

Amazon-Web-Services

Exam Questions SCS-C02

AWS Certified Security - Specialty



NEW QUESTION 1

A security engineer needs to develop a process to investigate and respond to potential security events on a company's Amazon EC2 instances. All the EC2 instances are backed by Amazon Elastic Block Store (Amazon EBS). The company uses AWS Systems Manager to manage all the EC2 instances and has installed Systems Manager Agent (SSM Agent) on all the EC2 instances.

The process that the security engineer is developing must comply with AWS security best practices and must meet the following requirements:

- A compromised EC2 instance's volatile memory and non-volatile memory must be preserved for forensic purposes.
- A compromised EC2 instance's metadata must be updated with corresponding incident ticket information.
- A compromised EC2 instance must remain online during the investigation but must be isolated to prevent the spread of malware.
- Any investigative activity during the collection of volatile data must be captured as part of the process. Which combination of steps should the security engineer take to meet these requirements with the LEAST operational overhead? (Select THREE.)

- A. Gather any relevant metadata for the compromised EC2 instance
- B. Enable termination protection
- C. Isolate the instance by updating the instance's security groups to restrict access
- D. Detach the instance from any Auto Scaling groups that the instance is a member of
- E. Deregister the instance from any Elastic Load Balancing (ELB) resources.
- F. Gather any relevant metadata for the compromised EC2 instance
- G. Enable termination protection
- H. Move the instance to an isolation subnet that denies all source and destination traffic
- I. Associate the instance with the subnet to restrict access
- J. Detach the instance from any Auto Scaling groups that the instance is a member of
- K. Deregister the instance from any Elastic Load Balancing (ELB) resources.
- L. Use Systems Manager Run Command to invoke scripts that collect volatile data.
- M. Establish a Linux SSH or Windows Remote Desktop Protocol (RDP) session to the compromised EC2 instance to invoke scripts that collect volatile data.
- N. Create a snapshot of the compromised EC2 instance's EBS volume for follow-up investigation
- O. Tag the instance with any relevant metadata and incident ticket information.
- P. Create a Systems Manager State Manager association to generate an EBS volume snapshot of the compromised EC2 instance
- Q. Tag the instance with any relevant metadata and incident ticket information.

Answer: ACE

NEW QUESTION 2

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 3

A company's security team is building a solution for logging and visualization. The solution will assist the company with the large variety and velocity of data that it receives from IAM across multiple accounts. The security team has enabled IAM CloudTrail and VPC Flow Logs in all of its accounts. In addition, the company has an organization in IAM Organizations and has an IAM Security Hub master account.

The security team wants to use Amazon Detective. However, the security team cannot enable Detective and is unsure why. What must the security team do to enable Detective?

- A. Enable Amazon Macie so that Security Hub will allow Detective to process findings from Macie.
- B. Disable IAM Key Management Service (IAM KMS) encryption on CloudTrail logs in every member account of the organization.
- C. Enable Amazon GuardDuty on all member accounts. Try to enable Detective in 48 hours.
- D. Ensure that the principal that launches Detective has the organizations ListAccounts permission.

Answer: D

NEW QUESTION 4

A company hosts an end-user application on AWS. Currently, the company deploys the application on Amazon EC2 instances behind an Elastic Load Balancer. The company wants to configure end-to-end encryption between the Elastic Load Balancer and the EC2 instances.

Which solution will meet this requirement with the LEAST operational effort?

- A. Use Amazon-issued AWS Certificate Manager (ACM) certificates on the EC2 instances and the Elastic Load Balancer to configure end-to-end encryption.
- B. Import a third-party SSL certificate to AWS Certificate Manager (ACM). Install the third-party certificate on the EC2 instances. Associate the ACM-imported third-party certificate with the Elastic Load Balancer.
- C. Deploy AWS CloudHSM. Import a third-party certificate. Configure the EC2 instances and the Elastic Load Balancer to use the CloudHSM-imported certificate.
- D. Import a third-party certificate bundle to AWS Certificate Manager (ACM). Install the third-party certificate on the EC2 instances. Associate the ACM-imported third-party certificate with the Elastic Load Balancer.

Answer: A

Explanation:

To configure end-to-end encryption between the Elastic Load Balancer and the EC2 instances with the least operational effort, the most appropriate solution would be to use Amazon-issued AWS Certificate Manager (ACM) certificates on the EC2 instances and the Elastic Load Balancer to configure end-to-end encryption.

AWS Certificate Manager - Amazon Web Services : Elastic Load Balancing - Amazon Web

Services : Amazon Elastic Compute Cloud - Amazon Web Services : AWS Certificate Manager - Amazon Web Services

NEW QUESTION 5

A Security Architect has been asked to review an existing security architecture and identify why the application servers cannot successfully initiate a connection to the database servers. The following summary describes the architecture:

- * 1. An Application Load Balancer, an internet gateway, and a NAT gateway are configured in the public subnet.
- * 2. Database, application, and web servers are configured on three different private subnets.
- * 3. The VPC has two route tables: one for the public subnet and one for all other subnets. The route table for the public subnet has a 0.0.0.0/0 route to the internet gateway. The route table for all other subnets has a 0.0.0.0/0 route to the NAT gateway. All private subnets can route to each other.
- * 4. Each subnet has a network ACL implemented that limits all inbound and outbound connectivity to only the required ports and protocols.
- * 5. There are 3 Security Groups (SGs): database, application, and web. Each group limits all inbound and outbound connectivity to the minimum required.

