

Amazon-Web-Services

Exam Questions SCS-C01

AWS Certified Security- Specialty



NEW QUESTION 1

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

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

Which the SIMPLEST change that would address this server issue?

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

Answer: A

Explanation:

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

NEW QUESTION 2

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

The company has not monitored account activity in the past.

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

Which solution will meet these requirements?

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

Answer: C

NEW QUESTION 3

A company needs to improve its ability to identify and prevent IAM policies that grant public access or cross-account access to resources. The company has implemented AWS Organizations and has started using AWS Identity and Access Management Access Analyzer to refine overly broad access to accounts in the organization.

A security engineer must automate a response in the company's organization for any newly created policies that are overly permissive. The automation must remediate external access and must notify the company's security team.

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

- A. Create an AWS Step Functions state machine that checks the resource type in the finding and adds an explicit Deny statement in the trust policy for the IAM rol
- B. Configure the state machine to publish a notification to an Amazon SimpleNotification Service (Amazon SNS) topic.
- C. Create an AWS Batch job that forwards any resource type findings to an AWS Lambda function. Configure the Lambda function to add an explicit Deny statement in the trust policy for the IAM rol
- D. Configure the AWS Batch job to publish a notification to an Amazon Simple Notification Service (Amazon SNS) topic.
- E. In Amazon EventBridge, create an event rule that matches active IAM Access Analyzer findings and invokes AWS Step Functions for resolution.
- F. In Amazon CloudWatch, create a metric filter that matches active IAM Access Analyzer findings and invokes AWS Batch for resolution.
- G. Create an Amazon Simple Queue Service (Amazon SQS) queu
- H. Configure the queue to forward a notification to the security team that an external principal has been granted access to the specific IAM role and has been blocked.
- I. Create an Amazon Simple Notification Service (Amazon SNS) topic for external or cross-account access notice
- J. Subscribe the security team's email addresses to the topic.

Answer: ACF

Explanation:

The correct answer is A, C, and F.

To automate a response for any newly created policies that are overly permissive, the security engineer needs to use a combination of services that can monitor, analyze, remediate, and notify the security incidents.

Option A is correct because creating an AWS Step Functions state machine that checks the resource type in the finding and adds an explicit Deny statement in the trust policy for the IAM role is a valid way to remediate external access. AWS Step Functions is a service that allows you to coordinate multiple AWS services into serverless workflows. You can use Step Functions to invoke AWS Lambda functions, which can modify the IAM policies programmatically. You can also use Step Functions to publish a notification to an Amazon SNS topic, which can send messages to subscribers such as email addresses.

Option B is incorrect because creating an AWS Batch job that forwards any resource type findings to an AWS Lambda function is not a suitable way to automate a response. AWS Batch is a service that enables you to run batch computing workloads on AWS. Batch is designed for large-scale and long-running jobs that can benefit from parallelization and dynamic provisioning of compute resources. Batch is not intended for event-driven and real-time workflows that require immediate response.

Option C is correct because creating an Amazon EventBridge event rule that matches active IAM Access Analyzer findings and invokes AWS Step Functions for

resolution is a valid way to monitor and analyze the security incidents. Amazon EventBridge is a serverless event bus service that allows you to connect your applications with data from various sources. EventBridge can use rules to match events and route them to targets for processing. You can use EventBridge to invoke AWS Step Functions state machines from the IAM Access Analyzer findings.

Option D is incorrect because creating an Amazon CloudWatch metric filter that matches active IAM Access Analyzer findings and invokes AWS Batch for resolution is not a suitable way to monitor and analyze the security incidents. Amazon CloudWatch is a service that provides monitoring and observability for your AWS resources and applications. CloudWatch can collect metrics, logs, and events from various sources and perform actions based on alarms or filters. However, CloudWatch cannot directly invoke AWS Batch jobs from the IAM Access Analyzer findings. You would need to use another service such as EventBridge or SNS to trigger the Batch job.

Option E is incorrect because creating an Amazon SQS queue that forwards a notification to the security team that an external principal has been granted access to the specific IAM role and has been blocked is not a valid way to notify the security incidents. Amazon SQS is a fully managed message queue service that enables you to decouple and scale microservices, distributed systems, and serverless applications. SQS can deliver messages to consumers that poll the queue for messages. However, SQS cannot directly forward a notification to the security team's email addresses. You would need to use another service such as SNS or SES to send email notifications.

Option F is correct because creating an Amazon SNS topic for external or cross-account access notices and subscribing the security team's email addresses to the topic is a valid way to notify the security incidents. Amazon SNS is a fully managed messaging service that enables you to decouple and scale microservices, distributed systems, and serverless applications. SNS can deliver messages to a variety of endpoints, such as email, SMS, or HTTP. You can use SNS to send email notifications to the security team when a critical security finding is detected.

References:

- > AWS Step Functions
- > AWS Batch
- > Amazon EventBridge
- > Amazon CloudWatch
- > Amazon SQS
- > Amazon SNS

NEW QUESTION 4

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

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

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

Answer: B

Explanation:

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

NEW QUESTION 5

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

Which solution will meet these requirements?

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

Answer: D

Explanation:

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

- Migrate the website to Amazon S3, which is a highly scalable and durable object storage service that can host static websites. To do this, create an S3 bucket with the same name as the domain name of the website, enable static website hosting for the bucket, upload the website files to the bucket, and configure the bucket policy to allow public read access to the objects. For more information, see [Hosting a static website on Amazon S3](#).
 - Import a public SSL certificate that is created by AWS Certificate Manager (ACM) to Amazon CloudFront, which is a global content delivery network (CDN) service that can improve the performance and security of web applications. To do this, request or import a public SSL certificate for the domain name of the website using ACM, create a CloudFront distribution with the S3 bucket as the origin, and associate the SSL certificate with the distribution. For more information, see [Using alternate domain names and HTTPS](#).
 - Configure CloudFront to block traffic from outside the US, which is one of the requirements. To do this, create a CloudFront web ACL using AWS WAF, which is a web application firewall service that lets you control access to your web applications. In the web ACL, create a rule that uses a geo match condition to block requests that originate from countries other than the US. Associate the web ACL with the CloudFront distribution. For more information, see [How AWS WAF works with Amazon CloudFront features](#).
 - Migrate DNS to Amazon Route 53, which is a highly available and scalable cloud DNS service that can route traffic to various AWS services. To do this, register or transfer your domain name to Route 53, create a hosted zone for your domain name, and create an alias record that points your domain name to your CloudFront distribution. For more information, see [Routing traffic to an Amazon CloudFront web distribution by using your domain name](#).
- The other options are incorrect because they either do not implement SSL/TLS encryption for the website (A), do not use managed services whenever possible (B), or do not block IP addresses from outside the US (C). **Verified References:**
- <https://docs.aws.amazon.com/AmazonS3/latest/userguide/HostingWebsiteOnS3Setup.html>
 - <https://docs.aws.amazon.com/waf/latest/developerguide/waf-cloudfront.html>
 - <https://docs.aws.amazon.com/Route53/latest/DeveloperGuide/routing-to-cloudfront-distribution.html>

