



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

An AWS account that is used for development projects has a VPC that contains two subnets. The first subnet is named public-subnet-1 and has the CIDR block 192.168.1.0/24 assigned. The other subnet is named private-subnet-2 and has the CIDR block 192.168.2.0/24 assigned. Each subnet contains Amazon EC2 instances.

Each subnet is currently using the VPC's default network ACL. The security groups that the EC2 instances in these subnets use have rules that allow traffic between each instance where required. Currently, all network traffic flow is working as expected between the EC2 instances that are using these subnets. A security engineer creates a new network ACL that is named subnet-2-NACL with default entries. The security engineer immediately configures private-subnet-2 to use the new network ACL and makes no other changes to the infrastructure. The security engineer starts to receive reports that the EC2 instances in public-subnet-1 and public-subnet-2 cannot communicate with each other.

Which combination of steps should the security engineer take to allow the EC2 instances that are running in these two subnets to communicate again? (Select TWO.)

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

Answer: CE

Explanation:

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

References: : Amazon VPC User Guide

NEW QUESTION 2

A company in France uses Amazon Cognito with the Cognito Hosted UI as an identity broker for sign-in and sign-up processes. The company is marketing an application and expects that all the application's users will come from France.

When the company launches the application the company's security team observes fraudulent sign-ups for the application. Most of the fraudulent registrations are from users outside of France.

The security team needs a solution to perform custom validation at sign-up. Based on the results of the validation the solution must accept or deny the registration request.

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

- A. Create a pre sign-up AWS Lambda trigger.
- B. Associate the Amazon Cognito function with the Amazon Cognito user pool.
- C. Use a geographic match rule statement to configure an AWS WAF web ACL.
- D. Associate the web ACL with the Amazon Cognito user pool.
- E. Configure an app client for the application's Amazon Cognito user pool.
- F. Use the app client ID to validate the requests in the hosted UI.
- G. Update the application's Amazon Cognito user pool to configure a geographic restriction setting.
- H. Use Amazon Cognito to configure a social identity provider (IdP) to validate the requests on the hosted UI.

Answer: B

Explanation:

<https://docs.aws.amazon.com/cognito/latest/developerguide/user-pool-lambda-post-authentication.html>

NEW QUESTION 3

A security engineer needs to create an Amazon S3 bucket policy to grant least privilege read access to IAM user accounts that are named User1, User2, and User3. These IAM user accounts are members of the AuthorizedPeople IAM group. The security engineer drafts the following S3 bucket policy:

```
{
  "Version": "2012-10-17",
  "Id": "AuthorizedPeoplePolicy",
  "Statement": [
    {
      "Sid": "Actions-Authorized-People",
      "Effect": "Allow",
      "Action": [
        "s3:GetObject"
      ],
      "Resource": "arn:aws:s3:::authorized-people-bucket/*"
    }
  ]
}
```

When the security engineer tries to add the policy to the S3 bucket, the following error message appears: "Missing required field Principal." The security engineer is adding a Principal element to the policy. The addition must provide read access to only User1, User2, and User3. Which solution meets these requirements?

A)

```
{
  "Principal": {
    "AWS": [
      "arn:aws:iam::1234567890:user/User1",
      "arn:aws:iam::1234567890:user/User2",
      "arn:aws:iam::1234567890:user/User3"
    ]
  }
}
```

B)

```
"Principal": {
  "AWS": [
    "arn:aws:iam::1234567890:root"
  ]
}
```

C)

```
"Principal": {
  "AWS": [
    "*"
  ]
}
```

D)

```
"Principal": {
  "AWS": "arn:aws:iam::1234567890:group/AuthorizedPeople"
}
```

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

Answer: A

NEW QUESTION 4

A company uses a third-party identity provider and SAML-based SSO for its AWS accounts. After the third-party identity provider renewed an expired signing certificate, users saw the following message when trying to log in:
 Error: Response Signature Invalid (Service: AWSSecurityTokenService; Status Code: 400; Error Code: InvalidIdentityToken)
 A security engineer needs to provide a solution that corrects the error and minimizes operational overhead.
 Which solution meets these requirements?

- A. Upload the third-party signing certificate's new private key to the AWS identity provider entity defined in AWS Identity and Access Management (IAM) by using the AWS Management Console.
- B. Sign the identity provider's metadata file with the new public key
- C. Upload the signature to the AWS identity provider entity defined in AWS Identity and Access Management (IAM) by using the AWS CLI.
- D. Download the updated SAML metadata file from the identity service provider
- E. Update the file in the AWS identity provider entity defined in AWS Identity and Access Management (IAM) by using the AWS CLI.
- F. Configure the AWS identity provider entity defined in AWS Identity and Access Management (IAM) to synchronously fetch the new public key by using the AWS Management Console.

Answer: C

Explanation:

This answer is correct because downloading the updated SAML metadata file from the identity service provider ensures that AWS has the latest information about the identity provider, including the new public key. Updating the file in the AWS identity provider entity defined in IAM by using the AWS CLI allows AWS to verify the signature of the SAML assertions sent by the identity provider. This solution also minimizes operational overhead because it can be automated with a script or a cron job.

NEW QUESTION 5

A company is implementing new compliance requirements to meet customer needs. According to the new requirements the company must not use any Amazon RDS DB instances or DB clusters that lack encryption of the underlying storage. The company needs a solution that will generate an email alert when an unencrypted DB instance or DB cluster is created. The solution also must terminate the unencrypted DB instance or DB cluster.
 Which solution will meet these requirements in the MOST operationally efficient manner?

