



Amazon-Web-Services

Exam Questions SCS-C02

AWS Certified Security - Specialty

About ExamBible

[Your Partner of IT Exam](#)

Found in 1998

ExamBible is a company specialized on providing high quality IT exam practice study materials, especially Cisco CCNA, CCDA, CCNP, CCIE, Checkpoint CCSE, CompTIA A+, Network+ certification practice exams and so on. We guarantee that the candidates will not only pass any IT exam at the first attempt but also get profound understanding about the certificates they have got. There are so many alike companies in this industry, however, ExamBible has its unique advantages that other companies could not achieve.

Our Advances

* 99.9% Uptime

All examinations will be up to date.

* 24/7 Quality Support

We will provide service round the clock.

* 100% Pass Rate

Our guarantee that you will pass the exam.

* Unique Gurantee

If you do not pass the exam at the first time, we will not only arrange FULL REFUND for you, but also provide you another exam of your claim, ABSOLUTELY FREE!

NEW QUESTION 1

A company is hosting a web application on Amazon EC2 instances behind an Application Load Balancer (ALB). The application has become the target of a DoS attack. Application logging shows that requests are coming from small number of client IP addresses, but the addresses change regularly. The company needs to block the malicious traffic with a solution that requires the least amount of ongoing effort. Which solution meets these requirements?

- A. Create an AWS WAF rate-based rule, and attach it to the ALB.
- B. Update the security group that is attached to the ALB to block the attacking IP addresses.
- C. Update the ALB subnet's network ACL to block the attacking client IP addresses.
- D. Create a AWS WAF rate-based rule, and attach it to the security group of the EC2 instances.

Answer: A

NEW QUESTION 2

A web application gives users the ability to log in verify their membership's validity and browse artifacts that are stored in an Amazon S3 bucket. When a user attempts to download an object, the application must verify the permission to access the object and allow the user to download the object from a custom domain name such as example.com.

What is the MOST secure way for a security engineer to implement this functionality?

- A. Configure read-only access to the object by using a bucket AC
- B. Remove the access after a set time has elapsed.
- C. Implement an IAM policy to give the user read access to the S3 bucket.
- D. Create an S3 presigned URL Provide the S3 presigned URL to the user through the application.
- E. Create an Amazon CloudFront signed UR
- F. Provide the CloudFront signed URL to the user through the application.

Answer: D

Explanation:

For this scenario you would need to set up static website hosting because a custom domain name is listed as a requirement. "Amazon S3 website endpoints do not support HTTPS or access points. If you want to use HTTPS, you can use Amazon CloudFront to serve a static website hosted on Amazon S3." This is not secure. <https://docs.aws.amazon.com/AmazonS3/latest/userguide/website-hosting-custom-domain-walkthrough.html> CloudFront signed URLs allow much more fine-grained control as well as HTTPS access with custom domain names:

<https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/private-content-signed-urls.html>

NEW QUESTION 3

A company has AWS accounts in an organization in AWS Organizations. The organization includes a dedicated security account.

All AWS account activity across all member accounts must be logged and reported to the dedicated security account. The company must retain all the activity logs in a secure storage location within the dedicated security account for 2 years. No changes or deletions of the logs are allowed.

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

- A. In the dedicated security account, create an Amazon S3 bucket
- B. Configure S3 Object Lock in compliance mode and a retention period of 2 years on the S3 bucket
- C. Set the bucket policy to allow the organization's management account to write to the S3 bucket.
- D. In the dedicated security account, create an Amazon S3 bucket
- E. Configure S3 Object Lock in compliance mode and a retention period of 2 years on the S3 bucket
- F. Set the bucket policy to allow the organization's member accounts to write to the S3 bucket.
- G. In the dedicated security account, create an Amazon S3 bucket that has an S3 Lifecycle configuration that expires objects after 2 year
- H. Set the bucket policy to allow the organization's member accounts to write to the S3 bucket.
- I. Create an AWS Cloud Trail trail for the organization
- J. Configure logs to be delivered to the logging Amazon S3 bucket in the dedicated security account.
- K. Turn on AWS CloudTrail in each account
- L. Configure logs to be delivered to an Amazon S3 bucket that is created in the organization's management account
- M. Forward the logs to the S3 bucket in the dedicated security account by using AWS Lambda and Amazon Kinesis Data Firehose.

Answer: BD

Explanation:

The correct answer is B and D. In the dedicated security account, create an Amazon S3 bucket. Configure S3 Object Lock in compliance mode and a retention period of 2 years on the S3 bucket. Set the bucket policy to allow the organization's member accounts to write to the S3 bucket. Create an AWS CloudTrail trail for the organization. Configure logs to be delivered to the logging Amazon S3 bucket in the dedicated security account.

According to the AWS documentation, AWS CloudTrail is a service that enables governance, compliance, operational auditing, and risk auditing of your AWS account. With CloudTrail, you can log, continuously monitor, and retain account activity related to actions across your AWS infrastructure. CloudTrail provides event history of your AWS account activity, including actions taken through the AWS Management Console, AWS SDKs, command line tools, and other AWS services.

To use CloudTrail 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 CloudTrail as a service principal for AWS Organizations, which lets you create an organization trail that applies to all accounts in your organization. An organization trail logs events for all AWS Regions and delivers the log files to an S3 bucket that you specify.

To create an organization trail, you need to use an administrator account, such as the organization's management account or a delegated administrator account. You can then configure the trail to deliver logs to an S3 bucket in the dedicated security account. This will ensure that all account activity across all member accounts and regions is logged and reported to the security account.

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 CloudTrail logs, you need to create an S3 bucket in the dedicated security account that will store the logs from the organization trail. You can then configure S3 Object Lock on the bucket to prevent objects from being deleted or overwritten for a fixed amount of time or indefinitely. You can also enable compliance mode on the bucket, which prevents any user, including the root user in your account, from deleting or modifying a locked object until it reaches its retention date.

