Amazon Web Services

Exam Questions SCS-C02

AWS Certified Security - Specialty

NEW QUESTION 1
An AWS account that is used for development projects has a VPC that contains two subnets. The first subnet is named public-subnet-1 and has the CIDR block 192.168.1.0/24 assigned. The other subnet is named private-subnet-2 and has the CIDR block 192.168.2.0/24 assigned. Each subnet contains Amazon EC2 instances.
Each subnet is currently using the VPC's default network ACL. The security groups that the EC2 instances in these subnets use have rules that allow traffic between each instance where required. Currently, all network traffic flow is working as expected between the EC2 instances that are using these subnets.
A security engineer creates a new network ACL that is named subnet-2-NACL with default entries. The security engineer immediately configures private-subnet-2 to use the new network ACL and makes no other changes to the infrastructure. The security engineer starts to receive reports that the EC2 instances in public-subnet-1 and public-subnet-2 cannot communicate with each other.
Which combination of steps should the security engineer take to allow the EC2 instances that are running in these two subnets to communicate again? (Select TWO.)

A. Add an outbound allow rule for 192.168.2.0/24 in the VPC's default network ACL.
B. Add an inbound allow rule for 192.168.2.0/24 in the VPC's default network ACL.
C. Add an outbound allow rule for 192.168.2.0/24 in subnet-2-NACL.
D. Add an inbound allow rule for 192.168.1.0/24 in subnet-2-NACL.
E. Add an outbound allow rule for 192.168.1.0/24 in subnet-2-NACL.

Answer: CE

Explanation:
The AWS documentation states that you can add an outbound allow rule for 192.168.2.0/24 in
subnet-2-NACL and add an outbound allow rule for 192.168.1.0/24 in subnet-2-NACL. This will allow the EC2 instances that are running in these two subnets to communicate again.
References: : Amazon VPC User Guide

NEW QUESTION 2
A company is using Amazon Macie, AWS Firewall Manager, Amazon Inspector, and AWS Shield Advanced in its AWS account. The company wants to receive alerts if a DDoS attack occurs against the account.
Which solution will meet this requirement?

A. Use Macie to detect an active DDoS even
B. Create Amazon CloudWatch alarms that respond to Macie findings.
C. Use Amazon Inspector to review resources and to invoke Amazon CloudWatch alarms for any resources that are vulnerable to DDoS attacks.
D. Create an Amazon CloudWatch alarm that monitors Firewall Manager metrics for an active DDoS event.
E. Create an Amazon CloudWatch alarm that monitors Shield Advanced metrics for an active DDoS event.

Answer: D

Explanation:
This answer is correct because AWS Shield Advanced is a service that provides comprehensive protection
against DDoS attacks of any size or duration. It also provides metrics and reports on the DDoS attack vectors, duration, and size. You can create an Amazon CloudWatch alarm that monitors Shield Advanced metrics such as DDoSAttackBitsPerSecond, DDoSAttackPacketsPerSecond, and DDoSAttackRequestsPerSecond to receive alerts if a DDoS attack occurs against your account.
For more information, see Monitoring AWS Shield Advanced with Amazon CloudWatch and AWS Shield Advanced metrics and alarms.

NEW QUESTION 3
A company stores sensitive documents in Amazon S3 by using server-side encryption with an IAM Key Management Service (IAM KMS) CMK. A new requirement mandates that the CMK that is used for these documents can be used only for S3 actions.
Which statement should the company add to the key policy to meet this requirement?
A)

```
{
  "Effect": "Deny",
  "Principal": "*",
  "Action": "kms:*",
  "Resource": "*",
  "Condition": {
    "StringNotEquals": {
      "kms:CallerAccount": "s3.amazonaws.com"
    }
  }
}
```

B)

```
{
  "Effect": "Deny",
  "Principal": "*",
  "Action": "s3:*",
  "Resource": "*",
  "Condition": {
    "StringNotEquals": {
      "kms:ViaService": "kms.*amazonaws.com"
    }
  }
}
```

A. Option A
B. Option B

Answer: A

NEW QUESTION 4
A security engineer must troubleshoot an administrator's inability to make an existing Amazon S3 bucket public in an account that is part of an organization n IAM Organizations. The administrator switched the role from the master account to a member account and then attempted to make one S3 bucket public. This action was immediately denied
Which actions should the security engineer take to troubleshoot the permissions issue? (Select TWO.)

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

Answer: DE

Explanation:
A is incorrect because reviewing the cross-account role permissions and the S3 bucket policy is not enough to troubleshoot the permissions issue. You also need to verify that the Amazon S3 block public access option in the member account is deactivated, as well as the permissions boundary and the SCPs of the role in the member account.
D is correct because evaluating the SCPs and the permissions boundary of the role in the member account can help you identify any missing permissions or explicit denies that could prevent the administrator from making the S3 bucket public.
E is correct because ensuring that the S3 bucket policy explicitly allows the s3 PutBucketPublicAccess action for the role in the member account can help you override any block public access settings that could prevent the administrator from making the S3 bucket public.

NEW QUESTION 5
An international company wants to combine AWS Security Hub findings across all the company's AWS Regions and from multiple accounts. In addition, the company
wants to create a centralized custom dashboard to correlate these findings with operational data for deeper
analysis and insights. The company needs an analytics tool to search and visualize Security Hub findings. Which combination of steps will meet these requirements? (Select THREE.)

