

SCS-C01 Dumps

AWS Certified Security- Specialty

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



NEW QUESTION 1

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

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

Which the SIMPLEST change that would address this server issue?

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

Answer: A

Explanation:

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

NEW QUESTION 2

A company wants to protect its website from man in-the-middle attacks by using Amazon CloudFront. Which solution will meet these requirements with the LEAST operational overhead?

- A. Use the SimpleCORS managed response headers policy.
- B. Use a Lambda@Edge function to add the Strict-Transport-Security response header.
- C. Use the SecurityHeadersPolicy managed response headers policy.
- D. Include the X-XSS-Protection header in a custom response headers policy.

Answer: C

Explanation:

<https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/using-managed-response-headers-poli> The SecurityHeadersPolicy is a managed policy provided by Amazon CloudFront that includes a set of recommended security headers to enhance the security of your website. These headers help protect against various types of attacks, including man-in-the-middle attacks. By applying the SecurityHeadersPolicy to your CloudFront distribution, the necessary security headers will be automatically added to the responses sent by CloudFront. This reduces operational overhead because you don't have to manually configure or manage the headers yourself.

NEW QUESTION 3

A company wants to prevent SSH access through the use of SSH key pairs for any Amazon Linux 2 Amazon EC2 instances in its AWS account. However, a system administrator occasionally will need to access these EC2 instances through SSH in an emergency. For auditing purposes, the company needs to record any commands that a user runs in an EC2 instance.

What should a security engineer do to configure access to these EC2 instances to meet these requirements?

- A. Use the EC2 serial console Configure the EC2 serial console to save all commands that are entered to an Amazon S3 bucket
- B. Provide the EC2 instances with an IAM role that allows the EC2 serial console to access Amazon S3. Configure an IAM account for the system administrator
- C. Provide an IAM policy that allows the IAM account to use the EC2 serial console.
- D. Use EC2 Instance Connect Configure EC2 Instance Connect to save all commands that are entered to Amazon CloudWatch Log
- E. Provide the EC2 instances with an IAM role that allows the EC2 instances to access CloudWatch Logs Configure an IAM account for the system administrator
- F. Provide an IAM policy that allows the IAM account to use EC2 Instance Connect.
- G. Use an EC2 key pair with an EC2 instance that needs SSH access Access the EC2 instance with this key pair by using SS
- H. Configure the EC2 instance to save all commands that are entered to Amazon CloudWatch Log
- I. Provide the EC2 instance with an IAM role that allows the EC2 instance to access Amazon S3 and CloudWatch Logs.
- J. Use AWS Systems Manager Session Manager Configure Session Manager to save all commands that are entered in a session to an Amazon S3 bucket
- K. Provide the EC2 instances with an IAM role that allows Systems Manager to manage the EC2 instance
- L. Configure an IAM account for the system administrator Provide an IAM policy that allows the IAM account to use Session Manager.

Answer: D

Explanation:

Open the AWS Systems Manager console at <https://console.aws.amazon.com/systems-manager/>. In the navigation pane, choose Session Manager. Choose the Preferences tab, and then choose Edit. Select the check box next to Enable under S3 logging. (Recommended) Select the check box next to Allow only encrypted S3 buckets. With this option turned on, log data is encrypted using the server-side encryption key specified for the bucket. If you don't want to encrypt the log data that is sent to Amazon S3, clear the check box. You must also clear the check box if encryption isn't allowed on the S3 bucket.

NEW QUESTION 4

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

Which solution will meet these requirements?

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

Answer: A

Explanation:

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

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

NEW QUESTION 5

A security engineer wants to use Amazon Simple Notification Service (Amazon SNS) to send email alerts to a company's security team for Amazon GuardDuty findings

that have a High severity level. The security engineer also wants to deliver these findings to a visualization tool for further examination.

Which solution will meet these requirements?