To set a retention period of 2 years on the S3 bucket, you need to create a default retention configuration for the bucket that specifies a retention mode (either governance or compliance) and a retention period (either a number of days or a date). You can then set the bucket policy to allow the organization's member accounts to write to the S3 bucket. This will ensure that all logs are retained in a secure storage location within the security account for 2 years and no changes or deletions are allowed.

Option A is incorrect because setting the bucket policy to allow the organization's management account to write to the S3 bucket is not sufficient, as it will not grant access to the other member accounts in the organization.

Option C is incorrect because using an S3 Lifecycle configuration that expires objects after 2 years is not secure, as it will allow users to delete or modify objects before they expire.

Option E is incorrect because using Lambda and Kinesis Data Firehose to forward logs from one S3 bucket to another is not necessary, as CloudTrail can directly deliver logs to an S3 bucket in another account. It also introduces additional operational overhead and complexity.

NEW QUESTION 4

An Incident Response team is investigating an IAM access key leak that resulted in Amazon EC2 instances being launched. The company did not discover the incident until many months later. The Director of Information Security wants to implement new controls that will alert when similar incidents happen in the future. Which controls should the company implement to achieve this? (Select TWO.)

- A. Enable VPC Flow Logs in all VPCs. Create a scheduled IAM Lambda function that downloads and parses the logs, and sends an Amazon SNS notification for violations.
- B. Use IAM CloudTrail to make a trail, and apply it to all Regions. Specify an Amazon S3 bucket to receive all the CloudTrail log files.
- C. Add the following bucket policy to the company's IAM CloudTrail bucket to prevent log tampering:

```
{ "Version": "2012-10-17", "Statement": [ { "Effect": "Deny", "Action": "s3:PutObject", "Principal": "*", "Resource": "arn:iam:s3:::cloudtrail/IAMLogs/111122223333/*" } ] }
```

 Create an Amazon S3 data event for an PutObject attempts, which sends notifications to an Amazon SNS topic.
- D. Create a Security Auditor role with permissions to access Amazon CloudWatch Logs in all Regions. Ship the logs to an Amazon S3 bucket and make a lifecycle policy to ship the logs to Amazon S3 Glacier.
- E. Verify that Amazon GuardDuty is enabled in all Regions, and create an Amazon CloudWatch Events rule for Amazon GuardDuty findings. Add an Amazon SNS topic as the rule's target.

Answer: AE

NEW QUESTION 5

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

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

Answer: D

Explanation:

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

NEW QUESTION 6

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

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

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

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

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

Answer: AD

NEW QUESTION 7

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

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

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

- A. Update the EC2 instance security group to add a rule that allows outbound traffic on port 443 for 0.0.0.0/0.
- B. Update the EC2 instance security group to add a rule that allows inbound traffic on port 443 to the VPC's CIDR range.
- C. Create an EC2 key pair
- D. Associate the key pair with the EC2 instance.
- E. Create a VPC interface endpoint for Systems Manager in the VPC where the EC2 instance is located.
- F. Attach a security group to the VPC interface endpoint
- G. Allow inbound traffic on port 443 to the VPC's CIDR range.
- H. Create a VPC interface endpoint for the EC2 instance in the VPC where the EC2 instance is located.

Answer: BCF

NEW QUESTION 8

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

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

Answer: D

Explanation:

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

NEW QUESTION 9

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

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

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

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

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

Answer: BE

NEW QUESTION 10

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

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

Answer: A

NEW QUESTION 10

A security engineer needs to implement a write-once-read-many (WORM) model for data that a company will store in Amazon S3 buckets. The company uses the S3 Standard storage class for all of its S3 buckets. The security engineer must en-sure that objects cannot be overwritten or deleted by any user, including the AWS account root user.

Which solution will meet these requirements?

- A. Create new S3 buckets with S3 Object Lock enabled in compliance mod
- B. Place objects in the S3 buckets.
- C. Use S3 Glacier Vault Lock to attach a Vault Lock policy to new S3 bucket
- D. Wait 24 hours to complete the Vault Lock proces

- E. Place objects in the S3 buckets.
- F. Create new S3 buckets with S3 Object Lock enabled in governance mod
- G. Place objects in the S3 buckets.
- H. Create new S3 buckets with S3 Object Lock enabled in governance mod
- I. Add a legal hold to the S3 bucket
- J. Place objects in the S3 buckets.

Answer: A

NEW QUESTION 13

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 14

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 16

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

- A. Assign an IAM policy to the instance profile to allow the EC2 instances to be managed by AWS Systems Manager. Provide the IAM user accounts with permission to use Systems Manager. Remove the SSH keys from the EC2 instances. Use Systems Manager Inventory to select the EC2 instance and connect.
- B. Assign an IAM policy to the IAM user accounts to provide permission to use AWS Systems Manager. Run Command. Remove the SSH keys from the EC2 instances. Use Run Command to open an SSH connection to the EC2 instance.
- C. Assign an IAM policy to the instance profile to allow the EC2 instances to be managed by AWS Systems Manager. Provide the IAM user accounts with permission to use Systems Manager. Remove the SSH keys from the EC2 instances. Use Systems Manager Session Manager to select the EC2 instance and connect.
- D. Assign an IAM policy to the IAM user accounts to provide permission to use the EC2 service in the AWS Management Console. Remove the SSH keys from the EC2 instances. Connect to the EC2 instance as the ec2-user through the AWS Management Console's EC2 SSH client method.

Answer: C

Explanation:

To provide access to the three individuals who have IAM user accounts to access the Amazon Linux 2 Amazon EC2 instances that are using the same IAM instance profile, the most appropriate solution would be to assign an IAM policy to the instance profile to allow the EC2 instances to be managed by AWS Systems Manager, provide the IAM user accounts with permission to use Systems Manager, remove the SSH keys from the EC2 instances, and use Systems Manager Session Manager to select the EC2 instance and connect.