- A. Create an AWS Config managed rule to detect unencrypted RDS storage
- B. Configure an automatic remediation action to publish messages to an Amazon Simple Notification Service (Amazon SNS) topic that includes an AWS Lambda function and an email delivery target as subscriber
- C. Configure the Lambda function to delete the unencrypted resource.
- D. Create an AWS Config managed rule to detect unencrypted RDS storage
- E. Configure a manual remediation action to invoke an AWS Lambda function
- F. Configure the Lambda function to publish messages to an Amazon Simple Notification Service (Amazon SNS) topic and to delete the unencrypted resource.
- G. Create an Amazon EventBridge rule that evaluates RDS event patterns and is initiated by the creation of DB instances or DB clusters. Configure the rule to publish messages to an Amazon Simple Notification Service (Amazon SNS) topic that includes an AWS Lambda function and an email delivery target as subscriber
- H. Configure the Lambda function to delete the unencrypted resource.
- I. Create an Amazon EventBridge rule that evaluates RDS event patterns and is initiated by the creation of DB instances or DB cluster
- J. Configure the rule to invoke an AWS Lambda function
- K. Configure the Lambda function to publish messages to an Amazon Simple Notification Service (Amazon SNS) topic and to delete the unencrypted resource.

Answer: A

Explanation:

<https://docs.aws.amazon.com/config/latest/developerguide/rds-storage-encrypted.html>

NEW QUESTION 6

A security engineer wants to forward custom application-security logs from an Amazon EC2 instance to Amazon CloudWatch. The security engineer installs the

CloudWatch agent on the EC2 instance and adds the path of the logs to the CloudWatch configuration file. However, CloudWatch does not receive the logs. The security engineer verifies that the awslogs service is running on the EC2 instance. What should the security engineer do next to resolve the issue?

- A. Add AWS CloudTrail to the trust policy of the EC2 instance
- B. Send the custom logs to CloudTrail instead of CloudWatch.
- C. Add Amazon S3 to the trust policy of the EC2 instance
- D. Configure the application to write the custom logs to an S3 bucket that CloudWatch can use to ingest the logs.
- E. Add Amazon Inspector to the trust policy of the EC2 instance
- F. Use Amazon Inspector instead of the CloudWatch agent to collect the custom logs.
- G. Attach the CloudWatchAgentServerPolicy AWS managed policy to the EC2 instance role.

Answer: D

Explanation:

The correct answer is D. Attach the CloudWatchAgentServerPolicy AWS managed policy to the EC2 instance role.

According to the AWS documentation¹, the CloudWatch agent is a software agent that you can install on your EC2 instances to collect system-level metrics and logs. To use the CloudWatch agent, you need to attach an IAM role or user to the EC2 instance that grants permissions for the agent to perform actions on your behalf. The CloudWatchAgentServerPolicy is an AWS managed policy that provides the necessary permissions for the agent to write metrics and logs to CloudWatch². By attaching this policy to the EC2 instance role, the security engineer can resolve the issue of CloudWatch not receiving the custom application-security logs.

The other options are incorrect for the following reasons:

- > A. Adding AWS CloudTrail to the trust policy of the EC2 instance is not relevant, because CloudTrail is a service that records API activity in your AWS account, not custom application logs³. Sending the custom logs to CloudTrail instead of CloudWatch would not meet the requirement of forwarding them to CloudWatch.
- > B. Adding Amazon S3 to the trust policy of the EC2 instance is not necessary, because S3 is a storage service that does not require any trust relationship with EC2 instances⁴. Configuring the application to write the custom logs to an S3 bucket that CloudWatch can use to ingest the logs would be an alternative solution, but it would be more complex and costly than using the CloudWatch agent directly.
- > C. Adding Amazon Inspector to the trust policy of the EC2 instance is not helpful, because Inspector is a service that scans EC2 instances for software vulnerabilities and unintended network exposure, not custom application logs⁵. Using Amazon Inspector instead of the CloudWatch agent would not meet the requirement of forwarding them to CloudWatch.

References:

- 1: Collect metrics, logs, and traces with the CloudWatch agent - Amazon CloudWatch
- 2: CloudWatchAgentServerPolicy - AWS Managed Policy
- 3: What Is AWS CloudTrail? - AWS CloudTrail
- 4: Amazon S3 FAQs - Amazon Web Services
- 5: Automated Software Vulnerability Management - Amazon Inspector - AWS

NEW QUESTION 7

A company uses an Amazon S3 bucket to store reports. Management has mandated that all new objects stored in this bucket must be encrypted at rest using server-side encryption with a client-specified IAM Key Management Service (IAM KMS) CMK owned by the same account as the S3 bucket. The IAM account number is 111122223333, and the bucket name is reportbucket. The company's security specialist must write the S3 bucket policy to ensure the mandate can be implemented.

Which statement should the security specialist include in the policy?

- A.

```
{
  "Effect": "Deny",
  "Principal": "*",
  "Action": "s3:PutObject",
  "Resource": "arn:aws:s3:::reportbucket/*",
  "Condition": {
    "StringEquals": {
      "s3:x-amz-server-side-encryption": "AES256"
    }
  }
}
```
- B.

```
{
  "Effect": "Deny",
  "Principal": "*",
  "Action": "s3:PutObject",
  "Resource": "arn:aws:s3:::reportbucket/*",
  "Condition": {
    "StringNotLike": {
      "s3:x-amz-server-side-encryption-aws-kms-key-id": "arn:aws:kms:*:111122223333:key/*"
    }
  }
}
```
- C.

```
{
  "Effect": "Deny",
  "Principal": "*",
  "Action": "s3:PutObject",
  "Resource": "arn:aws:s3:::reportbucket/*",
  "Condition": {
    "StringNotLike": {
      "s3:x-amz-server-side-encryption": "aws:kms"
    }
  }
}
```
- D.

```
{
  "Effect": "Deny",
  "Principal": "*",
  "Action": "s3:PutObject",
  "Resource": "arn:aws:s3:::reportbucket/*",
  "Condition": {
    "StringNotLike": {
      "s3:x-amz-server-side-encryption": "aws:kms"
    }
  }
}
```

```

{
  "Effect": "Deny",
  "Principal": "*",
  "Action": "s3:PutObject",
  "Resource": "arn:aws:s3:::reportbucket/*",
  "Condition": {
    "StringNotLikeIfExists": {
      "s3:x-amz-server-side-encryption-aws-kms-key-id": "arn:aws:kms:*:111122223333:key/*"
    }
  }
}

```

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

Answer: D

NEW QUESTION 8

A Security Engineer creates an Amazon S3 bucket policy that denies access to all users. A few days later, the Security Engineer adds an additional statement to the bucket policy to allow read-only access to one other employee. Even after updating the policy, the employee still receives an access denied message. What is the likely cause of this access denial?