- A. Set up GuardDuty to send notifications to an Amazon CloudWatch alarm with two targets in CloudWatc
- B. From CloudWatch, stream the findings through Amazon Kinesis Data Streams into an Amazon OpenSearch Service domain as the first target for deliver
- C. Use Amazon QuickSight to visualize the finding
- D. Use OpenSearch queries for further analysi
- E. Deliver email alerts to the security team by configuring an SNS topic as a second target for the CloudWatch alar
- F. Use event pattern matching with an Amazon EventBridge event rule to send only High severity findings in the alerts.
- G. Set up GuardDuty to send notifications to AWS CloudTrail with two targets in CloudTrai
- H. From CloudTrail, stream the findings through Amazon Kinesis Data Firehose into an Amazon OpenSearch Service domain as the first target for deliver
- I. Use OpenSearch Dashboards to visualize the finding
- J. Use OpenSearch queries for further analysi
- K. Deliver email alerts to the security team by configuring an SNS topic as a second target for CloudTrai
- L. Use event pattern matching with a CloudTrail event rule to send only High severity findings in the alerts.
- M. Set up GuardDuty to send notifications to Amazon EventBridge with two target
- N. From EventBridge, stream the findings through Amazon Kinesis Data Firehose into an Amazon OpenSearch Service domain as the first target for deliver
- O. Use OpenSearch Dashboards to visualize the finding
- P. Use OpenSearch queries for further analysi
- Q. Deliver email alerts to the security team by configuring an SNS topic as a second target for EventBridg
- R. Use event pattern matching with an EventBridge event rule to send only High severity findings in the alerts.
- S. Set up GuardDuty to send notifications to Amazon EventBridge with two target
- T. From EventBridge, stream the findings through Amazon Kinesis Data Streams into an Amazon OpenSearch Service domain as the first target for deliver
- . Use Amazon QuickSight to visualize the finding
- . Use OpenSearch queries for further analysi
- . Deliver email alerts to the security team by configuring an SNS topic as a second target for EventBridg
- . Use event pattern matching with an EventBridge event rule to send only High severity findings in the alerts.

Answer: C

NEW QUESTION 6

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 7

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

The company must create the break glass user. The company must log any activities of the break glass user and send the logs to a security team. Which combination of solutions will meet these requirements? (Select TWO.)

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

Answer: AE

Explanation:

The combination of solutions that will meet the requirements are:

- > A. Create a local individual break glass IAM user for the security team. Create a trail in AWS CloudTrail that has Amazon CloudWatch Logs turned on. Use Amazon EventBridge to monitor local user activities. This is a valid solution because it allows the security team to access the workload AWS account and instances using a local IAM user that does not depend on SAML federation. It also enables logging and monitoring of the break glass user activities using AWS CloudTrail, Amazon CloudWatch Logs, and Amazon EventBridge123.
- > E. Configure AWS Systems Manager Session Manager for Amazon EC2. Configure an AWS CloudTrail filter based on Session Manager. Send the results to an Amazon Simple Notification Service (Amazon SNS) topic. This is a valid solution because it allows the security team to access the workload instances without opening any inbound ports or managing SSH keys or bastion hosts. It also enables logging and notification of the break glass user activities using AWS CloudTrail, Session Manager, and Amazon SNS456.

The other options are incorrect because:

- > B. Creating a break glass EC2 key pair for the AWS account and providing it to the security team is not a valid solution, because it requires opening inbound ports on the instances and managing SSH keys, which increases the security risk and complexity7.
- > C. Creating a break glass IAM role for the account and allowing security team members to perform the AssumeRoleWithSAML operation is not a valid solution, because it still depends on SAML federation, which might not work in case of SAML errors8.
- > D. Creating a local individual break glass IAM user on the operating system level of each workload instance and configuring unrestricted security groups on the instances to grant access to the break glass IAM users is not a valid solution, because it requires opening inbound ports on the instances and managing multiple local users, which increases the security risk and complexity9.

References:

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

NEW QUESTION 8

A company's Chief Security Officer has requested that a Security Analyst review and improve the security posture of each company IAM account The Security Analyst decides to do this by Improving IAM account root user security.

Which actions should the Security Analyst take to meet these requirements? (Select THREE.)

- A. Delete the access keys for the account root user in every account.
- B. Create an admin IAM user with administrative privileges and delete the account root user in every account.
- C. Implement a strong password to help protect account-level access to the IAM Management Console by the account root user.
- D. Enable multi-factor authentication (MFA) on every account root user in all accounts.
- E. Create a custom IAM policy to limit permissions to required actions for the account root user and attach the policy to the account root user.
- F. Attach an IAM role to the account root user to make use of the automated credential rotation in IAM STS.