References: : AWS Systems Manager Session Manager - AWS Systems Manager : AWS Systems Manager AWS Management Console : AWS Identity and Access Management - AWS Management Console : Amazon Elastic Compute Cloud - Amazon Web Services : Amazon Linux 2 - Amazon Web Services : AWS Systems Manager - AWS Management Console : AWS Systems Manager - AWS Management Console : AWS Systems Manager - AWS Management Console

NEW QUESTION 19

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

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

Answer: C

Explanation:

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

NEW QUESTION 23

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

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

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

Answer: B

Explanation:

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

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

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

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

References:

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

NEW QUESTION 25

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

Which solution meets these criteria?

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

Answer: A

Explanation:

<https://awscli.amazonaws.com/v2/documentation/api/latest/reference/kms/import-key-material.html> `aws kms import-key-material \`

```
--key-id 1234abcd-12ab-34cd-56ef-1234567890ab \
--encrypted-key-material fileb://EncryptedKeyMaterial.bin \
--import-token fileb://ImportToken.bin \
--expiration-model KEY_MATERIAL_EXPIRES \
--valid-to 2021-09-21T19:00:00Z
```

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

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

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

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

The other options are incorrect for the following reasons:

* B. A customer managed CMK that uses AWS provided key material does not expire automatically. You can enable automatic rotation of the key material every year, but this does not prevent access to the data encrypted with the previous key material. You would need to manually delete the CMK and its backing key material to make the data inaccessible³.

* C. An AWS managed CMK is a KMS key that is created, owned, and managed by an AWS service on your behalf. You have limited control over the key configuration, permissions, rotation, and deletion. You cannot use an AWS managed CMK to encrypt data in other AWS services or applications. You also cannot set an expiration date for the key material of an AWS managed CMK⁴.

* D. Operation system-native encryption that uses GnuPG is not a solution that uses AWS KMS. GnuPG is a command line tool that implements the OpenPGP standard for encrypting and signing data. It does not integrate with Amazon EBS or other AWS services. It also does not provide a way to automatically expire the key material used for encryption⁵.

References:

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

NEW QUESTION 26

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

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

- A. Have a Database Administrator encrypt the credentials and store the ciphertext in Amazon S3. Grant permission to the instance role associated with the EC2 instance to read the object and decrypt the ciphertext.

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

Answer: CE

Explanation:

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

* E. Configure the Java application to catch a connection failure and make a call to AWS Secrets Manager to retrieve updated credentials when the password is rotated. Grant permission to the instance role associated with the EC2 instance to access Secrets Manager.

By configuring the Java application to catch a connection failure and make a call to AWS Secrets Manager to retrieve updated credentials, you can avoid hard-coding the credentials in your application code or configuration files. This way, your application can dynamically obtain the latest credentials from Secrets Manager whenever the password is rotated, without needing to restart or redeploy the application. To enable this, you need to grant permission to the instance role associated with the EC2 instance to access Secrets Manager using IAM policies². You can also use the AWS SDK for Java to integrate your application with Secrets Manager³.

NEW QUESTION 28

A company purchased a subscription to a third-party cloud security scanning solution that integrates with AWS Security Hub. A security engineer needs to implement a solution that will remediate the findings from the third-party scanning solution automatically. Which solution will meet this requirement?

- A. Set up an Amazon EventBridge rule that reacts to new Security Hub findings.
- B. Configure an AWS Lambda function as the target for the rule to remediate the findings.
- C. Set up a custom action in Security Hub.
- D. Configure the custom action to call AWS Systems Manager Automation runbooks to remediate the findings.
- E. Set up a custom action in Security Hub.
- F. Configure an AWS Lambda function as the target for the custom action to remediate the findings.
- G. Set up AWS Config rules to use AWS Systems Manager Automation runbooks to remediate the findings.

Answer: A

NEW QUESTION 30

A company deploys a distributed web application on a fleet of Amazon EC2 instances. The fleet is behind an Application Load Balancer (ALB) that will be configured to terminate the TLS connection. All TLS traffic to the ALB must stay secure, even if the certificate private key is compromised. How can a security engineer meet this requirement?

- A. Create an HTTPS listener that uses a certificate that is managed by IAM Certificate Manager (ACM).
- B. Create an HTTPS listener that uses a security policy that uses a cipher suite with perfect forward secrecy (PFS).
- C. Create an HTTPS listener that uses the Server Order Preference security feature.
- D. Create a TCP listener that uses a custom security policy that allows only cipher suites with perfect forward secrecy (PFS).

Answer: A

NEW QUESTION 32

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

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

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

Answer: BD

Explanation:

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

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

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

NEW QUESTION 36

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

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

- A. Use a dynamic reference in the CloudFormation template to reference the database credentials in Secrets Manager.

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

Answer: A

Explanation:

➤ Option A: This option meets the requirements of following security best practices and configuring sensitive database credentials in the CloudFormation template. A dynamic reference is a way to specify external values that are stored and managed in other services, such as Secrets Manager, in the stack templates¹. When using a dynamic reference, CloudFormation retrieves the value of the specified reference when necessary during stack and change set operations¹. Dynamic references can be used for certain resources that support them, such as AWS::RDS::DBInstance¹. By using a dynamic reference to reference the database credentials in Secrets Manager, the company can leverage the existing integration between these services and avoid hardcoding the secret information in the template. Secrets Manager is a service that helps you protect secrets needed to access your applications, services, and IT resources². Secrets Manager enables you to rotate, manage, and retrieve database credentials, API keys, and other secrets throughout their lifecycle².