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

Answer: BDF

Explanation:
The correct answer is B, D, and F. Designate an AWS account in an organization in AWS Organizations as a delegated administrator for Security Hub. Publish events to Amazon EventBridge from the delegated administrator account, all member accounts, and required Regions that are enabled for Security Hub findings. In each Region, create an Amazon EventBridge rule to deliver findings to an Amazon Kinesis Data Firehose delivery stream. Configure the Kinesis Data Firehose delivery streams to deliver the logs to a single Amazon S3 bucket. Partition the Amazon S3 data. Use AWS Glue to crawl the S3 bucket and build the schema. Use Amazon Athena to query the data and create views to flatten nested attributes. Build Amazon QuickSight dashboards that use the Athena views.
According to the AWS documentation, AWS Security Hub is a service that provides you with a comprehensive view of your security state across your AWS accounts, and helps you check your environment against security standards and best practices. You can use Security Hub to aggregate security findings from various sources, such as AWS services, partner products, or your own applications.
To use Security Hub with multiple AWS accounts and Regions, you need to enable AWS Organizations with all features enabled. This allows you to centrally manage your accounts and apply policies across your organization. You can also use Security Hub as a service principal for AWS Organizations, which lets you designate a delegated administrator account for Security Hub. The delegated administrator account can enable Security Hub automatically in all existing and future accounts in your organization, and can view and manage findings from all accounts.
According to the AWS documentation, Amazon EventBridge is a serverless event bus that makes it easy to connect applications using data from your own applications, integrated software as a service (SaaS) applications, and AWS services. You can use EventBridge to create rules that match events from various sources and route them to targets for processing.
To use EventBridge with Security Hub findings, you need to enable Security Hub as an event source in EventBridge. This will allow you to publish events from Security Hub to EventBridge in the same Region. You can then create EventBridge rules that match Security Hub findings based on criteria such as severity, type, or resource. You can also specify targets for your rules, such as Lambda functions, SNS topics, or Kinesis Data Firehose delivery streams.
According to the AWS documentation, Amazon Kinesis Data Firehose is a fully managed service that delivers real-time streaming data to destinations such as Amazon S3, Amazon Redshift, Amazon Elasticsearch Service (Amazon ES), and Splunk. You can use Kinesis Data Firehose to transform and enrich your data before delivering it to your destination.
To use Kinesis Data Firehose with Security Hub findings, you need to create a Kinesis Data Firehose delivery stream in each Region where you have enabled

Security Hub. You can then configure the delivery stream to receive events from EventBridge as a source, and deliver the logs to a single S3 bucket as a destination. You can also enable data transformation or compression on the delivery stream if needed.
According to the AWS documentation, Amazon S3 is an object storage service that offers scalability, data availability, security, and performance. You can use S3 to store and retrieve any amount of data from anywhere on the web. You can also use S3 features such as lifecycle management, encryption, versioning, and replication to optimize your storage.
To use S3 with Security Hub findings, you need to create an S3 bucket that will store the logs from Kinesis Data Firehose delivery streams. You can then partition the data in the bucket by using prefixes such as account ID or Region. This will improve the performance and cost-effectiveness of querying the data.
According to the AWS documentation, AWS Glue is a fully managed extract, transform, and load (ETL) service that makes it easy to prepare and load your data for analytics. You can use Glue to crawl your data sources, identify data formats, and suggest schemas and transformations. You can also use Glue Data Catalog as a central metadata repository for your data assets.
To use Glue with Security Hub findings, you need to create a Glue crawler that will crawl the S3 bucket and build the schema for the data. The crawler will create tables in the Glue Data Catalog that you can query using standard SQL.
According to the AWS documentation, Amazon Athena is an interactive query service that makes it easy to analyze data in Amazon S3 using standard SQL. Athena is serverless, so there is no infrastructure to manage, and you pay only for the queries that you run. You can use Athena with Glue Data Catalog as a metadata store for your tables.
To use Athena with Security Hub findings, you need to create views in Athena that will flatten nested attributes in the data. For example, you can create views that extract fields such as account ID, Region, resource type, resource ID, finding type, finding title, and finding description from the JSON data. You can then query the views using SQL and join them with other tables if needed.
According to the AWS documentation, Amazon QuickSight is a fast, cloud-powered business intelligence
service that makes it easy to deliver insights to everyone in your organization. You can use QuickSight to create and publish interactive dashboards that include machine learning insights. You can also use QuickSight to connect to various data sources, such as Athena, S3, or RDS.
To use QuickSight with Security Hub findings, you need to create QuickSight dashboards that use the Athena views as data sources. You can then visualize and analyze the findings using charts, graphs, maps, or tables. You can also apply filters, calculations, or aggregations to the data. You can then share the dashboards with your users or embed them in your applications.

NEW QUESTION 6
A company has recently recovered from a security incident that required the restoration of Amazon EC2 instances from snapshots.
After performing a gap analysis of its disaster recovery procedures and backup strategies, the company is concerned that, next time, it will not be able to recover the EC2 instances if the AWS account was compromised and Amazon EBS snapshots were deleted.
All EBS snapshots are encrypted using an AWS KMS CMK. Which solution would solve this problem?