Which of the following accurately reflects the access control mechanisms the Architect should verify?

- A. Outbound SG configuration on database servers. Inbound SG configuration on application servers. Inbound and outbound network ACL configuration on the database subnet. Inbound and outbound network ACL configuration on the application server subnet.
- B. Inbound SG configuration on database servers. Outbound SG configuration on application servers. Inbound and outbound network ACL configuration on the database subnet. Inbound and outbound network ACL configuration on the application server subnet.
- C. Inbound and outbound SG configuration on database servers. Inbound and outbound SG configuration on application servers. Inbound network ACL configuration on the database subnet. Outbound network ACL configuration on the application server subnet.
- D. Inbound SG configuration on database servers. Outbound SG configuration on application servers. Inbound network ACL configuration on the database subnet. Outbound network ACL configuration on the application server subnet.

Answer: A

Explanation:

This is the accurate reflection of the access control mechanisms that the Architect should verify. Access control mechanisms are methods that regulate who can access what resources and how. Security groups and network ACLs are two types of access control mechanisms that can be applied to EC2 instances and subnets. Security groups are stateful, meaning they remember and return traffic that was previously allowed. Network ACLs are stateless, meaning they do not remember or return traffic that was previously allowed. Security groups and network ACLs can have inbound and outbound rules that specify the source, destination, protocol, and port of the traffic. By verifying the outbound security group configuration on database servers, the inbound security group configuration on application servers, and the inbound and outbound network ACL configuration on both the database and application server subnets, the Architect can check if there are any misconfigurations or conflicts that prevent the application servers from initiating a connection to the database servers. The other options are either inaccurate or incomplete for verifying the access control mechanisms.

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 is using an AWS Key Management Service (AWS KMS) AWS owned key in its application to encrypt files in an AWS account. The company's security team wants the ability to change to new key material for new files whenever a potential key breach occurs. A security engineer must implement a solution that gives the security team the ability to change the key whenever the team wants to do so. Which solution will meet these requirements?

- A. Create a new customer managed key. Add a key rotation schedule to the key. Invoke the key rotation schedule every time the security team requests a key change.
- B. Create a new AWS managed key. Add a key rotation schedule to the key. Invoke the key rotation schedule every time the security team requests a key change.
- C. Create a key alias. Create a new customer managed key every time the security team requests a key change. Associate the alias with the new key.
- D. Create a key alias. Create a new AWS managed key every time the security team requests a key change. Associate the alias with the new key.

Answer: A

Explanation:

To meet the requirement of changing the key material for new files whenever a potential key breach occurs, the most appropriate solution would be to create a new customer managed key, add a key rotation schedule to the key, and invoke the key rotation schedule every time the security team requests a key change.

References: : Rotating AWS KMS keys - AWS Key Management Service

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 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 10

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

The company must create the break glass user. The company must log any activities of the break glass user and send the logs to a security team.

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

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

Answer: AE

Explanation:

The combination of solutions that will meet the requirements are:

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

The other options are incorrect because:

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

References:

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

NEW QUESTION 10

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

Which CMK-related problems possibly account for the error? (Select two.)

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

Answer: AD

Explanation:

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

NEW QUESTION 15

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

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

- A. Delete the access keys for the account root user in every account.
- B. Create an admin IAM user with administrative privileges and delete the account root user in every account.
- C. Implement a strong password to help protect account-level access to the IAM Management Console by the account root user.
- D. Enable multi-factor authentication (MFA) on every account root user in all accounts.

- E. Create a custom IAM policy to limit permissions to required actions for the account root user and attach the policy to the account root user.
- F. Attach an IAM role to the account root user to make use of the automated credential rotation in IAM STS.

Answer: ADE

Explanation:

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

NEW QUESTION 19

A security team is developing an application on an Amazon EC2 instance to get objects from an Amazon S3 bucket. All objects in the S3 bucket are encrypted with an AWS Key Management Service (AWS KMS) customer managed key. All network traffic for requests that are made within the VPC is restricted to the AWS infrastructure. This traffic does not traverse the public internet.

The security team is unable to get objects from the S3 bucket Which factors could cause this issue? (Select THREE.)

- A. The IAM instance profile that is attached to the EC2 instance does not allow the s3 ListBucket action to the S3: bucket in the AWS accounts.
- B. The IAM instance profile that is attached to the EC2 instance does not allow the s3 ListParts action to the S3; bucket in the AWS accounts.
- C. The KMS key policy that encrypts the object in the S3 bucket does not allow the kms; ListKeys action to the EC2 instance profile ARN.
- D. The KMS key policy that encrypts the object in the S3 bucket does not allow the kms Decrypt action to the EC2 instance profile ARN.
- E. The security group that is attached to the EC2 instance is missing an outbound rule to the S3 managed prefix list over port 443.
- F. The security group that is attached to the EC2 instance is missing an inbound rule from the S3 managed prefix list over port 443.