Answer: ADE

Explanation:

because these are the actions that can improve IAM account root user security. IAM account root user is a user that has complete access to all AWS resources and services in an account. IAM account root user security is a set of best practices that help protect the account root user from unauthorized or accidental use. Deleting the access keys for the account root user in every account can help prevent programmatic access by the account root user, which reduces the risk of compromise or misuse. Enabling MFA on every account root user in all accounts can help add an extra layer of security for console access by requiring a verification code in addition to a password. Creating a custom IAM policy to limit permissions to required actions for the account root user and attaching the policy to the account root user can help enforce the principle of least privilege and restrict the account root user from performing unnecessary or dangerous actions. The other options are either invalid or ineffective for improving IAM account root user security.

NEW QUESTION 9

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

specific AWS services.

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

- A. Use AWS SSO to set up service-linked roles with IAM policy statements that include the Condition, Resource, and NotAction elements to allow access to only the Regions and services that are needed.
- B. Deactivate AWS Security Token Service (AWS STS) in Regions that the developers are not allowed to use.
- C. Create SCPs that include the Condition, Resource, and NotAction elements to allow access to only the Regions and services that are needed.
- D. For each AWS account, create tailored identity-based policies for AWS 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 10

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

What should the security engineer recommend?

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

Answer: C

NEW QUESTION 10

A company hosts business-critical applications on Amazon EC2 instances in a VPC. The VPC uses default DHCP options sets. A security engineer needs to log all DNS queries that internal resources make in the VPC. The security engineer also must create a list of the most common DNS queries over time.

Which solution will meet these requirements?

- A. Install the Amazon CloudWatch agent on each EC2 instance in the VP
- B. Use the CloudWatch agent to stream the DNS query logs to an Amazon CloudWatch Logs log grou
- C. Use CloudWatch metric filters to automatically generate metrics that list the most common ONS queries.
- D. Install a BIND DNS server in the VP
- E. Create a bash script to list the DNS request number of common DNS queries from the BIND logs.
- F. Create VPC flow logs for all subnets in the VP
- G. Stream the flow logs to an Amazon CloudWatch Logs log grou
- H. Use CloudWatch Logs Insights to list the most common DNS queries for the log group in a custom dashboard.
- I. Configure Amazon Route 53 Resolver query loggin
- J. Add an Amazon CloudWatch Logs log group as the destinatio
- K. Use Amazon CloudWatch Contributor Insights to analyze the data and create time series that display the most common DNS queries.

Answer: D

Explanation:

<https://aws.amazon.com/blogs/aws/log-your-vpc-dns-queries-with-route-53-resolver-query-logs/>

NEW QUESTION 14

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 18

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 20

A company wants to monitor the deletion of AWS Key Management Service (AWS KMS) customer managed keys. A security engineer needs to create an alarm that will notify the company before a KMS key is deleted. The security engineer has configured the integration of AWS CloudTrail with Amazon CloudWatch. What should the security engineer do next to meet these requirements?

- A. Specify the deletion time of the key material during KMS key creation
- B. Create a custom AWS Config rule to assess the key's scheduled deletion
- C. Configure the rule to trigger upon a configuration change
- D. Send a message to an Amazon Simple Notification Service (Amazon SNS) topic if the key is scheduled for deletion.
- E. Create an Amazon EventBridge rule to detect KMS API calls of DeleteAlias
- F. Create an AWS Lambda function to send an Amazon Simple Notification Service (Amazon SNS) message to the company
- G. Add the Lambda function as the target of the EventBridge rule.
- H. Create an Amazon EventBridge rule to detect KMS API calls of DisableKey and ScheduleKeyDeletion. Create an AWS Lambda function to send an Amazon Simple Notification Service (Amazon SNS) message to the company
- I. Add the Lambda function as the target of the EventBridge rule.
- J. Create an Amazon Simple Notification Service (Amazon SNS) policy to detect KMS API calls of RevokeGrant and ScheduleKeyDeletion. Create an AWS Lambda function to generate the alarm and send the notification to the company
- K. Add the Lambda function as the target of the SNS policy.