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

Answer: C

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

NEW QUESTION 7
A company uses AWS Organizations and has production workloads across multiple AWS accounts. A security engineer needs to design a solution that will proactively monitor for suspicious behavior across all the accounts that contain production workloads.
The solution must automate remediation of incidents across the production accounts. The solution also must publish a notification to an Amazon Simple Notification Service (Amazon SNS) topic when a critical security finding is detected. In addition, the solution must send all security incident logs to a dedicated account.
Which solution will meet these requirements?

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

Answer: D

Explanation:
The correct answer is D.
To design a solution that will proactively monitor for suspicious behavior across all the accounts that contain production workloads, the security engineer needs to

use a service that can aggregate and analyze security findings from multiple sources. AWS Security Hub is a service that provides a comprehensive view of your security posture across your AWS accounts and enables you to check your environment against security standards and best practices. Security Hub also integrates with other AWS services, such as Amazon GuardDuty, AWS Config, and AWS Systems Manager, to collect and correlate security findings.
To automate remediation of incidents across the production accounts, the security engineer needs to use a service that can trigger actions based on events. Amazon EventBridge is a serverless event bus service that allows you to connect your applications with data from a variety of sources. EventBridge can use rules to match events and route them to targets for processing. You can use EventBridge to invoke a custom AWS Lambda function from the Security Hub findings. Lambda is a serverless compute service that lets you run code without provisioning or managing servers.
To publish a notification to an Amazon SNS topic when a critical security finding is detected, the security engineer needs to use a service that can send messages to subscribers. Amazon SNS is a fully managed messaging service that enables you to decouple and scale microservices, distributed systems, and serverless applications. SNS can deliver messages to a variety of endpoints, such as email, SMS, or HTTP. You can configure the Lambda function to also publish notifications to the SNS topic.
To send all security incident logs to a dedicated account, the security engineer needs to use a service that can aggregate and store log data from multiple sources. AWS Security Hub allows you to aggregate security findings from multiple accounts into a single account using the delegated administrator feature. This feature enables you to designate an AWS account as the administrator for Security Hub in an organization. The administrator account can then view and manage Security Hub findings from all member accounts.
Therefore, option D is correct because it meets all the requirements of the solution. Option A is incorrect because GuardDuty does not provide a comprehensive view of your security posture across your AWS accounts. GuardDuty is primarily a threat detection service that monitors for malicious or unauthorized behavior. Option B is incorrect because Config and Systems Manager are not designed to automate remediation of incidents based on Security Hub findings. Config is a service that enables you to assess, audit, and evaluate the configurations of your AWS resources, while Systems Manager is a service that allows you to manage your infrastructure on AWS at scale. Option C is incorrect because GuardDuty does not provide a comprehensive view of your security posture across your AWS accounts.
References:

- AWS Security Hub
- Amazon EventBridge
- AWS Lambda
- Amazon SNS
- Aggregating Security Hub findings across accounts

NEW QUESTION 8
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 Stor
B. In the template, replace all references to the value with {{resolve:ssm:MySSMParameterName:I}}.
C. Store the API key value in AWS Secrets Manage
D. In the template, replace all references to the value with { {resolve:secretsmanager:MySecretId:SecretString}}.
E. Store the API key value in Amazon DynamoD
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 bucke
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 9
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-easMb Each subnet is associate 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-1 b 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 inus-east-lb 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 inus-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 10
You need to create a policy and apply it for just an individual user. How could you accomplish this in the right way?
Please select:

A. Add an IAM managed policy for the user
B. Add a service policy for the user
C. Add an IAM role for the user
D. Add an inline policy for the user

Answer: D

Explanation:
Options A and B are incorrect since you need to add an inline policy just for the user Option C is invalid because you don't assign an IAM role to a user
The IAM Documentation mentions the following
An inline policy is a policy that's embedded in a principal entity (a user, group, or role)—that is, the policy is an inherent part of the principal entity. You can create a policy and embed it in a principal entity, either when you create the principal entity or later.
For more information on IAM Access and Inline policies, just browse to the below URL: https://docs.IAM.amazon.com/IAM/latest/UserGuide/access
The correct answer is: Add an inline policy for the user Submit your Feedback/Queries to our Experts

NEW QUESTION 11
An international company has established a new business entity in South Korea. The company also has established a new AWS account to contain the workload for the South Korean region. The company has set up the workload in the new account in the ap-northeast-2 Region. The workload consists of three Auto Scaling groups of Amazon EC2 instances. All workloads that operate in this Region must keep system logs and application logs for 7 years.
A security engineer must implement a solution to ensure that no logging data is lost for each instance during scaling activities. The solution also must keep the logs for only the required period of 7 years.
Which combination of steps should the security engineer take to meet these requirements? (Choose three.)

