

Amazon-Web-Services

Exam Questions SCS-C01

AWS Certified Security- Specialty



NEW QUESTION 1

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 2

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 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)

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

C)

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

D)

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

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

Answer: A**NEW QUESTION 4**

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 am aws kmsus-west-1 " as the principal.
- D. Create an IAM policy that allows access to the customer managed key in us-east-1. Specify arn aws rds 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 5

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 6

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 SS
- 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-eleme

NEW QUESTION 7

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

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

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

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

- A. Configure Amazon Macie to identify any objects in DOC-EXAMPLE-BUCKET1 that contain PII and that were available to the access key.
- B. Use Amazon CloudWatch Logs Insights to identify any objects in DOC-EXAMPLE-BUCKET1 that contain PII and that were available to the access key.
- C. Use Amazon OpenSearch Service (Amazon Elasticsearch Service) to query the CloudTrail logs in DOC-EXAMPLE-BUCKET2 for API calls that used the access key to access an object that contained PII.
- D. Use Amazon Athena to query the CloudTrail logs in DOC-EXAMPLE-BUCKET2 for any API calls that used the access key to access an object that contained PII.
- E. Use AWS Identity and Access Management Access Analyzer to identify any API calls that used the access key to access objects that contained PII in DOC-EXAMPLE-BUCKET1.

Answer: AD

NEW QUESTION 8

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 9

Which of the following are valid configurations for using SSL certificates with Amazon CloudFront? (Select THREE)

- A. Default AWS Certificate Manager certificate
- B. Custom SSL certificate stored in AWS KMS
- C. Default CloudFront certificate
- D. Custom SSL certificate stored in AWS Certificate Manager
- E. Default SSL certificate stored in AWS Secrets Manager
- F. Custom SSL certificate stored in AWS IAM

Answer: ABC

Explanation:

The key length for an RSA certificate that you use with CloudFront is 2048 bits, even though ACM supports larger keys. If you use an imported certificate with CloudFront, your key length must be 1024 or 2048 bits and cannot exceed 2048 bits. You must import the certificate in the US East (N. Virginia) Region. You must have permission to use and import the SSL/TLS certificate
<https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/cnames-and-https-requirements.html>

NEW QUESTION 10

A security engineer recently rotated all IAM access keys in an AWS account. The security engineer then configured AWS Config and enabled the following AWS Config managed rules; mfa-enabled-for-iam-console-access, iam-user-mfa-enabled, access-key-rotated, and iam-user-unused-credentials-check. The security engineer notices that all resources are displaying as noncompliant after the IAM GenerateCredentialReport API operation is invoked. What could be the reason for the noncompliant status?

- A. The IAM credential report was generated within the past 4 hours.
- B. The security engineer does not have the GenerateCredentialReport permission.
- C. The security engineer does not have the GetCredentialReport permission.
- D. The AWS Config rules have a MaximumExecutionFrequency value of 24 hours.

Answer: D

Explanation:

The correct answer is D. The AWS Config rules have a MaximumExecutionFrequency value of 24 hours. According to the AWS documentation¹, the MaximumExecutionFrequency parameter specifies the maximum frequency with which AWS Config runs evaluations for a rule. For AWS Config managed rules, this value can be one of the following:

- One_Hour
- Three_Hours
- Six_Hours
- Twelve_Hours
- TwentyFour_Hours

If the rule is triggered by configuration changes, it will still run evaluations when AWS Config delivers the configuration snapshot. However, if the rule is triggered periodically, it will not run evaluations more often than the specified frequency.

In this case, the security engineer enabled four AWS Config managed rules that are triggered periodically. Therefore, these rules will only run evaluations every 24 hours, regardless of when the IAM credential report is generated. This means that the resources will display as noncompliant until the next evaluation cycle, which could take up to 24 hours after the IAM access keys are rotated.

The other options are incorrect because:

- A. The IAM credential report can be generated at any time, but it will not affect the compliance status of the resources until the next evaluation cycle of the AWS Config rules.
-

B. The security engineer was able to invoke the IAM GenerateCredentialReport API operation, which means they have the GenerateCredentialReport permission. This permission is required to generate a credential report that lists all IAM users in an AWS account and their credential status².

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

References:

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

NEW QUESTION 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 backhsted IPs
- C. Configure Linux iptables or Windows Firewall to block any traffic from the malicious IP
- D. Modify the hosted zone in Amazon Route 53 and create a DNS sinkhole for the malicious IP

Answer: D

Explanation:

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

NEW QUESTION 12

A company uses AWS Organizations to manage a multi-accountAWS 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 administra-tor 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 organiza-tion. The solution must turn on AWS Config automatically during account crea-tion.

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

- A. Create an AWS CloudFormation template that contains the 10 required AVVS 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 Confi
- H. De-ploy the template by using CloudFormation StackSets in the security-01 ac-count.
- I. Create an AWS CloudFormation template that will activate AWS Confi
- J. De-ploy the template by using CloudFormation StackSets in the management-01 account.