NEW QUESTION 39

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

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

- A. Open a support case with the IAM Security team and ask them to remove the malicious code from the affected instance
- B. Respond to the notification and list the actions that have been taken to address the incident
- C. Delete all IAM users and resources in the account
- D. Detach the internet gateway from the VPC remove aft rules that contain 0.0.0.0V0 from the security groups, and create a NACL rule to deny all traffic Inbound from the internet
- E. Delete the identified compromised instances and delete any associated resources that the Security team did not create.

Answer: DE

Explanation:

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

NEW QUESTION 41

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 45

A company is using AWS Organizations to implement a multi-account strategy. The company does not have on-premises infrastructure. All workloads run on AWS. The company currently has eight member accounts. The company anticipates that it will have no more than 20 AWS accounts total at any time.

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

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

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

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

- A. Use a CloudFormation template in the member accounts to launch workload
- B. Configure the template to use the Fn::ImportValue function to obtain the subnet ID values.
- C. Use a transit gateway in the VPC within Account-
- D. Configure the member accounts to use the transit gateway to access the subnets in Account-A to launch workloads.
- E. Use AWS Resource Access Manager (AWS RAM) to share Account-A's VPC subnets with the remaining member account
- F. Configure the member accounts to use the shared subnets to launch workloads.
- G. Create a peering connection between Account-A and the remaining member account
- H. Configure the member accounts to use the subnets in Account-A through the VPC peering connection to launch workloads.

Answer: C

Explanation:

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

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

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

The other options are incorrect because:

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

References:

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

NEW QUESTION 48

A company has an AWS account that includes an Amazon S3 bucket. The S3 bucket uses server-side encryption with AWS KMS keys (SSE-KMS) to encrypt all the objects at rest by using a customer managed key. The S3 bucket does not have a bucket policy.

An IAM role in the same account has an IAM policy that allows s3 List* and s3 Get* permissions for the S3 bucket. When the IAM role attempts to access an object in the S3 bucket the role receives an access denied message.

Why does the IAM role not have access to the objects that are in the S3 bucket?

- A. The IAM role does not have permission to use the KMS CreateKey operation.
- B. The S3 bucket lacks a policy that allows access to the customer managed key that encrypts the objects.
- C. The IAM role does not have permission to use the customer managed key that encrypts the objects that are in the S3 bucket.
- D. The ACL of the S3 objects does not allow read access for the objects when the objects are encrypted at rest.

Answer: C

Explanation:

When using server-side encryption with AWS KMS keys (SSE-KMS), the requester must have both Amazon S3 permissions and AWS KMS permissions to access the objects. The Amazon S3 permissions are for the bucket and object operations, such as s3:ListBucket and s3:GetObject. The AWS KMS permissions are for the key operations, such as kms:GenerateDataKey and kms:Decrypt. In this case, the IAM role has the necessary Amazon S3 permissions, but not the AWS KMS permissions to use the customer managed key that encrypts the objects. Therefore, the IAM role receives an access denied message when trying to access the objects. Verified References:

- <https://docs.aws.amazon.com/AmazonS3/latest/userguide/troubleshoot-403-errors.html>
- <https://repost.aws/knowledge-center/s3-access-denied-error-kms>
- <https://repost.aws/knowledge-center/cross-account-access-denied-error-s3>

NEW QUESTION 52

A company is using AWS Organizations to manage multiple AWS accounts for its human resources, finance, software development, and production departments. All the company's developers are part of the software development AWS account.

The company discovers that developers have launched Amazon EC2 instances that were preconfigured with software that the company has not approved for use. The company wants to implement a solution to ensure that developers can launch EC2 instances with only approved software applications and only in the software development AWS account.

Which solution will meet these requirements?

- A. In the software development account, create AMIs of preconfigured instances that include only approved software
- B. Include the AMI IDs in the condition section of an AWS CloudFormation template to launch the appropriate AMI based on the AWS Region
- C. Provide the developers with the CloudFormation template to launch EC2 instances in the software development account.
- D. Create an Amazon EventBridge rule that runs when any EC2 RunInstances API event occurs in the software development account
- E. Specify AWS Systems Manager Run Command as a target of the rule
- F. Configure Run Command to run a script that will install all approved software onto the instances that the developers launch.
- G. Use an AWS Service Catalog portfolio that contains EC2 products with appropriate AMIs that include only approved software
- H. Grant the developers permission to portfolio access only the Service Catalog to launch a product in the software development account.
- I. In the management account, create AMIs of preconfigured instances that include only approved software
- J. Use AWS CloudFormation StackSets to launch the AMIs across any AWS account in the organization
- K. Grant the developers permission to launch the stack sets within the management account.

Answer: C

NEW QUESTION 57

A development team is using an IAM Key Management Service (IAM KMS) CMK to try to encrypt and decrypt a secure string parameter from IAM Systems Manager Parameter Store. However, the development team receives an error message on each attempt.

Which issues that are related to the CMK could be reasons for the error? (Select TWO.)

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

- D. The CMK that is used in the attempt is not enabled.
- E. The CMK that is used in the attempt is using an alias.

Answer: AD

NEW QUESTION 58

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

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

Answer: C

Explanation:

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

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

NEW QUESTION 61

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_eroups.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 64

Which of the following bucket policies will ensure that objects being uploaded to a bucket called 'demo' are encrypted. Please select:

- A.

```
C:\Users\wk\Desktop\mudassar\Untitled.jpg
{
  "Version": "2012-10-17",
  "Id": "PutObj",
  "Statement": [
    {
      "Sid": "DenyUploads",
      "Effect": "Deny",
      "Principal": "*",
      "Action": "s3:PutObject",
      "Resource": "arn:aws:s3:::demo/*",
      "Condition": {
        "StringNotEquals": {
          "s3:x-amz-server-side-encryption": "aws:kms"
        }
      }
    }
  ]
}
```
- B.

```
C:\Users\wk\Desktop\mudassar\Untitled.jpg
{
  "Version": "2012-10-17",
  "Id": "PutObj",
  "Statement": [
    {
      "Sid": "DenyUploads",
      "Effect": "Deny",
      "Principal": "*",
      "Action": "s3:PutObject",
      "Resource": "arn:aws:s3:::demo/*",
      "Condition": {
        "StringNotEquals": {
          "s3:x-amz-server-side-encryption": "aws:kms"
        }
      }
    }
  ]
}
```

```
"Version": "2012-10-17",
"Id": "PutObj",
"Statement": [
  {
    "Sid": "DenyUploads",
    "Effect": "Deny",
    "Principal": "*",
    "Action": "s3:PutObject",
    "Resource": "arn:aws:s3:::demo/*",
    "Condition": {
      "StringEquals": {
        "s3:x-amz-server-side-encryption": "aws:kms"
      }
    }
  }
]
```

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

```
"Version": "2012-10-17",
"Id": "PutObj",
"Statement": [
  {
    "Sid": "DenyUploads",
    "Effect": "Deny",
    "Principal": "*",
    "Action": "s3:PutObject",
    "Resource": "arn:aws:s3:::demo/*"
  }
]
```

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

```
"Version": "2012-10-17",
"Id": "PutObj",
"Statement": [
  {
    "Sid": "DenyUploads",
    "Effect": "Deny",
    "Principal": "*",
    "Action": "s3:PutObjectEncrypted",
    "Resource": "arn:aws:s3:::demo/*"
  }
]
```

Answer: A

Explanation:

The condition of "s3:x-amz-server-side-encryption":"IAM:kms" ensures that objects uploaded need to be encrypted.

Options B,C and D are invalid because you have to ensure the condition of ns3:x-amz-server-side-encryption":"IAM:kms" is present

For more information on IAM KMS best practices, just browse to the below URL: <https://dl.IAMstatic.com/whitepapers/IAM-kms-best-praaices.pdf>

Submit your Feedback/Queries to our Expert

NEW QUESTION 67

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

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

- A. Deploy an AWS Config managed rule to run on a periodic basis of 24 hour
- B. Select the access-keys-rotated managed rule, and set the maxAccessKeyAge parameter to 90 day
- C. Create an Amazon EventBridge (Amazon CloudWatch Events) rule with an event pattern that matches the compliance type of NON_COMPLIANT from AWS Config for the managed rule
- D. Configure EventBridge (CloudWatch Events) to send an Amazon Simple Notification Service (Amazon SNS) notification to the security team.
- E. Create a script to export a .csv file from the AWS Trusted Advisor check for IAM access key rotation. Load the script into an AWS Lambda function that will upload the .csv file to an Amazon S3 bucket
- F. Create an Amazon Athena table query that runs when the .csv file is uploaded to the S3 bucket
- G. Publish the results for any keys older than 90 days by using an invocation of an Amazon Simple Notification Service (Amazon SNS) notification to the security team.
- H. Create a script to download the IAM credentials report on a periodic basis
- I. Load the script into an AWS Lambda function that will run on a schedule through Amazon EventBridge (Amazon CloudWatch Events). Configure the Lambda

script to load the report into memory and to filter the report for records in which the key was last rotated at least 90 days ago.
J. If any records are detected, send an Amazon Simple Notification Service (Amazon SNS) notification to the security team.
K. Create an AWS Lambda function that queries the IAM API to list all the users.
L. Iterate through the users by using the ListAccessKeys operation.
M. Verify that the value in the CreateDate field is not at least 90 days old.
N. Send an Amazon Simple Notification Service (Amazon SNS) notification to the security team if the value is at least 90 days old.
O. Create an Amazon EventBridge (Amazon CloudWatch Events) rule to schedule the Lambda function to run each day.

Answer: A

NEW QUESTION 71

A security engineer needs to run an AWS CloudFormation script. The CloudFormation script builds AWS infrastructure to support a stack that includes web servers and a MySQL database. The stack has been deployed in pre-production environments and is ready for production. The production script must comply with the principle of least privilege. Additionally, separation of duties must exist between the security engineer's IAM account and CloudFormation. Which solution will meet these requirements?

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

Answer: A

Explanation:

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

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

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

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

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

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

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

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

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

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

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

NEW QUESTION 73

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

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

Answer: C

NEW QUESTION 76

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

Which S3 bucket policy will meet this requirement?

A. {

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

B. {

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

C. {

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

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

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

Answer: B

Explanation:

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

NEW QUESTION 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 80

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

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

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

- A. Create a service role that has a composite principal that contains each service that needs the necessary permission.
- B. Configure the role to allow the sts:AssumeRole action.
- C. Create a service role that has cloudformation.amazonaws.com as the service principal.
- D. Configure the role to allow the sts:AssumeRole action.
- E. For each required set of permissions, add a separate policy to the role to allow those permissions.
- F. Add the ARN of each CloudFormation stack in the resource field of each policy.
- G. For each required set of permissions, add a separate policy to the role to allow those permissions.
- H. Add the ARN of each service that needs the permissions in the resource field of the corresponding policy.
- I. Update each stack to use the service role.
- J. Add a policy to each member role to allow the iam:PassRole action.
- K. Set the policy's resource field to the ARN of the service role.

Answer: BDF

NEW QUESTION 82

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

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

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

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

How should the security engineer correct the error?

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

Answer: D

Explanation:

CloudWatchLogsReadOnlyAccess doesn't include "logs:CreateLogStream" but it includes "logs:Get*"

<https://docs.aws.amazon.com/AmazonCloudWatch/latest/logs/iam-identity-based-access-control-cwl.html#:~:te>

NEW QUESTION 87

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

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

Answer: B

Explanation:

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

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

NEW QUESTION 91

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

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

Answer: C

Explanation:

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

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

NEW QUESTION 95

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

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

Answer: C

Explanation:

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

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

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

- Create a vault in S3 Glacier using the AWS Management Console, AWS Command Line Interface (AWS CLI), or AWS SDKs.
- Create a Vault Lock policy in S3 Glacier that meets all the regulatory requirements using the IAM policy language. The policy can include conditions such as

aws:CurrentTime or aws:SecureTransport to further restrict access to the vault.