Answer: C

Explanation:

The AWS documentation states that you can create an Amazon EventBridge rule to detect KMS API calls of DisableKey and ScheduleKeyDeletion. You can then create an AWS Lambda function to send an Amazon Simple Notification Service (Amazon SNS) message to the company. You can add the Lambda function as the target of the EventBridge rule. This method will meet the requirements.

References: : AWS KMS Developer Guide

NEW QUESTION 21

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 status.
- 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 operation.

References:

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

NEW QUESTION 22

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

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

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

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

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

Answer: BE

NEW QUESTION 26

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's 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 in 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 in 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 in 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 28

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 the 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 29

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

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

Answer: BDF

Explanation:

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

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

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

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

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

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

before delivering it to your destination.

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

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

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

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

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

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

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

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

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

NEW QUESTION 30

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 34

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 36

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 38

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

How should a security engineer implement this solution?

- A. Add the file-system-id_efs_IAM-region_amazonIAM.com URL to the allow list for the data center firewall. Install the IAM CLI on the data center-based servers to

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

Answer: B

Explanation:

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

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

NEW QUESTION 41

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 42

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 43

A security engineer is checking an AWS CloudFormation template for vulnerabilities. The security engineer finds a parameter that has a default value that exposes an application's API key in plaintext. The parameter is referenced several times throughout the template. The security engineer must replace the parameter while maintaining the ability to reference the value in the template. Which solution will meet these requirements in the MOST secure way? `{resolve:s3:MyBucketName:MyObjectName}}`.

- A. Store the API key value as a SecureString parameter in AWS Systems Manager Parameter Store
- B. In the template, replace all references to the value with `{{resolve:ssm:MySSMParameterName:}}`.
- C. Store the API key value in AWS Secrets Manager
- D. In the template, replace all references to the value with `{{resolve:secretsmanager:MySecretId:SecretString}}`.
- E. Store the API key value in Amazon DynamoDB
- F. In the template, replace all references to the value with `{{resolve:dynamodb:MyTableName:MyPrimaryKey}}`.
- G. Store the API key value in a new Amazon S3 bucket
- H. In the template, replace all references to the value with `{`

Answer: B

Explanation:

The correct answer is B. Store the API key value in AWS Secrets Manager. In the template, replace all references to the value with

`{{resolve:secretsmanager:MySecretId:SecretString}}`.

This answer is correct because AWS Secrets Manager is a service that helps you protect secrets that are needed to access your applications, services, and IT resources. You can store and manage secrets such as database credentials, API keys, and other sensitive data in Secrets Manager. You can also use Secrets Manager to rotate, manage, and retrieve your secrets throughout their lifecycle¹. Secrets Manager integrates with AWS CloudFormation, which allows you to reference secrets from your templates using the

`{{resolve:secretsmanager:...}}` syntax². This way, you can avoid exposing your secrets in plaintext and still use them in your resources.

The other options are incorrect because:

- A. Storing the API key value as a SecureString parameter in AWS Systems Manager Parameter Store is not a solution, because AWS CloudFormation does not support references to SecureString parameters. This means that you cannot use the `{{resolve:ssm:...}}` syntax to retrieve encrypted parameter values from Parameter Store³. You would have to use a custom resource or a Lambda function to decrypt the parameter value, which adds complexity and overhead to your template.
- C. Storing the API key value in Amazon DynamoDB is not a solution, because AWS CloudFormation does not support references to DynamoDB items. This means that you cannot use the `{{resolve:dynamodb:...}}` syntax to retrieve item values from DynamoDB tables⁴. You would have to use a custom resource or a Lambda function to query the DynamoDB table, which adds complexity and overhead to your template.
- D. Storing the API key value in a new Amazon S3 bucket is not a solution, because AWS CloudFormation does not support references to S3 objects. This means that you cannot use the `{{resolve:s3:...}}` syntax to retrieve object values from S3 buckets⁵. You would have to use a custom resource or a Lambda function to download the object from S3, which adds complexity and overhead to your template.