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

Answer: ABC

Explanation:
The correct combination of steps that the security engineer should take to meet these requirements are A. Ensure that the Amazon CloudWatch agent is installed on all the EC2 instances that the Auto Scaling groups launch. Generate a CloudWatch agent configuration file to forward the required logs to Amazon CloudWatch Logs., B. Set the log retention for desired log groups to 7 years., and C. Attach an IAM role to the launch configuration or launch template that the Auto Scaling groups use. Configure the role to provide the necessary permissions to forward logs to Amazon CloudWatch Logs.
* A. This answer is correct because it meets the requirement of ensuring that no logging data is lost for each instance during scaling activities. By installing the CloudWatch agent on all the EC2 instances, the security engineer can collect and send system logs and application logs to CloudWatch Logs, which is a service that stores and monitors log data. By generating a CloudWatch agent configuration file, the security engineer can specify which logs to forward and how often.
* B. This answer is correct because it meets the requirement of keeping the logs for only the required period of 7 years. By setting the log retention for desired log groups, the security engineer can control how long
CloudWatch Logs retains log events before deleting them. The security engineer can choose a predefined retention period of 7 years, or use a custom value.
* C. This answer is correct because it meets the requirement of providing the necessary permissions to forward logs to CloudWatch Logs. By attaching an IAM role to the launch configuration or launch template that the Auto Scaling groups use, the security engineer can grant permissions to the EC2 instances that are launched by the Auto Scaling groups. By configuring the role to provide the necessary permissions, such as cloudwatch:PutLogEvents and cloudwatch:CreateLogStream, the security engineer can allow the EC2 instances to send log data to CloudWatch Logs.

NEW QUESTION 12
A Security Engineer has been tasked with enabling IAM Security Hub to monitor Amazon EC2 instances fix CVE in a single IAM account The Engineer has already enabled IAM Security Hub and Amazon Inspector m the IAM Management Console and has installed me Amazon Inspector agent on an EC2 instances that need

to be monitored.
Which additional steps should the Security Engineer lake 10 meet this requirement?

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

Answer: D

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

NEW QUESTION 13
A Development team has built an experimental environment to test a simple stale web application It has built an isolated VPC with a private and a public subnet. The public subnet holds only an Application Load Balancer a NAT gateway, and an internet gateway. The private subnet holds ail of the Amazon EC2 instances There are 3 different types of servers Each server type has its own Security Group that limits access lo only required connectivity. The Security Groups nave both inbound and outbound rules applied Each subnet has both inbound and outbound network ACIs applied to limit access to only required connectivity
Which of the following should the team check if a server cannot establish an outbound connection to the
internet? (Select THREE.)

A. The route tables and the outbound rules on the appropriate private subnet security group
B. The outbound network ACL rules on the private subnet and the Inbound network ACL rules on the public subnet
C. The outbound network ACL rules on the private subnet and both the inbound and outbound rules on the public subnet
D. The rules on any host-based firewall that may be applied on the Amazon EC2 instances
E. The Security Group applied to the Application Load Balancer and NAT gateway
F. That the 0.0.0./0 route in the private subnet route table points to the internet gateway in the public subnet

Answer: CEF

Explanation:
because these are the factors that could affect the outbound connection to the internet from a server in a private subnet. The outbound network ACL rules on the private subnet and both the inbound and outbound rules on the public subnet must allow the traffic to pass through8. The security group applied to the application load balancer and NAT gateway must also allow the traffic from the private subnet9. The 0.0.0.0/0 route in the private subnet route table must point to the NAT gateway in the public subnet, not the internet gateway10. The other options are either irrelevant or incorrect for troubleshooting the outbound connection issue.

NEW QUESTION 14
A company wants to ensure that its IAM resources can be launched only in the us-east-1 and us-west-2 Regions.
What is the MOST operationally efficient solution that will prevent developers from launching Amazon EC2 instances in other Regions?

A. Enable Amazon GuardDuty in all Region
B. Create alerts to detect unauthorized activity outside us-east-1 and us-west-2.
C. Use an organization in IAM Organization
D. Attach an SCP that allows all actions when the IAM: Requested Region condition key is either us-east-1 or us-west-2. Delete the FullIAMAccess policy.
E. Provision EC2 resources by using IAM Cloud Formation templates through IAM CodePipelin
F. Allow only the values of us-east-1 and us-west-2 in the IAM CloudFormation template's parameters.
G. Create an IAM Config rule to prevent unauthorized activity outside us-east-1 and us-west-2.

Answer: C

NEW QUESTION 15
A company is testing its incident response plan for compromised credentials. The company runs a database on an Amazon EC2 instance and stores the sensitive data-base credentials as a secret in AWS Secrets Manager. The secret has rotation configured with an AWS Lambda function that uses the generic rotation function template. The EC2 instance and the Lambda function are deployed in the same pri-vate subnet. The VPC has a Secrets Manager VPC endpoint.
A security engineer discovers that the secret cannot rotate. The security engi-neer determines that the VPC endpoint is working as intended. The Amazon Cloud-Watch logs contain the following error:
"setSecret: Unable to log into database". Which solution will resolve this error?

A. Use the AWS Management Console to edit the JSON structure of the secret in Secrets Manager so that the secret automatically conforms with the struc-ture that the database requires.
B. Ensure that the security group that is attached to the Lambda function al-lows outbound connections to the EC2 instanc
C. Ensure that the security group that is attached to the EC2 instance allows inbound connections from the security group that is attached to the Lambda function.
D. Use the Secrets Manager list-secrets command in the AWS CLI to list the secre
E. Identify the database credential
F. Use the Secrets Manager rotate-secret command in the AWS CLI to force the immediate rotation of the se-cret.
G. Add an internet gateway to the VP
H. Create a NAT gateway in a public sub-ne
I. Update the VPC route tables so that traffic from the Lambda function and traffic from the EC2 instance can reach the Secrets Manager public endpoint.