- A. The ACL in the bucket needs to be updated
- B. The IAM policy does not allow the user to access the bucket
- C. It takes a few minutes for a bucket policy to take effect
- D. The allow permission is being overridden by the deny

Answer: D

NEW QUESTION 9

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 10

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 10

A company has several workloads running on AWS. Employees are required to authenticate using on-premises ADFS and SSO to access the AWS Management Console. Developers migrated an existing legacy web application to an Amazon EC2 instance. Employees need to access this application from anywhere on the internet, but currently, there is no authentication system built into the application.

How should the Security Engineer implement employee-only access to this system without changing the application?

- A. Place the application behind an Application Load Balancer (ALB). Use Amazon Cognito as authentication for the ALB.
- B. Define a SAML-based Amazon Cognito user pool and connect it to ADFS.
- C. Implement AWS SSO in the master account and link it to ADFS as an identity provider.
- D. Define the EC2 instance as a managed resource, then apply an IAM policy on the resource.
- E. Define an Amazon Cognito identity pool, then install the connector on the Active Directory server.
- F. Use the Amazon Cognito SDK on the application instance to authenticate the employees using their Active Directory user names and passwords.
- G. Create an AWS Lambda custom authorizer as the authenticator for a reverse proxy on Amazon EC2. Ensure the security group on Amazon EC2 only allows access from the Lambda function.

Answer: A

Explanation:

<https://docs.aws.amazon.com/elasticloadbalancing/latest/application/listener-authenticate-users.html>

NEW QUESTION 13

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

What should the security engineer recommend?

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

Answer: C

NEW QUESTION 14

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 VP
- E. Create a bash script to list the DNS request number of common DNS queries from the BIND logs.
- F. Create VPC flow logs for all subnets in the VP
- G. Stream the flow logs to an Amazon CloudWatch Logs log 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 19

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 24

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 26

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

What should the security engineer do next to meet these requirements?

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

Answer: C

Explanation:

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

References: : AWS KMS Developer Guide

NEW QUESTION 31

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

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

Answer: D

Explanation:

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

- > One_Hour
- > Three_Hours
- > Six_Hours
- > Twelve_Hours
- > TwentyFour_Hours

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

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

The other options are incorrect because:

- > A. The IAM credential report can be generated at any time, but it will not affect the compliance status of the resources until the next evaluation cycle of the AWS Config rules.
- > B. The security engineer was able to invoke the IAM GenerateCredentialReport API operation, which means they have the GenerateCredentialReport permission. This permission is required to generate a credential report that lists all IAM users in an AWS account and their credential status².
- > C. The security engineer does not need the GetCredentialReport permission to enable or evaluate AWS Config rules. This permission is required to retrieve a credential report that was previously generated by using the GenerateCredentialReport operation².

References:

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

NEW QUESTION 34

A company uses AWS Organizations to manage several AWS accounts. The company processes a large volume of sensitive data. The company uses a serverless approach to microservices. The company stores all the data in either Amazon S3 or Amazon DynamoDB. The company reads the data by using either AWS Lambda functions or container-based services that the company hosts on Amazon Elastic Kubernetes Service (Amazon EKS) on AWS Fargate.

The company must implement a solution to encrypt all the data at rest and enforce least privilege data access controls. The company creates an AWS Key Management Service (AWS KMS) customer managed key.

What should the company do next to meet these requirements?

- A. Create a key policy that allows the kms:Decrypt action only for Amazon S3 and DynamoDB
- B. Create an SCP that denies the creation of S3 buckets and DynamoDB tables that are not encrypted with the key.
- C. Create an IAM policy that denies the kms:Decrypt action for the key
- D. Create a Lambda function that runs on a schedule to attach the policy to any new role
- E. Create an AWS Config rule to send alerts for resources that are not encrypted with the key.
- F. Create a key policy that allows the kms:Decrypt action only for Amazon S3, DynamoDB, Lambda, and Amazon EKS
- G. Create an SCP that denies the creation of S3 buckets and DynamoDB tables that are not encrypted with the key.
- H. Create a key policy that allows the kms:Decrypt action only for Amazon S3, DynamoDB, Lambda, and Amazon EKS
- I. Create an AWS Config rule to send alerts for resources that are not encrypted with the key.

Answer: B

NEW QUESTION 35

A company used a lift-and-shift approach to migrate from its on-premises data centers to the AWS Cloud. The company migrated on-premises VMS to Amazon EC2 instances. Now the company wants to replace some of the components that are running on the EC2 instances with managed AWS services that provide similar functionality.

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

Which solution will meet these requirements?

- A. Create an Amazon CloudWatch Logs log group
- B. Configure the load balancers to send logs to the log group
- C. Use the CloudWatch Logs console to search the log
- D. Create CloudWatch Logs filters on the logs for the required metrics.
- E. Create an Amazon S3 bucket
- F. Configure the load balancers to send logs to the S3 bucket
- G. Use Amazon Athena to search the logs that are in the S3 bucket
- H. Create Amazon CloudWatch filters on the S3 log files for the required metrics.
- I. Create an Amazon S3 bucket

- J. Configure the load balancers to send logs to the S3 bucket
- K. Use Amazon Athena to search the logs that are in the S3 bucket
- L. Create Athena queries for the required metric
- M. Publish the metrics to Amazon CloudWatch.
- N. Create an Amazon CloudWatch Logs log group
- O. Configure the load balancers to send logs to the log group
- P. Use the AWS Management Console to search the log
- Q. Create Amazon Athena queries for the required metric
- R. Publish the metrics to Amazon CloudWatch.

