

Amazon-Web-Services

Exam Questions SCS-C02

AWS Certified Security - Specialty



NEW QUESTION 1

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

Each subnet is currently using the VPC's default network ACL. The security groups that the EC2 instances in these subnets use have rules that allow traffic between each instance where required. Currently, all network traffic flow is working as expected between the EC2 instances that are using these subnets.

A security engineer creates a new network ACL that is named subnet-2-NACL with default entries. The security engineer immediately configures private-subnet-2 to use the new network ACL and makes no other changes to the infrastructure. The security engineer starts to receive reports that the EC2 instances in public-subnet-1 and public-subnet-2 cannot communicate with each other.

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

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

Answer: CE

Explanation:

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

References: : Amazon VPC User Guide

NEW QUESTION 2

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

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

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

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

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

Answer: B

Explanation:

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

NEW QUESTION 3

A company is using Amazon Macie, AWS Firewall Manager, Amazon Inspector, and AWS Shield Advanced in its AWS account. The company wants to receive alerts if a DDoS attack occurs against the account.

Which solution will meet this requirement?

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

Answer: D

Explanation:

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

For more information, see Monitoring AWS Shield Advanced with Amazon CloudWatch and AWS Shield Advanced metrics and alarms.

NEW QUESTION 4

A company has a relational database workload that runs on Amazon Aurora MySQL. According to new compliance standards the company must rotate all database credentials every 30 days. The company needs a solution that maximizes security and minimizes development effort.

Which solution will meet these requirements?

- A. Store the database credentials in AWS Secrets Manager
- B. Configure automatic credential rotation for every 30 days.
- C. Store the database credentials in AWS Systems Manager Parameter Store
- D. Create an AWS Lambda function to rotate the credentials every 30 days.

- E. Store the database credentials in an environment file or in a configuration file
- F. Modify the credentials every 30 days.
- G. Store the database credentials in an environment file or in a configuration file
- H. Create an AWS Lambda function to rotate the credentials every 30 days.

Answer: A

Explanation:

To rotate database credentials every 30 days, the most secure and efficient solution is to store the database credentials in AWS Secrets Manager and configure automatic credential rotation for every 30 days. Secrets Manager can handle the rotation of the credentials in both the secret and the database, and it can use AWS KMS to encrypt the credentials. Option B is incorrect because it requires creating a custom Lambda function to rotate the credentials, which is more effort than using Secrets Manager. Option C is incorrect because it stores the database credentials in an environment file or a configuration file, which is less secure than using Secrets Manager. Option D is incorrect because it combines the drawbacks of option B and option C. Verified References:

- <https://docs.aws.amazon.com/secretsmanager/latest/userguide/rotating-secrets.html>
- https://docs.aws.amazon.com/secretsmanager/latest/userguide/rotate-secrets_turn-on-for-other.html

NEW QUESTION 5

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 6

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 7

Your company has just set up a new central server in a VPC. There is a requirement for other teams who have their servers located in different VPC's in the same region to connect to the central server. Which of the below options is best suited to achieve this requirement.

Please select:

- A. Set up VPC peering between the central server VPC and each of the teams VPCs.
- B. Set up IAM DirectConnect between the central server VPC and each of the teams VPCs.
- C. Set up an IPSec Tunnel between the central server VPC and each of the teams VPCs.
- D. None of the above options will work.

Answer: A

Explanation:

A VPC peering connection is a networking connection between two VPCs that enables you to route traffic between them using private IPv4 addresses or IPv6 addresses. Instances in either VPC can communicate with each other as if they are within the same network. You can create a VPC peering connection between your own VPCs, or with a VPC in another IAM account within a single region.

Options B and C are invalid because you need to use VPC Peering Option D is invalid because VPC Peering is available

For more information on VPC Peering please see the below Link:

<http://docs.IAM.amazon.com/AmazonVPC/latest/UserGuide/vpc-peering.html>

The correct answer is: Set up VPC peering between the central server VPC and each of the teams VPCs. Submit your Feedback/Queries to our Experts

NEW QUESTION 8

A company developed an application by using AWS Lambda, Amazon S3, Amazon Simple Notification Service (Amazon SNS), and Amazon DynamoDB. An external application puts objects into the company's S3 bucket and tags the objects with date and time. A Lambda function periodically pulls data from the company's S3 bucket based on date and time tags and inserts specific values into a DynamoDB table for further processing.

The data includes personally identifiable information (PII). The company must remove data that is older than 30 days from the S3 bucket and the DynamoDB table.

Which solution will meet this requirement with the MOST operational efficiency?

- A. Update the Lambda function to add a TTL S3 flag to S3 object
- B. Create an S3 Lifecycle policy to expire objects that are older than 30 days by using the TTL S3 flag.
- C. Create an S3 Lifecycle policy to expire objects that are older than 30 day
- D. Update the Lambda function to add the TTL attribute in the DynamoDB tabl
- E. Enable TTL on the DynamoDB table to expire entires that are older than 30 days based on the TTL attribute.
- F. Create an S3 Lifecycle policy to expire objects that are older than 30 days and to add all prefixes to the S3 bucke
- G. Update the Lambda function to delete entries that are older than 30 days.
- H. Create an S3 Lifecycle policy to expire objects that are older than 30 days by using object tag
- I. Update the Lambda function to delete entries that are older than 30 days.