Answer: ADE

Explanation:

<https://docs.aws.amazon.com/vpc/latest/userguide/security-group-rules.html>

To get objects from an S3 bucket that are encrypted with a KMS customer managed key, the security team needs to have the following factors in place:

- The IAM instance profile that is attached to the EC2 instance must allow the s3:GetObject action to the S3 bucket or object in the AWS account. This permission is required to read the object from S3. Option A is incorrect because it specifies the s3:ListBucket action, which is only required to list the objects in the bucket, not to get them.
- The KMS key policy that encrypts the object in the S3 bucket must allow the kms:Decrypt action to the EC2 instance profile ARN. This permission is required to decrypt the object using the KMS key. Option D is correct.
- The security group that is attached to the EC2 instance must have an outbound rule to the S3 managed prefix list over port 443. This rule is required to allow HTTPS traffic from the EC2 instance to S3 within the AWS infrastructure. Option E is correct. Option B is incorrect because it specifies the s3:ListParts action, which is only required for multipart uploads, not for getting objects. Option C is incorrect because it specifies the kms:ListKeys action, which is not required for getting objects. Option F is incorrect because it specifies an inbound rule from the S3 managed prefix list, which is not required for getting objects. Verified References:
 - <https://docs.aws.amazon.com/kms/latest/developerguide/control-access.html>
 - <https://docs.aws.amazon.com/vpc/latest/userguide/vpc-endpoints-s3.html>

NEW QUESTION 21

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

- A. Install the Amazon CloudWatch agent on each EC2 instance in the VP
- B. Use the CloudWatch agent to stream the DNS query logs to an Amazon CloudWatch Logs log grou
- C. Use CloudWatch metric filters to automatically generate metrics that list the most common 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 loggin
- J. Add an Amazon CloudWatch Logs log group as the destinatio
- K. Use Amazon CloudWatch Contributor Insights to analyze the data and create time series that display the most common DNS queries.

Answer: D

Explanation:

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

NEW QUESTION 26

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 30

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 33

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 38

A company uses AWS Organizations to manage a small number of AWS accounts. However, the company plans to add 1,000 more accounts soon. The company allows only a centralized security team to create IAM roles for all AWS accounts and teams. Application teams submit requests for IAM roles to the security team. The security team has a backlog of IAM role requests and cannot review and provision the IAM roles quickly.

The security team must create a process that will allow application teams to provision their own IAM roles. The process must also limit the scope of IAM roles and prevent privilege escalation.

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

- A. Create an IAM group for each application team
- B. Associate policies with each IAM group
- C. Provision IAM users for each application team member
- D. Add the new IAM users to the appropriate IAM group by using role-based access control (RBAC).
- E. Delegate application team leads to provision IAM roles for each team

- F. Conduct a quarterly review of the IAM roles the team leads have provisioned
- G. Ensure that the application team leads have the appropriate training to review IAM roles.
- H. Put each AWS account in its own OU
- I. Add an SCP to each OU to grant access to only the AWS services that the teams plan to use
- J. Include conditions in the AWS account of each team.
- K. Create an SCP and a permissions boundary for IAM role
- L. Add the SCP to the root OU so that only roles that have the permissions boundary attached can create any new IAM roles.

Answer: D

Explanation:

To create a process that will allow application teams to provision their own IAM roles, while limiting the scope of IAM roles and preventing privilege escalation, the following steps are required:

➤ Create a service control policy (SCP) that defines the maximum permissions that can be granted to any IAM role in the organization. An SCP is a type of policy that you can use with AWS Organizations to manage permissions for all accounts in your organization. SCPs restrict permissions for entities in member accounts, including each AWS account root user, IAM users, and roles. For more information, see [Service control policies overview](#).

➤ Create a permissions boundary for IAM roles that matches the SCP. A permissions boundary is an advanced feature for using a managed policy to set the maximum permissions that an identity-based policy can grant to an IAM entity. A permissions boundary allows an entity to perform only the actions that are allowed by both its identity-based policies and its permissions boundaries. For more information, see [Permissions boundaries for IAM entities](#).

➤ Add the SCP to the root organizational unit (OU) so that it applies to all accounts in the organization.

This will ensure that no IAM role can exceed the permissions defined by the SCP, regardless of how it is created or modified.

➤ Instruct the application teams to attach the permissions boundary to any IAM role they create. This will prevent them from creating IAM roles that can escalate their own privileges or access resources they are not authorized to access.

This solution will meet the requirements with the least operational overhead, as it leverages AWS Organizations and IAM features to delegate and limit IAM role creation without requiring manual reviews or approvals.

The other options are incorrect because they either do not allow application teams to provision their own IAM roles (A), do not limit the scope of IAM roles or prevent privilege escalation (B), or do not take advantage of managed services whenever possible (C).

Verified References:

➤ https://docs.aws.amazon.com/IAM/latest/UserGuide/access_policies_boundaries.html

NEW QUESTION 40

A company created an IAM account for its developers to use for testing and learning purposes. Because the account will be shared among multiple teams of developers, the company wants to restrict the ability to stop and terminate Amazon EC2 instances so that a team can perform these actions only on the instances it owns.

Developers were instructed to tag all their instances with a Team tag key and use the team name in the tag value. One of the first teams to use this account is Business Intelligence. A security engineer needs to develop a highly scalable solution for providing developers with access to the appropriate resources within the account. The security engineer has already created individual IAM roles for each team.

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

A. For each team, create an IAM policy similar to the one that follows. Populate the `ec2:ResourceTag/Team` condition key with a proper team name. Attach resulting policies to the corresponding IAM roles.