Answer: BE

NEW QUESTION 14

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

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

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

Answer: DE

Explanation:

➤ A is incorrect because reviewing the cross-account role permissions and the S3 bucket policy is not enough to troubleshoot the permissions issue. You also need to verify that the Amazon S3 block public access option in the member account is deactivated, as well as the permissions boundary and the SCPs of the role in the member account.

➤ D is correct because evaluating the SCPs and the permissions boundary of the role in the member account can help you identify any missing permissions or explicit denies that could prevent the administrator from making the S3 bucket public.

➤ E is correct because ensuring that the S3 bucket policy explicitly allows the s3 PutBucketPublicAccess action for the role in the member account can help you override any block public access settings that could prevent the administrator from making the S3 bucket public.

NEW QUESTION 19

A company used a lift-and-shift approach to migrate from its on-premises data centers to the AWS Cloud. The company migrated on-premises VMS to Amazon

EC2 instances. Now the company wants to replace some of components that are running on the EC2 instances with managed AWS services that provide similar functionality.

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

Which solution will meet these requirements?

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

Answer: C

Explanation:

- Amazon S3 is a service that provides scalable, durable, and secure object storage. You can use Amazon S3 to store and retrieve any amount of data from anywhere on the web¹
- AWS Elastic Load Balancing is a service that distributes incoming application or network traffic across multiple targets, such as EC2 instances, containers, or IP addresses. You can use Elastic Load Balancing to increase the availability and fault tolerance of your applications²
- Elastic Load Balancing supports access logging, which captures detailed information about requests sent to your load balancer. Each log contains information such as the time the request was received, the client's IP address, latencies, request paths, and server responses. You can use access logs to analyze traffic patterns and troubleshoot issues³
- You can configure your load balancer to store access logs in an Amazon S3 bucket that you specify. You can also specify the interval for publishing the logs, which can be 5 or 60 minutes. The logs are stored in a hierarchical folder structure by load balancer name, IP address, year, month, day, and time.
- Amazon Athena is a service that allows you to analyze data in Amazon S3 using standard SQL. You can use Athena to run ad-hoc queries and get results in seconds. Athena is serverless, so there is no infrastructure to manage and you pay only for the queries that you run.
- You can use Athena to search the access logs that are stored in your S3 bucket. You can create a table in Athena that maps to your S3 bucket and then run SQL queries on the table. You can also use the Athena console or API to view and download the query results.
- You can also use Athena to create queries for the required metrics, such as the number of requests per cipher or protocol. You can then publish the metrics to Amazon CloudWatch, which is a service that monitors and manages your AWS resources and applications. You can use CloudWatch to collect and track metrics, create alarms, and automate actions based on the state of your resources.
- By using this solution, you can meet the requirements of ensuring that all the load balancer logs are centralized and searchable for auditing and that metrics are generated to show which ciphers are in use.

NEW QUESTION 22

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

Which solution will meet these requirements?

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

Answer: B

Explanation:

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

This answer is correct because it meets the requirements of sending an alert when three or more failed sign-in attempts to the AWS Management Console occur during a 5-minute period. By configuring CloudTrail to send events to CloudWatch Logs, the security engineer can create a metric filter that matches the desired pattern of failed sign-in events. Then, by creating a CloudWatch alarm based on the metric filter, the security engineer can set a threshold of 3 and a period of 5 minutes, and choose an action such as sending an email or an Amazon Simple Notification Service (Amazon SNS) message when the alarm is triggered¹².

The other options are incorrect because:

- > A. Turning on Insights events on the trail and configuring an alarm on the insight is not a solution, because Insights events are used to analyze unusual activity in management events, such as spikes in API call volume or error rates. Insights events do not capture failed sign-in attempts to the AWS Management Console3.
- > C. Creating an Amazon Athena table from the CloudTrail events and running a query for failed sign-in events is not a solution, because it does not provide a mechanism to send an alert based on the query results. Amazon Athena is an interactive query service that allows analyzing data in Amazon S3 using standard SQL, but it does not support creating notifications or alarms from queries4.
- > D. Creating an analyzer in AWS Identity and Access Management Access Analyzer and configuring it to send an Amazon SNS notification when a failed sign-in event occurs 3 times for any IAM user within a period of 5 minutes is not a solution, because IAM Access Analyzer is not a service that monitors sign-in events, but a service that helps identify resources that are shared with external entities. IAM Access Analyzer does not generate findings for failed sign-in attempts to the AWS Management Console5.