NEW QUESTION 7

A development team is attempting to encrypt and decode a secure string parameter from the IAM Systems Manager Parameter Store using an IAM Key Management Service (IAM KMS) CMK. However, each attempt results in an error message being sent to the development team. Which CMK-related problems possibly account for the error? (Select two.)

- A. The CMK is used in the attempt does not exist.
- B. The CMK is used in the attempt needs to be rotated.
- C. The CMK is used in the attempt is using the CMK's key ID instead of the CMK ARN.
- D. The CMK is used in the attempt is not enabled.
- E. The CMK is used in the attempt is using an alias.

Answer: AD

Explanation:

<https://docs.IAM.amazonaws.com/kms/latest/developerguide/services-parameter-store.html#parameter-store-cmk-fa>

NEW QUESTION 8

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

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

Answer: ABC

Explanation:

The key length for an RSA certificate that you use with CloudFront is 2048 bits, even though ACM supports larger keys. If you use an imported certificate with CloudFront, your key length must be 1024 or 2048 bits and cannot exceed 2048 bits. You must import the certificate in the US East (N. Virginia) Region. You must have permission to use and import the SSL/TLS certificate

<https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/cnames-and-https-requirements.html>

NEW QUESTION 9

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

- A. Use AWS IAM Access Analyzer to analyze the policies
- B. View the findings from policy validation checks.
- C. Review AWS Trusted Advisor checks for all accounts in the organization.

- D. Set up AWS Audit Manage
- E. Run an assessment for all AWS Regions for all accounts.
- F. Ensure that Amazon Inspector agents are installed on all Amazon EC2 in-stances in all accounts.

Answer: A

NEW QUESTION 10

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

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

Which solution will meet these requirements?

- A. Create an Amazon CloudWatch Logs log grou
- B. Configure the load balancers to send logs to the log grou
- C. Use the CloudWatch Logs console to search the log
- D. Create CloudWatch Logs filters on the logs for the required met-rics.
- E. Create an Amazon S3 bucke
- F. Configure the load balancers to send logs to the S3 bucke
- G. Use Amazon Athena to search the logs that are in the S3 bucke
- H. Create Amazon CloudWatch filters on the S3 log files for the re-quired metrics.
- I. Create an Amazon S3 bucke
- J. Configure the load balancers to send logs to the S3 bucke
- K. Use Amazon Athena to search the logs that are in the S3 bucke
- L. Create Athena queries for the required metric
- M. Publish the metrics to Amazon CloudWatch.
- N. Create an Amazon CloudWatch Logs log grou
- O. Configure the load balancers to send logs to the log grou
- 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 10

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

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

Answer: C

Explanation:

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

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

Security Status 3 out of 5 complete.

- ⚠ Delete your root access keys ▼
- ⚠ Activate MFA on your root account ▼
- ✅ Create individual IAM users ▼
- ✅ Use groups to assign permissions ▼
- ✅ Apply an IAM password policy ▼

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

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

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

NEW QUESTION 13

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 14

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

Please select:

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

Answer: C

Explanation:

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

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

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

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

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

NEW QUESTION 19

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

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

The security team has enabled CloudTrail for AWS management events with a basic configuration in the AWS Region in which EventBridge is being tested.

Verification of the EventBridge event pattern indicates that the pattern is set up correctly. The security team must implement a solution so that the s3:PutObjectAcl API call will invoke an EventBridge event. The solution must not generate false notifications.

Which solution will meet these requirements?

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

Answer: D

Explanation:

The correct answer is D. Enable CloudTrail to monitor data events for read and write operations to S3 buckets. According to the AWS documentation¹, CloudTrail data events are the resource operations performed on or within a resource. These are also known as data plane operations. Data events are often high-volume activities. For example, Amazon S3 object-level API activity (such as GetObject, DeleteObject, and PutObject) is a data event.

By default, trails do not log data events. To record CloudTrail data events, you must explicitly add the supported resources or resource types for which you want to collect activity. For more information, see Logging data events in the Amazon S3 User Guide².

In this case, the security team wants EventBridge to watch for the s3:PutObjectAcl API invocation logs from CloudTrail. This API uses the acl subresource to set the access control list (ACL) permissions for a new or existing object in an S3 bucket³. This is a data event that affects the S3 object resource type. Therefore, the security team must enable CloudTrail to monitor data events for read and write operations to S3 buckets in order to invoke an EventBridge event for this API call. The other options are incorrect because:

- > A. Modifying the EventBridge event pattern by selecting Amazon S3 and All Events as the event type will not capture the s3:PutObjectAcl API call, because this is a data event and not a management event. Management events provide information about management operations that are performed on resources in your AWS account. These are also known as control plane operations⁴.
- > B. Modifying the EventBridge event pattern by selecting Amazon S3 and Bucket Level Operations as the event type will not capture the s3:PutObjectAcl API call, because this is a data event that affects the S3 object resource type and not the S3 bucket resource type. Bucket level operations are management events that affect the configuration or metadata of an S3 bucket⁵.
- > C. Enabling CloudTrail Insights to identify unusual API activity will not help the security team monitor new S3 objects or changes to any S3 bucket policy or setting that result in public access. CloudTrail Insights helps AWS users identify and respond to unusual activity associated with API calls and API error rates by continuously analyzing CloudTrail management events⁶. It does not analyze data events or generate EventBridge events.