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "NotAction": [
        "ec2:StopInstances",
        "ec2:TerminateInstances"
      ],
      "Resource": "*"
    },
    {
      "Effect": "Allow",
      "Action": [
        "ec2:StopInstances",
        "ec2:TerminateInstances"
      ],
      "Resource": "*",
      "Condition": {
        "StringEquals": {
          "ec2:ResourceTag/Team": "BusinessIntelligence"
        }
      }
    }
  ]
}
```

B. For each team create an IAM policy similar to the one that follows. Populate the `IAM:TagKeys/Team` condition key with a proper team name.

C. Attach the resulting policies to the corresponding IAM roles.


```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "NotAction": [
        "ec2:StopInstances",
        "ec2:TerminateInstances"
      ],
      "Resource": "*"
    },
    {
      "Effect": "Allow",
      "Action": [
        "ec2:StopInstances",
        "ec2:TerminateInstances"
      ],
      "Resource": "*",
      "Condition": {
        "ForAnyValue:StringEquals": {
          "aws:TagKeys/Team": "BusinessIntelligence"
        }
      }
    }
  ]
}
```

- D. Tag each IAM role with a Team tag key
 E. and use the team name in the tag value
 F. Create an IAM policy similar to the one that follows, and attach it to all the IAM roles used by developers.

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "NotAction": [
        "ec2:StopInstances",
        "ec2:TerminateInstances"
      ],
      "Resource": "*"
    },
    {
      "Effect": "Allow",
      "Action": [
        "ec2:StopInstances",
        "ec2:TerminateInstances"
      ],
      "Resource": "*",
      "Condition": {
        "StringEquals": {
          "ec2:ResourceTag/Team": "${aws:PrincipalTag/Team}"
        }
      }
    }
  ]
}
```

- G. Tag each IAM role with the Team key, and use the team name in the tag value
 H. Create an IAM policy similar to the one that follows, and attach it to all the IAM roles used by developers.