References:

1: What is AWS Secrets Manager? 2: Referencing AWS Secrets Manager secrets from Parameter Store parameters 3: Using dynamic references to specify template values 4: Amazon DynamoDB 5: Amazon Simple Storage Service (S3)

NEW QUESTION 48

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

Which solution meets these criteria?

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

Answer: A

Explanation:

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

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

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

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

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

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

References:

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

NEW QUESTION 49

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

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

What should the Security Engineer do to troubleshoot this issue?

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

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

B)
Add the following statement to the CMK key policy:

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

C)
Add the following statement to the CMK key policy:

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

D)
Add the following statement to the CMK key policy:

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

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

Answer: D

NEW QUESTION 50

A security engineer receives a notice from the AWS Abuse team about suspicious activity from a Linux-based Amazon EC2 instance that uses Amazon Elastic Block Store (Amazon EBS)-based storage. The instance is making connections to known malicious addresses.

The instance is in a development account within a VPC that is in the us-east-1 Region. The VPC contains an internet gateway and has a subnet in us-east-1a and us-east-1b. Each subnet is associated with a route table that uses the internet gateway as a default route. Each subnet also uses the default network ACL. The suspicious EC2 instance runs within the us-east-1b subnet. During an initial investigation, a security engineer discovers that the suspicious instance is the only instance that runs in the subnet.

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

- A. Log in to the suspicious instance and use the netstat command to identify remote connections. Use the IP addresses from these remote connections to create deny rules in the security group of the instance. Install diagnostic tools on the instance for investigation. Update the outbound network ACL for the subnet in us-east-1b to explicitly deny all connections as the first rule during the investigation of the instance.
- B. Update the outbound network ACL for the subnet in us-east-1b to explicitly deny all connections as the first rule. Replace the security group with a new security

group that allows connections only from a diagnostics security group Update the outbound network ACL for the us-east-1b subnet to remove the deny all rule Launch a new EC2 instance that has diagnostic tools Assign the new security group to the new EC2 instance Use the new EC2 instance to investigate the suspicious instance

- C. Ensure that the Amazon Elastic Block Store (Amazon EBS) volumes that are attached to the suspicious EC2 instance will not delete upon termination Terminate the instance Launch a new EC2 instance in us-east-1a that has diagnostic tools Mount the EBS volumes from the terminated instance for investigation
- D. Create an AWS WAF web ACL that denies traffic to and from the suspicious instance Attach the AWS WAF web ACL to the instance to mitigate the attack Log in to the instance and install diagnostic tools to investigate the instance

Answer: B

Explanation:

This option suggests updating the outbound network ACL for the subnet in us-east-1b to explicitly deny all connections as the first rule, replacing the security group with a new one that only allows connections from a diagnostics security group, and launching a new EC2 instance with diagnostic tools to investigate the suspicious instance. This option will immediately mitigate the attack and provide the necessary tools for investigation.

NEW QUESTION 52

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

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

Answer: CE

NEW QUESTION 53

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

What should the security engineer recommend?

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

Answer: C

Explanation:

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

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

NEW QUESTION 57

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 role
- 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 action
- 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 60

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

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

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

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

Answer: BC

Explanation:

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

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

NEW QUESTION 62

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 67

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

The company is performing control tests on specific GuardDuty findings to make sure that the company's security team can detect and respond to security events.

The security team launched an Amazon EC2 instance and attempted to run DNS requests against a test domain, example.com, to generate a DNS finding.

However, the GuardDuty finding was never created in the Security Hub delegated administrator account.

Why was the finding was not created in the Security Hub delegated administrator account?

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

Answer: C

Explanation:

The correct answer is C. The GuardDuty integration with Security Hub was never activated in the AWS account where the finding was generated.

According to the AWS documentation¹, GuardDuty findings are automatically sent to Security Hub only if the GuardDuty integration with Security Hub is enabled in the same account and Region. This means that the security tooling account, which is the delegated administrator for both GuardDuty and Security Hub, must

enable the GuardDuty integration with Security Hub in each member account and Region where GuardDuty is enabled. Otherwise, the findings from GuardDuty will not be visible in Security Hub.