References:

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

NEW QUESTION 23

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

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

Answer: C

Explanation:

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

NEW QUESTION 26

To meet regulatory requirements, a Security Engineer needs to implement an IAM policy that restricts the use of AWS services to the us-east-1 Region. What policy should the Engineer implement?

A.

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": "*",
      "Resource": "*",
      "Condition": {
        "StringEquals": {
          "aws:RequestedRegion": "us-east-1"
        }
      }
    }
  ]
}
```

B. A computer code with black text Description automatically generated

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": "*",
      "Resource": "*",
      "Condition": {
        "StringEquals": {
          "ec2:Region": "us-east-1"
        }
      }
    }
  ]
}
```

C. A computer code with black text Description automatically generated

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Deny",
      "Action": "*",
      "Resource": "*",
      "Condition": {
        "StringNotEquals": {
          "aws:RequestedRegion": "us-east-1"
        }
      }
    }
  ]
}
```

D. A computer code with text Description automatically generated

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Deny",
      "NotAction": "*",
      "Resource": "*",
      "Condition": {
        "StringEquals": {
          "aws:RequestedRegion": "us-east-1"
        }
      }
    }
  ]
}
```

Answer: C

Explanation:

https://docs.aws.amazon.com/IAM/latest/UserGuide/reference_policies_examples_aws_deny-requested-region.h

NEW QUESTION 28

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 29

A company uses Amazon GuardDuty. The company's security team wants all High severity findings to automatically generate a ticket in a third-party ticketing

system through email integration.
 Which solution will meet this requirement?

- A. Create a verified identity for the third-party ticketing email system in Amazon Simple Email Service (Amazon SES). Create an Amazon EventBridge rule that includes an event pattern that matches High severity GuardDuty finding
- B. Specify the SES identity as the target for the EventBridge rule.
- C. Create an Amazon Simple Notification Service (Amazon SNS) topic
- D. Subscribe the third-party ticketing email system to the SNS topic
- E. Create an Amazon EventBridge rule that includes an event pattern that matches High severity GuardDuty finding
- F. Specify the SNS topic as the target for the EventBridge rule.
- G. Use the GuardDuty CreateFilter API operation to build a filter in GuardDuty to monitor for High severity finding
- H. Export the results of the filter to an Amazon Simple Notification Service (Amazon SNS) topic
- I. Subscribe the third-party ticketing email system to the SNS topic.
- J. Use the GuardDuty CreateFilter API operation to build a filter in GuardDuty to monitor for High severity finding
- K. Create an Amazon Simple Notification Service (Amazon SNS) topic
- L. Subscribe the third-party ticketing email system to the SNS topic
- M. Create an Amazon EventBridge rule that includes an event pattern that matches GuardDuty findings that are selected by the filter
- N. Specify the SNS topic as the target for the EventBridge rule.

Answer: B

Explanation:

The correct answer is B. Create an Amazon Simple Notification Service (Amazon SNS) topic. Subscribe the third-party ticketing email system to the SNS topic. Create an Amazon EventBridge rule that includes an event pattern that matches High severity GuardDuty findings. Specify the SNS topic as the target for the EventBridge rule.

According to the AWS documentation¹, you can use Amazon EventBridge to create rules that match events from GuardDuty and route them to targets such as Amazon SNS topics. You can use event patterns to filter events based on criteria such as severity, type, or resource. For example, you can create a rule that matches only High severity findings and sends them to an SNS topic that is subscribed by a third-party ticketing email system. This way, you can automate the creation of tickets for High severity findings and notify the security team.

NEW QUESTION 32

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

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

Answer: A

Explanation:

The IAM Documentation mentions the following

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

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

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

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

NEW QUESTION 33

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 36

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

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

A)

Apply the following SCP:

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

B)

Apply the following SCP:

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

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

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

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

Answer: ACE

NEW QUESTION 37

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

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

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

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

Answer: B

NEW QUESTION 40

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 43

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 database
- 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 48

A company is running an Amazon RDS for MySQL DB instance in a VPC. The VPC must not send or receive network traffic through the internet. A security engineer wants to use AWS Secrets Manager to rotate the DB instance credentials automatically. Because of a security policy, the security engineer cannot use the standard AWS Lambda function that Secrets Manager provides to rotate the credentials. The security engineer deploys a custom Lambda function in the VPC. The custom Lambda function will be responsible for rotating the secret in Secrets Manager. The security engineer edits the DB instance's security group to allow connections from this function. When the function is invoked, the function cannot communicate with Secrets Manager to rotate the secret properly. What should the security engineer do so that the function can rotate the secret?

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

Answer: D

Explanation:

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

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

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

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

NEW QUESTION 50

During a manual review of system logs from an Amazon Linux EC2 instance, a Security Engineer noticed that there are sudo commands that were never properly

alerted or reported on the Amazon CloudWatch Logs agent
Why were there no alerts on the sudo commands?

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

Answer: B

Explanation:

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

NEW QUESTION 53

A company uses AWS Organizations. The company wants to implement short-term credentials for third-party AWS accounts to use to access accounts within the company's organization. Access is for the AWS Management Console and third-party software-as-a-service (SaaS) applications. Trust must be enhanced to prevent two external accounts from using the same credentials. The solution must require the least possible operational effort.
Which solution will meet these requirements?

- A. Use a bearer token authentication with OAuth or SAML to manage and share a central Amazon Cognito user pool across multiple Amazon API Gateway APIs.
- B. Implement AWS IAM Identity Center (AWS Single Sign-On), and use an identity source of choice. Grant access to users and groups from other accounts by using permission sets that are assigned by account.
- C. Create a unique IAM role for each external account
- D. Create a trust policy
- E. Use AWS Secrets Manager to create a random external key.
- F. Create a unique IAM role for each external account
- G. Create a trust policy that includes a condition that uses the sts:ExternalId condition key.

Answer: D

Explanation:

The correct answer is D.

