



Amazon-Web-Services

Exam Questions SCS-C01

AWS Certified Security- Specialty

NEW QUESTION 1

A company is attempting to conduct forensic analysis on an Amazon EC2 instance, but the company is unable to connect to the instance by using AWS Systems Manager Session Manager. The company has installed AWS Systems Manager Agent (SSM Agent) on the EC2 instance. The EC2 instance is in a subnet in a VPC that does not have an internet gateway attached. The company has associated a security group with the EC2 instance. The security group does not have inbound or outbound rules. The subnet's network ACL allows all inbound and outbound traffic. Which combination of actions will allow the company to conduct forensic analysis on the EC2 instance without compromising forensic data? (Select THREE.)

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

Answer: BCF

NEW QUESTION 2

A company uses AWS Organizations to manage a small number of AWS accounts. However, the company plans to add 1 000 more accounts soon. The company allows only a centralized security team to create IAM roles for all AWS accounts and teams. Application teams submit requests for IAM roles to the security team. The security team has a backlog of IAM role requests and cannot review and provision the IAM roles quickly. The security team must create a process that will allow application teams to provision their own IAM roles. The process must also limit the scope of IAM roles and prevent privilege escalation. Which solution will meet these requirements with the LEAST operational overhead?

- A. Create an IAM group for each application team
- B. Associate policies with each IAM group
- C. Provision IAM users for each application team member
- D. Add the new IAM users to the appropriate IAM group by using role-based access control (RBAC).
- E. Delegate application team leads to provision IAM roles for each team
- F. Conduct a quarterly review of the IAM roles the team leads have provisioned
- G. Ensure that the application team leads have the appropriate training to review IAM roles.
- H. Put each AWS account in its own OU
- I. Add an SCP to each OU to grant access to only the AWS services that the teams plan to use
- J. Include conditions in the AWS account of each team.
- K. Create an SCP and a permissions boundary for IAM roles
- L. Add the SCP to the root OU so that only roles that have the permissions boundary attached can create any new IAM roles.

Answer: D

Explanation:

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

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

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

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

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

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

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

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

Verified References:

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

NEW QUESTION 3

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

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

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

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

```

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

```

- B. For each team create an IAM policy similar to the one that follows Populate the IAM TagKeys/Team condition key with a proper team nam
- C. Attach the resuming policies to the corresponding IAM roles.

```

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

```

- D. Tag each IAM role with a Team tag ke
- E. and use the team name in the tag valu
- F. Create an IAM policy similar to the one that follows, and attach 4 to all the IAM roles used by developers.

```

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

```

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

```

{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "NotAction": [
        "ec2:StopInstances",
        "ec2:TerminateInstances"
      ],
      "Resource": "*"
    },
    {
      "Effect": "Allow",
      "Action": [
        "ec2:StopInstances",
        "ec2:TerminateInstances"
      ],
      "Resource": "*",
      "Condition": {
        "ForAnyValue:StringEquals": {
          "aws:TagKeys/Team": "2 (aws:PrincipalTag/Team)"
        }
      }
    }
  ]
}

```

Answer: A

NEW QUESTION 4

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

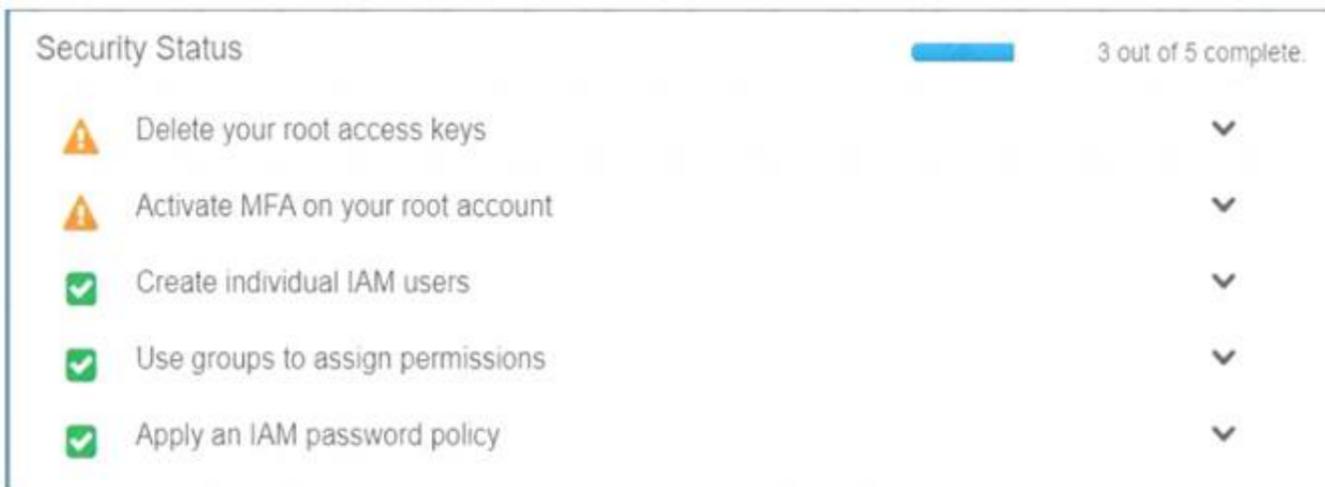
- A. Use short but complex password on the root account and any administrators.
- B. Use IAM IAM Geo-Lock and disallow anyone from logging in except for in your city.
- C. Use MFA on all users and accounts, especially on the root account.
- D. Don't write down or remember the root account password after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

Answer: C

Explanation:

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

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



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

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

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

NEW QUESTION 5

To meet regulatory requirements, a Security Engineer needs to implement an IAM policy that restricts the use of AWS services to the us-east-1 Region. What policy should the Engineer implement?

A.

```

{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": "*",
      "Resource": "*",
      "Condition": {
        "StringEquals": {
          "aws:RequestedRegion": "us-east-1"
        }
      }
    }
  ]
}

```

B. A computer code with black text Description automatically generated

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": "*",
      "Resource": "*",
      "Condition": {
        "StringEquals": {
          "ec2:Region": "us-east-1"
        }
      }
    }
  ]
}
```

C. A computer code with black text Description automatically generated

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Deny",
      "Action": "*",
      "Resource": "*",
      "Condition": {
        "StringNotEquals": {
          "aws:RequestedRegion": "us-east-1"
        }
      }
    }
  ]
}
```