The other options are incorrect because:

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

References:

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

NEW QUESTION 72

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

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

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

Answer: DE

Explanation:

these are the recommended actions to take when you receive an abuse notice from AWS8. You should review the abuse notice to see what content or activity was reported and detach the internet gateway from the VPC to isolate the affected instances from the internet. You should also remove any rules that allow inbound traffic from 0.0.0.0/0 from the security groups and create a network access control list (NACL) rule to deny all traffic inbound from the internet. You should then delete the compromised instances and any associated resources that you did not create. The other options are either inappropriate or unnecessary for responding to the abuse notice.

NEW QUESTION 77

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 Mirrorin
- 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 bucke
- 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 bucke
- 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 78

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 82

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

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

Which solution will meet these requirements?

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

Answer: C

Explanation:

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

NEW QUESTION 87

A company has thousands of AWS Lambda functions. While reviewing the Lambda functions, a security engineer discovers that sensitive information is being stored in environment variables and is viewable as plaintext in the Lambda console. The values of the sensitive information are only a few characters long.

What is the MOST cost-effective way to address this security issue?

- A. Set up IAM policies from the Lambda console to hide access to the environment variables.
- B. Use AWS Step Functions to store the environment variable
- C. Access the environment variables at runtime
- D. Use IAM permissions to restrict access to the environment variables to only the Lambda functions that require access.
- E. Store the environment variables in AWS Secrets Manager, and access them at runtime
- F. Use IAM permissions to restrict access to the secrets to only the Lambda functions that require access.
- G. Store the environment variables in AWS Systems Manager Parameter Store as secure string parameters, and access them at runtime
- H. Use IAM permissions to restrict access to the parameters to only the Lambda functions that require access.

Answer: D

Explanation:

Storing sensitive information in environment variables is not a secure practice, as anyone who has access to the Lambda console or the Lambda function code can view them as plaintext. To address this security issue, the security engineer needs to use a service that can store and encrypt the environment variables, and access them at runtime using IAM permissions. The most cost-effective way to do this is to use AWS Systems Manager Parameter Store, which is a service that provides secure, hierarchical storage for configuration data management and secrets management. Parameter Store allows you to store values as standard parameters (plaintext) or secure string parameters (encrypted). Secure string parameters use a AWS Key Management Service (AWS KMS) customer master key (CMK) to encrypt the parameter value. To access the parameter value at runtime, the Lambda function needs to have IAM permissions to decrypt the parameter using the KMS CMK.

The other options are incorrect because:

- Option A is incorrect because setting up IAM policies from the Lambda console to hide access to the environment variables will not prevent someone who has access to the Lambda function code from viewing them as plaintext. IAM policies can only control who can perform actions on AWS resources, not what they can see in the code or the console.
- Option B is incorrect because using AWS Step Functions to store the environment variables is not a secure or cost-effective solution. AWS Step Functions is a service that lets you coordinate multiple AWS services into serverless workflows. Step Functions does not provide any encryption or secrets management capabilities, and it will incur additional charges for each state transition in the workflow. Moreover, storing environment variables in Step Functions will make them visible in the execution history of the workflow, which can be accessed by anyone who has permission to view the Step Functions console or API.
- Option C is incorrect because storing the environment variables in AWS Secrets Manager and accessing them at runtime is not a cost-effective solution. AWS Secrets Manager is a service that helps you protect secrets needed to access your applications, services, and IT resources. Secrets Manager enables you to rotate, manage, and retrieve secrets throughout their lifecycle. While Secrets Manager can securely store and encrypt environment variables using KMS CMKs, it will incur higher charges than Parameter Store for storing and retrieving secrets. Unless the security engineer needs the advanced features of Secrets Manager, such as automatic rotation of secrets or integration with other AWS services, Parameter Store is a cheaper and simpler option.

NEW QUESTION 89

A company needs a forensic-logging solution for hundreds of applications running in Docker on Amazon EC2. The solution must perform real-time analytics on the logs, must support the replay of messages, and must persist the logs.