To implement short-term credentials for third-party AWS accounts, you can use IAM roles and trust policies. A trust policy is a JSON policy document that defines who can assume the role. You can specify the AWS account ID of the third-party account as a principal in the trust policy, and use the sts:ExternalId condition key to enhance the security of the role. The sts:ExternalId condition key is a unique identifier that is agreed upon by both parties and included in the AssumeRole request. This way, you can prevent the "confused deputy" problem, where an unauthorized party can use the same role as a legitimate party.

Option A is incorrect because bearer token authentication with OAuth or SAML is not suitable for granting access to AWS accounts and resources. Amazon Cognito and API Gateway are used for building web and mobile applications that require user authentication and authorization.

Option B is incorrect because AWS IAM Identity Center (AWS Single Sign-On) is a service that simplifies the management of access to multiple AWS accounts and cloud applications for your workforce users. It does not support granting access to third-party AWS accounts.

Option C is incorrect because using AWS Secrets Manager to create a random external key is not necessary and adds operational complexity. You can use the sts:ExternalId condition key instead to provide a unique identifier for each external account.

NEW QUESTION 57

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 58

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 and 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 61

A company is building a data processing application that uses AWS Lambda functions. The application's Lambda functions need to communicate with an Amazon RDS DB instance that is deployed within a VPC in the same AWS account. Which solution meets these requirements in the MOST secure way?

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

Answer: C

Explanation:

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

References: : AWS Lambda Developer Guide

NEW QUESTION 63

A company has a web server in the AWS Cloud. The company will store the content for the web server in an Amazon S3 bucket. A security engineer must use an Amazon CloudFront distribution to speed up delivery of the content. None of the files can be publicly accessible from the S3 bucket directly. Which solution will meet these requirements?

- A. Configure the permissions on the individual files in the S3 bucket so that only the CloudFront distribution has access to them.
- B. Create an origin access identity (OAI). Associate the OAI with the CloudFront distribution.
- C. Configure the S3 bucket permissions so that only the OAI can access the files in the S3 bucket.
- D. Create an S3 role in AWS Identity and Access Management (IAM). Allow only the CloudFront distribution to assume the role to access the files in the S3 bucket.
- E. Create an S3 bucket policy that uses only the CloudFront distribution ID as the principal and the Amazon Resource Name (ARN) as the target.

Answer: B

NEW QUESTION 66

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

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

Answer: C

Explanation:

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

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

This solution provides the following security benefits:

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

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

NEW QUESTION 68

A company's IAM account consists of approximately 300 IAM users. Now there is a mandate that an access change is required for 100 IAM users to have unlimited privileges to S3. As a system administrator, how can you implement this effectively so that there is no need to apply the policy at the individual user level? Please select:

- A. Create a new role and add each user to the IAM role
- B. Use the IAM groups and add users, based upon their role, to different groups and apply the policy to group
- C. Create a policy and apply it to multiple users using a JSON script
- D. Create an S3 bucket policy with unlimited access which includes each user's IAM account ID

Answer: B

Explanation:

Option A is incorrect since you don't add a user to the IAM Role Option C is incorrect since you don't assign multiple users to a policy Option D is incorrect since this is not an ideal approach

An IAM group is used to collectively manage users who need the same set of permissions. By having groups, it becomes easier to manage permissions. So if you change the permissions on the group scale, it will affect all the users in that group

For more information on IAM Groups, just browse to the below URL: https://docs.IAM.amazon.com/IAM/latest/UserGuide/id_groups.html

The correct answer is: Use the IAM groups and add users, based upon their role, to different groups and apply the policy to group

Submit your Feedback/Queries to our Experts

NEW QUESTION 71

A company is running its workloads in a single AWS Region and uses AWS Organizations. A security engineer must implement a solution to prevent users from launching resources in other Regions.

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

- A. Create an IAM policy that has an aws RequestedRegion condition that allows actions only in the designated Region Attach the policy to all users.
- B. Create an IAM policy that has an aws RequestedRegion condition that denies actions that are not in the designated Region Attach the policy to the AWS account in AWS Organizations.
- C. Create an IAM policy that has an aws RequestedRegion condition that allows the desired actions Attach the policy only to the users who are in the designated Region.
- D. Create an SCP that has an aws RequestedRegion condition that denies actions that are not in the designated Region
- E. Attach the SCP to the AWS account in AWS Organizations.

Answer: D

Explanation:

Although you can use a IAM policy to prevent users launching resources in other regions. The best practice is to use SCP when using AWS organizations.

https://docs.aws.amazon.com/organizations/latest/userguide/orgs_manage_policies_scps_examples_general.htm

NEW QUESTION 74

A company's security engineer is designing an isolation procedure for Amazon EC2 instances as part of an incident response plan. The security engineer needs to isolate a target instance to block any traffic to and from the target instance, except for traffic from the company's forensics team. Each of the company's EC2 instances has its own dedicated security group. The EC2 instances are deployed in subnets of a VPC. A subnet can contain multiple instances.

The security engineer is testing the procedure for EC2 isolation and opens an SSH session to the target instance. The procedure starts to simulate access to the target instance by an attacker. The security engineer removes the existing security group rules and adds security group rules to give the forensics team access to the target instance on port 22.

After these changes, the security engineer notices that the SSH connection is still active and usable. When the security engineer runs a ping command to the public IP address of the target instance, the ping command is blocked.

What should the security engineer do to isolate the target instance?

- A. Add an inbound rule to the security group to allow traffic from 0.0.0.0/0 for all port
- B. Add an outbound rule to the security group to allow traffic to 0.0.0.0/0 for all port
- C. Then immediately delete these rules.
- D. Remove the port 22 security group rule
- E. Attach an instance role policy that allows AWS Systems Manager Session Manager connections so that the forensics team can access the target instance.
- F. Create a network ACL that is associated with the target instance's subnet
- G. Add a rule at the top of the inbound rule set to deny all traffic from 0.0.0.0/0. Add a rule at the top of the outbound rule set to deny all traffic to 0.0.0.0/0.
- H. Create an AWS Systems Manager document that adds a host-level firewall rule to block all inbound traffic and outbound traffic
- I. Run the document on the target instance.