Answer: C

Explanation:

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

NEW QUESTION 39

A company that uses AWS Organizations wants to see AWS Security Hub findings for many AWS accounts and AWS Regions. Some of the accounts are in the company's organization, and some accounts are in organizations that the company manages for customers. Although the company can see findings in the Security Hub administrator account for accounts in the company's organization, there are no findings from accounts in other organizations. Which combination of steps should the company take to see findings from accounts that are outside the organization that includes the Security Hub administrator account? (Select TWO.)

- A. Use a designated administration account to automatically set up member accounts.
- B. Create the AWS Service Role For Security Hub service-linked role for Security Hub.
- C. Send an administration request from the member accounts.
- D. Enable Security Hub for all member accounts.
- E. Send invitations to accounts that are outside the company's organization from the Security Hub administrator account.

Answer: CE

Explanation:

To see Security Hub findings for accounts that are outside the organization that includes the Security Hub administrator account, the following steps are required:

- Send invitations to accounts that are outside the company's organization from the Security Hub administrator account. This will allow the administrator account to view and manage findings from those accounts. The administrator account can send invitations by using the Security Hub console, API, or CLI. For more information, see [Sending invitations to member accounts](#).
- Send an administration request from the member accounts. This will allow the member accounts to accept the invitation from the administrator account and establish a relationship with it. The member accounts can send administration requests by using the Security Hub console, API, or CLI. For more information, see [Sending administration requests](#).

This solution will enable the company to see Security Hub findings for many AWS accounts and AWS Regions, including accounts that are outside its own organization.

The other options are incorrect because they either do not establish a relationship between the administrator and member accounts (A, B), do not enable Security Hub for all member accounts (D), or do not use a valid service for Security Hub (F).

Verified References:

- <https://docs.aws.amazon.com/securityhub/latest/userguide/securityhub-member-accounts.html>

NEW QUESTION 43

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

- A. In CloudTrail, turn on Insights events on the trail
- B. Configure an alarm on the insight with eventName matching ConsoleLogin and errorMessage matching "Failed authentication". Configure a threshold of 3 and a period of 5 minutes.
- C. Configure CloudTrail to send events to Amazon CloudWatch Log
- D. Create a metric filter for the relevant log group
- E. Create a filter pattern with eventName matching ConsoleLogin and errorMessage matching "Failed authentication". Create a CloudWatch alarm with a threshold of 3 and a period of 5 minutes.
- F. Create an Amazon Athena table from the CloudTrail event
- G. Run a query for eventName matching ConsoleLogin and for errorMessage matching "Failed authentication". Create a notification action from the query to send

an Amazon Simple Notification Service (Amazon SNS) notification when the count equals 3 within a period of 5 minutes.

H. In AWS Identity and Access Management Access Analyzer, create a new analyze

I. Configure the analyzer to send an Amazon Simple Notification Service (Amazon SNS) notification when a failed sign-in event occurs 3 times for any IAM user within a period of 5 minutes.

Answer: B

Explanation:

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

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

The other options are incorrect because:

> A. Turning on Insights events on the trail and configuring an alarm on the insight is not a solution, because Insights events are used to analyze unusual activity in management events, such as spikes in API call volume or error rates. Insights events do not capture failed sign-in attempts to the AWS Management Console³.

> C. Creating an Amazon Athena table from the CloudTrail events and running a query for failed sign-in events is not a solution, because it does not provide a mechanism to send an alert based on the query results. Amazon Athena is an interactive query service that allows analyzing data in Amazon S3 using standard SQL, but it does not support creating notifications or alarms from queries⁴.

> D. Creating an analyzer in AWS Identity and Access Management Access Analyzer and configuring it to send an Amazon SNS notification when a failed sign-in event occurs 3 times for any IAM user within a period of 5 minutes is not a solution, because IAM Access Analyzer is not a service that monitors sign-in events, but a service that helps identify resources that are shared with external entities. IAM Access Analyzer does not generate findings for failed sign-in attempts to the AWS Management Console⁵.

References:

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

NEW QUESTION 44

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

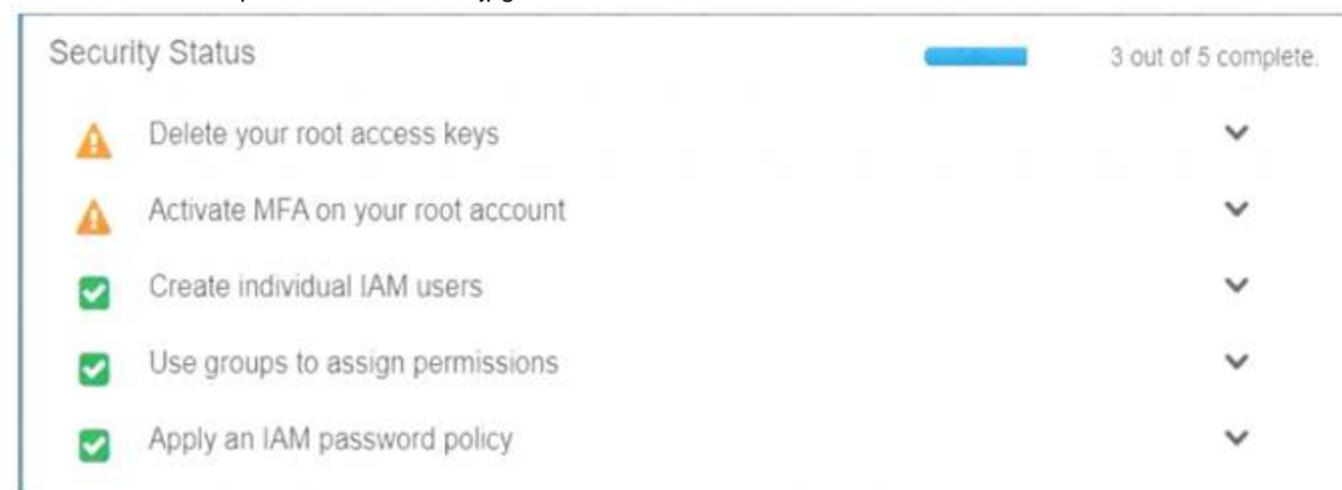
- A. Use short but complex password on the root account and any administrators.
- B. Use IAM Geo-Lock and disallow anyone from logging in except for in your city.
- C. Use MFA on all users and accounts, especially on the root account.
- D. Don't write down or remember the root account password after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