Which IAM services should be used to meet these requirements? (Select TWO)

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

Answer: BD

Explanation:

Amazon Kinesis and Amazon Elasticsearch are both suitable for forensic-logging solutions. Amazon Kinesis can collect, process, and analyze streaming data in real time. Amazon Elasticsearch can store, search, and analyze log data using the popular open-source tool Elasticsearch. The other options are not designed for forensic-logging purposes. Amazon Athena is a query service that can analyze data in S3, Amazon SQS is a message queue service that can decouple and scale microservices, and Amazon EMR is a big data platform that can run Apache Spark and Hadoop clusters.

NEW QUESTION 91

A security engineer needs to run an AWS CloudFormation script. The CloudFormation script builds AWS infrastructure to support a stack that includes web servers and a MySQL database. The stack has been deployed in pre-production environments and is ready for production.

The production script must comply with the principle of least privilege. Additionally, separation of duties must exist between the security engineer's IAM account and CloudFormation.

Which solution will meet these requirements?

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

Answer: A

Explanation:

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

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

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

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

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

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

➤ Run the CloudFormation script in the production environment using the new role.

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

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

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

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

NEW QUESTION 96

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

How should the Security team accomplish this?

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

Answer: C

NEW QUESTION 97

A company has a guideline that mandates the encryption of all Amazon S3 bucket data in transit. A security engineer must implement an S3 bucket policy that denies any S3 operations if data is not encrypted.

Which S3 bucket policy will meet this requirement?

A. {
 "Version": "2012-10-17",
 "Statement": [{
 "Sid": "AllowSSLRequestOnly",
 "Action": "s3:*",
 "Effect": "Deny",
 "Resource": [
 "arn:aws:s3:::DOC-EXAMPLE-BUCKET",
 "arn:aws:s3:::DOC-EXAMPLE-BUCKET/*"
],
 "Condition": {
 "Bool": {
 "aws:SecureTransport": "true"
 }
 },
 "Principal": "*"
 }]
 }

B. {
 "Version": "2012-10-17",
 "Statement": [{
 "Sid": "AllowSSLRequestOnly",
 "Action": "s3:*",
 "Effect": "Deny",
 "Resource": [
 "arn:aws:s3:::DOC-EXAMPLE-BUCKET",
 "arn:aws:s3:::DOC-EXAMPLE-BUCKET/*"
],
 "Condition": {
 "Bool": {
 "aws:SecureTransport": "false"
 }
 },
 "Principal": "*"
 }]
 }

C.

```
{
  "Version": "2012-10-17",
  "Statement": [{
    "Sid": "AllowSSLRequestOnly",
    "Action": "s3:*",
    "Effect": "Deny",
    "Resource": [
      "arn:aws:s3:::DOC-EXAMPLE-BUCKET",
      "arn:aws:s3:::DOC-EXAMPLE-BUCKET/*"
    ],
    "Condition": {
      "StringNotEquals": {
        "s3:x-amz-server-side-encryption": "AES256"
      }
    }
  ]
}
```

D. A screenshot of a computer code Description automatically generated {

```
"Version": "2012-10-17",
"Statement": [{
  "Sid": "AllowSSLRequestOnly",
  "Action": "s3:*",
  "Effect": "Deny",
  "Resource": [
    "arn:aws:s3:::DOC-EXAMPLE-BUCKET",
    "arn:aws:s3:::DOC-EXAMPLE-BUCKET/*"
  ],
  "Condition": {
    "StringNotEquals": {
      "s3:x-amz-server-side-encryption": true
    }
  }
}],
"Principal": "*"
}
```

Answer: B

Explanation:

<https://aws.amazon.com/blogs/security/how-to-use-bucket-policies-and-apply-defense-in-depth-to-help-secure-y>

NEW QUESTION 99

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 featur
- 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 104

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

Answer: BDF

NEW QUESTION 107

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

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

Answer: B

Explanation:

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

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

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

Verified References:

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

NEW QUESTION 109

A security engineer logs in to the AWS Lambda console with administrator permissions. The security engineer is trying to view logs in Amazon CloudWatch for a Lambda function that is named myFunction.

When the security engineer chooses the option in the Lambda console to view logs in CloudWatch, an "error loading Log Streams" message appears.