Answer: C

NEW QUESTION 78

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

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

Which approach should the Security Engineer use?

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

Answer: A

Explanation:

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

NEW QUESTION 82

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 86

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

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

Answer: C

Explanation:

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

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

NEW QUESTION 91

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

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

Answer: D

Explanation:

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

NEW QUESTION 95

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

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

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

Answer: B

Explanation:

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

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

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

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

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

Verified References:

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

NEW QUESTION 96

A company has deployed Amazon GuardDuty and now wants to implement automation for potential threats. The company has decided to start with RDP brute force attacks that come from Amazon EC2 instances in the company's AWS environment. A security engineer needs to implement a solution that blocks the detected communication from a suspicious instance until investigation and potential remediation can occur.

Which solution will meet these requirements?

- A. Configure GuardDuty to send the event to an Amazon Kinesis data stream
- B. Process the event with an Amazon Kinesis Data Analytics for Apache Flink application that sends a notification to the company through Amazon Simple Notification Service (Amazon SNS). Add rules to the network ACL to block traffic to and from the suspicious instance.
- C. Configure GuardDuty to send the event to Amazon EventBridge (Amazon CloudWatch Events). Deploy an AWS WAF web ACL
- D. Process the event with an AWS Lambda function that sends a notification to the company through Amazon Simple Notification Service (Amazon SNS) and adds a web ACL rule to block traffic to and from the suspicious instance.
- E. Enable AWS Security Hub to ingest GuardDuty findings and send the event to Amazon EventBridge (Amazon CloudWatch Events). Deploy AWS Network Firewall
- F. Process the event with an AWS Lambda function that adds a rule to a Network Firewall firewall policy to block traffic to and from the suspicious instance.
- G. Enable AWS Security Hub to ingest GuardDuty finding
- H. Configure an Amazon Kinesis data stream as an event destination for Security Hub
- I. Process the event with an AWS Lambda function that replaces the security group of the suspicious instance with a security group that does not allow any connections.

Answer: C

Explanation:

<https://aws.amazon.com/blogs/security/automatically-block-suspicious-traffic-with-aws-network-firewall-and-a>

NEW QUESTION 98

A company is operating a website using Amazon CloudFront. CloudFront servers some content from Amazon S3 and other from web servers running EC2 instances behind an Application Load Balancer (ALB). Amazon DynamoDB is used as the data store. The company already uses IAM Certificate Manager (ACM) to store a public TLS certificate that can optionally secure connections between the website users and CloudFront. The company has a new requirement to enforce end-to-end encryption in transit.

Which combination of steps should the company take to meet this requirement? (Select THREE.)

- A. Update the CloudFront distribution
- B. configuring it to optionally use HTTPS when connecting to origins on Amazon S3
- C. Update the web application configuration on the web servers to use HTTPS instead of HTTP when connecting to DynamoDB
- D. Update the CloudFront distribution to redirect HTTP connections to HTTPS
- E. Configure the web servers on the EC2 instances to listen using HTTPS using the public ACM TLS certificate Update the ALB to connect to the target group using HTTPS
- F. Update the ALB listen to listen using HTTPS using the public ACM TLS certificate
- G. Update the CloudFront distribution to connect to the HTTPS listener.
- H. Create a TLS certificate Configure the web servers on the EC2 instances to use HTTPS only with that certificate
- I. Update the ALB to connect to the target group using HTTPS.

Answer: BCE

Explanation:

To enforce end-to-end encryption in transit, the company should do the following:

- Update the web application configuration on the web servers to use HTTPS instead of HTTP when connecting to DynamoDB. This ensures that the data is encrypted when it travels from the web servers to the data store.
- Update the CloudFront distribution to redirect HTTP requests to HTTPS. This ensures that the viewers always use HTTPS when they access the website through CloudFront.
- Update the ALB to listen using HTTPS using the public ACM TLS certificate. Update the CloudFront distribution to connect to the HTTPS listener. This ensures that the data is encrypted when it travels from CloudFront to the ALB and from the ALB to the web servers.

NEW QUESTION 103

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

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

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

Answer: D

Explanation:

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

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

NEW QUESTION 108

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

Which solution meets these requirements?

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

Answer: A

NEW QUESTION 113

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

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

What should the security engineer do next?

- A. Place the network interface in promiscuous mode to capture the traffic.
- B. Configure VPC Flow Logs to send traffic to the monitoring EC2 instance using a Network Load Balancer.
- C. Configure VPC traffic mirroring to send traffic to the monitoring EC2 instance using a Network Load Balancer.
- D. Use Amazon Inspector to detect network-level attacks and trigger an IAM Lambda function to send the suspicious packets to the EC2 instance.

Answer: D

NEW QUESTION 115

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

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

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

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

Answer: C

NEW QUESTION 120

Your development team is using access keys to develop an application that has access to S3 and DynamoDB. A new security policy has outlined that the credentials should not be older than 2 months, and should be rotated. How can you achieve this? Please select:

- A. Use the application to rotate the keys in every 2 months via the SDK
- B. Use a script to query the creation date of the key
- C. If older than 2 months, create new access key and update all applications to use it inactivate the old key and delete it.
- D. Delete the user associated with the keys after every 2 month
- E. Then recreate the user again.
- F. Delete the IAM Role associated with the keys after every 2 month
- G. Then recreate the IAM Role again.

Answer: B

Explanation:

One can use the CLI command list-access-keys to get the access keys. This command also returns the "CreateDate" of the keys. If the CreateDate is older than 2 months, then the keys can be deleted.

The Returns list-access-keys CLI command returns information about the access key IDs associated with the specified IAM user. If there are none, the action returns an empty list

Option A is incorrect because you might as use a script for such maintenance activities Option C is incorrect because you would not rotate the users themselves

Option D is incorrect because you don't use IAM roles for such a purpose For more information on the CLI command, please refer to the below Link:

<http://docs.IAM.amazon.com/cli/latest/reference/iam/list-access-keys.html>

The correct answer is: Use a script to query the creation date of the keys. If older than 2 months, create new access key and update all applications to use it inactivate the old key and delete it.