Answer: B

NEW QUESTION 9

An AWS account administrator created an IAM group and applied the following managed policy to require that each individual user authenticate using multi-factor authentication:

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": "ec2:*",
      "Resource": "*"
    },
    {
      "Sid": "BlockAnyAccessUnlessSignedInWithMFA",
      "Effect": "Deny",
      "Action": "ec2:*",
      "Resource": "*",
      "Condition": {
        "BoolIfExists": {
          "aws:MultiFactorAuthPresent": false
        }
      }
    }
  ]
}
```

After implementing the policy, the administrator receives reports that users are unable to perform Amazon EC2 commands using the AWS CLI. What should the administrator do to resolve this problem while still enforcing multi-factor authentication?

- A. Change the value of aws:MultiFactorAuthPresent to true.
- B. Instruct users to run the aws sts get-session-token CLI command and pass the multi-factor authentication --serial-number and --token-code parameter
- C. Use these resulting values to make API/CLI calls.
- D. Implement federated API/CLI access using SAML 2.0, then configure the identity provider to enforce multi-factor authentication.
- E. Create a role and enforce multi-factor authentication in the role trust polic
- F. Instruct users to run the sts assume-role CLI command and pass --serial-number and --token-code parameter
- G. Store the resultingvalues in environment variable
- H. Add sts:AssumeRole to NotAction in the policy.

Answer: B

Explanation:

The correct answer is B. Instruct users to run the aws sts get-session-token CLI command and pass the multi-factor authentication --serial-number and --token-code parameters. Use these resulting values to make API/CLI calls.

According to the AWS documentation¹, the aws sts get-session-token CLI command returns a set of temporary credentials for an AWS account or IAM user. The credentials consist of an access key ID, a secret access key, and a security token. These credentials are valid for the specified duration only. The session duration for IAM users can be between 15 minutes and 36 hours, with a default of 12 hours.

You can use the --serial-number and --token-code parameters to provide the MFA device serial number and the MFA code from the device. The MFA device must be associated with the user who is making the

get-session-token call. If you do not provide these parameters when your IAM user or role has a policy that requires MFA, you will receive an Access Denied error. The temporary security credentials that are returned by the get-session-token command can then be used to make subsequent API or CLI calls that require MFA authentication. You can use environment variables or a profile in your AWS CLI configuration file to specify the temporary credentials.

Therefore, this solution will resolve the problem of users being unable to perform EC2 commands using the AWS CLI, while still enforcing MFA.

The other options are incorrect because:

- A. Changing the value of aws:MultiFactorAuthPresent to true will not work, because this is a condition key that is evaluated by AWS when a request is made. You cannot set this value manually in your policy or request. You must provide valid MFA information to AWS for this condition key to be true.
- C. Implementing federated API/CLI access using SAML 2.0 may work, but it requires more operational effort than using the get-session-token command. You would need to configure a SAML identity provider and trust relationship with AWS, and use a custom SAML client to request temporary credentials from AWS STS. This solution may also introduce additional security risks if the identity provider is compromised.
- D. Creating a role and enforcing MFA in the role trust policy may work, but it also requires more operational effort than using the get-session-token command. You would need to create a role for each user or group that needs to perform EC2 commands, and specify a trust policy that requires MFA. You would also need to grant the users permission to assume the role, and instruct them to use the sts assume-role command instead of the get-session-token command.

References:

1: get-session-token — AWS CLI Command Reference

NEW QUESTION 10

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 10

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

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

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

Answer: D

Explanation:

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

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

NEW QUESTION 14

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

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

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

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

Answer: BD

Explanation:

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

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

To use CloudTrail with multiple AWS accounts and regions, you need to enable AWS Organizations with all features enabled. This allows you to centrally manage your accounts and apply policies across your organization. You can also use CloudTrail as a service principal for AWS Organizations, which lets you create an organization trail that applies to all accounts in your organization. An organization trail logs events for all AWS Regions and delivers the log files to an S3 bucket that you specify.

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

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

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

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

Option A is incorrect because setting the bucket policy to allow the organization's management account to write to the S3 bucket is not sufficient, as it will not

grant access to the other member accounts in the organization.

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

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

NEW QUESTION 17

A corporation is preparing to acquire several companies. A Security Engineer must design a solution to ensure that newly acquired IAM accounts follow the corporation's security best practices. The solution should monitor each Amazon S3 bucket for unrestricted public write access and use IAM managed services. What should the Security Engineer do to meet these requirements?

- A. Configure Amazon Macie to continuously check the configuration of all S3 buckets.
- B. Enable IAM Config to check the configuration of each S3 bucket.
- C. Set up IAM Systems Manager to monitor S3 bucket policies for public write access.
- D. Configure an Amazon EC2 instance to have an IAM role and a cron job that checks the status of all S3 buckets.

Answer: C

Explanation:

because this is a solution that can monitor each S3 bucket for unrestricted public write access and use IAM managed services. S3 is a service that provides object storage in the cloud. Systems Manager is a service that helps you automate and manage your AWS resources. You can use Systems Manager to monitor S3 bucket policies for public write access by using a State Manager association that runs a predefined document called AWS-FindS3BucketWithPublicWriteAccess. This document checks each S3 bucket in an account and reports any bucket that has public write access enabled. The other options are either not suitable or not feasible for meeting the requirements.