Answer: C

Explanation:

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

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



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

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

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

NEW QUESTION 47

A security engineer is designing an IAM policy to protect AWS API operations. The policy must enforce multi-factor authentication (MFA) for IAM users to access certain services in the AWS production account. Each session must remain valid for only 2 hours. The current version of the IAM policy is as follows:

```
{
  "Version": "2012-10-17",
  "Statement": [{
    "Effect": "Allow",
    "Action": [
      "ec2:DescribeInstances",
      "ec2:StopInstances",
      "ec2:TerminateInstances"
    ],
    "Resource": ["*"]
  }]
}
```

Which combination of conditions must the security engineer add to the IAM policy to meet these requirements? (Select TWO.)

- A. "Bool" : { "aws : Multi FactorAuthPresent": "true" }
- B. "B001" : { "aws : MultiFactorAuthPresent": "false" }
- C. "NumericLessThan" : { "aws : Multi FactorAuthAge" : "7200" }
- D. "NumericGreaterThan" : { "aws : MultiFactorAuthAge" : "7200" }
- E. "NumericLessThan" : { "MaxSessionDuration" : "7200" }

Answer: AC

Explanation:

The correct combination of conditions to add to the IAM policy is A and C. These conditions will ensure that IAM users must use MFA to access certain services in the AWS production account, and that each session will expire after 2 hours.

- > Option A: "Bool" : { "aws:MultiFactorAuthPresent" : "true" } is a valid condition that checks if the principal (the IAM user) has authenticated with MFA before making the request. This condition will enforce MFA for the IAM users to access the specified services. This condition key is supported by all AWS services that support IAM policies1.
- > Option B: "Bool" : { "aws:MultiFactorAuthPresent" : "false" } is the opposite of option A. This condition will allow access only if the principal has not authenticated with MFA, which is not the desired requirement. This condition key is supported by all AWS services that support IAM policies1.
- > Option C: "NumericLessThan" : { "aws:MultiFactorAuthAge" : "7200" } is a valid condition that checks if the time since the principal authenticated with MFA is less than 7200 seconds (2 hours). This condition will enforce the session duration limit for the IAM users. This condition key is supported by all AWS services that support IAM policies1.
- > Option D: "NumericGreaterThan" : { "aws:MultiFactorAuthAge" : "7200" } is the opposite of option C. This condition will allow access only if the time since the principal authenticated with MFA is more than 7200 seconds (2 hours), which is not the desired requirement. This condition key is supported by all AWS services that support IAM policies1.
- > Option E: "NumericLessThan" : { "MaxSessionDuration" : "7200" } is not a valid condition key.

MaxSessionDuration is a property of an IAM role, not a condition key. It specifies the maximum session duration (in seconds) for the role, which can be between 3600 and 43200 seconds (1 to 12 hours). This property can be set when creating or modifying a role, but it cannot be used as a condition in a policy2.

NEW QUESTION 52

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

Please select:

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

Answer: C

Explanation:

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

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

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

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

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

NEW QUESTION 53

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

Specifically, the security team wants EventBridge to watch for the s3:PutObjectAcl, s3>DeleteBucketPolicy, and s3:PutBucketPolicy API invocation logs from

CloudTrail. While developing the solution in a single account, the security team discovers that the s3:PutObjectAcl API call does not invoke an EventBridge event. However, the s3:DeleteBucketPolicy API call and the s3:PutBucketPolicy API call do invoke an event.

The security team has enabled CloudTrail for AWS management events with a basic configuration in the AWS Region in which EventBridge is being tested. Verification of the EventBridge event pattern indicates that the pattern is set up correctly. The security team must implement a solution so that the s3:PutObjectAcl API call will invoke an EventBridge event. The solution must not generate false notifications. Which solution will meet these requirements?

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

Answer: D

Explanation:

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

By default, trails do not log data events. To record CloudTrail data events, you must explicitly add the supported resources or resource types for which you want to collect activity. For more information, see Logging data events in the Amazon S3 User Guide². In this case, the security team wants EventBridge to watch for the s3:PutObjectAcl API invocation logs from CloudTrail. This API uses the acl subresource to set the access control list (ACL) permissions for a new or existing object in an S3 bucket³. This is a data event that affects the S3 object resource type. Therefore, the security team must enable CloudTrail to monitor data events for read and write operations to S3 buckets in order to invoke an EventBridge event for this API call. The other options are incorrect because:

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

References:

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

NEW QUESTION 55

A Network Load Balancer (NLB) target instance is not entering the InService state. A security engineer determines that health checks are failing. Which factors could cause the health check failures? (Select THREE.)

- A. The target instance's security group does not allow traffic from the NLB.
- B. The target instance's security group is not attached to the NLB.
- C. The NLB's security group is not attached to the target instance.
- D. The target instance's subnet network ACL does not allow traffic from the NLB.
- E. The target instance's security group is not using IP addresses to allow traffic from the NLB.
- F. The target network ACL is not attached to the NLB.

Answer: ACD

NEW QUESTION 58

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

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

Answer: C

Explanation:

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

NEW QUESTION 60

A company is implementing a new application in a new IAM account. A VPC and subnets have been created for the application. The application has been peered to an existing VPC in another account in the same IAM Region for database access. Amazon EC2 instances will regularly be created and terminated in the application VPC, but only some of them will need access to the databases in the peered VPC over TCP port 1521. A security engineer must ensure that only the EC2 instances that need access to the databases can access them through the network.

How can the security engineer implement this solution?