Submit your Feedback/Queries to our Experts

NEW QUESTION 124

A security engineer recently rotated the host keys for an Amazon EC2 instance. The security engineer is trying to access the EC2 instance by using the EC2 Instance Connect feature. However, the security engineer receives an error (or failed host key validation. Before the rotation of the host keys EC2 Instance Connect worked correctly with this EC2 instance.

What should the security engineer do to resolve this error?

- A. Import the key material into AWS Key Management Service (AWS KMS).
- B. Manually upload the new host key to the AWS trusted host keys database.
- C. Ensure that the AmazonSSMManagedInstanceCore policy is attached to the EC2 instance profile.
- D. Create a new SSH key pair for the EC2 instance.

Answer: B

Explanation:

To set up a CloudFront distribution for an S3 bucket that hosts a static website, and to allow only specified IP addresses to access the website, the following steps are required:

➤ Create a CloudFront origin access identity (OAI), which is a special CloudFront user that you can associate with your distribution. An OAI allows you to restrict access to your S3 content by using signed URLs or signed cookies. For more information, see Using an origin access identity to restrict access to your Amazon S3 content.

➤ Create the S3 bucket policy so that only the OAI has access. This will prevent users from accessing the website directly by using S3 URLs, as they will receive an Access Denied error. To do this, use the AWS Policy Generator to create a bucket policy that grants s3:GetObject permission to the OAI, and attach it to the S3 bucket. For more information, see Restricting access to Amazon S3 content by using an origin access identity.

➤ Create an AWS WAF web ACL and add an IP set rule. AWS WAF is a web application firewall service that lets you control access to your web applications. An IP set is a condition that specifies a list of IP addresses or IP address ranges that requests originate from. You can use an IP set rule to allow or block requests based on the IP addresses of the requesters. For more information, see Working with IP match conditions.

➤ Associate the web ACL with the CloudFront distribution. This will ensure that the web ACL filters all requests for your website before they reach your origin. You can do this by using the AWS WAF console, API, or CLI. For more information, see Associating or disassociating a web ACL with a CloudFront distribution.

This solution will meet the requirements of allowing only specified IP addresses to access the website and preventing direct access by using S3 URLs. The other options are incorrect because they either do not create a CloudFront distribution for the S3 bucket (A), do not use an OAI to restrict access to the S3 bucket ©, or do not use AWS WAF to block traffic from outside the specified IP addresses (D).

Verified References:

➤ <https://docs.aws.amazon.com/waf/latest/developerguide/web-acl-ip-conditions.html>

NEW QUESTION 125

A Security Engineer is working with a Product team building a web application on AWS. The application uses Amazon S3 to host the static content, Amazon API Gateway to provide RESTful services; and Amazon DynamoDB as the backend data store. The users already exist in a directory that is exposed through a SAML identity provider.

Which combination of the following actions should the Engineer take to enable users to be authenticated into the web application and call APIs? (Choose three.)

- A. Create a custom authorization service using AWS Lambda.
- B. Configure a SAML identity provider in Amazon Cognito to map attributes to the Amazon Cognito user pool attributes.
- C. Configure the SAML identity provider to add the Amazon Cognito user pool as a relying party.
- D. Configure an Amazon Cognito identity pool to integrate with social login providers.
- E. Update DynamoDB to store the user email addresses and passwords.
- F. Update API Gateway to use a COGNITO_USER_POOLS authorizer.

Answer: BCF

Explanation:

The combination of the following actions should the Engineer take to enable users to be authenticated into the web application and call APIs are:

- B. Configure a SAML identity provider in Amazon Cognito to map attributes to the Amazon Cognito user pool attributes. This is a necessary step to federate the existing users from the SAML identity provider to the Amazon Cognito user pool, which will be used for authentication and authorization¹.
- C. Configure the SAML identity provider to add the Amazon Cognito user pool as a relying party. This is a necessary step to establish a trust relationship between the SAML identity provider and the Amazon Cognito user pool, which will allow the users to sign in using their existing credentials².
- F. Update API Gateway to use a COGNITO_USER_POOLS authorizer. This is a necessary step to enable API Gateway to use the Amazon Cognito user pool as an authorizer for the RESTful services, which will validate the identity or access tokens that are issued by Amazon Cognito when a user signs in successfully³. The other options are incorrect because:
 - A. Creating a custom authorization service using AWS Lambda is not a necessary step, because Amazon Cognito user pools can provide built-in authorization features, such as scopes and groups, that can be used to control access to API resources⁴.
 - D. Configuring an Amazon Cognito identity pool to integrate with social login providers is not a necessary step, because the users already exist in a directory that is exposed through a SAML identity provider, and there is no requirement to support social login providers⁵.
 - E. Updating DynamoDB to store the user email addresses and passwords is not a necessary step, because the user credentials are already stored in the SAML identity provider, and there is no need to duplicate them in DynamoDB⁶.

References:

1: Using Tokens with User Pools 2: Adding SAML Identity Providers to a User Pool 3: Control Access to a REST API Using Amazon Cognito User Pools as Authorizer 4: API Authorization with Resource Servers and OAuth 2.0 Scopes 5: Using Identity Pools (Federated Identities) 6: Amazon DynamoDB

NEW QUESTION 127

A company receives a notification from the AWS Abuse team about an AWS account. The notification indicates that a resource in the account is compromised. The company determines that the compromised resource is an Amazon EC2 instance that hosts a web application. The compromised EC2 instance is part of an EC2 Auto Scaling group.

The EC2 instance accesses Amazon S3 and Amazon DynamoDB resources by using an IAM access key and secret key. The IAM access key and secret key are stored inside the AMI that is specified in the Auto Scaling group's launch configuration. The company is concerned that the credentials that are stored in the AMI might also have been exposed.

The company must implement a solution that remediates the security concerns without causing downtime for the application. The solution must comply with security best practices. Which solution will meet these requirements?