References:

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

NEW QUESTION 24

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

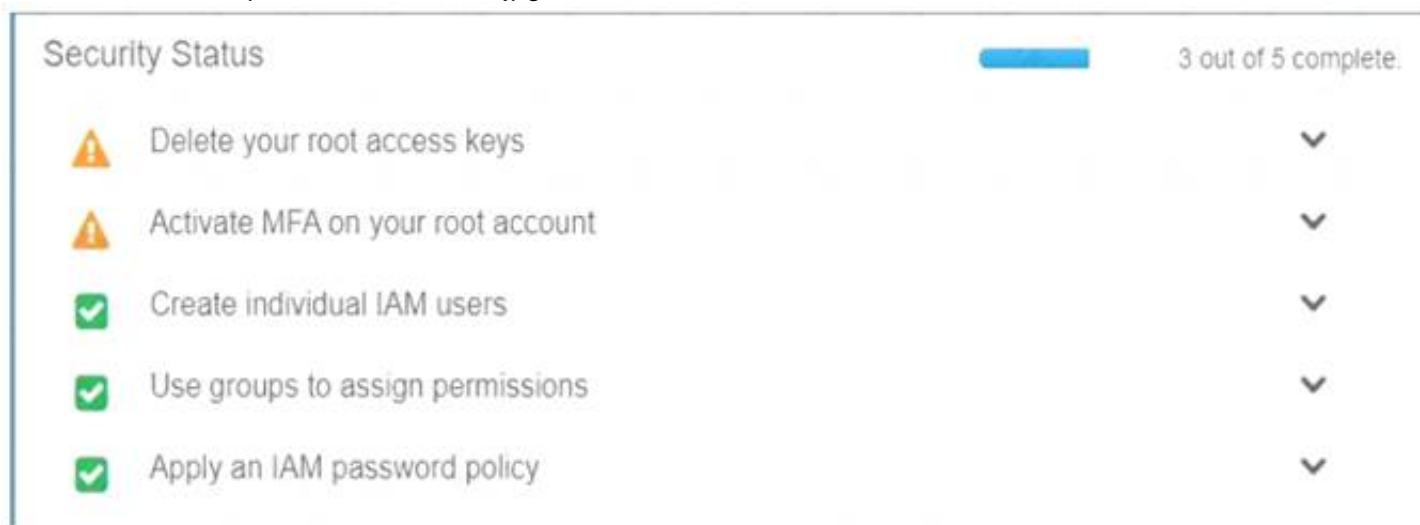
- A. Use short but complex password on the root account and any administrators.
- B. Use IAM IAM Geo-Lock and disallow anyone from logging in except for in your city.
- C. Use MFA on all users and accounts, especially on the root account.
- D. Don't write down or remember the root account password after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

Answer: C

Explanation:

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

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



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

http://docs.IAM.amazon.com/IAM/latest/UserGuide/id_credentials_mfa.html

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

NEW QUESTION 27

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

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

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

Answer: B

NEW QUESTION 32

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

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

Answer: C

Explanation:

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

NEW QUESTION 33

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

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

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

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

Answer: C

Explanation:

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

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

The other options are incorrect for the following reasons:

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

NEW QUESTION 35

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 36

A company has a single AWS account and uses an Amazon EC2 instance to test application code. The company recently discovered that the instance was compromised. The instance was serving up malware. The analysis of the instance showed that the instance was compromised 35 days ago.

A security engineer must implement a continuous monitoring solution that automatically notifies the company's security team about compromised instances through an email distribution list for high severity findings. The security engineer must implement the solution as soon as possible.

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

- A. Enable AWS Security Hub in the AWS account.
- B. Enable Amazon GuardDuty in the AWS account.
- C. Create an Amazon Simple Notification Service (Amazon SNS) topic.
- D. Subscribe the security team's email distribution list to the topic.

- E. Create an Amazon Simple Queue Service (Amazon SQS) queue
- F. Subscribe the security team's email distribution list to the queue.
- G. Create an Amazon EventBridge (Amazon CloudWatch Events) rule for GuardDuty findings of high severity
- H. Configure the rule to publish a message to the topic.
- I. Create an Amazon EventBridge (Amazon CloudWatch Events) rule for Security Hub findings of high severity
- J. Configure the rule to publish a message to the queue.

Answer: BCE

NEW QUESTION 40

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

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

Answer: C

Explanation:

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

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

NEW QUESTION 45

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

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

Answer: C

Explanation:

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

NEW QUESTION 50

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

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

Answer: A

Explanation:

The IAM Documentation mentions the following

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

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

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

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

NEW QUESTION 54

A company has multiple departments. Each department has its own IAM account. All these accounts belong to the same organization in IAM Organizations.

A large .csv file is stored in an Amazon S3 bucket in the sales department's IAM account. The company wants to allow users from the other accounts to access the .csv file's content through the combination of IAM Glue and Amazon Athena. However, the company does not want to allow users from the other accounts to access other files in the same folder.

Which solution will meet these requirements?

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

Answer: A

NEW QUESTION 56

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

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

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

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

Answer: D

Explanation:

The correct answer is D. Change the existing CloudTrail trail so that it applies to all Regions.

According to the AWS documentation¹, you can configure CloudTrail to deliver log files from multiple Regions to a single S3 bucket for a single account. To change an existing single-Region trail to log in all Regions, you must use the AWS CLI and add the `--is-multi-region-trail` option to the `update-trail` command². This will ensure that you log global service events and capture all management event activity in your account.

Option A is incorrect because creating a new CloudTrail trail for each Region will incur additional costs and increase operational overhead. Option B is incorrect because changing the S3 bucket to receive notifications will not affect the delivery of log files from other Regions. Option C is incorrect because creating a new CloudTrail trail that applies to all Regions will result in duplicate log files for the original Region and also incur additional costs.