D. A computer code with text Description automatically generated

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Deny",
      "NotAction": "*",
      "Resource": "*",
      "Condition": {
        "StringEquals": {
          "aws:RequestedRegion": "us-east-1"
        }
      }
    }
  ]
}
```

Answer: C

Explanation:

https://docs.aws.amazon.com/IAM/latest/UserGuide/reference_policies_examples_aws_deny-requested-region.h

NEW QUESTION 6

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 7

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

Please select:

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

Answer: A

Explanation:

The IAM Documentation mentions the following

To determine whether a log file was modified, deleted, or unchanged after CloudTrail delivered it you can use CloudTrail log file integrity validation. This feature is built using industry standard algorithms: SHA-256 for hashing and SHA-256 with RSA for digital signing. This makes it computationally infeasible to modify, delete or forge CloudTrail log files without detection. You can use the IAM CLI to validate the files in the location where CloudTrail delivered them

Validated log files are invaluable in security and forensic investigations. For example, a validated log file enables you to assert positively that the log file itself has not changed, or that particular user credentials performed specific API activity. The CloudTrail log file integrity validation process also lets you know if a log file has been deleted or changed, or assert positively that no log files were delivered to your account during a given period of time.

Options B.C and D is invalid because you need to check for log File Integrity Validation for cloudtrail logs For more information on Cloudtrail log file validation, please visit the below URL: <http://docs.IAM.amazon.com/IAMcloudtrail/latest/userguide/cloudtrail-log-file-validation-intro.html>

The correct answer is: Use CloudTrail Log File Integrity Validation. omit your Feedback/Queries to our Expert

NEW QUESTION 8

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

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

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

Answer: B

Explanation:

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

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

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

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

References:

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

NEW QUESTION 9

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

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

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

- A. Create a new CloudTrail trail
- B. Select the new Regions where the company added resources.
- C. Change the S3 bucket to receive notifications to track all actions from all Regions.
- D. Create a new CloudTrail trail that applies to all Regions.
- E. Change the existing CloudTrail trail so that it applies to all Regions.

Answer: D

Explanation:

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

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

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

NEW QUESTION 10

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

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

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

Answer: B

Explanation:

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

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

`{{resolve:secretsmanager:...}}` syntax². This way, you can avoid exposing your secrets in plaintext and still use them in your resources.

The other options are incorrect because:

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

References:

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

NEW QUESTION 10

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

Which solution meets these criteria?

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

Answer: A

Explanation:

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

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

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

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

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

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

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

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

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

References:

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

NEW QUESTION 11

A security engineer is defining the controls required to protect the IAM account root user credentials in an IAM Organizations hierarchy. The controls should also limit the impact in case these credentials have been compromised.

Which combination of controls should the security engineer propose? (Select THREE.)

A)

Apply the following SCP:

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "GRRESTRICTROOTUSER",
      "Effect": "Deny",
      "Action": "*",
      "Resource": [
        "*"
      ],
      "Condition": {
        "StringLike": {
          "aws:PrincipalArn": [
            "arn:aws:iam::*:root"
          ]
        }
      }
    }
  ]
}
```

B)

Apply the following SCP:

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "GRRESTRICTROOTUSER",
      "Effect": "Deny",
      "Principal": "arn:aws:iam::*:root",
      "Action": "*",
      "Resource": [
        "*"
      ]
    }
  ]
}
```

C) Enable multi-factor authentication (MFA) for the root user.

D) Set a strong randomized password and store it in a secure location.

E) Create an access key ID and secret access key, and store them in a secure location.

F) Apply the following permissions boundary to the root user:

```

{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "GRRESTRICTROOTUSER",
      "Effect": "Deny",
      "Action": "*",
      "Resource": [
        "*"
      ],
      "Condition": {
        "StringLike": {
          "aws:PrincipalArn": [
            "arn:aws:iam::*:root"
          ]
        }
      }
    }
  ]
}

```

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

Answer: ACE

NEW QUESTION 16

A company is running an application in The eu-west-1 Region. The application uses an IAM Key Management Service (IAM KMS) CMK to encrypt sensitive data. The company plans to deploy the application in the eu-north-1 Region.

A security engineer needs to implement a key management solution for the application deployment in the new Region. The security engineer must minimize changes to the application code.

Which change should the security engineer make to the IAM KMS configuration to meet these requirements?

- A. Update the key policies in eu-west-1. Point the application in eu-north-1 to use the same CMK as the application in eu-west-1.
- B. Allocate a new CMK to eu-north-1 to be used by the application that is deployed in that Region.
- C. Allocate a new CMK to eu-north-1. Create the same alias name for both key
- D. Configure the application deployment to use the key alias.
- E. Allocate a new CMK to eu-north-1. Create an alias for eu-'-1. Change the application code to point to the alias for eu-'-1.

Answer: B

NEW QUESTION 21

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

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

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

Answer: BD

Explanation:

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

NEW QUESTION 22

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

- A. Use the AWS account root user access keys instead of the AWS Management Console.
- B. Enable multi-factor authentication for the AWS IAM users with the AdministratorAccess managed policy attached to them.

- C. Enable multi-factor authentication for the AWS account root user.
- D. Use AWS KMS to encrypt all AWS account root user and AWS IAM access keys and set automatic rotation to 30 days.
- E. Do not create access keys for the AWS account root user; instead, create AWS IAM users.

Answer: CE

NEW QUESTION 25

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 instanc
- N. Start the instance.
- O. When the volume is detached from the original instance, attach the volume to another instance as a data volum
- P. Modify the authorized_keys file with a new public ke
- 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 28

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

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

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

Answer: A

Explanation:

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

NEW QUESTION 33

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

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

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

Answer: DE

Explanation:

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

NEW QUESTION 34

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

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

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

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

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

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

Answer: C

Explanation:

The correct answer is C. Use AWS Resource Access Manager (AWS RAM) to share Account-A's VPC subnets with the remaining member accounts. Configure the member accounts to use the shared subnets to launch workloads.

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

The other options are incorrect because:

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

References:

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

NEW QUESTION 37

A security engineer is designing an IAM policy for a script that will use the AWS CLI. The script currently assumes an IAM role that is attached to three AWS managed IAM policies: AmazonEC2FullAccess, AmazonDynamoDBFullAccess, and AmazonVPCFullAccess.

The security engineer needs to construct a least privilege IAM policy that will replace the AWS managed IAM policies that are attached to this role.

Which solution will meet these requirements in the MOST operationally efficient way?

- A. In AWS CloudTrail, create a trail for management event
- B. Run the script with the existing AWS managed IAM policies
- C. Use IAM Access Analyzer to generate a new IAM policy that is based on access activity in the trail
- D. Replace the existing AWS managed IAM policies with the generated IAM policy for the role.
- E. Remove the existing AWS managed IAM policies from the role
- F. Attach the IAM Access Analyzer Role Policy Generator to the role
- G. Run the script
- H. Return to IAM Access Analyzer and generate a least privilege IAM policy
- I. Attach the new IAM policy to the role.
- J. Create an account analyzer in IAM Access Analyzer
- K. Create an archive rule that has a filter that checks whether the PrincipalArn value matches the ARN of the role
- L. Run the script
- M. Remove the existing AWS managed IAM policies from the role.
- N. In AWS CloudTrail, create a trail for management event
- O. Remove the existing AWS managed IAM policies from the role
- P. Run the script
- Q. Find the authorization failure in the trail event that is associated with the script
- R. Create a new IAM policy that includes the action and resource that caused the authorization failure
- S. Repeat the process until the script succeeds
- T. Attach the new IAM policy to the role.

Answer: A

NEW QUESTION 40

A company maintains an open-source application that is hosted on a public GitHub repository. While creating a new commit to the repository, an engineer uploaded their IAM access key and secret access key. The engineer reported the mistake to a manager, and the manager immediately disabled the access key. The company needs to assess the impact of the exposed access key. A security engineer must recommend a solution that requires the least possible managerial

overhead.

Which solution meets these requirements?

- A. Analyze an IAM Identity and Access Management (IAM) use report from IAM Trusted Advisor to see when the access key was last used.
- B. Analyze Amazon CloudWatch Logs for activity by searching for the access key.
- C. Analyze VPC flow logs for activity by searching for the access key
- D. Analyze a credential report in IAM Identity and Access Management (IAM) to see when the access key was last used.

Answer: A

Explanation:

To assess the impact of the exposed access key, the security engineer should recommend the following solution:

➤ Analyze an IAM use report from AWS Trusted Advisor to see when the access key was last used. This allows the security engineer to use a tool that provides information about IAM entities and credentials in their account, and check if there was any unauthorized activity with the exposed access key.

NEW QUESTION 44

A security engineer is working with a company to design an ecommerce application. The application will run on Amazon EC2 instances that run in an Auto Scaling group behind an Application Load Balancer (ALB). The application will use an Amazon RDS DB instance for its database. The only required connectivity from the internet is for HTTP and HTTPS traffic to the application. The application must communicate with an external payment provider that allows traffic only from a preconfigured allow list of IP addresses. The company must ensure that communications with the external payment provider are not interrupted as the environment scales.

Which combination of actions should the security engineer recommend to meet these requirements? (Select THREE.)

- A. Deploy a NAT gateway in each private subnet for every Availability Zone that is in use.
- B. Place the DB instance in a public subnet.
- C. Place the DB instance in a private subnet.
- D. Configure the Auto Scaling group to place the EC2 instances in a public subnet.
- E. Configure the Auto Scaling group to place the EC2 instances in a private subnet.
- F. Deploy the ALB in a private subnet.

Answer: ACE

NEW QUESTION 46

A company's IAM account consists of approximately 300 IAM users. Now there is a mandate that an access change is required for 100 IAM users to have unlimited privileges to S3. As a system administrator, how can you implement this effectively so that there is no need to apply the policy at the individual user level? Please select:

- A. Create a new role and add each user to the IAM role
- B. Use the IAM groups and add users, based upon their role, to different groups and apply the policy to group
- C. Create a policy and apply it to multiple users using a JSON script
- D. Create an S3 bucket policy with unlimited access which includes each user's IAM account ID

Answer: B

Explanation:

Option A is incorrect since you don't add a user to the IAM Role Option C is incorrect since you don't assign multiple users to a policy Option D is incorrect since this is not an ideal approach

An IAM group is used to collectively manage users who need the same set of permissions. By having groups, it becomes easier to manage permissions. So if you change the permissions on the group scale, it will affect all the users in that group

For more information on IAM Groups, just browse to the below URL: https://docs.IAM.amazon.com/IAM/latest/UserGuide/id_eroups.html

The correct answer is: Use the IAM groups and add users, based upon their role, to different groups and apply the policy to group

Submit your Feedback/Queries to our Experts

NEW QUESTION 48

A business stores website images in an Amazon S3 bucket. The firm serves the photos to end users through Amazon CloudFront. The firm learned lately that the photographs are being accessible from nations in which it does not have a distribution license.

Which steps should the business take to safeguard the photographs and restrict their distribution? (Select two.)

- A. Update the S3 bucket policy to restrict access to a CloudFront origin access identity (OAI).
- B. Update the website DNS record to use an Amazon Route 53 geolocation record deny list of countries where the company lacks a license.
- C. Add a CloudFront geo restriction deny list of countries where the company lacks a license.
- D. Update the S3 bucket policy with a deny list of countries where the company lacks a license.
- E. Enable the Restrict Viewer Access option in CloudFront to create a deny list of countries where the company lacks a license.

Answer: AC

Explanation:

For Enable Geo-Restriction, choose Yes. For Restriction Type, choose Whitelist to allow access to certain countries, or choose Blacklist to block access from certain countries. <https://IAM.amazon.com/premiumsupport/knowledge-center/cloudfront-geo-restriction/>

NEW QUESTION 49

A company uses Amazon EC2 Linux instances in the AWS Cloud. A member of the company's security team recently received a report about common vulnerability identifiers on the instances.

A security engineer needs to verify patching and perform remediation if the instances do not have the correct patches installed. The security engineer must determine which EC2 instances are at risk and must implement a solution to automatically update those instances with the applicable patches.

What should the security engineer do to meet these requirements?

- A. Use AWS Systems Manager Patch Manager to view vulnerability identifiers for missing patches on the instance

- B. Use Patch Manager also to automate the patching process.
- C. Use AWS Shield Advanced to view vulnerability identifiers for missing patches on the instance
- D. Use AWS Systems Manager Patch Manager to automate the patching process.
- E. Use Amazon GuardDuty to view vulnerability identifiers for missing patches on the instance
- F. Use Amazon Inspector to automate the patching process.
- G. Use Amazon Inspector to view vulnerability identifiers for missing patches on the instance
- H. Use Amazon Inspector also to automate the patching process.

Answer: A

Explanation:

<https://aws.amazon.com/about-aws/whats-new/2020/10/now-use-aws-systems-manager-to-view-vulnerability-id>