- A. Rotate the potentially compromised access key that the EC2 instance uses. Create a new AMI without the potentially compromised credentials. Perform an EC2 Auto Scaling instance refresh.
- B. Delete or deactivate the potentially compromised access key. Create an EC2 Auto Scaling linked IAM role that includes a custom policy that matches the potentially compromised access key permission. Associate the new IAM role with the Auto Scaling group. Perform an EC2 Auto Scaling instance refresh.
- C. Delete or deactivate the potentially compromised access key. Create a new AMI without the potentially compromised credentials. Create an IAM role that includes the correct permissions. Create a launch template for the Auto Scaling group to reference the new AMI and IAM role. Perform an EC2 Auto Scaling instance refresh.
- D. Rotate the potentially compromised access key. Create a new AMI without the potentially compromised access key. Use a user data script to supply the new access key as environmental variables in the Auto Scaling group's launch configuration. Perform an EC2 Auto Scaling instance refresh.

Answer: C

Explanation:

The AWS documentation states that you can create a new AMI without the potentially compromised credentials and create an IAM role that includes the correct permissions. You can then create a launch template for the Auto Scaling group to reference the new AMI and IAM role. This method is the most secure way to remediate the security concerns without causing downtime for the application.

References: : AWS Security Best Practices

NEW QUESTION 128

A company hosts a web application on an Apache web server. The application runs on Amazon EC2 instances that are in an Auto Scaling group. The company configured the EC2 instances to send the Apache web server logs to an Amazon CloudWatch Logs group that the company has configured to expire after 1 year. Recently, the company discovered in the Apache web server logs that a specific IP address is sending suspicious requests to the web application. A security engineer wants to analyze the past week of Apache web server logs to determine how many requests that the IP address sent and the corresponding URLs that the IP address requested.

What should the security engineer do to meet these requirements with the LEAST effort?

- A. Export the CloudWatch Logs group data to Amazon S3. Use Amazon Macie to query the logs for the specific IP address and the requested URLs.

- B. Configure a CloudWatch Logs subscription to stream the log group to an Amazon OpenSearch Service cluster.
- C. Use OpenSearch Service to analyze the logs for the specific IP address and the requested URLs.
- D. Use CloudWatch Logs Insights and a custom query syntax to analyze the CloudWatch logs for the specific IP address and the requested URLs.
- E. Export the CloudWatch Logs group data to Amazon S3. Use AWS Glue to crawl the S3 bucket for only the log entries that contain the specific IP address.
- F. Use AWS Glue to view the results.

Answer: C

NEW QUESTION 132

A security engineer is configuring a new website that is named example.com. The security engineer wants to secure communications with the website by requiring users to connect to example.com through HTTPS.

Which of the following is a valid option for storing SSL/TLS certificates?

- A. Custom SSL certificate that is stored in AWS Key Management Service (AWS KMS)
- B. Default SSL certificate that is stored in Amazon CloudFront.
- C. Custom SSL certificate that is stored in AWS Certificate Manager (ACM)
- D. Default SSL certificate that is stored in Amazon S3

Answer: C

NEW QUESTION 133

A company is migrating one of its legacy systems from an on-premises data center to AWS. The application server will run on AWS, but the database must remain in the on-premises data center for compliance reasons. The database is sensitive to network latency. Additionally, the data that travels between the on-premises data center and AWS must have IPsec encryption.

Which combination of AWS solutions will meet these requirements? (Choose two.)

- A. AWS Site-to-Site VPN
- B. AWS Direct Connect
- C. AWS VPN CloudHub
- D. VPC peering
- E. NAT gateway

Answer: AB

Explanation:

The correct combination of AWS solutions that will meet these requirements is A. AWS Site-to-Site VPN and B. AWS Direct Connect.

* A. AWS Site-to-Site VPN is a service that allows you to securely connect your on-premises data center to your AWS VPC over the internet using IPsec encryption. This solution meets the requirement of encrypting the data in transit between the on-premises data center and AWS.

* B. AWS Direct Connect is a service that allows you to establish a dedicated network connection between your on-premises data center and your AWS VPC. This solution meets the requirement of reducing network latency between the on-premises data center and AWS.

* C. AWS VPN CloudHub is a service that allows you to connect multiple VPN connections from different locations to the same virtual private gateway in your AWS VPC. This solution is not relevant for this scenario, as there is only one on-premises data center involved.

* D. VPC peering is a service that allows you to connect two or more VPCs in the same or different regions using private IP addresses. This solution does not meet the requirement of connecting an on-premises data center to AWS, as it only works for VPCs.

* E. NAT gateway is a service that allows you to enable internet access for instances in a private subnet in your AWS VPC. This solution does not meet the requirement of connecting an on-premises data center to AWS, as it only works for outbound traffic from your VPC.

NEW QUESTION 135

A company has a group of Amazon EC2 instances in a single private subnet of a VPC with no internet gateway attached. A security engineer has installed the Amazon CloudWatch agent on all instances in that subnet to capture logs from a specific application. To ensure that the logs flow securely, the company's networking team has created VPC endpoints for CloudWatch monitoring and CloudWatch logs. The networking team has attached the endpoints to the VPC. The application is generating logs. However, when the security engineer queries CloudWatch, the logs do not appear.

Which combination of steps should the security engineer take to troubleshoot this issue? (Choose three.)

- A. Ensure that the EC2 instance profile that is attached to the EC2 instances has permissions to create log streams and write logs.
- B. Create a metric filter on the logs so that they can be viewed in the AWS Management Console.
- C. Check the CloudWatch agent configuration file on each EC2 instance to make sure that the CloudWatch agent is collecting the proper log files.
- D. Check the VPC endpoint policies of both VPC endpoints to ensure that the EC2 instances have permissions to use them.
- E. Create a NAT gateway in the subnet so that the EC2 instances can communicate with CloudWatch.
- F. Ensure that the security groups allow all the EC2 instances to communicate with each other to aggregate logs before sending.

Answer: ACD

Explanation:

The possible steps to troubleshoot this issue are:

➤ A. Ensure that the EC2 instance profile that is attached to the EC2 instances has permissions to create log streams and write logs. This is a necessary step because the CloudWatch agent uses the credentials from the instance profile to communicate with CloudWatch1.

➤ C. Check the CloudWatch agent configuration file on each EC2 instance to make sure that the CloudWatch agent is collecting the proper log files. This is a necessary step because the CloudWatch agent needs to know which log files to monitor and send to CloudWatch2.