NEW QUESTION 19

A company is running workloads in a single IAM account on Amazon EC2 instances and Amazon EMR clusters a recent security audit revealed that multiple Amazon Elastic Block Store (Amazon EBS) volumes and snapshots are not encrypted

The company's security engineer is working on a solution that will allow users to deploy EC2 Instances and EMR clusters while ensuring that all new EBS volumes and EBS snapshots are encrypted at rest. The solution must also minimize operational overhead

Which steps should the security engineer take to meet these requirements?

- A. Create an Amazon Event Bridge (Amazon Cloud watch Events) event with an EC2 instance as the source and create volume as the event trigger
- B. When the event is triggered invoke an IAM Lambda function to evaluate and notify the security engineer if the EBS volume that was created is not encrypted.
- C. Use a customer managed IAM policy that will verify that the encryption tag of the CreateVolume context is set to true
- D. Apply this rule to all users.
- E. Create an IAM Config rule to evaluate the configuration of each EC2 instance on creation or modification. Have the IAM Config rule trigger an IAM Lambda function to alert the security team and terminate the instance if the EBS volume is not encrypted
- F. 5
- G. Use the IAM Management Console or IAM CLI to enable encryption by default for EBS volumes in each IAM Region where the company operates.

Answer: D

Explanation:

To ensure that all new EBS volumes and EBS snapshots are encrypted at rest and minimize operational overhead, the security engineer should do the following:

➤ Use the AWS Management Console or AWS CLI to enable encryption by default for EBS volumes in each AWS Region where the company operates. This allows the security engineer to automatically encrypt any new EBS volumes and snapshots created from those volumes, without requiring any additional actions from users.

NEW QUESTION 21

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 team
- B. Create a trail in AWS CloudTrail that has Amazon CloudWatch Logs turned on
- C. Use Amazon EventBridge to monitor local user activities.
- D. Create a break glass EC2 key pair for the AWS account
- E. Provide the key pair to the security team
- F. Use AWS CloudTrail to monitor key pair activities
- G. Send notifications to the security team by using Amazon Simple Notification Service (Amazon SNS).
- H. Create a break glass IAM role for the account
- I. Allow security team members to perform the AssumeRoleWithSAML operation
- J. Create an AWS Cloud Trail trail that has Amazon CloudWatch Logs turned on
- 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 Manager
- 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 23

A company is designing a multi-account structure for its development teams. The company is using AWS Organizations and AWS Single Sign-On (AWS SSO). The company must implement a solution so that the development teams can use only specific AWS Regions and so that each AWS account allows access to only specific AWS services.

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

- A. Use AWS SSO to set up service-linked roles with IAM policy statements that include the Condition, Resource, and NotAction elements to allow access to only the Regions and services that are needed.
- B. Deactivate AWS Security Token Service (AWS STS) in Regions that the developers are not allowed to use.
- C. Create SCPs that include the Condition, Resource, and NotAction elements to allow access to only the Regions and services that are needed.
- D. For each AWS account, create tailored identity-based policies for AWS SS
- E. Use statements that include the Condition, Resource, and NotAction elements to allow access to only the Regions and services that are needed.

Answer: C

Explanation:

https://docs.aws.amazon.com/organizations/latest/userguide/orgs_manage_policies_scps_syntax.html#scp-eleme

NEW QUESTION 28

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

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

Answer: ABC

Explanation:

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

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

NEW QUESTION 33

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

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

Answer: D

Explanation:

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

NEW QUESTION 37

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

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

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

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

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

Answer: BE

NEW QUESTION 40

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

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

Which solution will meet these requirements?

- A. Create an Amazon CloudWatch Logs log group
- B. Configure the load balancers to send logs to the log group
- C. Use the CloudWatch Logs console to search the log
- D. Create CloudWatch Logs filters on the logs for the required metrics.
- E. Create an Amazon S3 bucket
- F. Configure the load balancers to send logs to the S3 bucket
- G. Use Amazon Athena to search the logs that are in the S3 bucket
- H. Create Amazon CloudWatch filters on the S3 log files for the required metrics.
- I. Create an Amazon S3 bucket
- J. Configure the load balancers to send logs to the S3 bucket
- K. Use Amazon Athena to search the logs that are in the S3 bucket
- L. Create Athena queries for the required metric
- M. Publish the metrics to Amazon CloudWatch.
- N. Create an Amazon CloudWatch Logs log group
- O. Configure the load balancers to send logs to the log group
- P. Use the AWS Management Console to search the log
- Q. Create Amazon Athena queries for the required metric
- R. Publish the metrics to Amazon CloudWatch.

Answer: C

Explanation:

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

NEW QUESTION 44

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

Which solution will meet these requirements?

- A. Create new S3 buckets with S3 Object Lock enabled in compliance mod
- B. Place objects in the S3 buckets.
- C. Use S3 Glacier Vault Lock to attach a Vault Lock policy to new S3 bucket
- D. Wait 24 hours to complete the Vault Lock proces
- E. Place objects in the S3 buckets.
- F. Create new S3 buckets with S3 Object Lock enabled in governance mod
- G. Place objects in the S3 buckets.
- H. Create new S3 buckets with S3 Object Lock enabled in governance mod
- I. Add a legal hold to the S3 bucket
- J. Place objects in the S3 buckets.

Answer: A

NEW QUESTION 46

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

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

Answer: BDF

Explanation:

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

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

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

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

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

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

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

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

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

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

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

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