- A. Create a new security group in the database VPC and create an inbound rule that allows all traffic from the IP address range of the application VP
- B. Add a new network ACL rule on the database subnet
- C. Configure the rule to TCP port 1521 from the IP address range of the application VP
- D. Attach the new security group to the database instances that the application instances need to access.
- E. Create a new security group in the application VPC with an inbound rule that allows the IP address range of the database VPC over TCP port 1521. Create a new security group in the database VPC with an inbound rule that allows the IP address range of the application VPC over port 1521. Attach the new security group to the database instances and the application instances that need database access.
- F. Create a new security group in the application VPC with no inbound rule
- G. Create a new security group in the database VPC with an inbound rule that allows TCP port 1521 from the new application security group in the application VP
- H. Attach the application security group to the application instances that need database access, and attach the database security group to the database instances.
- I. Create a new security group in the application VPC with an inbound rule that allows the IP address range of the database VPC over TCP port 1521. Add a new network ACL rule on the database subnet
- J. Configure the rule to allow all traffic from the IP address range of the application VP
- K. Attach the new security group to the application instances that need database access.

Answer: C

NEW QUESTION 61

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 64

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

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

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

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

Answer: C

Explanation:

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

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

The other options are incorrect for the following reasons:

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

NEW QUESTION 67

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

Which solution will meet this requirement?

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

Answer: B

Explanation:

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

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

NEW QUESTION 72

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 74

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

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

Which solution will meet these requirements?

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

Answer: A

NEW QUESTION 79

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

Which solution will meet these requirements?

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

Answer: D

Explanation:

The correct answer is D.

To design a solution that will proactively monitor for suspicious behavior across all the accounts that contain production workloads, the security engineer needs to use a service that can aggregate and analyze security findings from multiple sources. AWS Security Hub is a service that provides a comprehensive view of your security posture across your AWS accounts and enables you to check your environment against security standards and best practices. Security Hub also integrates with other AWS services, such as Amazon GuardDuty, AWS Config, and AWS Systems Manager, to collect and correlate security findings.

To automate remediation of incidents across the production accounts, the security engineer needs to use a service that can trigger actions based on events. Amazon EventBridge is a serverless event bus service that allows you to connect your applications with data from a variety of sources. EventBridge can use rules to match events and route them to targets for processing. You can use EventBridge to invoke a custom AWS Lambda function from the Security Hub findings.

Lambda is a serverless compute service that lets you run code without provisioning or managing servers.

To publish a notification to an Amazon SNS topic when a critical security finding is detected, the security engineer needs to use a service that can send messages to subscribers. Amazon SNS is a fully managed messaging service that enables you to decouple and scale microservices, distributed systems, and serverless applications. SNS can deliver messages to a variety of endpoints, such as email, SMS, or HTTP. You can configure the Lambda function to also publish notifications to the SNS topic.

To send all security incident logs to a dedicated account, the security engineer needs to use a service that can aggregate and store log data from multiple sources. AWS Security Hub allows you to aggregate security findings from multiple accounts into a single account using the delegated administrator feature. This feature enables you to designate an AWS account as the administrator for Security Hub in an organization. The administrator account can then view and manage Security Hub findings from all member accounts.

Therefore, option D is correct because it meets all the requirements of the solution. Option A is incorrect because GuardDuty does not provide a comprehensive view of your security posture across your AWS accounts. GuardDuty is primarily a threat detection service that monitors for malicious or unauthorized behavior. Option B is incorrect because Config and Systems Manager are not designed to automate remediation of incidents based on Security Hub findings. Config is a service that enables you to assess, audit, and evaluate the configurations of your AWS resources, while Systems Manager is a service that allows you to manage your infrastructure on AWS at scale. Option C is incorrect because GuardDuty does not provide a comprehensive view of your security posture across your AWS accounts.

References:

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

NEW QUESTION 80

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

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

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

- A. Create a new CloudTrail trail
- B. Select the new Regions where the company added resources.

- C. Change the S3 bucket to receive notifications to track all actions from all Regions.
- D. Create a new CloudTrail trail that applies to all Regions.
- E. Change the existing CloudTrail trail so that it applies to all Regions.

Answer: D

Explanation:

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

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

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

NEW QUESTION 81

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

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

Answer: B

Explanation:

The correct answer is B. Store the API key value in AWS Secrets Manager. In the template, replace all references to the value with `{{resolve:secretsmanager:MySecretId:SecretString}}`.

This answer is correct because AWS Secrets Manager is a service that helps you protect secrets that are needed to access your applications, services, and IT resources. You can store and manage secrets such as database credentials, API keys, and other sensitive data in Secrets Manager. You can also use Secrets Manager to rotate, manage, and retrieve your secrets throughout their lifecycle¹. Secrets Manager integrates with AWS CloudFormation, which allows you to reference secrets from your templates using the `{{resolve:secretsmanager:...}}` syntax². This way, you can avoid exposing your secrets in plaintext and still use them in your resources.

The other options are incorrect because:

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

References:

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

NEW QUESTION 84

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

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

What should the Security Engineer do to troubleshoot this issue?

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

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

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

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

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

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

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

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

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

Answer: D

NEW QUESTION 87

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 Manager.

* 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 90

A company needs a security engineer to implement a scalable solution for multi-account authentication and authorization. The solution should not introduce additional user-managed architectural components. Native IAM features should be used as much as possible. The security engineer has set up IAM Organizations with all features activated and IAM SSO enabled.

Which additional steps should the security engineer take to complete the task?

- A. Use AD Connector to create users and groups for all employees that require access to IAM accounts. Assign AD Connector groups to IAM accounts and link to the IAM roles in accordance with the employees' job functions and access requirements. Instruct employees to access IAM accounts by using the IAM Directory Service user portal.
- B. Use an IAM SSO default directory to create users and groups for all employees that require access to IAM account.
- C. Assign groups to IAM accounts and link to permission sets in accordance with the employees' job functions and access requirement.
- D. Instruct employees to access IAM accounts by using the IAM SSO user portal.
- E. Use an IAM SSO default directory to create users and groups for all employees that require access to IAM account.
- F. Link IAM SSO groups to the IAM users present in all accounts to inherit existing permission.
- G. Instruct employees to access IAM accounts by using the IAM SSO user portal.
- H. Use IAM Directory Service for Microsoft Active Directory to create users and groups for all employees that require access to IAM accounts. Enable IAM Management Console access in the created directory and specify IAM SSO as a source of information for integrated accounts and permission set.
- I. Instruct employees to access IAM accounts by using the IAM Directory Service user portal.