Answer: B

Explanation:
This answer is correct because ensuring that the security groups allow bidirectional communication between the Lambda function and the EC2 instance will resolve the error. The error indicates that the Lambda function cannot connect to the database, which might be due to firewall rules blocking the traffic. By allowing outbound connections from the Lambda function and inbound connections to the EC2 instance, the security engineer can enable the rotation function to access and update the database credentials.

NEW QUESTION 16
A company's application team needs to host a MySQL database on IAM. According to the company's security policy, all data that is stored on IAM must be encrypted at rest. In addition, all cryptographic material must be compliant with FIPS 140-2 Level 3 validation.
The application team needs a solution that satisfies the company's security requirements and minimizes
operational overhead.
Which solution will meet these requirements?

A. Host the database on Amazon RD
B. Use Amazon Elastic Block Store (Amazon EBS) for encryption.Use an IAM Key Management Service (IAM KMS) custom key store that is backed by IAM CloudHSM for key management.
C. Host the database on Amazon RD
D. Use Amazon Elastic Block Store (Amazon EBS) for encryption.Use an IAM managed CMK in IAM Key Management Service (IAM KMS) for key management.
E. Host the database on an Amazon EC2 instanc
F. Use Amazon Elastic Block Store (Amazon EBS) for encryptio
G. Use a customer managed CMK in IAM Key Management Service (IAM KMS) for key management.
H. Host the database on an Amazon EC2 instanc
I. Use Transparent Data Encryption (TDE) for encryption and key management.

Answer: B

NEW QUESTION 17
A company is evaluating its security posture. In the past, the company has observed issues with specific hosts and host header combinations that affected the company's business. The company has configured AWS WAF web ACLs as an initial step to mitigate these issues.
The company must create a log analysis solution for the AWS WAF web ACLs to monitor problematic activity. The company wants to process all the AWS WAF logs in a central location. The company must have the ability to filter out requests based on specific hosts.
A security engineer starts to enable access logging for the AWS WAF web ACLs.
What should the security engineer do next to meet these requirements with the MOST operational efficiency?

A. Specify Amazon Redshift as the destination for the access log
B. Deploy the Amazon Athena Redshift connecto
C. Use Athena to query the data from Amazon Redshift and to filter the logs by host.
D. Specify Amazon CloudWatch as the destination for the access log
E. Use Amazon CloudWatch Logs Insights to design a query to filter the logs by host.
F. Specify Amazon CloudWatch as the destination for the access log
G. Export the CloudWatch logs to an Amazon S3 bucke
H. Use Amazon Athena to query the logs and to filter the logs by host.
I. Specify Amazon CloudWatch as the destination for the access log
J. Use Amazon Redshift Spectrum to query the logs and to filter the logs by host.

Answer: C

Explanation:
The correct answer is C. Specify Amazon CloudWatch as the destination for the access logs. Export the CloudWatch logs to an Amazon S3 bucket. Use Amazon Athena to query the logs and to filter the logs by host.
According to the AWS documentation1, AWS WAF offers logging for the traffic that your web ACLs analyze. The logs include information such as the time that AWS WAF received the request from your protected AWS resource, detailed information about the request, and the action setting for the rule that the request matched. You can send your logs to an Amazon CloudWatch Logs log group, an Amazon Simple Storage Service (Amazon S3) bucket, or an Amazon Kinesis Data Firehose.
To create a log analysis solution for the AWS WAF web ACLs, you can use Amazon Athena, which is an interactive query service that makes it easy to analyze data in Amazon S3 using standard SQL2. You can use Athena to query and filter the AWS WAF logs by host or any other criteria. Athena is serverless, so there is no infrastructure to manage, and you pay only for the queries that you run.
To use Athena with AWS WAF logs, you need to export the CloudWatch logs to an S3 bucket. You can do this by creating a subscription filter that sends your log events to a Kinesis Data Firehose delivery stream, which then delivers the data to an S3 bucket3. Alternatively, you can use AWS DMS to migrate your CloudWatch logs to S34.
After you have exported your CloudWatch logs to S3, you can create a table in Athena that points to your S3 bucket and use the AWS service log format that matches your log schema5. For example, if you are using JSON format for your AWS WAF logs, you can use the AWSJSONSerDe serde. Then you can run SQL queries on your Athena table and filter the results by host or any other field in your log data.
Therefore, this solution meets the requirements of creating a log analysis solution for the AWS WAF web ACLs with the most operational efficiency. This solution does not require setting up any additional infrastructure or services, and it leverages the existing capabilities of CloudWatch, S3, and Athena.
The other options are incorrect because:

- A. Specifying Amazon Redshift as the destination for the access logs is not possible, because AWS WAF does not support sending logs directly to Redshift. You would need to use an intermediate service such as Kinesis Data Firehose or AWS DMS to load the data from CloudWatch or S3 to Redshift. Deploying the Amazon Athena Redshift connector is not necessary, because you can query Redshift data directly from Athena without using a connector6. This solution would also incur additional costs and operational overhead of managing a Redshift cluster.
- B. Specifying Amazon CloudWatch as the destination for the access logs is possible, but using Amazon CloudWatch Logs Insights to design a query to filter the logs by host is not efficient or scalable. CloudWatch Logs Insights is a feature that enables you to interactively search and analyze your log data in CloudWatch Logs7. However, CloudWatch Logs Insights has some limitations, such as a maximum query duration of 20 minutes, a maximum of 20 log groups per query, and a maximum retention period of 24 months8. These limitations may affect your ability to perform complex and long-running analysis on your AWS WAF logs.
- D. Specifying Amazon CloudWatch as the destination for the access logs is possible, but using Amazon Redshift Spectrum to query the logs and filter them by host is not efficient or cost-effective. Redshift Spectrum is a feature of Amazon Redshift that enables you to run queries against exabytes of data in S3 without loading or transforming any data9. However, Redshift Spectrum requires a Redshift cluster to process the queries, which adds additional costs and operational overhead. Redshift Spectrum also charges you based on the number of bytes scanned by each query, which can be expensive if you have large volumes of log data10.

References:
1: Logging AWS WAF web ACL traffic - Amazon Web Services 2: What Is Amazon Athena? - Amazon Athena 3: Streaming CloudWatch Logs Data to Amazon S3 - Amazon CloudWatch Logs 4: Migrate data from CloudWatch Logs using AWS Database Migration Service - AWS Database Migration Service 5: Querying AWS service logs - Amazon Athena 6: Querying data from Amazon Redshift - Amazon Athena 7: Analyzing log data with CloudWatch Logs Insights - Amazon CloudWatch Logs 8: CloudWatch Logs Insights quotas - Amazon CloudWatch 9: Querying external data using Amazon Redshift Spectrum - Amazon Redshift 10: Amazon Redshift Spectrum pricing - Amazon Redshift

NEW QUESTION 18
A company has recently recovered from a security incident that required the restoration of Amazon EC2 instances from snapshots. The company uses an AWS Key
Management Service (AWS KMS) customer managed key to encrypt all Amazon Elastic Block Store (Amazon EBS) snapshots.
The company performs a gap analysis of its disaster recovery procedures and backup strategies. A security engineer needs to implement a solution so that the company can recover the EC2 instances if the AWS account is compromised and the EBS snapshots are deleted.
Which solution will meet this requirement?

A. Create a new Amazon S3 bucke
B. Use EBS lifecycle policies to move EBS snapshots to the new S3 bucke
C. Use lifecycle policies to move snapshots to the S3 Glacier Instant Retrieval storage clas
D. Use S3 Object Lock to prevent deletion of the snapshots.
E. Use AWS Systems Manager to distribute a configuration that backs up all attached disks to Amazon S3.
F. Create a new AWS account that has limited privilege
G. Allow the new account to access the KMS key that encrypts the EBS snapshot
H. Copy the encrypted snapshots to the new account on a recurring basis.
I. Use AWS Backup to copy EBS snapshots to Amazon S3. Use S3 Object Lock to prevent deletion of the snapshots.

Answer: C

Explanation:
This solution meets the requirement of recovering the EC2 instances if the AWS account is compromised and the EBS snapshots are deleted. By creating a new AWS account with limited privileges, the company can isolate the backup snapshots from the main account and reduce the risk of accidental or malicious deletion. By allowing the new account to access the KMS key that encrypts the EBS snapshots, the company can ensure that the snapshots are copied in an encrypted form and can be decrypted when needed. By copying the encrypted snapshots to the new account on a recurring basis, the company can maintain a consistent backup schedule and minimize data loss.

NEW QUESTION 19
A company has two VPCs in the same AWS Region and in the same AWS account Each VPC uses a CIDR block that does not overlap with the CIDR block of the other VPC One VPC contains AWS Lambda functions that run inside a subnet that accesses the internet through a NAT gateway. The Lambda functions require access to a publicly accessible Amazon Aurora MySQL database that is running in the other VPC
A security engineer determines that the Aurora database uses a security group rule that allows connections from the NAT gateway IP address that the Lambda functions use. The company's security policy states that no database should be publicly accessible.
What is the MOST secure way that the security engineer can provide the Lambda functions with access to the Aurora database?

A. Move the Aurora database into a private subnet that has no internet access routes in the database's current VPC Configure the Lambda functions to use the Auroradatabase's new private IP address to access the database Configure the Aurora databases security group to allow access from the private IP addresses of the Lambda functions
B. Establish a VPC endpoint between the two VPCs in the Aurora database's VPC configure a service VPC endpoint for Amazon RDS In the Lambda functions' VPC.configure an interface VPC endpoint that uses the service endpoint in the Aurora database's VPC Configure the service endpoint to allow connections from the Lambda functions.
C. Establish an AWS Direct Connect interface between the VPCs Configure the Lambda functions to use a new route table that accesses the Aurora database through the Direct Connect interface Configure the Aurora database's security group to allow access from the Direct Connect interface IP address
D. Move the Lambda functions into a public subnet in their VPC Move the Aurora database into a private subnet in its VPC Configure the Lambda functions to use the Aurora database's new private IP address to access the database Configure the Aurora database to allow access from the public IP addresses of the Lambda functions

Answer: B

Explanation:
This option involves creating a VPC Endpoint between the two VPCs that allows private communication between them without going through the internet or exposing any public IP addresses. In this option, a VPC endpoint for Amazon RDS will be established, and an interface VPC endpoint will be created that points to the service endpoint in the Aurora database's VPC. This way, the Lambda functions can use the private IP address of the Aurora database to access it through the VPC endpoint without exposing any public IP addresses or allowing public internet access to the database.