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

According to the AWS documentation, Amazon QuickSight is a fast, cloud-powered business intelligence

service that makes it easy to deliver insights to everyone in your organization. You can use QuickSight to create and publish interactive dashboards that include machine learning insights. You can also use QuickSight to connect to various data sources, such as Athena, S3, or RDS. To use QuickSight with Security Hub findings, you need to create QuickSight dashboards that use the Athena views as data sources. You can then visualize and analyze the findings using charts, graphs, maps, or tables. You can also apply filters, calculations, or aggregations to the data. You can then share the dashboards with your users or embed them in your applications.

NEW QUESTION 50

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

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

Answer: C

Explanation:

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

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

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

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

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

NEW QUESTION 53

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

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

Answer: ACD

NEW QUESTION 54

A security engineer needs to build a solution to turn IAM CloudTrail back on in multiple IAM Regions in case it is ever turned off. What is the MOST efficient way to implement this solution?

- A. Use IAM Config with a managed rule to trigger the IAM-EnableCloudTrail remediation.
- B. Create an Amazon EventBridge (Amazon CloudWatch Events) event with a cloudtrail.amazonaws.com event source and a StartLogging event name to trigger an IAM Lambda function to call the StartLogging API.
- C. Create an Amazon CloudWatch alarm with a cloudtrail.amazonaws.com event source and a StopLogging event name to trigger an IAM Lambda function to call the StartLogging API.
- D. Monitor IAM Trusted Advisor to ensure CloudTrail logging is enabled.

Answer: B

NEW QUESTION 55

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

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

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

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

Answer: C

Explanation:

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

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

The other options are incorrect for the following reasons:

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

NEW QUESTION 58

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

Which solution will meet this requirement?

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

Answer: B

Explanation:

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

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

NEW QUESTION 62

A company is planning to use Amazon Elastic File System (Amazon EFS) with its on-premises servers. The company has an existing IAM Direct Connect connection established between its on-premises data center and an IAM Region. A security policy states that the company's on-premises firewall should only have specific IP addresses added to the allow list and not a CIDR range. The company also wants to restrict access so that only certain data center-based servers have access to Amazon EFS.

How should a security engineer implement this solution?

- A. Add the file-system-id efs IAM-region amazonIAM.com URL to the allow list for the data center firewall. Install the IAM CLI on the data center-based servers to mount the EFS file system in the EFS security group. Add the data center IP range to the allow list. Mount the EFS using the EFS file system name.
- B. Assign an Elastic IP address to Amazon EFS and add the Elastic IP address to the allow list for the data center firewall. Install the IAM CLI on the data center-based servers to mount the EFS file system. In the EFS security group, add the IP addresses of the data center servers to the allow list. Mount the EFS using the Elastic IP address.
- C. Add the EFS file system mount target IP addresses to the allow list for the data center firewall. In the EFS security group, add the data center server IP addresses to the allow list. Use the Linux terminal to mount the EFS file system using the IP address of one of the mount targets.
- D. Assign a static range of IP addresses for the EFS file system by contacting IAM Support. In the EFS security group, add the data center server IP addresses to the allow list. Use the Linux terminal to mount the EFS file system using one of the static IP addresses.

Answer: B

Explanation:

To implement the solution, the security engineer should do the following:

- Assign an Elastic IP address to Amazon EFS and add the Elastic IP address to the allow list for the data center firewall. This allows the security engineer to use a specific IP address for the EFS file system that can be added to the firewall rules, instead of a CIDR range or a URL.
- Install the AWS CLI on the data center-based servers to mount the EFS file system. This allows the security engineer to use the mount helper provided by AWS CLI to mount the EFS file system with encryption in transit.
- In the EFS security group, add the IP addresses of the data center servers to the allow list. This allows the security engineer to restrict access to the EFS file system to only certain data center-based servers.
- Mount the EFS using the Elastic IP address. This allows the security engineer to use the Elastic IP address as the DNS name for mounting the EFS file system.

NEW QUESTION 67

A company has multiple departments. Each department has its own IAM account. All these accounts belong to the same organization in IAM Organizations. A large .csv file is stored in an Amazon S3 bucket in the sales department's IAM account. The company wants to allow users from the other accounts to access the .csv file's content through the combination of IAM Glue and Amazon Athena. However, the company does not want to allow users from the other accounts to access other files in the same folder. Which solution will meet these requirements?

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

Answer: A

NEW QUESTION 69

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

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

Answer: D

Explanation:

The correct answer is D.

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

To automate remediation of incidents across the production accounts, the security engineer needs to use a service that can trigger actions based on events. Amazon EventBridge is a serverless event bus service that allows you to connect your applications with data from a variety of sources. EventBridge can use rules to match events and route them to targets for processing. You can use EventBridge to invoke a custom AWS Lambda function from the Security Hub findings. Lambda is a serverless compute service that lets you run code without provisioning or managing servers.

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

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

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

References:

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

NEW QUESTION 70

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 75

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