NEW QUESTION 61

A security engineer has enabled IAM Security Hub in their IAM account, and has enabled the Center for Internet Security (CIS) IAM Foundations compliance standard. No evaluation results on compliance are returned in the Security Hub console after several hours. The engineer wants to ensure that Security Hub can evaluate their resources for CIS IAM Foundations compliance.

Which steps should the security engineer take to meet these requirements?

- A. Add full Amazon Inspector IAM permissions to the Security Hub service role to allow it to perform the CIS compliance evaluation
- B. Ensure that IAM Trusted Advisor is enabled in the account and that the Security Hub service role has permissions to retrieve the Trusted Advisor security-related recommended actions
- C. Ensure that IAM Config is enabled in the account, and that the required IAM Config rules have been created for the CIS compliance evaluation
- D. is enabled in the account, and that the required IAM Config rules have been created for the CIS compliance evaluation
- E. Ensure that the correct trail in IAM CloudTrail has been configured for monitoring by Security Hub and that the Security Hub service role has permissions to perform the `GetObject` operation on CloudTrail's Amazon S3 bucket

Answer: C

Explanation:

To ensure that Security Hub can evaluate their resources for CIS AWS Foundations compliance, the security engineer should do the following:

- Ensure that AWS Config is enabled in the account. This is a service that enables continuous assessment and audit of your AWS resources for compliance.
- Ensure that the required AWS Config rules have been created for the CIS compliance evaluation. These are rules that represent your desired configuration settings for specific AWS resources or for an entire AWS account.

NEW QUESTION 62

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

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

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

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

Answer: B

NEW QUESTION 66

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 69

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

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

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

Answer: BC

Explanation:

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

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

Verified References:

- <https://repost.aws/knowledge-center/cross-account-access-iam>
- https://docs.aws.amazon.com/organizations/latest/userguide/orgs_manage_accounts_access.html
- https://docs.aws.amazon.com/IAM/latest/UserGuide/tutorial_cross-account-with-roles.html

NEW QUESTION 72

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 76

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

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

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

Answer: A

Explanation:

➤ Option A: This option meets the requirements of following security best practices and configuring sensitive database credentials in the CloudFormation template. A dynamic reference is a way to specify external values that are stored and managed in other services, such as Secrets Manager, in the stack templates¹. When using a dynamic reference, CloudFormation retrieves the value of the specified reference when necessary during stack and change set operations¹. Dynamic references can be used for certain resources that support them, such as `AWS::RDS::DBInstance`¹. By using a dynamic reference to reference the database credentials in Secrets Manager, the company can leverage the existing integration between these services and avoid hardcoding the secret information in the template. Secrets Manager is a service that helps you protect secrets needed to access your applications, services, and IT resources². Secrets Manager enables you to rotate, manage, and retrieve database credentials, API keys, and other secrets throughout their lifecycle².

NEW QUESTION 79

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

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

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

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

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

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

Answer: D

Explanation:

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

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

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

A Secrets Manager interface VPC endpoint is a private connection between the VPC and Secrets Manager that does not require an internet gateway, NAT device, VPN connection, or AWS Direct Connect connection¹. By configuring a Secrets Manager interface VPC endpoint, the security engineer can enable the custom Lambda function to communicate with Secrets Manager without sending or receiving network traffic through the internet. The security engineer must include the Lambda function's private subnet during the configuration process to allow the function to use the endpoint².

The other options are incorrect for the following reasons:

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

NEW QUESTION 81

A company is using AWS Organizations to implement a multi-account strategy. The company does not have on-premises infrastructure. All workloads run on AWS.

The company currently has eight member accounts. The company anticipates that it will have no more than 20 AWS accounts total at any time.

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

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

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

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

- A. Use a CloudFormation template in the member accounts to launch workload
- B. Configure the template to use the `Fn::ImportValue` function to obtain the subnet ID values.
- C. Use a transit gateway in the VPC within Account-
- D. Configure the member accounts to use the transit gateway to access the subnets in Account-A to launch workloads.

- E. Use AWS Resource Access Manager (AWS RAM) to share Account-A's VPC subnets with the remaining member account
- F. Configure the member accounts to use the shared subnets to launch workloads.
- G. Create a peering connection between Account-A and the remaining member account
- H. Configure the member accounts to use the subnets in Account-A through the VPC peering connection to launch workloads.

Answer: C

Explanation:

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

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

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

The other options are incorrect because:

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

References:

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

NEW QUESTION 83

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 86

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

Please select:

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

Answer: A

Explanation:

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

Options B and C are invalid because these options will not reduce network latency Options D is invalid because this is only used to connect 2 VPC's

For more information on IAM direct connect, just browse to the below URL: <https://IAM.amazon.com/directconnect>

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

NEW QUESTION 88

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

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

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

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

Answer: A

NEW QUESTION 93

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 96

An IAM user receives an Access Denied message when the user attempts to access objects in an Amazon S3 bucket. The user and the S3 bucket are in the same AWS account. The S3 bucket is configured to use server-side encryption with AWS KMS keys (SSE-KMS) to encrypt all of its objects at rest by using a customer managed key from the same AWS account. The S3 bucket has no bucket policy defined. The IAM user has been granted permissions through an IAM policy that allows the kms:Decrypt permission to the customer managed key. The IAM policy also allows the s3:List* and s3:Get* permissions for the S3 bucket and its objects.