The IAM policy for the Lambda function's execution role contains the following:

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": "logs:CreateLogGroup",
      "Resource": "arn:aws:logs:us-east-1:111111111111:*"
    },
    {
      "Effect": "Allow",
      "Action": ["logs:PutLogEvents"],
      "Resource": ["arn:aws:logs:us-east-1:111111111111:log-
group:/aws/Lambda/myFunction:*"]
    }
  ]
}
```

How should the security engineer correct the error?

- A. Move the `logs:CreateLogGroup` action to the second Allow statement.
- B. Add the `logs:PutDestination` action to the second Allow statement.
- C. Add the `logs:GetLogEvents` action to the second Allow statement.
- D. Add the `logs:CreateLogStream` action to the second Allow statement.

Answer: D

Explanation:

`CloudWatchLogsReadOnlyAccess` doesn't include `"logs:CreateLogStream"` but it includes `"logs:Get*"`
<https://docs.aws.amazon.com/AmazonCloudWatch/latest/logs/iam-identity-based-access-control-cwl.html#:~:te>

NEW QUESTION 110

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 115

A company needs to retain log data archives for several years to be compliant with regulations. The log data is no longer used but it must be retained. What is the MOST secure and cost-effective solution to meet these requirements?

- A. Archive the data to Amazon S3 and apply a restrictive bucket policy to deny the s3 DeleteObject API
- B. Archive the data to Amazon S3 Glacier and apply a Vault Lock policy
- C. Archive the data to Amazon S3 and replicate it to a second bucket in a second IAM Region. Choose the S3 Standard-Infrequent Access (S3 Standard-1A) storage class and apply a restrictive bucket policy to deny the s3 DeleteObject API
- D. Migrate the log data to a 16 TB Amazon Elastic Block Store (Amazon EBS) volume. Create a snapshot of the EBS volume

Answer: B

Explanation:

To securely and cost-effectively retain log data archives for several years, the company should do the following:

- Archive the data to Amazon S3 Glacier and apply a Vault Lock policy. This allows the company to use a low-cost storage class that is designed for long-term archival of data that is rarely accessed. It also allows the company to enforce compliance controls on their S3 Glacier vault by locking a vault access policy that cannot be changed.

NEW QUESTION 117

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

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

Answer: C

Explanation:

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

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

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

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

The other options are incorrect because:

- Option A is incorrect because creating an Amazon S3 bucket and configuring it to use S3 Object Lock in compliance mode will not prevent anyone from changing or deleting the data. S3 Object Lock is a feature that allows you to store objects using a WORM model in S3. You can apply two types of object locks: retention periods and legal holds. A retention period specifies a fixed period of time during which an object remains locked. A legal hold is an indefinite lock on an object until it is removed. However, S3 Object Lock only prevents objects from being overwritten or deleted by any user, including the root user in your AWS account. It does not prevent objects from being modified by other means, such as changing their metadata or encryption settings. Moreover, S3 Object Lock requires that you enable versioning on your bucket, which will incur additional storage costs for storing multiple versions of an object.
- Option B is incorrect because creating an Amazon S3 bucket and configuring it to use S3 Object Lock in governance mode will not prevent anyone from changing or deleting the data. S3 Object Lock in governance mode works similarly to compliance mode, except that users with specific IAM permissions can change or delete objects that are locked. This means that users who have `s3:BypassGovernanceRetention` permission can remove retention periods or legal holds.

from objects and overwrite or delete them before they expire. This option does not provide strong enforcement for compliance controls as required by the regulatory requirements.

➤ Option D is incorrect because creating an Amazon S3 bucket and using a lifecycle rule to transition the data to a vault in S3 Glacier will not prevent anyone from changing or deleting the data. Lifecycle rules are actions that Amazon S3 automatically performs on objects during their lifetime. You can use lifecycle rules to transition objects between storage classes or expire them after a certain period of time. However, lifecycle rules do not apply any compliance controls on objects or prevent them from being modified or deleted by users. Moreover, transitioning objects from S3 to S3 Glacier using lifecycle rules will incur additional charges for retrieval requests and data transfers.

NEW QUESTION 122

.....

Thank You for Trying Our Product

* 100% Pass or Money Back

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

* One year free update

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

* Trusted by Millions

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

* Shop Securely

All transactions are protected by VeriSign!

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

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