Answer: B

NEW QUESTION 91

A Systems Engineer is troubleshooting the connectivity of a test environment that includes a virtual security appliance deployed inline. In addition to using the virtual security appliance, the Development team wants to use security groups and network ACLs to accomplish various security requirements in the environment. What configuration is necessary to allow the virtual security appliance to route the traffic?

- A. Disable network ACLs.
- B. Configure the security appliance's elastic network interface for promiscuous mode.
- C. Disable the Network Source/Destination check on the security appliance's elastic network interface.
- D. Place the security appliance in the public subnet with the internet gateway.

Answer: C

Explanation:

<https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/using-eni.html#eni-basics> Source/destination checking "You must disable source/destination checks if the instance runs services such as network address translation, routing, or firewalls."

The correct answer is C. Disable the Network Source/Destination check on the security appliance's elastic network interface.

This answer is correct because disabling the Network Source/Destination check allows the virtual security appliance to route traffic that is not addressed to or from itself. By default, this check is enabled on all EC2 instances, and it prevents them from forwarding traffic that does not match their own IP or MAC addresses. However, for a virtual security appliance that acts as a router or a firewall, this check needs to be disabled, otherwise it will drop the traffic that it is supposed to route¹².

The other options are incorrect because:

- > A. Disabling network ACLs is not a solution, because network ACLs are optional layers of security for the subnets in a VPC. They can be used to allow or deny traffic based on IP addresses and ports, but they do not affect the routing behavior of the virtual security appliance³.
- > B. Configuring the security appliance's elastic network interface for promiscuous mode is not a solution, because promiscuous mode is a mode for a network interface that causes it to pass all traffic it receives to the CPU, rather than passing only the frames that it is programmed to receive. Promiscuous mode is normally used for packet sniffing or monitoring, but it does not enable the network interface to route traffic⁴.
- > D. Placing the security appliance in the public subnet with the internet gateway is not a solution, because it does not address the routing issue of the virtual security appliance. The security appliance can be placed in either a public or a private subnet, depending on the network design and security requirements, but it still needs to have the Network Source/Destination check disabled to route traffic properly⁵.

References:

1: Enabling or disabling source/destination checks - Amazon Elastic Compute Cloud 2: Virtual security appliance - Wikipedia 3: Network ACLs - Amazon Virtual Private Cloud 4: Promiscuous mode - Wikipedia 5: NAT instances - Amazon Virtual Private Cloud

NEW QUESTION 95

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 98

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 100

A business requires a forensic logging solution for hundreds of Docker-based apps running on Amazon EC2. The solution must analyze logs in real time, provide message replay, and persist logs.

Which Amazon Web Offerings (IAM) services should be employed to satisfy these requirements? (Select two.)

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

Answer: BD

NEW QUESTION 103

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

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

Answer: CE

NEW QUESTION 106

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

What should the security engineer recommend?

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

Answer: C

Explanation:

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

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

NEW QUESTION 107

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

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

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

- F. Attach the trust policy to the role.
- G. Establish a trust relationship between the IAM user and the AWS account that contains the resources.
- H. Create a role in the IAM user's AWS account
- I. Create an identity policy that allows the sts: AssumeRole action
- J. Attach the identity policy to the role.

Answer: BC

Explanation:

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

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

Verified References:

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

NEW QUESTION 112

A security engineer needs to set up an Amazon CloudFront distribution for an Amazon S3 bucket that hosts a static website. The security engineer must allow only specified IP addresses to access the website. The security engineer also must prevent users from accessing the website directly by using S3 URLs. Which solution will meet these requirements?

- A. Generate an S3 bucket policy
- B. Specify cloudfront.amazonaws.com as the principal
- C. Use the aws:SourceIp condition key to allow access only if the request comes from the specified IP addresses.
- D. Create a CloudFront origin access identity (OAI). Create the S3 bucket policy so that only the OAI has access
- E. Create an AWS WAF web ACL and add an IP set rule
- F. Associate the web ACL with the CloudFront distribution.
- G. Implement security groups to allow only the specified IP addresses access and to restrict S3 bucket access by using the CloudFront distribution.
- H. Create an S3 bucket access point to allow access from only the CloudFront distribution
- I. Create an AWS WAF web ACL and add an IP set rule
- J. Associate the web ACL with the CloudFront distribution.

Answer: B

NEW QUESTION 115

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 groups.
- 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 packages.

NEW QUESTION 119

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 templates. When using a dynamic reference, CloudFormation retrieves the value of the specified reference when necessary during stack and change set

operations¹. Dynamic references can be used for certain resources that support them, such as `AWS::RDS::DBInstance`¹. By using a dynamic reference to reference the database credentials in Secrets Manager, the company can leverage the existing integration between these services and avoid hardcoding the secret information in the template. Secrets Manager is a service that helps you protect secrets needed to access your applications, services, and IT resources². Secrets Manager enables you to rotate, manage, and retrieve database credentials, API keys, and other secrets throughout their lifecycle².

NEW QUESTION 124

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

The company is performing control tests on specific GuardDuty findings to make sure that the company's security team can detect and respond to security events. The security team launched an Amazon EC2 instance and attempted to run DNS requests against a test domain, `example.com`, to generate a DNS finding.

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

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

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

Answer: C

Explanation:

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

According to the AWS documentation¹, GuardDuty findings are automatically sent to Security Hub only if the GuardDuty integration with Security Hub is enabled in the same account and Region. This means that the security tooling account, which is the delegated administrator for both GuardDuty and Security Hub, must enable the GuardDuty integration with Security Hub in each member account and Region where GuardDuty is enabled. Otherwise, the findings from GuardDuty will not be visible in Security Hub.