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

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

Answer: D

Explanation:

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

This answer is correct because the KMS key policy is the primary way to control access to the KMS key, and it must explicitly allow the AWS account to have full access to the key. If the KMS key policy has been edited to remove this permission, then the IAM policy that grants kms:Decrypt permission to the IAM user has no effect, and the IAM user cannot decrypt the objects in the S3 bucket¹².

The other options are incorrect because:

- > A. The IAM policy does not need to allow the kms:DescribeKey permission, because this permission is not required for decrypting objects in S3 using SSE-KMS. The kms:DescribeKey permission allows getting information about a KMS key, such as its creation date, description, and key state³.
- > B. The S3 bucket has not been changed to use the AWS managed key to encrypt objects at rest, because this would not cause an Access Denied message for the IAM user. The AWS managed key is a default KMS key that is created and managed by AWS for each AWS account and Region. The IAM user does not need any permissions on this key to use it for SSE-KMS⁴.
- > C. An S3 bucket policy does not need to be added to allow the IAM user to access the objects, because the IAM user already has s3:List* and s3:Get* permissions for the S3 bucket and its objects through an IAM policy. An S3 bucket policy is an optional way to grant cross-account access or public access to an S3 bucket⁵.

References:

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

NEW QUESTION 98

A company is using IAM Organizations. The company wants to restrict IAM usage to the eu-west-1 Region for all accounts under an OU that is named "development." The solution must persist restrictions to existing and new IAM accounts under the development OU.

- ☐ A. Include the following SCP on the development OU:

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "DenyNonDefaultRegions",
      "Effect": "Deny",
      "NotAction": [
        <Desired Global Services> ],
      "Resource": "*",
      "Condition": {
        "StringNotEquals": {
          "aws:RequestedRegion": [
            "eu-west-1"
          ]
        }
      },
      "ArnNotLike": {
        "aws:PrincipalARN": "arn:aws:iam::*:role/AWSExecution"
      }
    }
  ]
}
```

- ☐ B. Include the following SCP on the development account:

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "DenyNonDefaultRegions",
      "Effect": "Deny",
      "NotAction": [
        <Desired Global Services> ],
      "Resource": "*",
      "Condition": {
        "StringNotEquals": {
          "aws:RequestedRegion": [
            "eu-west-1"
          ]
        },
        "ArnNotLike": {
          "aws:PrincipalARN": "arn:aws:iam::*:role/AWSExecution"
        }
      }
    }
  ]
}
```

- ☐ C. Include the following SCP on the development OU:

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "DenyNonDefaultRegions",
      "Effect": "Deny",
      "NotAction": [
        <Desired Global Services> ],
      "Resource": "*",
      "Condition": {
        "StringEquals": {
          "aws:RequestedRegion": [
            "eu-west-1"
          ]
        },
        "ArnNotLike": {
          "aws:PrincipalARN": "arn:aws:iam::*:role/AWSExecution"
        }
      }
    }
  ]
}
```

☐ D. Include the following SCP on the development OU

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "DenyNonDefaultRegions",
      "Effect": "Allow",
      "NotAction": [
        <Desired Global Services> ],
      "Resource": "*",
      "Condition": {
        "StringNotEquals": {
          "aws:RequestedRegion": [
            "us-east-1"
          ]
        },
        "ArnNotLike": {
          "aws:PrincipalARN": "arn:aws:iam::*:role/AWSExecution"
        }
      }
    }
  ]
}
```

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

Answer: A

NEW QUESTION 101

A company is using AWS WAF to protect a customized public API service that is based on Amazon EC2 instances. The API uses an Application Load Balancer. The AWS WAF web ACL is configured with an AWS Managed Rules rule group. After a software upgrade to the API and the client application, some types of requests are no longer working and are causing application stability issues. A security engineer discovers that AWS WAF logging is not turned on for the web ACL. The security engineer needs to immediately return the application to service, resolve the issue, and ensure that logging is not turned off in the future. The security engineer turns on logging for the web ACL and specifies Amazon Cloud-Watch Logs as the destination. Which additional set of steps should the security engineer take to meet the re-quirements?