NEW QUESTION 50

A company became aware that one of its access keys was exposed on a code sharing website 11 days ago. A Security Engineer must review all use of the exposed access keys to determine the extent of the exposure. The company enabled IAM CloudTrail in all regions when it opened the account. Which of the following will allow the Security Engineer to complete the task?

- A. Filter the event history on the exposed access key in the CloudTrail console. Examine the data from the past 11 days.
- B. Use the IAM CLI to generate an IAM credential report. Extract all the data from the past 11 days.
- C. Use Amazon Athena to query the CloudTrail logs from Amazon S3. Retrieve the rows for the exposed access key for the past 11 days.
- D. Use the Access Advisor tab in the IAM console to view all of the access key activity for the past 11 days.

Answer: C

Explanation:

Amazon Athena is a service that enables you to analyze data in Amazon S3 using standard SQL¹. You can use Athena to query the CloudTrail logs that are stored in S3 and filter them by the exposed access key and the date range². The other options are not effective ways to review the use of the exposed access key.

NEW QUESTION 55

A company is deploying an Amazon EC2-based application. The application will include a custom health-checking component that produces health status data in JSON format. A Security Engineer must implement a secure solution to monitor application availability in near-real time by analyzing the health status data. Which approach should the Security Engineer use?

- A. Use Amazon CloudWatch monitoring to capture Amazon EC2 and networking metrics. Visualize metrics using Amazon CloudWatch dashboards.
- B. Run the Amazon Kinesis Agent to write the status data to Amazon Kinesis Data Firehose. Store the streaming data from Kinesis Data Firehose in Amazon Redshift.
- C. Run a script on the pool data and analyze the data in Amazon Redshift.
- D. Write the status data directly to a public Amazon S3 bucket from the health-checking component. Configure S3 events to invoke an IAM Lambda function that analyzes the data.
- E. Generate events from the health-checking component and send them to Amazon CloudWatch Events. Include the status data as event payload.
- F. Use CloudWatch Events rules to invoke an IAM Lambda function that analyzes the data.

Answer: A

Explanation:

Amazon CloudWatch monitoring is a service that collects and tracks metrics from AWS resources and applications, and provides visualization tools and alarms to monitor performance and availability¹. The health status data in JSON format can be sent to CloudWatch as custom metrics², and then displayed in CloudWatch dashboards³. The other options are either inefficient or insecure for monitoring application availability in near-real time.

NEW QUESTION 56

You have an S3 bucket defined in IAM. You want to ensure that you encrypt the data before sending it across the wire. What is the best way to achieve this. Please select:

- A. Enable server side encryption for the S3 bucket.
- B. This request will ensure that the data is encrypted first.
- C. Use the IAM Encryption CLI to encrypt the data first.
- D. Use a Lambda function to encrypt the data before sending it to the S3 bucket.
- E. Enable client encryption for the bucket.

Answer: B

Explanation:

One can use the IAM Encryption CLI to encrypt the data before sending it across to the S3 bucket. Options A and C are invalid because this would still mean that data is transferred in plain text. Option D is invalid because you cannot just enable client side encryption for the S3 bucket. For more information on Encrypting and Decrypting data, please visit the below URL:

<https://IAM.amazonaws.com/blogs/security/how-to-encrypt-and-decrypt-your-data-with-the-IAM-encryption-cli> The correct answer is: Use the IAM Encryption CLI to encrypt the data first. Submit your Feedback/Queries to our Experts

NEW QUESTION 60

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

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

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

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

How should the security engineer correct the error?

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

Answer: D

Explanation:

CloudWatchLogsReadOnlyAccess doesn't include "logs:CreateLogStream" but it includes "logs:Get*" <https://docs.aws.amazon.com/AmazonCloudWatch/latest/logs/iam-identity-based-access-control-cwl.html#:~:te>

NEW QUESTION 65

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

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

Answer: C

Explanation:

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

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

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

- > Create a vault in S3 Glacier using the AWS Management Console, AWS Command Line Interface (AWS CLI), or AWS SDKs.
- > Create a Vault Lock policy in S3 Glacier that meets all the regulatory requirements using the IAM policy language. The policy can include conditions such as aws:CurrentTime or aws:SecureTransport to further restrict access to the vault.
- > Initiate the lock by attaching the Vault Lock policy to the vault, which sets the lock to an in-progress state and returns a lock ID. While the policy is in the in-progress state, you have 24 hours to validate your Vault Lock policy before the lock ID expires. To prevent your vault from exiting the in-progress state, you must complete the Vault Lock process within these 24 hours. Otherwise, your Vault Lock policy will be deleted.
- > Use the lock ID to complete the lock process. If the Vault Lock policy doesn't work as expected, you can stop the Vault Lock process and restart from the beginning.
- > Upload the data to the vault using either direct upload or multipart upload methods. For more information about S3 Glacier and Vault Lock, see S3 Glacier Vault Lock.

The other options are incorrect because:

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

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

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

NEW QUESTION 69

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

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

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

Answer: A

Explanation:

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

- Configure AWS KMS and use a custom key store. This allows the security engineer to use a key manager outside of AWS KMS that they own and manage, such as an AWS CloudHSM cluster or an external key manager.
- Create a customer managed CMK with no key material. Import the company's keys and key material into the CMK. This allows the security engineer to use their own key material for encryption and decryption operations, and to specify an expiration date for it.

NEW QUESTION 73

A security engineer wants to evaluate configuration changes to a specific AWS resource to ensure that the resource meets compliance standards. However, the security engineer is concerned about a situation in which several configuration changes are made to the resource in quick succession. The security engineer wants to record only the latest configuration of that resource to indicate the cumulative impact of the set of changes.

Which solution will meet this requirement in the MOST operationally efficient way?

- A. Use AWS CloudTrail to detect the configuration changes by filtering API calls to monitor the changes.Use the most recent API call to indicate the cumulative impact of multiple calls
- B. Use AWS Config to detect the configuration changes and to record the latest configuration in case of multiple configuration changes.
- C. Use Amazon CloudWatch to detect the configuration changes by filtering API calls to monitor the change
- D. Use the most recent API call to indicate the cumulative impact of multiple calls.
- E. Use AWS Cloud Map to detect the configuration change
- F. Generate a report of configuration changes from AWS Cloud Map to track the latest state by using a sliding time window.

Answer: B

Explanation:

AWS Config is a service that enables you to assess, audit, and evaluate the configurations of your AWS resources. AWS Config continuously monitors and records your AWS resource configurations and allows you to automate the evaluation of recorded configurations against desired configurations.

To evaluate configuration changes to a specific AWS resource and ensure that it meets compliance standards, the security engineer should use AWS Config to detect the configuration changes and to record the latest configuration in case of multiple configuration changes. This will allow the security engineer to view the current state of the resource and its compliance status, as well as its configuration history and timeline.

AWS Config records configuration changes as ConfigurationItems, which are point-in-time snapshots of the resource's attributes, relationships, and metadata. If multiple configuration changes occur within a short period of time, AWS Config records only the latest ConfigurationItem for that resource. This indicates the cumulative impact of the set of changes on the resource's configuration.

This solution will meet the requirement in the most operationally efficient way, as it leverages AWS Config's features to monitor, record, and evaluate resource configurations without requiring additional tools or services.

The other options are incorrect because they either do not record the latest configuration in case of multiple configuration changes (A, C), or do not use a valid service for evaluating resource configurations (D).

Verified References:

- <https://docs.aws.amazon.com/config/latest/developerguide/WhatIsConfig.html>
- <https://docs.aws.amazon.com/config/latest/developerguide/config-item-table.html>

NEW QUESTION 76

A company wants to configure DNS Security Extensions (DNSSEC) for the company's primary domain. The company registers the domain with Amazon Route 53. The company hosts the domain on Amazon EC2 instances by using BIND.

What is the MOST operationally efficient solution that meets this requirement?

- A. Set the dnssec-enable option to yes in the BIND configuratio
- B. Create a zone-signing key (ZSK) and a key-signing key (KSK) Restart the BIND service.
- C. Migrate the zone to Route 53 with DNSSEC signing enable
- D. Create a zone-signing key (ZSK) and a key-signing key (KSK) that are based on an AW
- E. Key Management Service (AWS KMS) customer managed key.
- F. Set the dnssec-enable option to yes in the BIND configuratio

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

Answer: D

Explanation:

To configure DNSSEC for a domain registered with Route 53, the most operationally efficient solution is to migrate the zone to Route 53 with DNSSEC signing enabled, create a key-signing key (KSK) that is based on an AWS Key Management Service (AWS KMS) customer managed key, and add a delegation signer (DS) record to the parent zone. This way, Route 53 handles the zone-signing key (ZSK) and the signing of the records in the hosted zone, and the customer only needs to manage the KSK in AWS KMS and provide the DS record to the domain registrar. Option A is incorrect because it does not involve migrating the zone to Route 53, which would simplify the DNSSEC configuration. Option B is incorrect because it creates both a ZSK and a KSK based on AWS KMS customer managed keys, which is unnecessary and less efficient than letting Route 53 manage the ZSK. Option C is incorrect because it does not involve migrating the zone to Route 53, and it requires running the dnssec-signzone command manually, which is less efficient than letting Route 53 sign the zone automatically. Verified

References:

- > <https://docs.aws.amazon.com/Route53/latest/DeveloperGuide/domain-configure-dnssec.html>
- > <https://aws.amazon.com/about-aws/whats-new/2020/12/announcing-amazon-route-53-support-dnssec/>

NEW QUESTION 79

A developer signed in to a new account within an IAM Organization organizational unit (OU) containing multiple accounts. Access to the Amazon S3 service is restricted with the following SCP.

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Deny",
      "Action": "s3:*",
      "Resource": "*"
    }
  ]
}
```