```
{resolve:s3:MyBucketName:MyObjectName}}.
```

- A. Store the API key value as a SecureString parameter in AWS Systems Manager Parameter Store
- B. In the template, replace all references to the value with `{{resolve:ssm:MySSMParameterName:1}}`.
- C. Store the API key value in AWS Secrets Manager
- D. In the template, replace all references to the value with `{{resolve:secretsmanager:MySecretId:SecretString}}`.
- E. Store the API key value in Amazon DynamoDB
- F. In the template, replace all references to the value with `{{resolve:dynamodb:MyTableName:MyPrimaryKey}}`.
- G. Store the API key value in a new Amazon S3 bucket
- H. In the template, replace all references to the value with `{{resolve:s3:MyBucketName:MyObjectName}}`.

Answer: B

Explanation:

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

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

The other options are incorrect because:

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

References:

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

NEW QUESTION 80

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

- A. Set up an Amazon EventBridge rule that reacts to new Security Hub findings

- B. Configure an AWS Lambda function as the target for the rule to reme-diate the findings.
- C. Set up a custom action in Security Hu
- D. Configure the custom action to call AWS Systems Manager Automation runbooks to remediate the findings.
- E. Set up a custom action in Security Hu
- F. Configure an AWS Lambda function as the target for the custom action to remediate the findings.
- G. Set up AWS Config rules to use AWS Systems Manager Automation runbooks to remediate the findings.

Answer: A

NEW QUESTION 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 secre
- B. Update the IAM principals in the role trust policy as required.
- C. Deploy a VPC endpoint for Secrets Manage
- D. Create and attach an endpoint policy that specifies the IAM principals that are allowed to access the secre
- E. Update the list of IAM principals as required.
- F. Use a tag-based approach by attaching a resource policy to the secre
- 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 explicitl
- J. Attach the policies to an IAM grou
- K. Add all IAM principals to the IAM grou
- L. Remove principals from the group when they need acces
- M. Add the principals to the group again when access is no longer allowed.

Answer: C

NEW QUESTION 85

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 89

A company accidentally deleted the private key for an Amazon Elastic Block Store (Amazon EBS)-backed Amazon EC2 instance. A security engineer needs to regain access to the instance.

Which combination of steps will meet this requirement? (Choose two.)

- A. Stop the instanc
- B. Detach the root volum
- C. Generate a new key pair.
- D. Keep the instance runnin
- E. Detach the root volum
- F. Generate a new key pair.
- G. When the volume is detached from the original instance, attach the volume to another instance as a data volum
- H. Modify the authorized_keys file with a new public ke
- I. Move the volume back to the original instanc
- J. Start the instance.
- K. When the volume is detached from the original instance, attach the volume to another instance as a data volum
- L. Modify the authorized_keys file with a new private ke

- M. Move the volume back to the original instance
- N. Start the instance.
- O. When the volume is detached from the original instance, attach the volume to another instance as a data volume
- P. Modify the authorized_keys file with a new public key
- Q. Move the volume back to the original instance that is running.

Answer: AC

Explanation:

If you lose the private key for an EBS-backed instance, you can regain access to your instance. You must stop the instance, detach its root volume and attach it to another instance as a data volume, modify the authorized_keys file with a new public key, move the volume back to the original instance, and restart the instance.
<https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/TroubleshootingInstancesConnecting.html#replacing>

NEW QUESTION 94

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

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

Answer: B

NEW QUESTION 96

A company's Security Team received an email notification from the Amazon EC2 Abuse team that one or more of the company's Amazon EC2 instances may have been compromised. Which combination of actions should the Security team take to respond to the notification? (Select TWO.)

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

Answer: DE

Explanation:

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

NEW QUESTION 97

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

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

Answer: D

Explanation:

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

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

NEW QUESTION 101

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 Tower
- 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 table
- C. Create an SCP that denies the `CreateInternetGateway` action
- D. Attach the SCP to all accounts except the security inspection account.
- E. Create a centrally managed VPC in the security inspection account
- 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 peer
- H. Create an SCP that denies the `AttachInternetGateway` action
- I. Attach the SCP to all accounts except the security inspection account.
- J. Use IAM Control Tower
- K. Modify the default Account Factory networking template to automatically associate new accounts with a centrally managed transit gateway and to create a default route to the transit gateway in the default route table
- L. Create an SCP that denies the `AttachInternetGateway` action
- 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 share
- P. Create an SCP that denies the `CreateInternetGateway` action
- Q. Attach the SCP to all accounts except the security inspection account
- R. Create routes in the route tables of all accounts that point to the shared transit gateway.

Answer: C

NEW QUESTION 106

A security engineer needs to configure an Amazon S3 bucket policy to restrict access to an S3 bucket that is named `DOC-EXAMPLE-BUCKET`. The policy must allow access to only `DOC-EXAMPLE-BUCKET` from only the following endpoint: `vpce-1a2b3c4d`. The policy must deny all access to `DOC-EXAMPLE-BUCKET` if the specified endpoint is not used.

Which bucket policy statement meets these requirements?

- A. A computer code with black text Description automatically generated