➤ Initiate the lock by attaching the Vault Lock policy to the vault, which sets the lock to an in-progress state and returns a lock ID. While the policy is in the in-progress state, you have 24 hours to validate your Vault Lock policy before the lock ID expires. To prevent your vault from exiting the in-progress state, you must complete the Vault Lock process within these 24 hours. Otherwise, your Vault Lock policy will be deleted.

➤ Use the lock ID to complete the lock process. If the Vault Lock policy doesn't work as expected, you can stop the Vault Lock process and restart from the beginning.

➤ Upload the data to the vault using either direct upload or multipart upload methods. For more information about S3 Glacier and Vault Lock, see S3 Glacier Vault Lock.

The other options are incorrect because:

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

➤ Option B is incorrect because creating an Amazon S3 bucket and configuring it to use S3 Object Lock in governance mode will not prevent anyone from changing or deleting the data. S3 Object Lock in governance mode works similarly to compliance mode, except that users with specific IAM permissions can change or delete objects that are locked. This means that users who have s3:BypassGovernanceRetention permission can remove retention periods or legal holds from objects and overwrite or delete them before they expire. This option does not provide strong enforcement for compliance controls as required by the regulatory requirements.

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

NEW QUESTION 98

A company needs to encrypt all of its data stored in Amazon S3. The company wants to use IAM Key Management Service (IAM KMS) to create and manage its encryption keys. The company's security policies require the ability to Import the company's own key material for the keys, set an expiration date on the keys, and delete keys immediately, if needed.

How should a security engineer set up IAM KMS to meet these requirements?

- A. Configure IAM KMS and use a custom key stor
- B. Create a customer managed CMK with no key material Import the company's keys and key material into the CMK
- C. Configure IAM KMS and use the default Key store Create an IAM managed CMK with no key material Import the company's key material into the CMK
- D. Configure IAM KMS and use the default key store Create a customer managed CMK with no key material import the company's key material into the CMK
- E. Configure IAM KMS and use a custom key stor
- F. Create an IAM managed CMK with no key material.Import the company's key material into the CMK.

Answer: A

Explanation:

To meet the requirements of importing their own key material, setting an expiration date on the keys, and deleting keys immediately, the security engineer should do the following:

➤ Configure AWS KMS and use a custom key store. This allows the security engineer to use a key manager outside of AWS KMS that they own and manage, such as an AWS CloudHSM cluster or an external key manager.

➤ Create a customer managed CMK with no key material. Import the company's keys and key material into the CMK. This allows the security engineer to use their own key material for encryption and decryption operations, and to specify an expiration date for it.

NEW QUESTION 101

A security engineer needs to implement a solution to create and control the keys that a company uses for cryptographic operations. The security engineer must create symmetric keys in which the key material is generated and used within a custom key store that is backed by an AWS CloudHSM cluster.

The security engineer will use symmetric and asymmetric data key pairs for local use within applications. The security engineer also must audit the use of the keys. How can the security engineer meet these requirements?

- A. To create the keys use AWS Key Management Service (AWS KMS) and the custom key stores with the CloudHSM cluste
- B. For auditing, use Amazon Athena
- C. To create the keys use Amazon S3 and the custom key stores with the CloudHSM cluste
- D. For auditing use AWS CloudTrail.
- E. To create the keys use AWS Key Management Service (AWS KMS) and the custom key stores with the CloudHSM cluste
- F. For auditing, use Amazon GuardDuty.
- G. To create the keys use AWS Key Management Service (AWS KMS) and the custom key stores with the CloudHSM cluste
- H. For auditing, use AWS CloudTrail.

Answer: D

Explanation:

AWS KMS supports asymmetric KMS keys that represent a mathematically related RSA, elliptic curve (ECC), or SM2 (China Regions only) public and private key pair. These key pairs are generated in AWS KMS hardware security modules certified under the FIPS 140-2 Cryptographic Module Validation Program, except in the China (Beijing) and China (Ningxia) Regions. The private key never leaves the AWS KMS HSMs unencrypted.

<https://docs.aws.amazon.com/kms/latest/developerguide/symmetric-asymmetric.html>

NEW QUESTION 106

A security engineer is designing a cloud architecture to support an application. The application runs on Amazon EC2 instances and processes sensitive information, including credit card numbers.

The application will send the credit card numbers to a component that is running in an isolated environment. The component will encrypt, store, and decrypt the

numbers.

The component then will issue tokens to replace the numbers in other parts of the application.

The component of the application that manages the tokenization process will be deployed on a separate set of EC2 instances. Other components of the application must not be able to store or access the credit card numbers.

Which solution will meet these requirements?

- A. Use EC2 Dedicated Instances for the tokenization component of the application.
- B. Place the EC2 instances that manage the tokenization process into a partition placement group.
- C. Create a separate VP
- D. Deploy new EC2 instances into the separate VPC to support the data tokenization.
- E. Deploy the tokenization code onto AWS Nitro Enclaves that are hosted on EC2 instances.

Answer: D

Explanation:

AWS Nitro Enclaves are isolated and hardened virtual machines that run on EC2 instances and provide a secure environment for processing sensitive data. Nitro Enclaves have no persistent storage, interactive access, or external networking, and they can only communicate with the parent instance through a secure local channel. Nitro Enclaves also support cryptographic attestation, which allows verifying the identity and integrity of the enclave and its code. Nitro Enclaves are ideal for implementing data protection solutions such as tokenization, encryption, and key management.

Using Nitro Enclaves for the tokenization component of the application meets the requirements of isolating the sensitive data from other parts of the application, encrypting and storing the credit card numbers securely, and issuing tokens to replace the numbers. Other components of the application will not be able to access or store the credit card numbers, as they are only available within the enclave.