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "NotAction": [
        "ec2:StopInstances",
        "ec2:TerminateInstances"
      ],
      "Resource": "*"
    },
    {
      "Effect": "Allow",
      "Action": [
        "ec2:StopInstances",
        "ec2:TerminateInstances"
      ],
      "Resource": "*",
      "Condition": {
        "ForAnyValue:StringEquals": {
          "aws:TagKeys/Team": "${aws:PrincipalTag/Team}"
        }
      }
    }
  ]
}
```

Answer: A

NEW QUESTION 42

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

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

Answer: BDF

Explanation:

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

NEW QUESTION 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 50

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 53

A developer is building a serverless application hosted on AWS that uses Amazon Redshift as a data store. The application has separate modules for readwrite and read-only functionality. The modules need their own database users for compliance reasons. Which combination of steps should a security engineer implement to grant appropriate access? (Select TWO.)

- A. Configure cluster security groups for each application module to control access to database users that are required for read-only and readwrite
- B. Configure a VPC endpoint for Amazon Redshift. Configure an endpoint policy that maps database users to each application module, and allow access to the tables that are required for read-only and read/write
- C. Configure an IAM policy for each module. Specify the ARN of an Amazon Redshift database user that allows the GetClusterCredentials API call
- D. Create local database users for each module
- E. Configure an IAM policy for each module. Specify the ARN of an IAM user that allows the GetClusterCredentials API call

Answer: A

Explanation:

To grant appropriate access to separate modules for read-write and read-only functionality in a serverless application hosted on AWS that uses Amazon Redshift as a data store, a security engineer should configure cluster security groups for each application module to control access to database users that are required for read-only and readwrite, and configure an IAM policy for each module specifying the ARN of an IAM user that allows the GetClusterCredentials API call.

References: : Amazon Redshift - Amazon Web Services : Amazon Redshift - Amazon Web Services : Identity and Access Management - AWS Management Console : AWS Identity and Access Management - AWS Management Console

NEW QUESTION 55

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

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

Answer: A

Explanation:

The IAM Documentation mentions the following

To determine whether a log file was modified, deleted, or unchanged after CloudTrail delivered it, you can use CloudTrail log file integrity validation. This feature is

built using industry standard algorithms: SHA-256 for hashing and SHA-256 with RSA for digital signing. This makes it computationally infeasible to modify, delete or forge CloudTrail log files without detection. You can use the IAM CLI to validate the files in the location where CloudTrail delivered them. Validated log files are invaluable in security and forensic investigations. For example, a validated log file enables you to assert positively that the log file itself has not changed, or that particular user credentials performed specific API activity. The CloudTrail log file integrity validation process also lets you know if a log file has been deleted or changed, or assert positively that no log files were delivered to your account during a given period of time. Options B, C and D are invalid because you need to check for log File Integrity Validation for CloudTrail logs. For more information on CloudTrail log file validation, please visit the below URL: <http://docs.IAM.amazonaws.com/IAMcloudtrail/latest/userguide/cloudtrail-log-file-validation-intro.html>. The correct answer is: Use CloudTrail Log File Integrity Validation. omit your Feedback/Queries to our Expert

NEW QUESTION 59

Within a VPC, a corporation runs an Amazon RDS Multi-AZ DB instance. The database instance is connected to the internet through a NAT gateway via two subnets.

Additionally, the organization has application servers that are hosted on Amazon EC2 instances and use the RDS database. These EC2 instances have been deployed onto two more private subnets inside the same VPC. These EC2 instances connect to the internet through a default route via the same NAT gateway. Each VPC subnet has its own route table.

The organization implemented a new security requirement after a recent security examination. Never allow the database instance to connect to the internet. A security engineer must perform this update promptly without interfering with the network traffic of the application servers. How will the security engineer be able to comply with these requirements?

- A. Remove the existing NAT gateway
- B. Create a new NAT gateway that only the application server subnets can use.
- C. Configure the DB instance's inbound network ACL to deny traffic from the security group ID of the NAT gateway.
- D. Modify the route tables of the DB instance subnets to remove the default route to the NAT gateway.
- E. Configure the route table of the NAT gateway to deny connections to the DB instance subnets.

Answer: C

Explanation:

Each subnet has a route table, so modify the routing associated with DB instance subnets to prevent internet access.

NEW QUESTION 62

A startup company is using a single AWS account that has resources in a single AWS Region. A security engineer configures an AWS Cloud Trail 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 66

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

Which solution meets these criteria?

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

Answer: A

Explanation:

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

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

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

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

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

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

References:

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

NEW QUESTION 68

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

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

- A. Have a Database Administrator encrypt the credentials and store the ciphertext in Amazon S3. Grant permission to the instance role associated with the EC2 instance to read the object and decrypt the ciphertext.
- B. Configure a scheduled job that updates the credential in AWS Systems Manager Parameter Store and notifies the Engineer that the application needs to be restarted.
- C. Configure automatic rotation of credentials in AWS Secrets Manager.
- D. Store the credential in an encrypted string parameter in AWS Systems Manager Parameter Store.
- E. Grant permission to the instance role associated with the EC2 instance to access the parameter and the AWS KMS key that is used to encrypt it.
- F. Configure the Java application to catch a connection failure and make a call to AWS Secrets Manager to retrieve updated credentials when the password is rotated.
- G. Grant permission to the instance role associated with the EC2 instance to access Secrets Manager.

Answer: CE

Explanation:

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

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

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

NEW QUESTION 71

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 72

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

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

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

- A. Log in to the suspicious instance and use the netstat command to identify remote connections. Use the IP addresses from these remote connections to create deny rules in the security group of the instance. Install diagnostic tools on the instance for investigation. Update the outbound network ACL for the subnet in us-east-1b.