```
"Statement": [
  {
    "Sid": "Access-to-specific-VPCE-only",
    "Principal": "*",
    "Action": "s3:*",
    "Effect": "Allow",
    "Resource": ["arn:aws:s3::DOC-EXAMPLE-BUCKET",
                 "arn:aws:s3::DOC-EXAMPLE-BUCKET/*"],
    "Condition": {
      "StringNotEquals": {
        "aws:sourceVpce": "vpce-1a2b3c4d"
      }
    }
  }
]
```

- B. A computer code with black text Description automatically generated

```
"Statement": [
  {
    "Sid": "Access-to-specific-VPCE-only",
    "Principal": "*",
    "Action": "s3:*",
    "Effect": "Deny",
    "Resource": ["arn:aws:s3::DOC-EXAMPLE-BUCKET",
                 "arn:aws:s3::DOC-EXAMPLE-BUCKET/*"],
    "Condition": {
      "StringNotEquals": {
        "aws:sourceVpce": "vpce-1a2b3c4d"
      }
    }
  }
]
```

- C. A computer code with black text Description automatically generated

```
"Statement": [
  {
    "Sid": "Access-to-specific-VPCE-only",
    "Principal": "*",
    "Action": "s3:*",
    "Effect": "Deny",
    "Resource": ["arn:aws:s3::DOC-EXAMPLE-BUCKET",
      "arn:aws:s3::DOC-EXAMPLE-BUCKET/*"],
    "Condition": {
      "StringEquals": {
        "aws:sourceVpce": "vpce-1a2b3c4d"
      }
    }
  }
]
```

D. A computer code with black text Description automatically generated