- A. Edit the rules in the web ACL to include rules with Count action
- B. Review the logs to determine which rule is blocking the reques
- C. Modify the IAM policy of all AWS WAF administrators so that they cannot remove the log-ging configuration for any AWS WAF web ACLs.
- D. Edit the rules in the web ACL to include rules with Count action
- E. Review the logs to determine which rule is blocking the reques
- F. Modify the AWS WAF resource policy so that AWS WAF administrators cannot remove the log-ging configuration for any AWS WAF web ACLs.
- G. Edit the rules in the web ACL to include rules with Count and Challenge action
- H. Review the logs to determine which rule is blocking the reques
- I. Modify the AWS WAF resource policy so that AWS WAF administrators cannot remove the logging configuration for any AWS WAF web ACLs.
- J. Edit the rules in the web ACL to include rules with Count and Challenge action
- K. Review the logs to determine which rule is blocking the reques
- L. Modify the IAM policy of all AWS WAF administrators so that they cannot remove the logging configuration for any AWS WAF web ACLs.

Answer: A

Explanation:

This answer is correct because it meets the requirements of returning the application to service, resolving the issue, and ensuring that logging is not turned off in the future. By editing the rules in the web ACL to include rules with Count actions, the security engineer can test the effect of each rule without blocking or allowing requests. By reviewing the logs, the security engineer can identify which rule is causing the problem and modify or delete it accordingly. By modifying the IAM policy of all AWS WAF administrators, the security engineer can restrict their permissions to prevent them from removing the logging configuration for any AWS WAF web ACLs.

NEW QUESTION 103

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 107

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

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

- A. Configure cluster security groups for each application module to control access to database users that are required for read-only and read/write.

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

Answer: CD

Explanation:

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

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

NEW QUESTION 109

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 112

A company has implemented IAM WAF and Amazon CloudFront for an application. The application runs on Amazon EC2 instances that are part of an Auto Scaling group. The Auto Scaling group is behind an Application Load Balancer (ALB).

The IAM WAF web ACL uses an IAM Managed Rules rule group and is associated with the CloudFront distribution. CloudFront receives the request from IAM WAF and then uses the ALB as the distribution's origin.

During a security review, a security engineer discovers that the infrastructure is susceptible to a large, layer 7 DDoS attack.

How can the security engineer improve the security at the edge of the solution to defend against this type of attack?

- A. Configure the CloudFront distribution to use the Lambda@Edge feature
- B. Create an IAM Lambda function that imposes a rate limit on CloudFront viewer request
- C. Block the request if the rate limit is exceeded.
- D. Configure the IAM WAF web ACL so that the web ACL has more capacity units to process all IAM WAF rules faster.
- E. Configure IAM WAF with a rate-based rule that imposes a rate limit that automatically blocks requests when the rate limit is exceeded.
- F. Configure the CloudFront distribution to use IAM WAF as its origin instead of the ALB.

Answer: C

Explanation:

To improve the security at the edge of the solution to defend against a large, layer 7 DDoS attack, the security engineer should do the following:

- Configure AWS WAF with a rate-based rule that imposes a rate limit that automatically blocks requests when the rate limit is exceeded. This allows the security engineer to use a rule that tracks the number of requests from a single IP address and blocks subsequent requests if they exceed a specified threshold within a specified time period.

NEW QUESTION 115

An international company has established a new business entity in South Korea. The company also has established a new AWS account to contain the workload for the South Korean region. The company has set up the workload in the new account in the ap-northeast-2 Region. The workload consists of three Auto Scaling groups of Amazon EC2 instances. All workloads that operate in this Region must keep system logs and application logs for 7 years.

A security engineer must implement a solution to ensure that no logging data is lost for each instance during scaling activities. The solution also must keep the logs for only the required period of 7 years.

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

- A. Ensure that the Amazon CloudWatch agent is installed on all the EC2 instances that the Auto Scaling groups launch
- B. Generate a CloudWatch agent configuration file to forward the required logs to Amazon CloudWatch Logs.
- C. Set the log retention for desired log groups to 7 years.
- D. Attach an IAM role to the launch configuration or launch template that the Auto Scaling groups use. Configure the role to provide the necessary permissions to forward logs to Amazon CloudWatch Logs.
- E. Attach an IAM role to the launch configuration or launch template that the Auto Scaling groups use. Configure the role to provide the necessary permissions to forward logs to Amazon S3.
- F. Ensure that a log forwarding application is installed on all the EC2 instances that the Auto Scaling groups launch
- G. Configure the log forwarding application to periodically bundle the logs and forward the logs to Amazon S3.
- H. Configure an Amazon S3 Lifecycle policy on the target S3 bucket to expire objects after 7 years.

Answer: ABC

Explanation:

The correct combination of steps that the security engineer should take to meet these requirements are A. Ensure that the Amazon CloudWatch agent is installed on all the EC2 instances that the Auto Scaling groups launch. Generate a CloudWatch agent configuration file to forward the required logs to Amazon CloudWatch

Logs., B. Set the log retention for desired log groups to 7 years., and C. Attach an IAM role to the launch configuration or launch template that the Auto Scaling groups use. Configure the role to provide the necessary permissions to forward logs to Amazon CloudWatch Logs.

* A. This answer is correct because it meets the requirement of ensuring that no logging data is lost for each instance during scaling activities. By installing the CloudWatch agent on all the EC2 instances, the security engineer can collect and send system logs and application logs to CloudWatch Logs, which is a service that stores and monitors log data. By generating a CloudWatch agent configuration file, the security engineer can specify which logs to forward and how often.