How can the security engineer provide the developer with Amazon S3 access without affecting other account?

- A. Move the SCP to the root OU of organization to remove the restriction to access Amazon S3.
- B. Add an IAM policy for the developer, which grants S3 access.
- C. Create a new OU without applying the SCP restricting S3 access
- D. Move the developer account to this new OU.
- E. Add an allow list for the developer account for the S3 service.

Answer: C

NEW QUESTION 81

A company's Security Engineer has been tasked with restricting a contractor's IAM account access to the company's Amazon EC2 console without providing access to any other AWS services. The contractor's IAM account must not be able to gain access to any other AWS service, even if the IAM account is assigned additional permissions based on IAM group membership.

What should the Security Engineer do to meet these requirements?

- A. Create an Inline IAM user policy that allows for Amazon EC2 access for the contractor's IAM user.
- B. Create an IAM permissions boundary policy that allows Amazon EC2 access
- C. Associate the contractor's IAM account with the IAM permissions boundary policy.
- D. Create an IAM group with an attached policy that allows for Amazon EC2 access
- E. Associate the contractor's IAM account with the IAM group.
- F. Create an IAM role that allows for EC2 and explicitly denies all other service
- G. Instruct the contractor to always assume this role.

Answer: B

NEW QUESTION 84

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

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

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

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

L. Schedule the compromised key for deletion.

Answer: D

NEW QUESTION 85

A company is running internal microservices on Amazon Elastic Container Service (Amazon ECS) with the Amazon EC2 launch type. The company is using Amazon Elastic Container Registry (Amazon ECR) private repositories. A security engineer needs to encrypt the private repositories by using AWS Key Management Service (AWS KMS). The security engineer also needs to analyze the container images for any common vulnerabilities and exposures (CVEs). Which solution will meet these requirements?

- A. Enable KMS encryption on the existing ECR repository
- B. Install Amazon Inspector Agent from the ECS container instances' user data
- C. Run an assessment with the CVE rules.
- D. Recreate the ECR repositories with KMS encryption and ECR scanning enable
- E. Analyze the scan report after the next push of images.
- F. Recreate the ECR repositories with KMS encryption and ECR scanning enable
- G. Install AWS Systems Manager Agent on the ECS container instance
- H. Run an inventory report.
- I. Enable KMS encryption on the existing ECR repository
- J. Use AWS Trusted Advisor to check the ECS container instances and to verify the findings against a list of current CVEs.

Answer: B

NEW QUESTION 89

A company wants to receive an email notification about critical findings in AWS Security Hub. The company does not have an existing architecture that supports this functionality. Which solution will meet the requirement?

- A. Create an AWS Lambda function to identify critical Security Hub finding
- B. Create an Amazon Simple Notification Service (Amazon SNS) topic as the target of the Lambda function
- C. Subscribe an email endpoint to the SNS topic to receive published messages.
- D. Create an Amazon Kinesis Data Firehose delivery stream
- E. Integrate the delivery stream with Amazon EventBridge
- F. Create an EventBridge rule that has a filter to detect critical Security Hub finding
- G. Configure the delivery stream to send the findings to an email address.
- H. Create an Amazon EventBridge rule to detect critical Security Hub finding
- I. Create an Amazon Simple Notification Service (Amazon SNS) topic as the target of the EventBridge rule
- J. Subscribe an email endpoint to the SNS topic to receive published messages.
- K. Create an Amazon EventBridge rule to detect critical Security Hub finding
- L. Create an Amazon Simple Email Service (Amazon SES) topic as the target of the EventBridge rule
- M. Use the Amazon SES API to format the message
- N. Choose an email address to be the recipient of the message.

Answer: C

Explanation:

This solution meets the requirement of receiving an email notification about critical findings in AWS Security Hub. Amazon EventBridge is a serverless event bus that can receive events from AWS services and third-party sources, and route them to targets based on rules and filters. Amazon SNS is a fully managed pub/sub service that can send messages to various endpoints, such as email, SMS, mobile push, and HTTP. By creating an EventBridge rule that detects critical Security Hub findings and sends them to an SNS topic, the company can leverage the existing integration between these services and avoid writing custom code or managing servers. By subscribing an email endpoint to the SNS topic, the company can receive published messages in their inbox.

NEW QUESTION 92

An audit determined that a company's Amazon EC2 instance security group violated company policy by allowing unrestricted incoming SSH traffic. A security engineer must implement a near-real-time monitoring and alerting solution that will notify administrators of such violations.

Which solution meets these requirements with the MOST operational efficiency?

- A. Create a recurring Amazon Inspector assessment run that runs every day and uses the Network Reachability package
- B. Create an Amazon CloudWatch rule that invokes an IAM Lambda function when an assessment run starts
- C. Configure the Lambda function to retrieve and evaluate the assessment run report when it completes
- D. Configure the Lambda function also to publish an Amazon Simple Notification Service (Amazon SNS) notification if there are any violations for unrestricted incoming SSH traffic.
- E. Use the restricted-ssh IAM Config managed rule that is invoked by security group configuration changes that are not compliant
- F. Use the IAM Config remediation feature to publish a message to an Amazon Simple Notification Service (Amazon SNS) topic.
- G. Configure VPC Flow Logs for the VPC
- H. and specify an Amazon CloudWatch Logs group
- I. Subscribe the CloudWatch Logs group to an IAM Lambda function that parses new log entries, detects successful connections on port 22, and publishes a notification through Amazon Simple Notification Service (Amazon SNS).
- J. Create a recurring Amazon Inspector assessment run that runs every day and uses the Security Best Practices package
- K. Create an Amazon CloudWatch rule that invokes an IAM Lambda function when an assessment run starts
- L. Configure the Lambda function to retrieve and evaluate the assessment run report when it completes
- M. Configure the Lambda function also to publish an Amazon Simple Notification Service (Amazon SNS) notification if there are any violations for unrestricted incoming SSH traffic.

Answer: B

Explanation:

The most operationally efficient solution to implement a near-real-time monitoring and alerting solution that will notify administrators of security group violations is

to use the restricted-ssh AWS Config managed rule that is invoked by security group configuration changes that are not compliant. This rule checks whether security groups that are in use have inbound rules that allow unrestricted SSH traffic. If a violation is detected, AWS Config can use the remediation feature to publish a message to an Amazon Simple Notification Service (Amazon SNS) topic.

Option A is incorrect because creating a recurring Amazon Inspector assessment run that uses the Network Reachability package is not operationally efficient, as it requires setting up an assessment target and template, running the assessment every day, and invoking a Lambda function to retrieve and evaluate the assessment report. It also does not provide near-real-time monitoring and alerting, as it depends on the frequency and duration of the assessment run.

Option C is incorrect because configuring VPC Flow Logs for the VPC and specifying an Amazon CloudWatch Logs group is not operationally efficient, as it requires creating a log group and stream, enabling VPC Flow Logs for each subnet or network interface, and subscribing a Lambda function to parse and analyze the log entries. It also does not provide proactive monitoring and alerting, as it only detects successful connections on port 22 after they have occurred.

Option D is incorrect because creating a recurring Amazon Inspector assessment run that uses the Security

Best Practices package is not operationally efficient, for the same reasons as option A. It also does not provide specific monitoring and alerting for security group violations, as it covers a broader range of security issues. References:

- > [AWS Config Rules]
- > [AWS Config Remediation]
- > [Amazon Inspector]
- > [VPC Flow Logs]

NEW QUESTION 93

A company hosts a web application on an Apache web server. The application runs on Amazon EC2 instances that are in an Auto Scaling group. The company configured the EC2 instances to send the Apache web server logs to an Amazon CloudWatch Logs group that the company has configured to expire after 1 year. Recently, the company discovered in the Apache web server logs that a specific IP address is sending suspicious requests to the web application. A security engineer wants to analyze the past week of Apache web server logs to determine how many requests that the IP address sent and the corresponding URLs that the IP address requested.