```
"Statement": [
  {
    "Sid": "Access-to-specific-VPCE-only",
    "Principal": "*",
    "Action": "s3:*",
    "Effect": "Allow",
    "Resource": ["arn:aws:s3::DOC-EXAMPLE-BUCKET",
      "arn:aws:s3::DOC-EXAMPLE-BUCKET/*"],
    "Condition": {
      "StringEquals": {
        "aws:sourceVpce": "vpce-1a2b3c4d"
      }
    }
  }
]
```

Answer: B

Explanation:

<https://docs.aws.amazon.com/AmazonS3/latest/userguide/example-bucket-policies-vpc-endpoint.html>

NEW QUESTION 109

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 runtime
- 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 runtime
- 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 runtime
- 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 113

A company uses Amazon API Gateway to present REST APIs to users. An API developer wants to analyze API access patterns without the need to parse the log files.

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

- A. Configure access logging for the required API stage.
- B. Configure an AWS CloudTrail trail destination for API Gateway event
- C. Configure filters on the userIdentity, userAgent, and sourceIPAddress fields.
- D. Configure an Amazon S3 destination for API Gateway log
- E. Run Amazon Athena queries to analyze API access information.
- F. Use Amazon CloudWatch Logs Insights to analyze API access information.
- G. Select the Enable Detailed CloudWatch Metrics option on the required API stage.

Answer: CD

NEW QUESTION 115

A company is building a data processing application that uses AWS Lambda functions. The application's Lambda functions need to communicate with an Amazon RDS DB instance that is deployed within a VPC in the same AWS account.

Which solution meets these requirements in the MOST secure way?

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

Answer: C

Explanation:

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

References: : AWS Lambda Developer Guide

NEW QUESTION 119

A company needs a forensic-logging solution for hundreds of applications running in Docker on Amazon EC2. The solution must perform real-time analytics on the logs. The logs must support the replay of messages and must persist the logs.

Which IAM services should be used to meet these requirements? (Select TWO.)

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

Answer: BD

Explanation:

Amazon Kinesis and Amazon Elasticsearch are both suitable for forensic-logging solutions. Amazon Kinesis can collect, process, and analyze streaming data in real time. Amazon Elasticsearch can store, search, and analyze log data using the popular open-source tool Elasticsearch. The other options are not designed for forensic-logging purposes. Amazon Athena is a query service that can analyze data in S3, Amazon SQS is a message queue service that can decouple and scale microservices, and Amazon EMR is a big data platform that can run Apache Spark and Hadoop clusters.

NEW QUESTION 124

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

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

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

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

Answer: C

Explanation:

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

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

NEW QUESTION 129

A company is using AWS Organizations to manage multiple AWS accounts for its human resources, finance, software development, and production departments.

All the company's developers are part of the software development AWS account.

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

Which solution will meet these requirements?

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

Answer: C

NEW QUESTION 133

A company has launched an Amazon EC2 instance with an Amazon Elastic Block Store (Amazon EBS) volume in the us-east-1 Region. The volume is encrypted with an AWS Key Management Service (AWS KMS) customer managed key that the company's security team created. The security team has created an IAM key policy and has assigned the policy to the key. The security team has also created an IAM instance profile and has assigned the profile to the instance. The EC2 instance will not start and transitions from the pending state to the shutting-down state to the terminated state.

Which combination of steps should a security engineer take to troubleshoot this issue? (Select TWO.)

- A. Verify that the KMS key policy specifies a deny statement that prevents access to the key by using the aws:SourceIP condition key. Check that the range includes the EC2 instance IP address that is associated with the EBS volume.
- B. Verify that the KMS key that is associated with the EBS volume is set to the Symmetric key type.
- C. Verify that the KMS key that is associated with the EBS volume is in the Enabled state.
- D. Verify that the EC2 role that is associated with the instance profile has the correct IAM instance policy to launch an EC2 instance with the EBS volume.
- E. Verify that the key that is associated with the EBS volume has not expired and needs to be rotated.

Answer: CD

Explanation:

To troubleshoot the issue of an EC2 instance failing to start and transitioning to a terminated state when it has an EBS volume encrypted with an AWS KMS customer managed key, a security engineer should take the following steps:

* C. Verify that the KMS key that is associated with the EBS volume is in the Enabled state. If the key is not enabled, it will not function properly and could cause the EC2 instance to fail.

* D. Verify that the EC2 role that is associated with the instance profile has the correct IAM instance policy to launch an EC2 instance with the EBS volume. If the instance does not have the necessary permissions, it may not be able to mount the volume and could cause the instance to fail.

Therefore, options C and D are the correct answers.

NEW QUESTION 135

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 137

An IAM user receives an Access Denied message when the user attempts to access objects in an Amazon S3 bucket. The user and the S3 bucket are in the same AWS account. The S3 bucket is configured to use server-side encryption with AWS KMS keys (SSE-KMS) to encrypt all of its objects at rest by using a customer managed key from the same AWS account. The S3 bucket has no bucket policy defined. The IAM user has been granted permissions through an IAM policy that allows the kms:Decrypt permission to the customer managed key. The IAM policy also allows the s3:List* and s3:Get* permissions for the S3 bucket and its objects.

Which of the following is a possible reason that the IAM user cannot access the objects in the S3 bucket?

- A. The IAM policy needs to allow the kms:DescribeKey permission.
- B. The S3 bucket has been changed to use the AWS managed key to encrypt objects at rest.
- C. An S3 bucket policy needs to be added to allow the IAM user to access the objects.
- D. The KMS key policy has been edited to remove the ability for the AWS account to have full access to the key.

Answer: D

Explanation:

The possible reason that the IAM user cannot access the objects in the S3 bucket is D. The KMS key policy has been edited to remove the ability for the AWS account to have full access to the key.

This answer is correct because the KMS key policy is the primary way to control access to the KMS key, and it must explicitly allow the AWS account to have full access to the key. If the KMS key policy has been edited to remove this permission, then the IAM policy that grants kms:Decrypt permission to the IAM user has no effect, and the IAM user cannot decrypt the objects in the S3 bucket.

The other options are incorrect because:

- A. The IAM policy does not need to allow the kms:DescribeKey permission, because this permission is not required for decrypting objects in S3 using SSE-KMS. The kms:DescribeKey permission allows getting information about a KMS key, such as its creation date, description, and key state3.