* B. This answer is correct because it meets the requirement of keeping the logs for only the required period of 7 years. By setting the log retention for desired log groups, the security engineer can control how long

CloudWatch Logs retains log events before deleting them. The security engineer can choose a predefined retention period of 7 years, or use a custom value.

* C. This answer is correct because it meets the requirement of providing the necessary permissions to forward logs to CloudWatch Logs. By attaching an IAM role to the launch configuration or launch template that the Auto Scaling groups use, the security engineer can grant permissions to the EC2 instances that are launched by the Auto Scaling groups. By configuring the role to provide the necessary permissions, such as `cloudwatch:PutLogEvents` and `cloudwatch:CreateLogStream`, the security engineer can allow the EC2 instances to send log data to CloudWatch Logs.

NEW QUESTION 119

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 permission
- 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 permissions in the resource field of the corresponding policy.
- I. Update each stack to use the service role.
- J. Add a policy to each member role to allow the `iam:PassRole` action
- K. Set the policy's resource field to the ARN of the service role.

Answer: BDF

NEW QUESTION 124

An organization wants to log all IAM API calls made within all of its IAM accounts, and must have a central place to analyze these logs. What steps should be taken to meet these requirements in the MOST secure manner? (Select TWO)

- A. Turn on IAM CloudTrail in each IAM account
- B. Turn on CloudTrail in only the account that will be storing the logs
- C. Update the bucket ACL of the bucket in the account that will be storing the logs so that other accounts can log to it
- D. Create a service-based role for CloudTrail and associate it with CloudTrail in each account
- E. Update the bucket policy of the bucket in the account that will be storing the logs so that other accounts can log to it

Answer: AE

Explanation:

these are the steps that can meet the requirements in the most secure manner. CloudTrail is a service that records AWS API calls and delivers log files to an S3 bucket. Turning on CloudTrail in each IAM account can help capture all IAM API calls made within those accounts. Updating the bucket policy of the bucket in the account that will be storing the logs can help grant other accounts permission to write log files to that bucket. The other options are either unnecessary or insecure for logging and analyzing IAM API calls.

NEW QUESTION 128

A company's policy requires that all API keys be encrypted and stored separately from source code in a centralized security account. This security account is managed by the company's security team. However, an audit revealed that an API key is stored with the source code of an IAM Lambda function in an IAM CodeCommit repository in the DevOps account.

How should the security team securely store the API key?

- A. Create a CodeCommit repository in the security account using IAM Key Management Service (IAMKMS) for encryption. Require the development team to migrate the Lambda source code to this repository.
- B. Store the API key in an Amazon S3 bucket in the security account using server-side encryption with Amazon S3 managed encryption keys (SSE-S3) to encrypt the key. Create a signed URL for the S3 key.
- C. and specify the URL in a Lambda environmental variable in the IAM CloudFormation template. Update the Lambda function code to retrieve the key using the URL and call the API.
- D. Create a secret in IAM Secrets Manager in the security account to store the API key using IAM Key Management Service (IAM KMS) for encryption. Grant access to the IAM role used by the Lambda function so that the function can retrieve the key from Secrets Manager and call the API.
- E. Create an encrypted environment variable for the Lambda function to store the API key using IAM Key Management Service (IAM KMS) for encryption. Grant access to the IAM role used by the Lambda function so that the function can decrypt the key at runtime.

Answer: C

Explanation:

To securely store the API key, the security team should do the following:

- Create a secret in AWS Secrets Manager in the security account to store the API key using AWS Key Management Service (AWS KMS) for encryption. This allows the security team to encrypt and manage the API key centrally, and to configure automatic rotation schedules for it.
- Grant access to the IAM role used by the Lambda function so that the function can retrieve the key from Secrets Manager and call the API. This allows the security team to avoid storing the API key with the source code, and to use IAM policies to control access to the secret.

NEW QUESTION 133

A company has multiple accounts in the AWS Cloud. Users in the developer account need to have access to specific resources in the production account. What is the MOST secure way to provide this access?

- A. Create one IAM user in the production account
- B. Grant the appropriate permissions to the resources that are needed
- C. Share the password only with the users that need access.
- D. Create cross-account access with an IAM role in the developer account
- E. Grant the appropriate permissions to this role
- F. Allow users in the developer account to assume this role to access the production resources.
- G. Create cross-account access with an IAM user account in the production account
- H. Grant the appropriate permissions to this user account
- I. Allow users in the developer account to use this user account to access the production resources.
- J. Create cross-account access with an IAM role in the production account
- K. Grant the appropriate permissions to this role
- L. Allow users in the developer account to assume this role to access the production resources.

Answer: D

Explanation:

https://docs.aws.amazon.com/IAM/latest/UserGuide/tutorial_cross-account-with-roles.html

NEW QUESTION 136

A Security Engineer has been tasked with enabling IAM Security Hub to monitor Amazon EC2 instances for CVE in a single IAM account. The Engineer has already enabled IAM Security Hub and Amazon Inspector in the IAM Management Console and has installed the Amazon Inspector agent on an EC2 instances that need to be monitored.