What should the security engineer do to meet these requirements with the LEAST effort?

- A. Export the CloudWatch Logs group data to Amazon S3. Use Amazon Macie to query the logs for the specific IP address and the requested URLs.
- B. Configure a CloudWatch Logs subscription to stream the log group to an Amazon OpenSearch Service cluster.
- C. Use OpenSearch Service to analyze the logs for the specific IP address and the requested URLs.
- D. Use CloudWatch Logs Insights and a custom query syntax to analyze the CloudWatch logs for the specific IP address and the requested URLs.
- E. Export the CloudWatch Logs group data to Amazon S3. Use AWS Glue to crawl the S3 bucket for only the log entries that contain the specific IP address.
- F. Use AWS Glue to view the results.

Answer: C

NEW QUESTION 97

A company is testing its incident response plan for compromised credentials. The company runs a database on an Amazon EC2 instance and stores the sensitive database credentials as a secret in AWS Secrets Manager. The secret has rotation configured with an AWS Lambda function that uses the generic rotation function template. The EC2 instance and the Lambda function are deployed in the same private subnet. The VPC has a Secrets Manager VPC endpoint.

A security engineer discovers that the secret cannot rotate. The security engineer determines that the VPC endpoint is working as intended. The Amazon CloudWatch logs contain the following error:

"setSecret: Unable to log into database". Which solution will resolve this error?

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

Answer: B

Explanation:

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

NEW QUESTION 102

A company needs complete encryption of the traffic between external users and an application. The company hosts the application on a fleet of Amazon EC2 instances that run in an Auto Scaling group behind an Application Load Balancer (ALB).

How can a security engineer meet these requirements?

- A. Create a new Amazon-issued certificate in AWS Secrets Manager.
- B. Export the certificate from Secrets Manager.
- C. Import the certificate into the ALB and the EC2 instances.
- D. Create a new Amazon-issued certificate in AWS Certificate Manager (ACM). Associate the certificate with the ALB.
- E. Export the certificate from ACM.
- F. Install the certificate on the EC2 instances.
- G. Import a new third-party certificate into AWS Identity and Access Management (IAM). Export the certificate from IAM.
- H. Associate the certificate with the ALB and the EC2 instances.
- I. Import a new third-party certificate into AWS Certificate Manager (ACM). Associate the certificate with the ALB.
- J. Install the certificate on the EC2 instances.

Answer: D

Explanation:

The correct answer is D. Import a new third-party certificate into AWS Certificate Manager (ACM). Associate the certificate with the ALB. Install the certificate on the EC2 instances.

This answer is correct because it meets the requirements of complete encryption of the traffic between external users and the application. By importing a third-party certificate into ACM, the security engineer can use it to secure the communication between the ALB and the clients. By installing the same certificate on the EC2 instances, the security engineer can also secure the communication between the ALB and the instances. This way, both the front-end and back-end connections are encrypted with SSL/TLS1.

The other options are incorrect because:

- A. Creating a new Amazon-issued certificate in AWS Secrets Manager is not a solution, because AWS Secrets Manager is not a service for issuing certificates, but for storing and managing secrets such as database credentials and API keys2. AWS Secrets Manager does not integrate with ALB or EC2 for certificate deployment.
- B. Creating a new Amazon-issued certificate in AWS Certificate Manager (ACM) and exporting it from ACM is not a solution, because ACM does not allow exporting Amazon-issued certificates3. ACM only allows exporting private certificates that are issued by an AWS Private Certificate Authority (CA)4.
- C. Importing a new third-party certificate into AWS Identity and Access Management (IAM) is not a solution, because IAM is not a service for managing certificates, but for controlling access to AWS resources5. IAM does not integrate with ALB or EC2 for certificate deployment.

References:

1: How SSL/TLS works 2: What is AWS Secrets Manager? 3: Exporting an ACM Certificate 4: Exporting Private Certificates from ACM 5: What is IAM?

NEW QUESTION 103

A company is using AWS to run a long-running analysis process on data that is stored in Amazon S3 buckets. The process runs on a fleet of Amazon EC2 instances that are in an Auto Scaling group. The EC2 instances are deployed in a private subnet Of a VPC that does not have internet access. The EC2 instances and the S3 buckets are in the same AWS account

The EC2 instances access the S3 buckets through an S3 gateway endpoint that has the default access policy. Each EC2 instance is associated With an instance profile role that has a policy that explicitly allows the s3:GetObject action and the s3:PutObject action for only the required S3 buckets.

The company learns that one or more of the EC2 instances are compromised and are exfiltrating data to an S3 bucket that is outside the companys organization in AWS Organizations. A security engtneer must implement a solution to stop this exfiltration of data and to keep the EC2 processing job functional.

Which solution will meet these requirements?

- A. Update the policy on the S3 gateway endpoint to allow the S3 actions CY11y if the values of the aws:ResourceOrgID and aws:PrincipalOrgID condition keys match the companys values.
- B. Update the policy on the instance profile role to allow the S3 actions only if the value of the aws:ResourceOrgID condition key matches the company's value.
- C. Add a network ACL rule to the subnet of the EC2 instances to block outgoing connections on port 443.
- D. Apply an SCP on the AWS account to allow the \$3 actions only if the values of the aws:ResourceOrgID and aws:PrincipalOrgID condition keys match the company's values.

Answer: D

Explanation:

The correct answer is D.

To stop the data exfiltration from the compromised EC2 instances, the security engineer needs to implement a solution that can deny access to any S3 bucket that is outside the company's organization. The solution should also allow the EC2 instances to access the required S3 buckets within the company's organization for the analysis process.

Option A is incorrect because updating the policy on the S3 gateway endpoint will not affect the access to S3 buckets that are outside the company's organization. The S3 gateway endpoint only applies to S3 buckets that are in the same AWS Region as the VPC. The compromised EC2 instances can still access S3 buckets in other Regions or other AWS accounts through the internet gateway or NAT device.

Option B is incorrect because updating the policy on the instance profile role will not prevent the compromised EC2 instances from using other credentials or methods to access S3 buckets outside the company's organization. The instance profile role only applies to requests that are made using the credentials of that role. The compromised EC2 instances can still use other IAM users, roles, or access keys to access S3 buckets outside the company's organization.

Option C is incorrect because adding a network ACL rule to block outgoing connections on port 443 will also block legitimate connections to S3 buckets within the company's organization. The network ACL rule will prevent the EC2 instances from accessing any S3 bucket through HTTPS, regardless of whether it is inside or outside the company's organization.

Option D is correct because applying an SCP on the AWS account will effectively deny access to any S3 bucket that is outside the company's organization. The SCP will apply to all IAM users, roles, and resources in the AWS account, regardless of how they access S3. The SCP will use the aws:ResourceOrgID and aws:PrincipalOrgID condition keys to check whether the S3 bucket and the principal belong to the same organization as the AWS account. If they do not match, the SCP will deny the S3 actions.

References:

- Using service control policies
- AWS Organizations service control policy examples

NEW QUESTION 107

A company has an organization in AWS Organizations. The company wants to use AWS CloudFormation StackSets in the organization to deploy various AWS design patterns into environments. These patterns consist of Amazon EC2 instances, Elastic Load Balancing (ELB) load balancers, Amazon RDS databases, and Amazon Elastic Kubernetes Service (Amazon EKS) clusters or Amazon Elastic Container Service (Amazon ECS) clusters.