- B. The S3 bucket has not been changed to use the AWS managed key to encrypt objects at rest, because this would not cause an Access Denied message for the IAM user. The AWS managed key is a default KMS key that is created and managed by AWS for each AWS account and Region. The IAM user does not need any permissions on this key to use it for SSE-KMS4.
- C. An S3 bucket policy does not need to be added to allow the IAM user to access the objects, because the IAM user already has s3:List* and s3:Get* permissions for the S3 bucket and its objects through an IAM policy. An S3 bucket policy is an optional way to grant cross-account access or public access to an S3 bucket5.

References:

1: Key policies in AWS KMS 2: Using server-side encryption with AWS KMS keys (SSE-KMS) 3: AWS KMS API Permissions Reference 4: Using server-side encryption with Amazon S3 managed keys (SSE-S3) 5: Bucket policy examples

NEW QUESTION 139

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 143

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 145

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

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


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

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

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

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

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

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

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

Answer: A

Explanation:

The condition of "s3:x-amz-server-side-encryption": "IAM:kms" ensures that objects uploaded need to be encrypted.
Options B,C and D are invalid because you have to ensure the condition of "s3:x-amz-server-side-encryption": "IAM:kms" is present

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

NEW QUESTION 147

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

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

- A. Deploy an AWS Config managed rule to run on a periodic basis of 24 hour
- B. Select the access-keys-rotated managed rule, and set the maxAccessKeyAge parameter to 90 day
- C. Create an Amazon EventBridge (Amazon CloudWatch Events) rule with an event pattern that matches the compliance type of NON_COMPLIANT from AWS Config for the managed rule
- D. Configure EventBridge (CloudWatch Events) to send an Amazon Simple Notification Service (Amazon SNS) notification to the security team.
- E. Create a script to export a .csv file from the AWS Trusted Advisor check for IAM access key rotation. Load the script into an AWS Lambda function that will upload the .csv file to an Amazon S3 bucket
- F. Create an Amazon Athena table query that runs when the .csv file is uploaded to the S3 bucket
- G. Publish the results for any keys older than 90 days by using an invocation of an Amazon Simple Notification Service (Amazon SNS) notification to the security team.
- H. Create a script to download the IAM credentials report on a periodic basis
- I. Load the script into an AWS Lambda function that will run on a schedule through Amazon EventBridge (Amazon CloudWatch Events). Configure the Lambda script to load the report into memory and to filter the report for records in which the key was last rotated at least 90 days ago
- J. If any records are detected, send an Amazon Simple Notification Service (Amazon SNS) notification to the security team.
- K. Create an AWS Lambda function that queries the IAM API to list all the users
- L. Iterate through the users by using the ListAccessKeys operation
- M. Verify that the value in the CreateDate field is not at least 90 days old
- N. Send an Amazon Simple Notification Service (Amazon SNS) notification to the security team if the value is at least 90 days old
- O. Create an Amazon EventBridge (Amazon CloudWatch Events) rule to schedule the Lambda function to run each day.

Answer: A

NEW QUESTION 150

A security engineer needs to run an AWS CloudFormation script. The CloudFormation script builds AWS infrastructure to support a stack that includes web servers and a MySQL database. The stack has been deployed in pre-production environments and is ready for production.

The production script must comply with the principle of least privilege. Additionally, separation of duties must exist between the security engineer's IAM account and CloudFormation.

Which solution will meet these requirements?

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

Answer: A

Explanation:

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

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

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

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

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

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

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

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

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

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

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

NEW QUESTION 152

A company uses AWS Organizations. The company has teams that use an AWS CloudHSM hardware security module (HSM) that is hosted in a central AWS account. One of the teams creates its own new dedicated AWS account and wants to use the HSM that is hosted in the central account. How should a security engineer share the HSM that is hosted in the central account with the new dedicated account?

- A. Use AWS Resource Access Manager (AWS RAM) to share the VPC subnet ID of the HSM that is hosted in the central account with the new dedicated account.
- B. Configure the CloudHSM security group to accept inbound traffic from the private IP addresses of client instances in the new dedicated account.
- C. Use AWS Identity and Access Management (IAM) to create a cross-account role to access the CloudHSM cluster that is in the central account. Create a new IAM user in the new dedicated account. Assign the cross-account role to the new IAM user.
- D. Use AWS IAM Identity Center (AWS Single Sign-On) to create an AWS Security Token Service (AWS STS) token to authenticate from the new dedicated account to the central account.
- E. Use the cross-account permissions that are assigned to the STS token to invoke an operation on the HSM in the central account.
- F. Use AWS Resource Access Manager (AWS RAM) to share the ID of the HSM that is hosted in the central account with the new dedicated account.
- G. Configure the CloudHSM security group to accept inbound traffic from the private IP addresses of client instances in the new dedicated account.

Answer: A

Explanation:

<https://aws.amazon.com/premiumsupport/knowledge-center/cloudhsm-share-clusters/#:~:text=In%20the%20nav>

NEW QUESTION 156

You need to create a policy and apply it for just an individual user. How could you accomplish this in the right way? Please select:

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

Answer: D

Explanation:

Options A and B are incorrect since you need to add an inline policy just for the user. Option C is invalid because you don't assign an IAM role to a user.

The IAM Documentation mentions the following:

An inline policy is a policy that's embedded in a principal entity (a user, group, or role)—that is, the policy is an inherent part of the principal entity. You can create a policy and embed it in a principal entity, either when you create the principal entity or later.

For more information on IAM Access and Inline policies, just browse to the below URL: <https://docs.IAM.amazon.com/IAM/latest/UserGuide/access>

The correct answer is: Add an inline policy for the user. Submit your Feedback/Queries to our Experts

NEW QUESTION 159

.....

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](#)