NEW QUESTION 20
A company is using Amazon Elastic Container Service (Amazon ECS) to run its container-based application on AWS. The company needs to ensure that the container images contain no severe vulnerabilities. The company also must ensure that only specific IAM roles and specific AWS accounts can access the container images.
Which solution will meet these requirements with the LEAST management overhead?

A. Pull images from the public container registr
B. Publish the images to Amazon Elastic Container Registry (Amazon ECR) repositories with scan on push configured in a centralized AWS accoun
C. Use a CI/CD pipeline to deploy the images to different AWS account
D. Use identity-based policies to restrict access to which IAM principals can access the images.
E. Pull images from the public container registr
F. Publish the images to a private container registry that is hosted on Amazon EC2 instances in a centralized AWS accoun
G. Deploy host-based container scanning tools to EC2 instances that run Amazon EC
H. Restrict access to the container images by using basic authentication over HTTPS.
I. Pull images from the public container registr
J. Publish the images to Amazon Elastic Container Registry (Amazon ECR) repositories with scan on push configured in a centralized AWS accoun
K. Use a CI/CD pipeline to deploy the images to different AWS account
L. Use repository policies and identity-based policies to restrict access to which IAM principals and accounts can access the images.
M. Pull images from the public container registr
N. Publish the images to AWS CodeArtifact repositories in a centralized AWS accoun
O. Use a CI/CD pipeline to deploy the images to different AWS account
P. Use repository policies and identity-based policies to restrict access to which IAM principals and accounts can access the images.

Answer: C

Explanation:

The correct answer is C. Pull images from the public container registry. Publish the images to Amazon Elastic Container Registry (Amazon ECR) repositories with scan on push configured in a centralized AWS account.
Use a CI/CD pipeline to deploy the images to different AWS accounts. Use repository policies and identity-based policies to restrict access to which IAM principals and accounts can access the images.
This solution meets the requirements because:

- Amazon ECR is a fully managed container registry service that supports Docker and OCI images and artifacts1. It integrates with Amazon ECS and other AWS services to simplify the development and deployment of container-based applications.
- Amazon ECR provides image scanning on push, which uses the Common Vulnerabilities and Exposures (CVEs) database from the open-source Clair project to detect software vulnerabilities in container images2. The scan results are available in the AWS Management Console, AWS CLI, or AWS SDKs2.
- Amazon ECR supports cross-account access to repositories, which allows sharing images across multiple AWS accounts3. This can be achieved by using repository policies, which are resource-based policies that specify which IAM principals and accounts can access the repositories and what actions they can perform4. Additionally, identity-based policies can be used to control which IAM roles in each account can access the repositories5.

The other options are incorrect because:

- A. This option does not use repository policies to restrict cross-account access to the images, which is a requirement. Identity-based policies alone are not sufficient to control access to Amazon ECR repositories5.
- B. This option does not use Amazon ECR, which is a fully managed service that provides image scanning and cross-account access features. Hosting a private container registry on EC2 instances would require more management overhead and additional security measures.
- D. This option uses AWS CodeArtifact, which is a fully managed artifact repository service that supports Maven, npm, NuGet, PyPI, and generic package formats6. However, AWS CodeArtifact does not support Docker or OCI container images, which are required for Amazon ECS applications.

NEW QUESTION 21
A company's cloud operations team is responsible for building effective security for IAM cross-account access. The team asks a security engineer to help troubleshoot why some developers in the developer account (123456789012) in the developers group are not able to assume a cross-account role (ReadS3) into a production account (999999999999) to read the contents of an Amazon S3 bucket (productionapp). The two account policies are as follows:

Developer account 123456789012:
Developer group permissions:

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

Production account 999999999999:
Production account ReadS3 role policy:

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

Production account ReadS3 role policy - trust relationship:

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

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

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

Answer: AD

NEW QUESTION 22
A company has an organization in AWS Organizations that includes dedicated accounts for each of its business units. The company is collecting all AWS CloudTrail logs from the accounts in a single Amazon S3 bucket in the top-level account. The company's IT governance team has access to the top-level account. A security engineer needs to allow each business unit to access its own CloudTrail logs.
The security engineer creates an IAM role in the top-level account for each of the other accounts. For each role the security engineer creates an IAM policy to allow read-only permissions to objects in the S3 bucket with the prefix of the respective logs.
Which action must the security engineer take in each business unit account to allow an IAM user in that account to read the logs?

A. Attach a policy to the IAM user to allow the user to assume the role that was created in the top-level accoun
B. Specify the role's ARN in the policy.
C. Create an SCP that grants permissions to the top-level account.
D. Use the root account of the business unit account to assume the role that was created in the top-level accoun
E. Specify the role's ARN in the policy.