NEW QUESTION 107

A security administrator is setting up a new AWS account. The security administrator wants to secure the data that a company stores in an Amazon S3 bucket.

The security administrator also wants to reduce the chance of unintended data exposure and the potential for misconfiguration of objects that are in the S3 bucket.

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

- A. Configure the S3 Block Public Access feature for the AWS account.
- B. Configure the S3 Block Public Access feature for all objects that are in the bucket.
- C. Deactivate ACLs for objects that are in the bucket.
- D. Use AWS PrivateLink for Amazon S3 to access the bucket.

Answer: D

NEW QUESTION 111

An organization must establish the ability to delete an IAM KMS Customer Master Key (CMK) within a

24- hour timeframe to keep it from being used for encrypt or decrypt operations Which of the following actions will address this requirement?

- A. Manually rotate a key within KMS to create a new CMK immediately
- B. Use the KMS import key functionality to execute a delete key operation
- C. Use the schedule key deletion function within KMS to specify the minimum wait period for deletion
- D. Change the KMS CMK alias to immediately prevent any services from using the CMK.

Answer: C

Explanation:

the schedule key deletion function within KMS allows you to specify a waiting period before deleting a customer master key (CMK)⁴. The minimum waiting period is 7 days and the maximum is 30 days⁵. This function prevents the CMK from being used for encryption or decryption operations during the waiting period⁴. The other options are either invalid or ineffective for deleting a CMK within a 24-hour timeframe.

NEW QUESTION 112

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

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

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

Answer: ADF

NEW QUESTION 117

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 configuratio
- 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 AW
- E. Key Management Service (AWS KMS) customer managed key.

- F. Set the dnssec-enable option to yes in the BIND configuratio
- G. Create a zone-signing key (ZSK) and a key-signing key (KSK). Run the dnssec-signzone command to generate a delegation signer (DS) record Use AW
- H. Key Management Service (AWS KMS) to secure the keys.
- I. Migrate the zone to Route 53 with DNSSEC signing enable
- J. Create a key-signing key (KSK) that is based on an AWS Key Management Service (AWS KMS) customer managed ke
- 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 122

A company is hosting multiple applications within a single VPC in its IAM account. The applications are running behind an Application Load Balancer that is associated with an IAM WAF web ACL. The company's security team has identified that multiple port scans are originating from a specific range of IP addresses on the internet.

A security engineer needs to deny access from the offending IP addresses. Which solution will meet these requirements?

- A. Modify the IAM WAF web ACL with an IP set match rule statement to deny incoming requests from the IP address range.
- B. Add a rule to all security groups to deny the incoming requests from the IP address range.
- C. Modify the IAM WAF web ACL with a rate-based rule statement to deny the incoming requests from the IP address range.
- D. Configure the IAM WAF web ACL with regex match condition
- E. Specify a pattern set to deny the incoming requests based on the match condition

Answer: A

Explanation:

Note that the IP is known and the question wants us to deny access from that particular address and so we can use IP set match policy of WAF to block access.

NEW QUESTION 127

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

The application must

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

A Security Engineer is evaluating options for encryption key management

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

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

Answer: B

Explanation:

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

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

NEW QUESTION 132

A Development team has built an experimental environment to test a simple state 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 all of the Amazon EC2 instances There are 3 different types of servers Each server type has its own Security Group that limits access to only required connectivity. The Security Groups have both inbound and outbound rules applied Each subnet has both inbound and outbound network ACLs 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 137

An application is running on an Amazon EC2 instance that has an IAM role attached. The IAM role provides access to an AWS Key Management Service (AWS KMS) customer managed key and an Amazon S3 bucket. The key is used to access 2 TB of sensitive data that is stored in the S3 bucket.

A security engineer discovers a potential vulnerability on the EC2 instance that could result in the compromise of the sensitive data. Due to other critical operations, the security engineer cannot immediately shut down the EC2 instance for vulnerability patching.

What is the FASTEST way to prevent the sensitive data from being exposed?

- A. Download the data from the existing S3 bucket to a new EC2 instance
- B. Then delete the data from the S3 bucket
- C. Re-encrypt the data with a client-based key
- D. Upload the data to a new S3 bucket.
- E. Block access to the public range of S3 endpoint IP addresses by using a host-based firewall
- F. Ensure that internet-bound traffic from the affected EC2 instance is routed through the host-based firewall.
- G. Revoke the IAM role's active session permission
- H. Update the S3 bucket policy to deny access to the IAM role
- I. Remove the IAM role from the EC2 instance profile.
- J. Disable the current key
- K. Create a new KMS key that the IAM role does not have access to, and re-encrypt all the data with the new key
- L. Schedule the compromised key for deletion.

Answer: D

NEW QUESTION 141

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 months
- E. Then recreate the user again.
- F. Delete the IAM Role associated with the keys after every 2 months
- 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 143

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 145

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 148

Example.com is hosted on Amazon EC2 instances behind an Application Load Balancer (ALB). Third-party host intrusion detection system (HIDS) agents that capture the traffic of the EC2 instance are running on each host. The company must ensure they are using privacy enhancing technologies for users, without losing the assurance the third-party solution offers.

What is the MOST secure way to meet these requirements?

- A. Enable TLS pass through on the ALB, and handle decryption at the server using Elliptic Curve Diffie-Hellman (ECDHE) cipher suites.
- B. Create a listener on the ALB that uses encrypted connections with Elliptic Curve Diffie-Hellman (ECDHE) cipher suites, and pass the traffic in the clear to the server.
- C. Create a listener on the ALB that uses encrypted connections with Elliptic Curve Diffie-Hellman (ECDHE) cipher suites, and use encrypted connections to the servers that do not enable Perfect Forward Secrecy (PFS).
- D. Create a listener on the ALB that does not enable Perfect Forward Secrecy (PFS) cipher suites, and use encrypted connections to the servers using Elliptic Curve Diffie-Hellman (ECDHE) cipher suites.

