key file - Permissions for 'ssh/my_private_key.pem' are too open". Which command should the security administrator use to modify the private key file permissions to resolve this error?

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

Answer: B

Explanation:

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

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

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

Verified References:

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

NEW QUESTION 132

A company has several petabytes of data. The company must preserve this data for 7 years to comply with regulatory requirements. The company's compliance team asks a security officer to develop a strategy that will prevent anyone from changing or deleting the data. Which solution will meet this requirement MOST cost-effectively?

- A. Create an Amazon S3 bucket
- B. Configure the bucket to use S3 Object Lock in compliance mode
- C. Upload the data to the bucket
- D. Create a resource-based bucket policy that meets all the regulatory requirements.
- E. Create an Amazon S3 bucket
- F. Configure the bucket to use S3 Object Lock in governance mode
- G. Upload the data to the bucket
- H. Create a user-based IAM policy that meets all the regulatory requirements.
- I. Create a vault in Amazon S3 Glacier
- J. Create a Vault Lock policy in S3 Glacier that meets all the regulatory requirements
- K. Upload the data to the vault.
- L. Create an Amazon S3 bucket
- M. Upload the data to the bucket
- N. Use a lifecycle rule to transition the data to a vault in S3 Glacier
- O. Create a Vault Lock policy that meets all the regulatory requirements.

Answer: C

Explanation:

To preserve the data for 7 years and prevent anyone from changing or deleting it, the security officer needs to use a service that can store the data securely and enforce compliance controls. The most cost-effective way to do this is to use Amazon S3 Glacier, which is a low-cost storage service for data archiving and long-term backup. S3 Glacier allows you to create a vault, which is a container for storing archives. Archives are any data such as photos, videos, or documents that you want to store durably and reliably.

S3 Glacier also offers a feature called Vault Lock, which helps you to easily deploy and enforce compliance controls for individual vaults with a Vault Lock policy. You can specify controls such as "write once read many" (WORM) in a Vault Lock policy and lock the policy from future edits. Once a Vault Lock policy is locked, the policy can no longer be changed or deleted. S3 Glacier enforces the controls set in the Vault Lock policy to help achieve your compliance objectives. For example, you can use Vault Lock policies to enforce data retention by denying deletes for a specified period of time.

To use S3 Glacier and Vault Lock, the security officer needs to follow these steps:

- Create a vault in S3 Glacier using the AWS Management Console, AWS Command Line Interface (AWS CLI), or AWS SDKs.
- Create a Vault Lock policy in S3 Glacier that meets all the regulatory requirements using the IAM policy language. The policy can include conditions such as `aws:CurrentTime` or `aws:SecureTransport` to further restrict access to the vault.
- Initiate the lock by attaching the Vault Lock policy to the vault, which sets the lock to an in-progress state and returns a lock ID. While the policy is in the in-progress state, you have 24 hours to validate your Vault Lock policy before the lock ID expires. To prevent your vault from exiting the in-progress state, you must complete the Vault Lock process within these 24 hours. Otherwise, your Vault Lock policy will be deleted.
- Use the lock ID to complete the lock process. If the Vault Lock policy doesn't work as expected, you can stop the Vault Lock process and restart from the beginning.
- Upload the data to the vault using either direct upload or multipart upload methods. For more information about S3 Glacier and Vault Lock, see S3 Glacier

Vault Lock.

The other options are incorrect because:

- Option A is incorrect because creating an Amazon S3 bucket and configuring it to use S3 Object Lock in compliance mode will not prevent anyone from changing or deleting the data. S3 Object Lock is a feature that allows you to store objects using a WORM model in S3. You can apply two types of object locks: retention periods and legal holds. A retention period specifies a fixed period of time during which an object remains locked. A legal hold is an indefinite lock on an object until it is removed. However, S3 Object Lock only prevents objects from being overwritten or deleted by any user, including the root user in your AWS account. It does not prevent objects from being modified by other means, such as changing their metadata or encryption settings. Moreover, S3 Object Lock requires that you enable versioning on your bucket, which will incur additional storage costs for storing multiple versions of an object.
- Option B is incorrect because creating an Amazon S3 bucket and configuring it to use S3 Object Lock in governance mode will not prevent anyone from changing or deleting the data. S3 Object Lock in governance mode works similarly to compliance mode, except that users with specific IAM permissions can change or delete objects that are locked. This means that users who have `s3:BypassGovernanceRetention` permission can remove retention periods or legal holds from objects and overwrite or delete them before they expire. This option does not provide strong enforcement for compliance controls as required by the regulatory requirements.
- Option D is incorrect because creating an Amazon S3 bucket and using a lifecycle rule to transition the data to a vault in S3 Glacier will not prevent anyone from changing or deleting the data. Lifecycle rules are actions that Amazon S3 automatically performs on objects during their lifetime. You can use lifecycle rules to transition objects between storage classes or expire them after a certain period of time. However, lifecycle rules do not apply any compliance controls on objects or prevent them from being modified or deleted by users. Moreover, transitioning objects from S3 to S3 Glacier using lifecycle rules will incur additional charges for retrieval requests and data transfers.

NEW QUESTION 134

A security engineer wants to evaluate configuration changes to a specific AWS resource to ensure that the resource meets compliance standards. However, the security engineer is concerned about a situation in which several configuration changes are made to the resource in quick succession. The security engineer wants to record only the latest configuration of that resource to indicate the cumulative impact of the set of changes.

Which solution will meet this requirement in the MOST operationally efficient way?

- A. Use AWS CloudTrail to detect the configuration changes by filtering API calls to monitor the changes. Use the most recent API call to indicate the cumulative impact of multiple calls
- B. Use AWS Config to detect the configuration changes and to record the latest configuration in case of multiple configuration changes.
- C. Use Amazon CloudWatch to detect the configuration changes by filtering API calls to monitor the change
- D. Use the most recent API call to indicate the cumulative impact of multiple calls.
- E. Use AWS Cloud Map to detect the configuration change
- F. Generate a report of configuration changes from AWS Cloud Map to track the latest state by using a sliding time window.

Answer: B

Explanation:

AWS Config is a service that enables you to assess, audit, and evaluate the configurations of your AWS resources. AWS Config continuously monitors and records your AWS resource configurations and allows you to automate the evaluation of recorded configurations against desired configurations.

To evaluate configuration changes to a specific AWS resource and ensure that it meets compliance standards, the security engineer should use AWS Config to detect the configuration changes and to record the latest configuration in case of multiple configuration changes. This will allow the security engineer to view the current state of the resource and its compliance status, as well as its configuration history and timeline.

AWS Config records configuration changes as ConfigurationItems, which are point-in-time snapshots of the resource's attributes, relationships, and metadata. If multiple configuration changes occur within a short period of time, AWS Config records only the latest ConfigurationItem for that resource. This indicates the cumulative impact of the set of changes on the resource's configuration.

This solution will meet the requirement in the most operationally efficient way, as it leverages AWS Config's features to monitor, record, and evaluate resource configurations without requiring additional tools or services.

The other options are incorrect because they either do not record the latest configuration in case of multiple configuration changes (A, C), or do not use a valid service for evaluating resource configurations (D).

Verified References:

- <https://docs.aws.amazon.com/config/latest/developerguide/WhatIsConfig.html>
- <https://docs.aws.amazon.com/config/latest/developerguide/config-item-table.html>

NEW QUESTION 139

A company wants to configure DNS Security Extensions (DNSSEC) for the company's primary domain. The company registers the domain with Amazon Route 53. The company hosts the domain on Amazon EC2 instances by using BIND.

What is the MOST operationally efficient solution that meets this requirement?