Which additional steps should the Security Engineer take to meet this requirement?

- A. Configure the Amazon Inspector agent to use the CVE rule package
- B. Configure the Amazon Inspector agent to use the CVE rule package. Configure Security Hub to ingest from IAM Inspector by writing a custom resource policy
- C. Configure the Security Hub agent to use the CVE rule package. Configure IAM Inspector to ingest from Security Hub by writing a custom resource policy
- D. Configure the Amazon Inspector agent to use the CVE rule package. Install an additional Integration library. Allow the Amazon Inspector agent to communicate with Security Hub

Answer: D

Explanation:

You need to configure the Amazon Inspector agent to use the CVE rule package, which is a set of rules that check for vulnerabilities and exposures on your EC2 instances⁵. You also need to install an additional integration library that enables communication between the Amazon Inspector agent and Security Hub⁶. Security Hub is a service that provides you with a comprehensive view of your security state in AWS and helps you check your environment against security industry standards and best practices⁷. The other options are either incorrect or incomplete for meeting the requirement.

NEW QUESTION 140

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 144

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

of internet gateways. NAT gateways, and egress-only gateways.

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

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

Answer: ADF

NEW QUESTION 148

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 enabled
- D. Create a zone-signing key (ZSK) and a key-signing key (KSK) that are based on an AWS
- E. Key Management Service (AWS KMS) customer managed key.
- F. Set the dnssec-enable option to yes in the BIND configuration
- G. 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
- H. Key Management Service (AWS KMS) to secure the keys.
- I. Migrate the zone to Route 53 with DNSSEC signing enabled
- J. Create a key-signing key (KSK) that is based on an AWS Key Management Service (AWS KMS) customer managed key
- K. 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 149

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 RDS
- C. Enable encryption of the RDS DB instance
- 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 Balancer
- F. Install the same certificates on the EC2 instance
- G. Ensure that the database client software uses a TLS connection to Amazon RDS
- 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 RDS
- L. Use the encryption keys from CloudHSM for client-side encryption of application data.
- M. Use Amazon CloudFront with AWS WAF
- N. Send HTTP connections to the origin EC2 instance
- O. Ensure that the database client software uses a TLS connection to Amazon RDS
- 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 153

A company hosts an application on Amazon EC2 that is subject to specific rules for regulatory compliance. One rule states that traffic to and from the workload must be inspected for network-level attacks. This involves inspecting the whole packet.

To comply with this regulatory rule, a security engineer must install intrusion detection software on a c5n.4xlarge EC2 instance. The engineer must then configure the software to monitor traffic to and from the application instances.

What should the security engineer do next?

- A. Place the network interface in promiscuous mode to capture the traffic.
- B. Configure VPC Flow Logs to send traffic to the monitoring EC2 instance using a Network Load Balancer.
- C. Configure VPC traffic mirroring to send traffic to the monitoring EC2 instance using a Network Load Balancer.

D. Use Amazon Inspector to detect network-level attacks and trigger an IAM Lambda function to send the suspicious packets to the EC2 instance.

Answer: D

NEW QUESTION 158

A company is developing a highly resilient application to be hosted on multiple Amazon EC2 instances . The application will store highly sensitive user data in Amazon RDS tables

The application must

- Include migration to a different IAM Region in the application disaster recovery plan.
- Provide a full audit trail of encryption key administration events
- Allow only company administrators to administer keys.
- Protect data at rest using application layer encryption

A Security Engineer is evaluating options for encryption key management

Why should the Security Engineer choose IAM CloudHSM over IAM KMS for encryption key management in this situation?

- A. The key administration event logging generated by CloudHSM is significantly more extensive than IAM KMS.
- B. CloudHSM ensures that only company support staff can administer encryption keys, whereas IAM KMS allows IAM staff to administer keys
- C. The ciphertext produced by CloudHSM provides more robust protection against brute force decryption attacks than the ciphertext produced by IAM KMS
- D. CloudHSM provides the ability to copy keys to a different Region, whereas IAM KMS does not

Answer: B

Explanation:

CloudHSM allows full control of your keys such including Symmetric (AES), Asymmetric (RSA), Sha-256, SHA 512, Hash Based, Digital Signatures (RSA). On the other hand, AWS Key Management Service is a multi-tenant key storage that is owned and managed by AWS1.

References: 1: What are the differences between AWS Cloud HSM and KMS?

NEW QUESTION 161

.....

Thank You for Trying Our Product

We offer two products:

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

2nd - Questions and Answers in PDF Format

SCS-C01 Practice Exam Features:

- * SCS-C01 Questions and Answers Updated Frequently
- * SCS-C01 Practice Questions Verified by Expert Senior Certified Staff
- * SCS-C01 Most Realistic Questions that Guarantee you a Pass on Your FirstTry
- * SCS-C01 Practice Test Questions in Multiple Choice Formats and Updatesfor 1 Year

100% Actual & Verified — Instant Download, Please Click
[Order The SCS-C01 Practice Test Here](#)