lb to explicitly deny all connections as the first rule during the investigation of the instance

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

C. Ensure that the Amazon Elastic Block Store (Amazon EBS) volumes that are attached to the suspicious EC2 instance will not delete upon termination

Terminate the instance Launch a new EC2 instance in us-east-1a that has diagnostic tools Mount the EBS volumes from the terminated instance for investigation

D. Create an AWS WAF web ACL that denies traffic to and from the suspicious instance Attach the AWS WAF web ACL to the instance to mitigate the attack Log in to the instance and install diagnostic tools to investigate the instance

Answer: B

Explanation:

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

NEW QUESTION 76

A company is undergoing a layer 3 and layer 4 DDoS attack on its web servers running on IAM.

Which combination of IAM services and features will provide protection in this scenario? (Select THREE).

- A. Amazon Route 53
- B. IAM Certificate Manager (ACM)
- C. Amazon S3
- D. IAM Shield
- E. Elastic Load Balancer
- F. Amazon GuardDuty

Answer: DEF

NEW QUESTION 80

A company's security engineer is developing an incident response plan to detect suspicious activity in an AWS account for VPC hosted resources. The security engineer needs to provide visibility for as many AWS Regions as possible.

Which combination of steps will meet these requirements MOST cost-effectively? (Select TWO.)

- A. Turn on VPC Flow Logs for all VPCs in the account.
- B. Activate Amazon GuardDuty across all AWS Regions.
- C. Activate Amazon Detective across all AWS Regions.
- D. Create an Amazon Simple Notification Service (Amazon SNS) topic
- E. Create an Amazon EventBridge rule that responds to findings and publishes the findings to the SNS topic.
- F. Create an AWS Lambda function
- G. Create an Amazon EventBridge rule that invokes the Lambda function to publish findings to Amazon Simple Email Service (Amazon SES).

Answer: BD

Explanation:

To detect suspicious activity in an AWS account for VPC hosted resources, the security engineer needs to use a service that can monitor network traffic and API calls across all AWS Regions. Amazon GuardDuty is a threat detection service that can do this by analyzing VPC Flow Logs, AWS CloudTrail event logs, and DNS logs. By activating GuardDuty across all AWS Regions, the security engineer can provide visibility for as many regions as possible. GuardDuty generates findings that contain details about the potential threats detected in the account. To respond to these findings, the security engineer needs to create a mechanism that can notify the relevant stakeholders or take remedial actions. One way to do this is to use Amazon EventBridge, which is a serverless event bus service that can connect AWS services and third-party applications. By creating an EventBridge rule that responds to GuardDuty findings and publishes them to an Amazon Simple Notification Service (Amazon SNS) topic, the security engineer can enable subscribers of the topic to receive notifications via email, SMS, or other methods. This is a cost-effective solution that does not require any additional infrastructure or code.

NEW QUESTION 83

A team is using AWS Secrets Manager to store an application database password. Only a limited number of IAM principals within the account can have access to the secret. The principals who require access to the secret change frequently. A security engineer must create a solution that maximizes flexibility and scalability. Which solution will meet these requirements?

- A. Use a role-based approach by creating an IAM role with an inline permissions policy that allows access to the secret
- B. Update the IAM principals in the role trust policy as required.
- C. Deploy a VPC endpoint for Secrets Manager
- D. Create and attach an endpoint policy that specifies the IAM principals that are allowed to access the secret
- E. Update the list of IAM principals as required.
- F. Use a tag-based approach by attaching a resource policy to the secret
- G. Apply tags to the secret and the IAM principal
- H. Use the aws:PrincipalTag and aws:ResourceTag IAM condition keys to control access.
- I. Use a deny-by-default approach by using IAM policies to deny access to the secret explicitly
- J. Attach the policies to an IAM group
- K. Add all IAM principals to the IAM group
- L. Remove principals from the group when they need access
- M. Add the principals to the group again when access is no longer allowed.

Answer: C

NEW QUESTION 88

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 rol
- 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 accoun
- I. Create an identity policy that allows the sts: AssumeRole actio
- 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 91

A company is using IAM Secrets Manager to store secrets for its production Amazon RDS database. The Security Officer has asked that secrets be rotated every 3 months. Which solution would allow the company to securely rotate the secrets? (Select TWO.)

- A. Place the RDS instance in a public subnet and an IAM Lambda function outside the VP
- B. Schedule the Lambda function to run every 3 months to rotate the secrets.
- C. Place the RDS instance in a private subnet and an IAM Lambda function inside the VPC in the private subne
- D. Configure the private subnet to use a NAT gatewa
- E. Schedule the Lambda function to run every 3 months to rotate the secrets.
- F. Place the RDS instance in a private subnet and an IAM Lambda function outside the VP
- G. Configure the private subnet to use an internet gatewa
- H. Schedule the Lambda function to run every 3 months lo rotate the secrets.
- I. Place the RDS instance in a private subnet and an IAM Lambda function inside the VPC in the private subne
- J. Schedule the Lambda function to run quarterly to rotate the secrets.
- K. Place the RDS instance in a private subnet and an IAM Lambda function inside the VPC in the private subne
- L. Configure a Secrets Manager interface endpoin
- M. Schedule the Lambda function to run every 3 months to rotate the secrets.

Answer: BE

Explanation:

these are the solutions that can securely rotate the secrets for the production RDS database using Secrets Manager. Secrets Manager is a service that helps you manage secrets such as database credentials, API keys, and passwords. You can use Secrets Manager to rotate secrets automatically by using a Lambda function that runs on a schedule. The Lambda function needs to have access to both the RDS instance and the Secrets Manager service. Option B places the RDS instance in a private subnet and the Lambda function in the same VPC in another private subnet. The private subnet with the Lambda function needs to use a NAT gateway to access Secrets Manager over the internet. Option E places the RDS instance and the Lambda function in the same private subnet and configures a Secrets Manager interface endpoint, which is a private connection between the VPC and Secrets Manager. The other options are either insecure or incorrect for rotating secrets using Secrets Manager.

NEW QUESTION 94

A security engineer needs to see 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 polic
- B. Specify cloudfront amazonaws com as the principa
- 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 acces
- E. Create an AWS WAF web ACL and add an IP set rul
- 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 distributio
- I. Create an AWS WAF web ACL and add an IP set rul
- J. Associate the web ACL with the CloudFront distribution.

Answer: B

NEW QUESTION 96

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 98

A company is running an Amazon RDS for MySQL DB instance in a VPC. The VPC must not send or receive network traffic through the internet.

A security engineer wants to use AWS Secrets Manager to rotate the DB instance credentials automatically. Because of a security policy, the security engineer cannot use the standard AWS Lambda function that Secrets Manager provides to rotate the credentials.