The other options are incorrect because:

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

References:

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

NEW QUESTION 127

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

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

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

Answer: DE

Explanation:

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

NEW QUESTION 129

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 130

A company is using Amazon Route 53 Resolver for its hybrid DNS infrastructure. The company has set up Route 53 Resolver forwarding rules for authoritative domains that are hosted on on-premises DNS servers.

A new security mandate requires the company to implement a solution to log and query DNS traffic that goes to the on-premises DNS servers. The logs must show details of the source IP address of the instance from which the query originated. The logs also must show the DNS name that was requested in Route 53 Resolver. Which solution will meet these requirements?

- A. Use VPC Traffic Mirroring
- B. Configure all relevant elastic network interfaces as the traffic source, include amazon-dns in the mirror filter, and set Amazon CloudWatch Logs as the mirror target
- C. Use CloudWatch Insights on the mirror session logs to run queries on the source IP address and DNS name.
- D. Configure VPC flow logs on all relevant VPC
- E. Send the logs to an Amazon S3 bucket
- F. Use Amazon Athena to run SQL queries on the source IP address and DNS name.
- G. Configure Route 53 Resolver query logging on all relevant VPC
- H. Send the logs to Amazon CloudWatch Log
- I. Use CloudWatch Insights to run queries on the source IP address and DNS name.
- J. Modify the Route 53 Resolver rules on the authoritative domains that forward to the on-premises DNS server
- K. Send the logs to an Amazon S3 bucket
- L. Use Amazon Athena to run SQL queries on the source IP address and DNS name.

Answer: C

Explanation:

The correct answer is C. Configure Route 53 Resolver query logging on all relevant VPCs. Send the logs to Amazon CloudWatch Logs. Use CloudWatch Insights to run queries on the source IP address and DNS name.

According to the AWS documentation¹, Route 53 Resolver query logging lets you log the DNS queries that Route 53 Resolver handles for your VPCs. You can send the logs to CloudWatch Logs, Amazon S3, or Kinesis Data Firehose. The logs include information such as the following:

- The AWS Region where the VPC was created
- The ID of the VPC that the query originated from
- The IP address of the instance that the query originated from
- The instance ID of the resource that the query originated from
- The date and time that the query was first made
- The DNS name requested (such as prod.example.com)
- The DNS record type (such as A or AAAA)
- The DNS response code, such as NoError or ServFail
- The DNS response data, such as the IP address that is returned in response to the DNS query

You can use CloudWatch Insights to run queries on your log data and analyze the results using graphs and statistics². You can filter and aggregate the log data based on any field, and use operators and functions to perform calculations and transformations. For example, you can use CloudWatch Insights to find out how many queries were made for a specific domain name, or which instances made the most queries.

Therefore, this solution meets the requirements of logging and querying DNS traffic that goes to the on-premises DNS servers, showing details of the source IP address of the instance from which the query originated, and the DNS name that was requested in Route 53 Resolver.

The other options are incorrect because:

- A. Using VPC Traffic Mirroring would not capture the DNS queries that go to the on-premises DNS servers, because Traffic Mirroring only copies network traffic from an elastic network interface of an EC2 instance to a target for analysis³. Traffic Mirroring does not include traffic that goes through a Route 53 Resolver outbound endpoint, which is used to forward queries to on-premises DNS servers⁴. Therefore, this solution would not meet the requirements.
- B. Configuring VPC flow logs on all relevant VPCs would not capture the DNS name that was requested in Route 53 Resolver, because flow logs only record information about the IP traffic going to and from network interfaces in a VPC⁵. Flow logs do not include any information about the content or payload of a packet,

such as a DNS query or response. Therefore, this solution would not meet the requirements.

➤ D. Modifying the Route 53 Resolver rules on the authoritative domains that forward to the on-premises DNS servers would not enable logging of DNS queries, because Resolver rules only specify how to forward queries for specified domain names to your network. Resolver rules do not have any logging functionality by themselves. Therefore, this solution would not meet the requirements. References:

1: Resolver query logging - Amazon Route 53 2: Analyzing log data with CloudWatch Logs Insights - Amazon CloudWatch 3: What is Traffic Mirroring? - Amazon Virtual Private Cloud 4: Outbound Resolver endpoints - Amazon Route 53 5: Logging IP traffic using VPC Flow Logs - Amazon Virtual Private Cloud 6: Managing forwarding rules - Amazon Route 53

NEW QUESTION 134

A company is using Amazon Elastic Container Service (Amazon ECS) to deploy an application that deals with sensitive data. During a recent security audit, the company identified a security issue in which Amazon RDS credentials were stored with the application code in the company's source code repository. A security engineer needs to develop a solution to ensure that database credentials are stored securely and rotated periodically. The credentials should be accessible to the application only. The engineer also needs to prevent database administrators from sharing database credentials as plaintext with other teammates. The solution must also minimize administrative overhead. Which solution meets these requirements?

- A. Use the IAM Systems Manager Parameter Store to generate database credentials.
- B. Use an IAM profile for ECS tasks to restrict access to database credentials to specific containers only.
- C. Use IAM Secrets Manager to store database credentials.
- D. Use an IAM inline policy for ECS tasks to restrict access to database credentials to specific containers only.
- E. Use the IAM Systems Manager Parameter Store to store database credentials.
- F. Use IAM roles for ECS tasks to restrict access to database credentials to specific containers only.
- G. Use IAM Secrets Manager to store database credentials.
- H. Use IAM roles for ECS tasks to restrict access to database credentials to specific containers only.

Answer: D

Explanation:

To ensure that database credentials are stored securely and rotated periodically, the security engineer should do the following:

- Use AWS Secrets Manager to store database credentials. This allows the security engineer to encrypt and manage secrets centrally, and to configure automatic rotation schedules for them.
- Use IAM roles for ECS tasks to restrict access to database credentials to specific containers only. This allows the security engineer to grant fine-grained permissions to ECS tasks based on their roles, and to avoid sharing credentials as plaintext with other teammates.

NEW QUESTION 137

.....

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/>