F. Forward the credentials of the IAM role in the top-level account to the IAM user in the business unit account.

Answer: A

Explanation:
To allow an IAM user in one AWS account to access resources in another AWS account using IAM roles, the following steps are required:

- Create a role in the AWS account that contains the resources (the trusting account) and specify the AWS account that contains the IAM user (the trusted account) as a trusted entity in the role's trust policy. This allows users from the trusted account to assume the role and access resources in the trusting account.
- Attach a policy to the IAM user in the trusted account that allows the user to assume the role in the trusting account. The policy must specify the ARN of the role that was created in the trusting account.
- The IAM user can then switch roles or use temporary credentials to access the resources in the trusting account.

Verified References:

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

NEW QUESTION 23
Auditors for a health care company have mandated that all data volumes be encrypted at rest Infrastructure is deployed mainly via IAM CloudFormation however third-party frameworks and manual deployment are required on some legacy systems
What is the BEST way to monitor, on a recurring basis, whether all EBS volumes are encrypted?

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

Answer: B

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

NEW QUESTION 24
A company is expanding its group of stores. On the day that each new store opens, the company wants to launch a customized web application for that store. Each store's application will have a non-production environment and a production environment. Each environment will be deployed in a separate AWS account. The company uses AWS Organizations and has an OU that is used only for these accounts.
The company distributes most of the development work to third-party development teams. A security engineer needs to ensure that each team follows the company's
deployment plan for AWS resources. The security engineer also must limit access to the deployment plan to only the developers who need access. The security engineer already has created an AWS CloudFormation template that implements the deployment plan.
What should the security engineer do next to meet the requirements in the MOST secure way?

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

Answer: A

Explanation:
The correct answer is A. Create an AWS Service Catalog portfolio in the organization's management account. Upload the CloudFormation template. Add the template to the portfolio's product list. Share the portfolio with the OU.
According to the AWS documentation, AWS Service Catalog is a service that allows you to create and manage catalogs of IT services that are approved for use on AWS. You can use Service Catalog to centrally manage commonly deployed IT services and help achieve consistent governance and compliance requirements, while enabling users to quickly deploy only the approved IT services they need.
To use Service Catalog with multiple AWS accounts, you need to enable AWS Organizations with all features enabled. This allows you to centrally manage your accounts and apply policies across your organization. You can also use Service Catalog as a service principal for AWS Organizations, which lets you share your portfolios with organizational units (OUs) or accounts in your organization.
To create a Service Catalog portfolio, you need to use an administrator account, such as the organization's management account. You can upload your CloudFormation template as a product in your portfolio, and define constraints and tags for it. You can then share your portfolio with the OU that contains the accounts for the web applications. This will allow the developers in those accounts to launch products from the shared portfolio using the Service Catalog end user console.
Option B is incorrect because CloudFormation modules are reusable components that encapsulate one or more resources and their configurations. They are not meant to be used as templates for deploying entire stacks of resources. Moreover, sharing a module with an OU does not grant access to launch stacks from it.
Option C is incorrect because creating an IAM role that has a trust policy that allows cross-account access to the portfolio is not secure. It would allow any user in the OU accounts to assume the role and access the portfolio, regardless of their job function or access requirements.

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

NEW QUESTION 25
A company wants to monitor the deletion of customer managed CMKs A security engineer must create an alarm that will notify the company before a CMK is deleted The security engineer has configured the integration of IAM CloudTrail with Amazon CloudWatch
What should the security engineer do next to meet this requirement?

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

Answer: A

NEW QUESTION 26
A security engineer is creating an AWS Lambda function. The Lambda function needs to use a role that is named LambdaAuditRole to assume a role that is named AcmeAuditFactoryRole in a different AWS account.
When the code is processed, the following error message appears: "An error oc-curred (AccessDenied) when calling the AssumeRole operation."
Which combination of steps should the security engineer take to resolve this er-ror? (Select TWO.)

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

Answer: AC

NEW QUESTION 27
A company deploys a set of standard IAM roles in AWS accounts. The IAM roles are based on job functions within the company. To balance operational efficiency and security, a security engineer implemented AWS Organizations SCPs to restrict access to critical security services in all company accounts.
All of the company's accounts and OUs within AWS Organizations have a default FullAWSAccess SCP that is attached. The security engineer needs to ensure that no one can disable Amazon GuardDuty and AWS Security Hub. The security engineer also must not override other permissions that are granted by IAM policies that are defined in the accounts.
Which SCP should the security engineer attach to the root of the organization to meet these requirements? A)

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

B)

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

C)

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

D)

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

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

Answer: A

NEW QUESTION 28
A company wants to protect its website from man in-the-middle attacks by using Amazon CloudFront. Which solution will meet these requirements with the LEAST operational overhead?

A. Use the SimpleCORS managed response headers policy.
B. Use a Lambda@Edge function to add the Strict-Transport-Security response header.

C. Use the SecurityHeadersPolicy managed response headers policy.
D. Include the X-XSS-Protection header in a custom response headers policy.

Answer: C

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

NEW QUESTION 29
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 30
.....

Thank You for Trying Our Product

We offer two products:

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

2nd - Questons and Answers in PDF Format

SCS-C02 Practice Exam Features:

* SCS-C02 Questions and Answers Updated Frequently

* SCS-C02 Practice Questions Verified by Expert Senior Certified Staff

* SCS-C02 Most Realistic Questions that Guarantee you a Pass on Your FirstTry

* SCS-C02 Practice Test Questions in Multiple Choice Formats and Updatesfor 1 Year