The security engineer deploys a custom Lambda function in the VPC. The custom Lambda function will be responsible for rotating the secret in Secrets Manager. The security engineer edits the DB instance's security group to allow connections from this function. When the function is invoked, the function cannot communicate with Secrets Manager to rotate the secret properly.

What should the security engineer do so that the function can rotate the secret?

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

Answer: D

Explanation:

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

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

The correct answer is D. Configure a Secrets Manager interface VPC endpoint. Include the Lambda function's private subnet during the configuration process.

A Secrets Manager interface VPC endpoint is a private connection between the VPC and Secrets Manager that does not require an internet gateway, NAT device, VPN connection, or AWS Direct Connect connection¹. By configuring a Secrets Manager interface VPC endpoint, the security engineer can enable the custom Lambda function to communicate with Secrets Manager without sending or receiving network traffic through the internet. The security engineer must include the Lambda function's private subnet during the configuration process to allow the function to use the endpoint².

The other options are incorrect for the following reasons:

- A. An egress-only internet gateway is a VPC component that allows outbound communication over IPv6 from instances in the VPC to the internet, and prevents the internet from initiating an IPv6 connection with the instances³. However, this option does not meet the requirement that the VPC must not send or receive network traffic through the internet. Moreover, an egress-only internet gateway is for use with IPv6 traffic only, and Secrets Manager does not support IPv6 addresses².
- B. A NAT gateway is a VPC component that enables instances in a private subnet to connect to the internet or other AWS services, but prevents the internet from initiating connections with those instances⁴. However, this option does not meet the requirement that the VPC must not send or receive network traffic through the internet. Additionally, a NAT gateway requires an elastic IP address, which is a public IPv4 address⁴.
- C. A VPC peering connection is a networking connection between two VPCs that enables you to route traffic between them using private IPv4 addresses or IPv6 addresses⁵. However, this option does not work because Secrets Manager does not have a default VPC that can be peered with. Furthermore, a VPC peering connection does not provide a private connection to Secrets Manager APIs without an internet gateway or other devices².

NEW QUESTION 103

During a manual review of system logs from an Amazon Linux EC2 instance, a Security Engineer noticed that there are sudo commands that were never properly alerted or reported on the Amazon CloudWatch Logs agent

Why were there no alerts on the sudo commands?

- A. There is a security group blocking outbound port 80 traffic that is preventing the agent from sending the logs

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

Answer: B

Explanation:

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

NEW QUESTION 108

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 111

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

Which solution will meet these requirements?

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

Answer: D

Explanation:

The correct answer is D.

To implement short-term credentials for third-party AWS accounts, you can use IAM roles and trust policies. A trust policy is a JSON policy document that defines who can assume the role. You can specify the AWS account ID of the third-party account as a principal in the trust policy, and use the sts:ExternalId condition key to enhance the security of the role. The sts:ExternalId condition key is a unique identifier that is agreed upon by both parties and included in the AssumeRole request. This way, you can prevent the “confused deputy” problem, where an unauthorized party can use the same role as a legitimate party. Option A is incorrect because bearer token authentication with OAuth or SAML is not suitable for granting access to AWS accounts and resources. Amazon Cognito and API Gateway are used for building web and mobile applications that require user authentication and authorization. Option B is incorrect because AWS IAM Identity Center (AWS Single Sign-On) is a service that simplifies the management of access to multiple AWS accounts and cloud applications for your workforce users. It does not support granting access to third-party AWS accounts. Option C is incorrect because using AWS Secrets Manager to create a random external key is not necessary and adds operational complexity. You can use the sts:ExternalId condition key instead to provide a unique identifier for each external account.

NEW QUESTION 115

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 network⁶. Resolver rules do not have any logging functionality by themselves. Therefore, this solution would not meet the requirements. References:

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

NEW QUESTION 119

There is a requirement for a company to transfer large amounts of data between IAM and an on-premise location. There is an additional requirement for low latency and high consistency traffic to IAM. Given these requirements how would you design a hybrid architecture? Choose the correct answer from the options below

Please select:

- A. Provision a Direct Connect connection to an IAM region using a Direct Connect partner.
- B. Create a VPN tunnel for private connectivity, which increases network consistency and reduces latency.
- C. Create an iPSec tunnel for private connectivity, which increases network consistency and reduces latency.

D. Create a VPC peering connection between IAM and the Customer gateway.

Answer: A

Explanation:

IAM Direct Connect makes it easy to establish a dedicated network connection from your premises to IAM. Using IAM Direct Connect you can establish private connectivity between IAM and your datacenter, office, or colocation environment which in many cases can reduce your network costs, increase bandwidth throughput and provide a more consistent network experience than Internet-based connections.

Options B and C are invalid because these options will not reduce network latency Options D is invalid because this is only used to connect 2 VPC's

For more information on IAM direct connect, just browse to the below URL: <https://IAM.amazon.com/directconnect>

The correct answer is: Provision a Direct Connect connection to an IAM region using a Direct Connect partner. omit your Feedback/Queries to our Experts

NEW QUESTION 121

A company has multiple Amazon S3 buckets encrypted with customer-managed CMKs Due to regulatory requirements the keys must be rotated every year. The company's Security Engineer has enabled automatic key rotation for the CMKs; however the company wants to verify that the rotation has occurred.

What should the Security Engineer do to accomplish this?

- A. Filter IAM CloudTrail logs for KeyRotaton events
- B. Monitor Amazon CloudWatch Events for any IAM KMS CMK rotation events
- C. Using the IAM CL
- D. run the IAM kms get-key-rotation-status operation with the --key-id parameter to check the CMK rotation date
- E. Use Amazon Athena to query IAM CloudTrail logs saved in an S3 bucket to filter Generate New Key events

Answer: C

Explanation:

the aws kms get-key-rotation-status command returns a boolean value that indicates whether automatic rotation of the customer master key (CMK) is enabled¹.

This command also shows the date and time when the CMK was last rotated². The other options are not valid ways to check the CMK rotation status.

NEW QUESTION 122