- A. Set the `dnssec-enable` option to `yes` in the BIND configuration
- B. Create a zone-signing key (ZSK) and a key-signing key (KSK). Restart the BIND service.
- C. Migrate the zone to Route 53 with DNSSEC signing enable
- D. Create a zone-signing key (ZSK) and a key-signing key (KSK) that are based on an AWS Key Management Service (AWS KMS) customer managed key.
- E. Set the `dnssec-enable` option to `yes` in the BIND configuration
- F. Create a zone-signing key (ZSK) and a key-signing key (KSK). Run the `dnssec-signzone` command to generate a delegation signer (DS) record. Use AWS Key Management Service (AWS KMS) to secure the keys.
- G. Migrate the zone to Route 53 with DNSSEC signing enable
- H. Create a key-signing key (KSK) that is based on an AWS Key Management Service (AWS KMS) customer managed key.
- I. Add a delegation signer (DS) record to the parent zone.

Answer: D

Explanation:

To configure DNSSEC for a domain registered with Route 53, the most operationally efficient solution is to migrate the zone to Route 53 with DNSSEC signing enabled, create a key-signing key (KSK) that is based on an AWS Key Management Service (AWS KMS) customer managed key, and add a delegation signer (DS) record to the parent zone. This way, Route 53 handles the zone-signing key (ZSK) and the signing of the records in the hosted zone, and the customer only needs to manage the KSK in AWS KMS and provide the DS record to the domain registrar. Option A is incorrect because it does not involve migrating the zone to Route 53, which would simplify the DNSSEC configuration. Option B is incorrect because it creates both a ZSK and a KSK based on AWS KMS customer managed keys, which is unnecessary and less efficient than letting Route 53 manage the ZSK. Option C is incorrect because it does not involve migrating the zone to Route 53, and it requires running the `dnssec-signzone` command manually, which is less efficient than letting Route 53 sign the zone automatically. Verified

References:

- <https://docs.aws.amazon.com/Route53/latest/DeveloperGuide/domain-configure-dnssec.html>
- <https://aws.amazon.com/about-aws/whats-new/2020/12/announcing-amazon-route-53-support-dnssec/>

NEW QUESTION 143

A company is developing an ecommerce application. The application uses Amazon EC2 instances and an Amazon RDS MySQL database. For compliance reasons, data must be secured in transit and at rest. The company needs a solution that minimizes operational overhead and minimizes cost. Which solution meets these requirements?

- A. Use TLS certificates from AWS Certificate Manager (ACM) with an Application Load Balancer. Deploy self-signed certificates on the EC2 instance
- B. Ensure that the database client software uses a TLS connection to Amazon RD
- C. Enable encryption of the RDS DB instanc
- D. Enable encryption on the Amazon Elastic Block Store (Amazon EBS) volumes that support the EC2 instances.
- E. Use TLS certificates from a third-party vendor with an Application Load Balance
- F. Install the same certificates on the EC2 instance
- G. Ensure that the database client software uses a TLS connection to Amazon RD
- H. Use AWS Secrets Manager for client-side encryption of application data.
- I. Use AWS CloudHSM to generate TLS certificates for the EC2 instance
- J. Install the TLS certificates on the EC2 instance
- K. Ensure that the database client software uses a TLS connection to Amazon RD
- L. Use the encryption keys form CloudHSM for client-side encryption of application data.
- M. Use Amazon CloudFront with AWS WA
- N. Send HTTP connections to the origin EC2 instance
- O. Ensure that the database client software uses a TLS connection to Amazon RD
- P. Use AWS Key Management Service (AWS KMS) for client-side encryption of application data before the data is stored in the RDS database.

Answer: A

NEW QUESTION 146

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

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

Answer: C

NEW QUESTION 147

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

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

Answer: A

Explanation:

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

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

to the EC2 instances.

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

Verified References:

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

NEW QUESTION 152

.....

Thank You for Trying Our Product

* 100% Pass or Money Back

All our products come with a 90-day Money Back Guarantee.

* One year free update

You can enjoy free update one year. 24x7 online support.

* Trusted by Millions

We currently serve more than 30,000,000 customers.

* Shop Securely

All transactions are protected by VeriSign!

100% Pass Your SCS-C01 Exam with Our Prep Materials Via below:

<https://www.certleader.com/SCS-C01-dumps.html>