Answer: D

Explanation:

the most secure way to meet the requirements. TLS is a protocol that provides encryption and authentication for data in transit. ALB is a service that distributes incoming traffic across multiple EC2 instances. HIDS is a system that monitors and detects malicious activity on a host. ECDHE is a type of cipher suite that supports perfect forward secrecy, which is a property that ensures that past and current TLS traffic stays secure even if the certificate private key is leaked. By creating a listener on the ALB that does not enable PFS cipher suites, and using encrypted connections to the servers using ECDHE cipher suites, you can ensure that the HIDS agents can capture the traffic of the EC2 instance without compromising the privacy of the users. The other options are either less secure or less compatible with the third-party solution.

NEW QUESTION 149

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 RDS
- 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 RDS
- 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 instance
- F. Use Amazon Elastic Block Store (Amazon EBS) for encryption
- G. Use a customer managed CMK in IAM Key Management Service (IAM KMS) for key management.
- H. Host the database on an Amazon EC2 instance
- I. Use Transparent Data Encryption (TDE) for encryption and key management.

Answer: B

NEW QUESTION 154

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 159

An ecommerce company has a web application architecture that runs primarily on containers. The application containers are deployed on Amazon Elastic Container Service (Amazon ECS). The container images for the application are stored in Amazon Elastic Container Registry (Amazon ECR).

The company's security team is performing an audit of components of the application architecture. The security team identifies issues with some container images that are stored in the container repositories.

The security team wants to address these issues by implementing continual scanning and on-push scanning of the container images. The security team needs to implement a solution that makes any findings from these scans visible in a centralized dashboard. The security team plans to use the dashboard to view these findings along with other security-related findings that they intend to generate in the future.

There are specific repositories that the security team needs to exclude from the scanning process. Which solution will meet these requirements?

- A. Use Amazon Inspector
- B. Create inclusion rules in Amazon ECR to match repos-itories that need to be scanne
- C. Push Amazon Inspector findings to AWS Se-curity Hub.
- D. Use ECR basic scanning of container image
- E. Create inclusion rules in Ama-zon ECR to match repositories that need to be scanne
- F. Push findings to AWS Security Hub.
- G. Use ECR basic scanning of container image
- H. Create inclusion rules in Ama-zon ECR to match repositories that need to be scanne
- I. Push findings to Amazon Inspector.
- J. Use Amazon Inspector
- K. Create inclusion rules in Amazon Inspector to match repositories that need to be scanne
- L. Push Amazon Inspector findings to AWS Config.

Answer: A

NEW QUESTION 163

A company has an organization with SCPs in AWS Organizations. The root SCP for the organization is as follows:


```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "AllowsAllActions",
      "Effect": "Allow",
      "Action": "*",
      "Resource": "*"
    },
    {
      "Sid": "DenySES",
      "Effect": "Deny",
      "Action": "ses:*",
      "Resource": "*"
    }
  ]
}
```

The company's developers are members of a group that has an IAM policy that allows access to Amazon Simple Email Service (Amazon SES) by allowing ses:* actions. The account is a child to an OU that has an SCP that allows Amazon SES. The developers are receiving a not-authorized error when they try to access Amazon SES through the AWS Management Console.

Which change must a security engineer implement so that the developers can access Amazon SES?

- A. Add a resource policy that allows each member of the group to access Amazon SES.
- B. Add a resource policy that allows "Principal": {"AWS": "arn:aws:iam::account-number:group/Dev"}.
- C. Remove the AWS Control Tower control (guardrail) that restricts access to Amazon SES.
- D. Remove Amazon SES from the root SCP.

Answer: D

Explanation:

The correct answer is D. Remove Amazon SES from the root SCP.

This answer is correct because the root SCP is the most restrictive policy that applies to all accounts in the organization. The root SCP explicitly denies access to Amazon SES by using the NotAction element, which means that any action that is not listed in the element is denied. Therefore, removing Amazon SES from the root SCP will allow the developers to access it, as long as there are no other SCPs or IAM policies that deny it.

The other options are incorrect because:

➤ A. Adding a resource policy that allows each member of the group to access Amazon SES is not a solution, because resource policies are not supported by Amazon SES¹. Resource policies are policies that are attached to AWS resources, such as S3 buckets or SNS topics, to control access to those resources². Amazon SES does not have any resources that can have resource policies attached to them.

➤ B. Adding a resource policy that allows "Principal": {"AWS": "arn:aws:iam::account-number:group/Dev"} is not a solution, because resource policies do not support IAM groups as principals³. Principals are entities that can perform actions on AWS resources, such as IAM users, roles, or AWS accounts⁴. IAM groups are not principals, but collections of IAM users that share the same permissions⁵.

➤ C. Removing the AWS Control Tower control (guardrail) that restricts access to Amazon SES is not a solution, because AWS Control Tower does not have any guardrails that restrict access to Amazon SES⁶. Guardrails are high-level rules that govern the overall behavior of an organization's accounts⁷. AWS Control Tower provides a set of predefined guardrails that cover security, compliance, and operations domains⁸.

References:

1: Amazon Simple Email Service endpoints and quotas 2: Resource-based policies and IAM policies 3: Specifying a principal in a policy 4: Policy elements: Principal 5: IAM groups 6: AWS Control Tower guardrails reference 7: AWS Control Tower concepts 8: AWS Control Tower guardrails

NEW QUESTION 167

.....

Relate Links

100% Pass Your SCS-C02 Exam with ExamBible Prep Materials

<https://www.exambible.com/SCS-C02-exam/>

Contact us

We are proud of our high-quality customer service, which serves you around the clock 24/7.

Viste - <https://www.exambible.com/>