A company is using IAM Organizations to develop a multi-account secure networking strategy. The company plans to use separate centrally managed accounts for shared services, auditing, and security inspection. The company plans to provide dozens of additional accounts to application owners for production and development environments.

Company security policy requires that all internet traffic be routed through a centrally managed security inspection layer in the security inspection account. A security engineer must recommend a solution that minimizes administrative overhead and complexity.

Which solution meets these requirements?

- A. Use IAM Control Towe
- B. Modify the default Account Factory networking template to automatically associate new accounts with a centrally managed VPC through a VPC peering connection and to create a default route to the VPC peer in the default route tabl
- C. Create an SCP that denies the CreateInternetGateway actio
- D. Attach the SCP to all accounts except the security inspection account.
- E. Create a centrally managed VPC in the security inspection accoun
- F. Establish VPC peering connections between the security inspection account and other account
- G. Instruct account owners to create default routes in their account route tables that point to the VPC pee
- H. Create an SCP that denies theAttach InternetGateway actio
- I. Attach the SCP to all accounts except the security inspection account.
- J. Use IAM Control Towe
- K. Modify the default Account Factory networking template to automatically associate new accounts with a centrally managed transitgateway and to create a default route to the transit gateway in the default route tabl
- L. Create an SCP that denies the AttachInternetGateway actio
- M. Attach the SCP to all accounts except the security inspection account.
- N. Enable IAM Resource Access Manager (IAM RAM) for IAM Organization
- O. Create a shared transit gateway, and make it available by using an IAM RAM resource shar
- P. Create an SCP that denies the CreateInternetGateway actio
- Q. Attach the SCP to all accounts except the security inspection accoun
- R. Create routes in the route tables of all accounts that point to the shared transit gateway.

Answer: C

NEW QUESTION 123

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

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

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

Answer: D

Explanation:

Storing sensitive information in environment variables is not a secure practice, as anyone who has access to the Lambda console or the Lambda function code can view them as plaintext. To address this security issue, the security engineer needs to use a service that can store and encrypt the environment variables, and access them at runtime using IAM permissions. The most cost-effective way to do this is to use AWS Systems Manager Parameter Store, which is a service that

provides secure, hierarchical storage for configuration data management and secrets management. Parameter Store allows you to store values as standard parameters (plaintext) or secure string parameters (encrypted). Secure string parameters use a AWS Key Management Service (AWS KMS) customer master key (CMK) to encrypt the parameter value. To access the parameter value at runtime, the Lambda function needs to have IAM permissions to decrypt the parameter using the KMS CMK.

The other options are incorrect because:

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

NEW QUESTION 128

A company needs to store multiple years of financial records. The company wants to use Amazon S3 to store copies of these documents. The company must implement a solution to prevent the documents from being edited, replaced, or deleted for 7 years after the documents are stored in Amazon S3. The solution must also encrypt the documents at rest.

A security engineer creates a new S3 bucket to store the documents. What should the security engineer do next to meet these requirements?

- A. Configure S3 server-side encryptio
- B. Create an S3 bucket policy that has an explicit deny rule for all users for s3:DeleteObject and s3:PutObject API call
- C. Configure S3 Object Lock to use governance mode with a retention period of 7 years.
- D. Configure S3 server-side encryptio
- E. Configure S3 Versioning on the S3 bucke
- F. Configure S3 ObjectLock to use compliance mode with a retention period of 7 years.
- G. Configure S3 Versionin
- H. Configure S3 Intelligent-Tiering on the S3 bucket to move the documents to S3 Glacier Deep Archive storag
- I. Use S3 server-side encryption immediatel
- J. Expire the objects after 7 years.
- K. Set up S3 Event Notifications and use S3 server-side encryptio
- L. Configure S3 Event Notifications to target an AWS Lambda function that will review any S3 API call to the S3 bucket and deny the s3:DeleteObject and s3:PutObject API call
- M. Remove the S3 event notification after 7 years.

Answer: B

NEW QUESTION 133

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

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

Which solution will meet these requirements?

- A. In the software development account, create AMIS of preconfigured instanc-es that include only approved softwar
- B. Include the AMI IDs in the condi-tion section of an AWS CloudFormation template to launch the appropriate AMI based on the AWS Regio
- C. Provide the developers with theCloudFor-mation template to launch EC2 instances in the software development ac-count.
- D. Create an Amazon EventBridge rule that runs when any EC2 RunInstances API event occurs in the software development accoun
- E. Specify AWS Systems Man-ager Run Command as a target of the rul
- F. Configure Run Command to run a script that will install all approved software onto the instances that the developers launch.
- G. Use an AWS Service Catalog portfolio that contains EC2 products with ap-propriate AMIS that include only approved softwar
- H. Grant the developers permission to portfolio access only the Service Catalog to launch a prod-uct in the software development account.
- I. In the management account, create AMIS of preconfigured instances that in-clude only approved softwar
- J. Use AWS CloudFormation StackSets to launch the AMIS across any AWS account in the organizatio
- K. Grant the developers permission to launch the stack sets within the management account.

Answer: C

NEW QUESTION 138

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

Which solution will meet these requirements?

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

Answer: B

NEW QUESTION 141

A development team is using an IAM Key Management Service (IAM KMS) CMK to try to encrypt and decrypt a secure string parameter from IAM Systems

Manager Parameter Store. However, the development team receives an error message on each attempt. Which issues that are related to the CMK could be reasons for the error? (Select TWO.)

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

Answer: AD

NEW QUESTION 146

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

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

Answer: C

Explanation:

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

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

This solution provides the following security benefits:

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

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

NEW QUESTION 148

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

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

Answer: C

Explanation:

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

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

NEW QUESTION 151

.....

Thank You for Trying Our Product

We offer two products:

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

2nd - Questions and Answers in PDF Format

SCS-C02 Practice Exam Features:

- * SCS-C02 Questions and Answers Updated Frequently
- * SCS-C02 Practice Questions Verified by Expert Senior Certified Staff
- * SCS-C02 Most Realistic Questions that Guarantee you a Pass on Your FirstTry
- * SCS-C02 Practice Test Questions in Multiple Choice Formats and Updatesfor 1 Year

100% Actual & Verified — Instant Download, Please Click
[Order The SCS-C02 Practice Test Here](#)