➤ D. Check the VPC endpoint policies of both VPC endpoints to ensure that the EC2 instances have permissions to use them. This is a necessary step because the VPC endpoint policies control which principals can access the AWS services through the endpoints3.

The other options are incorrect because:

➤ B. Creating a metric filter on the logs is not a troubleshooting step, but a way to extract metric data from the logs. Metric filters do not affect the visibility of the logs in the AWS Management Console.

➤ E. Creating a NAT gateway in the subnet is not a solution, because the EC2 instances do not need internet access to communicate with CloudWatch through the VPC endpoints. A NAT gateway would also incur additional costs.

➤ F. Ensuring that the security groups allow all the EC2 instances to communicate with each other is not a necessary step, because the CloudWatch agent does not require log aggregation before sending. Each EC2 instance can send its own logs independently to CloudWatch.

References:

1: IAM Roles for Amazon EC2 2: CloudWatch Agent Configuration File: Logs Section 3: Using Amazon VPC Endpoints : Metric Filters : NAT Gateways : CloudWatch Agent Reference: Log Aggregation

NEW QUESTION 138

Your company uses IAM to host its resources. They have the following requirements

- 1) Record all API calls and Transitions
- 2) Help in understanding what resources are there in the account
- 3) Facility to allow auditing credentials and logins

Which services would suffice the above requirements Please select:

- A. IAM Inspector, CloudTrail, IAM Credential Reports
- B. CloudTrai
- C. IAM Credential Reports, IAM SNS
- D. CloudTrail, IAM Config, IAM Credential Reports
- E. IAM SQS, IAM Credential Reports, CloudTrail

Answer: C

Explanation:

You can use IAM CloudTrail to get a history of IAM API calls and related events for your account. This history includes calls made with the IAM Management Console, IAM Command Line Interface, IAM SDKs, and other IAM services.

Options A,B and D are invalid because you need to ensure that you use the services of CloudTrail, IAM Config, IAM Credential Reports

For more information on Cloudtrail, please visit the below URL:

<http://docs.IAM.amazon.com/IAMcloudtrail/latest/userguide/cloudtrail-user-guide.html>

IAM Config is a service that enables you to assess, audit and evaluate the configurations of your IAM resources. Config continuously monitors and records your IAM resource configurations and allows you to automate the evaluation of recorded configurations against desired configurations. With Config, you can review changes in configurations and relationships between IAM resources, dive into detailed resource configuration histories, and determine your overall compliance against the configurations specified in your internal guidelines. This enables you to simplify compliance auditing, security analysis, char management and operational troubleshooting.

For more information on the config service, please visit the below URL <https://IAM.amazon.com/config/>

You can generate and download a credential report that lists all users in your account and the status of their various credentials, including passwords, access keys, and MFA devices. You can get a credential report from the IAM Management Console, the IAM SDKs and Command Line Tools, or the IAM API.

For more information on Credentials Report, please visit the below URL: http://docs.IAM.amazon.com/IAM/latest/UserGuide/id_credentials_getting-report.html

The correct answer is: CloudTrail, IAM Config, IAM Credential Reports Submit your Feedback/Queries to our Experts

NEW QUESTION 142

A security engineer is trying to use Amazon EC2 Image Builder to create an image of an EC2 instance. The security engineer has configured the pipeline to send logs to an Amazon S3 bucket. When the security engineer runs the pipeline, the build fails with the following error: "AccessDenied: Access Denied status code: 403".

The security engineer must resolve the error by implementing a solution that complies with best practices for least privilege access.

Which combination of steps will meet these requirements? (Choose two.)

- A. Ensure that the following policies are attached to the IAM role that the security engineer is using: EC2InstanceProfileForImageBuilder, EC2InstanceProfileForImageBuilderECRContainerBuilds, and AmazonSSMManagedInstanceCore.
- B. Ensure that the following policies are attached to the instance profile for the EC2 instance: EC2InstanceProfileForImageBuilder, EC2InstanceProfileForImageBuilderECRContainerBuilds, and AmazonSSMManagedInstanceCore.
- C. Ensure that the AWSImageBuilderFullAccess policy is attached to the instance profile for the EC2 instance.
- D. Ensure that the security engineer's IAM role has the s3:PutObject permission for the S3 bucket.
- E. Ensure that the instance profile for the EC2 instance has the s3:PutObject permission for the S3 bucket.

Answer: BE

Explanation:

The most likely cause of the error is that the instance profile for the EC2 instance does not have the s3:PutObject permission for the S3 bucket. This permission is needed to upload logs to the bucket. Therefore, the security engineer should ensure that the instance profile has this permission.

One possible solution is to attach the AWSImageBuilderFullAccess policy to the instance profile for the EC2 instance. This policy grants full access to Image Builder resources and related AWS services, including the s3:PutObject permission for any bucket with "imagebuilder" in its name. However, this policy may grant more permissions than necessary, which violates the principle of least privilege.

Another possible solution is to create a custom policy that only grants the s3:PutObject permission for the specific S3 bucket that is used for logging. This policy can be attached to the instance profile along with the other policies that are required for Image Builder functionality: EC2InstanceProfileForImageBuilder, EC2InstanceProfileForImageBuilderECRContainerBuilds, and AmazonSSMManagedInstanceCore. This solution follows the principle of least privilege more closely than the previous one.

➤ Ensure that the following policies are attached to the instance profile for the EC2 instance: EC2InstanceProfileForImageBuilder, EC2InstanceProfileForImageBuilderECRContainerBuilds, and AmazonSSMManagedInstanceCore.

➤ Ensure that the instance profile for the EC2 instance has the s3:PutObject permission for the S3 bucket.

This can be done by either attaching the AWSImageBuilderFullAccess policy or creating a custom policy with this permission.

1: Using managed policies for EC2 Image Builder - EC2 Image Builder 2: PutObject - Amazon Simple Storage Service 3: AWSImageBuilderFullAccess - AWS Managed Policy

NEW QUESTION 145

.....

Thank You for Trying Our Product

We offer two products:

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

2nd - Questions and Answers in PDF Format

SCS-C01 Practice Exam Features:

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

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