Currently, the company's developers can create their own CloudFormation stacks to increase the overall speed of delivery. A centralized CI/CD pipeline in a shared services AWS account deploys each CloudFormation stack.

The company's security team has already provided requirements for each service in accordance with internal standards. If there are any resources that do not comply with the internal standards, the security team must receive notification to take appropriate action. The security team must implement a notification solution that gives developers the ability to maintain the same overall delivery speed that they currently have.

Which solution will meet these requirements in the MOST operationally efficient way?

- A. Create an Amazon Simple Notification Service (Amazon SNS) topic
- B. Subscribe the security team's email addresses to the SNS topic
- C. Create a custom AWS Lambda function that will run the aws cloudformation validate-template AWS CLI command on all CloudFormation templates before the build stage in the CI/CD pipeline
- D. Configure the CI/CD pipeline to publish a notification to the SNS topic if any issues are found.
- E. Create an Amazon Simple Notification Service (Amazon SNS) topic
- F. Subscribe the security team's email addresses to the SNS topic
- G. Create custom rules in CloudFormation Guard for each resource configuratio
- H. In the CI/CD pipeline, before the build stage, configure a Docker image to run the cfn-guard command on the CloudFormation template

- I. Configure the CI/CD pipeline to publish a notification to the SNS topic if any issues are found.
- J. Create an Amazon Simple Notification Service (Amazon SNS) topic and an Amazon Simple Queue Service (Amazon SQS) queue
- K. Subscribe the security team's email addresses to the SNS topic
- L. Create an Amazon S3 bucket in the shared services AWS account
- M. Include an event notification to publish to the SQS queue when new objects are added to the S3 bucket
- N. Require the developers to put their CloudFormation templates in the S3 bucket
- O. Launch EC2 instances that automatically scale based on the SQS queue depth
- P. Configure the EC2 instances to use CloudFormation Guard to scan the templates and deploy the templates if there are no issues
- Q. Configure the CI/CD pipeline to publish a notification to the SNS topic if any issues are found.
- R. Create a centralized CloudFormation stack set that includes a standard set of resources that the developers can deploy in each AWS account
- S. Configure each CloudFormation template to meet the security requirements
- T. For any new resources or configurations, update the CloudFormation template and send the template to the security team for review
- U. When the review is completed, add the new CloudFormation stack to the repository for the developers to use.

Answer: B

NEW QUESTION 109

A recent security audit found that IAM CloudTrail logs are insufficiently protected from tampering and unauthorized access. Which actions must the Security Engineer take to address these audit findings? (Select THREE)

- A. Ensure CloudTrail log file validation is turned on
- B. Configure an S3 lifecycle rule to periodically archive CloudTrail logs into Glacier for long-term storage
- C. Use an S3 bucket with tight access controls that exists in a separate account
- D. Use Amazon Inspector to monitor the file integrity of CloudTrail log files.
- E. Request a certificate through ACM and use a generated certificate private key to encrypt CloudTrail log files
- F. Encrypt the CloudTrail log files with server-side encryption with IAM KMS-managed keys (SSE-KMS)

Answer: ADE

NEW QUESTION 110

An application team wants to use IAM Certificate Manager (ACM) to request public certificates to ensure that data is secured in transit. The domains that are being used are not currently hosted on Amazon Route 53

The application team wants to use an IAM managed distribution and caching solution to optimize requests to its systems and provide better points of presence to customers. The distribution solution will use a primary domain name that is customized. The distribution solution also will use several alternative domain names. The certificates must renew automatically over an indefinite period of time

Which combination of steps should the application team take to deploy this architecture? (Select THREE.)

- A. Request a certificate from ACM in the us-west-2 Region. Add the domain names that the certificate will secure
- B. Send an email message to the domain administrators to request vacation of the domains for ACM
- C. Request validation of the domains for ACM through DNS. Insert CNAME records into each domain's DNS zone
- D. Create an Application Load Balancer for the caching solution. Select the newly requested certificate from ACM to be used for secure connections
- E. Create an Amazon CloudFront distribution for the caching solution. Enter the main CNAME record as the Origin Name. Enter the subdomain names or alternate names in the Alternate Domain Names Distribution Settings. Select the newly requested certificate from ACM to be used for secure connections
- F. Request a certificate from ACM in the us-east-1 Region. Add the domain names that the certificate will secure

Answer: CDE

NEW QUESTION 112

A company's Security Engineer is copying all application logs to centralized Amazon S3 buckets. Currently, each of the company's applications is in its own IAM account, and logs are pushed into S3 buckets associated with each account. The Engineer will deploy an IAM Lambda function into each account that copies the relevant log files to the centralized S3 bucket.

The Security Engineer is unable to access the log files in the centralized S3 bucket. The Engineer's IAM user policy from the centralized account looks like this:

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Action": "s3:Put*",
      "Resource": "arn:aws:s3:::centralizedbucket/*",
      "Effect": "Deny"
    },
    {
      "Action": ["s3:Get*", "s3:List*"],
      "Resource": [
        "arn:aws:s3:::centralizedbucket/*",
        "arn:aws:s3:::centralizedbucket/"
      ],
      "Effect": "Allow"
    }
  ]
}
```

The centralized S3 bucket policy looks like this:

```

{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Principal": {
        "AWS": [
          "arn:aws:iam::111122223333:role/LogCopier",
          "arn:aws:iam::444455556666:role/LogCopier"
        ]
      },
      "Action": ["s3:PutObject", "s3:PutObjectAcl"],
      "Resource": "arn:aws:s3:::centralizedbucket/*"
    }
  ]
}

```

Why is the Security Engineer unable to access the log files?

- A. The S3 bucket policy does not explicitly allow the Security Engineer access to the objects in the bucket.
- B. The object ACLs are not being updated to allow the users within the centralized account to access the objects
- C. The Security Engineers IAM policy does not grant permissions to read objects in the S3 bucket
- D. The s3:PutObject and s3:PutObjectAcl permissions should be applied at the S3 bucket level

Answer: C

NEW QUESTION 115

A company has deployed servers on Amazon EC2 instances in a VPC. External vendors access these servers over the internet. Recently, the company deployed a new application on EC2 instances in a new CIDR range. The company needs to make the application available to the vendors.

A security engineer verified that the associated security groups and network ACLs are allowing the required ports in the inbound direction. However, the vendors cannot connect to the application.

Which solution will provide the vendors access to the application?

- A. Modify the security group that is associated with the EC2 instances to have the same outbound rules as inbound rules.
- B. Modify the network ACL that is associated with the CIDR range to allow outbound traffic to ephemeral ports.
- C. Modify the inbound rules on the internet gateway to allow the required ports.
- D. Modify the network ACL that is associated with the CIDR range to have the same outbound rules as inbound rules.

Answer: B

Explanation:

The correct answer is B. Modify the network ACL that is associated with the CIDR range to allow outbound traffic to ephemeral ports.

This answer is correct because network ACLs are stateless, which means that they do not automatically allow return traffic for inbound connections. Therefore, the network ACL that is associated with the CIDR range of the new application must have outbound rules that allow traffic to ephemeral ports, which are the temporary ports used by the vendors' machines to communicate with the application servers. Ephemeral ports are typically in the range of 1024-65535. If the network ACL does not have such rules, the vendors will not be able to connect to the application.

The other options are incorrect because:

- > A. Modifying the security group that is associated with the EC2 instances to have the same outbound rules as inbound rules is not a solution, because security groups are stateful, which means that they automatically allow return traffic for inbound connections. Therefore, there is no need to add outbound rules to the security group for the vendors to access the application.
- > C. Modifying the inbound rules on the internet gateway to allow the required ports is not a solution, because internet gateways do not have inbound or outbound rules. Internet gateways are VPC components that enable communication between instances in a VPC and the internet. They do not filter traffic based on ports or protocols.
- > D. Modifying the network ACL that is associated with the CIDR range to have the same outbound rules as inbound rules is not a solution, because it does not address the issue of ephemeral ports. The outbound rules of the network ACL must match the ephemeral port range of the vendors' machines, not necessarily the inbound rules of the network ACL.

References:

1: Ephemeral port - Wikipedia 2: Security groups for your VPC - Amazon Virtual Private Cloud 3: Internet gateways - Amazon Virtual Private Cloud 4: Network ACLs - Amazon Virtual Private Cloud

NEW QUESTION 120

A company has an application that uses dozens of Amazon DynamoDB tables to store data. Auditors find that the tables do not comply with the company's data protection policy.

The company's retention policy states that all data must be backed up twice each month: once at midnight on the 15th day of the month and again at midnight on the 25th day of the month. The company must retain the backups for 3 months.

Which combination of steps should a security engineer take to meet these requirements? (Select TWO.)

- A. Use the DynamoDB on-demand backup capability to create a backup plan.
- B. Configure a lifecycle policy to expire backups after 3 months.
- C. Use AWS DataSync to create a backup plan.
- D. Add a backup rule that includes a retention period of 3 months.
- E. Use AWS Backup to create a backup plan.
- F. Add a backup rule that includes a retention period of 3 months.
- G. Set the backup frequency by using a cron schedule expression.
- H. Assign each DynamoDB table to the backup plan.
- I. Set the backup frequency by using a rate schedule expression.
- J. Assign each DynamoDB table to the backup plan.

Answer: AD

NEW QUESTION 121

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:

This solution ensures that the Lambda functions are deployed inside the VPC and can communicate with the Amazon RDS DB instance securely. The security group attached to the Lambda functions only allows outbound traffic to the VPC CIDR range, and the DB instance security group only allows traffic from the Lambda security group. This solution ensures that the Lambda functions can communicate with the DB instance securely and that the DB instance is not exposed to the public internet.

NEW QUESTION 122

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 private-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 inbound 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 123

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

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

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

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

Answer: ACE

NEW QUESTION 127

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

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

Answer: A

NEW QUESTION 128

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

- * 1 An Application Load Balancer, an internet gateway, and a NAT gateway are configured in the public subnet
 - * 2. Database, application, and web servers are configured on three different private subnets.
 - * 3 The VPC has two route tables: one for the public subnet and one for all other subnets The route table for the public subnet has a 0 0 0 0/0 route to the internet gateway The route table for all other subnets has a 0 0.0.0/0 route to the NAT gateway. All private subnets can route to each other
 - * 4 Each subnet has a network ACL implemented that limits all inbound and outbound connectivity to only the required ports and protocols
 - * 5 There are 3 Security Groups (SGs) database application and web Each group limits all inbound and outbound connectivity to the minimum required
- Which of the following accurately reflects the access control mechanisms the Architect should verify?

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

Answer: A

Explanation:

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

NEW QUESTION 130

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

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

Answer: A

Explanation:

To meet the requirement of changing the key material for new files whenever a potential key breach occurs, the most appropriate solution would be to create a new customer managed key, add a key rotation schedule to the key, and invoke the key rotation schedule every time the security team requests a key change. References: : Rotating AWS KMS keys - AWS Key Management Service

NEW QUESTION 131

A company plans to use AWS Key Management Service (AWS KMS) to implement an encryption strategy to protect data at rest. The company requires client-side encryption for company projects. The company is currently conducting multiple projects to test the company's use of AWS KMS. These tests have led to a sudden increase in the company's AWS resource consumption. The test projects include applications that issue multiple requests each second to KMS endpoints for encryption activities.

The company needs to develop a solution that does not throttle the company's ability to use AWS KMS. The solution must improve key usage for client-side encryption and must be cost optimized. Which solution will meet these requirements?

- A. Use keyrings with the AWS Encryption SD
- B. Use each keyring individually or combine keyrings into a multi-keyring
- C. Decrypt the data by using a keyring that has the primary key in the multi-keyring.
- D. Use data key caching
- E. Use the local cache that the AWS Encryption SDK provides with a caching cryptographic materials manager.
- F. Use KMS key rotation
- G. Use a local cache in the AWS Encryption SDK with a caching cryptographic materials manager.

- H. Use keyrings with the AWS Encryption SD
- I. Use each keyring individually or combine keyrings into a multi-keyrin
- J. Use any of the wrapping keys in the multi-keyring to decrypt the data.

Answer: B

Explanation:

The correct answer is B. Use data key caching. Use the local cache that the AWS Encryption SDK provides with a caching cryptographic materials manager. This answer is correct because data key caching can improve performance, reduce cost, and help the company stay within the service limits of AWS KMS. Data key caching stores data keys and related cryptographic material in a cache, and reuses them for encryption and decryption operations. This reduces the number of requests to AWS KMS endpoints and avoids throttling. The AWS Encryption SDK provides a local cache and a caching cryptographic materials manager (caching CMM) that interacts with the cache and enforces security thresholds that the company can set1.

The other options are incorrect because:

- A. Using keyrings with the AWS Encryption SDK does not address the problem of throttling or cost optimization. Keyrings are used to generate, encrypt, and decrypt data keys, but they do not cache or reuse them. Using each keyring individually or combining them into a multi-keyring does not reduce the number of requests to AWS KMS endpoints2.
- C. Using KMS key rotation does not address the problem of throttling or cost optimization. Key rotation is a security practice that creates new cryptographic material for a KMS key every year, but it does not affect the data that the KMS key protects. Key rotation does not reduce the number of requests to AWS KMS endpoints, and it might incur additional costs for storing multiple versions of key material3.
- D. Using keyrings with the AWS Encryption SDK does not address the problem of throttling or cost optimization, as explained in option A. Moreover, using any of the wrapping keys in the multi-keyring to decrypt the data is not a valid option, because only one of the wrapping keys can decrypt a given data key. The wrapping key that encrypts a data key is stored in the encrypted data key structure, and only that wrapping key can decrypt it4.

References:

1: Data key caching - AWS Encryption SDK 2: Using keyrings - AWS Encryption SDK 3: Rotating AWS KMS keys - AWS Key Management Service 4: How keyrings work - AWS Encryption SDK

NEW QUESTION 135

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 136

.....

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-C01 Practice Exam Features:

- * SCS-C01 Questions and Answers Updated Frequently
- * SCS-C01 Practice Questions Verified by Expert Senior Certified Staff
- * SCS-C01 Most Realistic Questions that Guarantee you a Pass on Your FirstTry
- * SCS-C01 Practice Test Questions in Multiple Choice Formats and Updatesfor 1 Year

100% Actual & Verified — Instant Download, Please Click
[Order The SCS-C01 Practice Test Here](#)