



# Amazon-Web-Services

## Exam Questions SCS-C02

AWS Certified Security - Specialty

### NEW QUESTION 1

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

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

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

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

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

**Answer:** CE

#### Explanation:

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

References: : Amazon VPC User Guide

### NEW QUESTION 2

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

Which solution will meet these requirements?

- A. Store the database credentials in AWS Secrets Manager
- B. Configure automatic credential rotation for every 30 days.
- C. Store the database credentials in AWS Systems Manager Parameter Store
- D. Create an AWS Lambda function to rotate the credentials every 30 days.
- E. Store the database credentials in an environment file or in a configuration file
- F. Modify the credentials every 30 days.
- G. Store the database credentials in an environment file or in a configuration file
- H. Create an AWS Lambda function to rotate the credentials every 30 days.

**Answer:** A

#### Explanation:

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

- > <https://docs.aws.amazon.com/secretsmanager/latest/userguide/rotating-secrets.html>
- > [https://docs.aws.amazon.com/secretsmanager/latest/userguide/rotate-secrets\\_turn-on-for-other.html](https://docs.aws.amazon.com/secretsmanager/latest/userguide/rotate-secrets_turn-on-for-other.html)

### NEW QUESTION 3

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

The data includes personally identifiable information (PII). The company must remove data that is older than 30 days from the S3 bucket and the DynamoDB table. Which solution will meet this requirement with the MOST operational efficiency?

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

**Answer:** B

### NEW QUESTION 4

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 5**

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-keyring
- 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 set.

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 endpoints.
- 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 material.
- 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 it.

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 6**

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

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

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

- A. In the dedicated security account, create an Amazon S3 bucket
- B. Configure S3 Object Lock in compliance mode and a retention period of 2 years on the S3 bucket
- C. Set the bucket policy to allow the organization's management account to write to the S3 bucket.
- D. In the dedicated security account, create an Amazon S3 bucket
- E. Configure S3 Object Lock in compliance mode and a retention period of 2 years on the S3 bucket
- F. Set the bucket policy to allow the organization's member accounts to write to the S3 bucket.
- G. In the dedicated security account, create an Amazon S3 bucket that has an S3 Lifecycle configuration that expires objects after 2 year
- H. Set the bucket policy to allow the organization's member accounts to write to the S3 bucket.



- I. Create an AWS Cloud Trail trail for the organization
- J. Configure logs to be delivered to the logging Amazon S3 bucket in the dedicated security account.
- K. Turn on AWS CloudTrail in each account
- L. Configure logs to be delivered to an Amazon S3 bucket that is created in the organization's management account
- M. Forward the logs to the S3 bucket in the dedicated security account by using AWS Lambda and Amazon Kinesis Data Firehose.

**Answer:** BD

**Explanation:**

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

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

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

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

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

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

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

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

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

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

**NEW QUESTION 7**

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

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

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

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

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

**Answer:** AD

**NEW QUESTION 8**

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 9

A security engineer must troubleshoot an administrator's inability to make an existing Amazon S3 bucket public in an account that is part of an organization n IAM Organizations. The administrator switched the role from the master account to a member account and then attempted to make one S3 bucket public. This action was immediately denied

Which actions should the security engineer take to troubleshoot the permissions issue? (Select TWO.)

- A. Review the cross-account role permissions and the S3 bucket policy Verify that the Amazon S3 block public access option in the member account is deactivated.
- B. Review the role permissions m the master account and ensure it has sufficient privileges to perform S3 operations
- C. Filter IAM CloudTrail logs for the master account to find the original deny event and update the cross-account role m the member account accordingly Verify that the Amazon S3 block public access option in the master account is deactivated.
- D. Evaluate the SCPs covering the member account and the permissions boundary of the role in the member account for missing permissions and explicit denies.
- E. Ensure the S3 bucket policy explicitly allows the s3 PutBucketPublicAccess action for the role m the member account

**Answer:** DE

### Explanation:

- A is incorrect because reviewing the cross-account role permissions and the S3 bucket policy is not enough to troubleshoot the permissions issue. You also need to verify that the Amazon S3 block public access option in the member account is deactivated, as well as the permissions boundary and the SCPs of the role in the member account.
- D is correct because evaluating the SCPs and the permissions boundary of the role in the member account can help you identify any missing permissions or explicit denies that could prevent the administrator from making the S3 bucket public.
- E is correct because ensuring that the S3 bucket policy explicitly allows the s3 PutBucketPublicAccess action for the role in the member account can help you override any block public access settings that could prevent the administrator from making the S3 bucket public.

### NEW QUESTION 10

A company created an IAM account for its developers to use for testing and learning purposes Because MM 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 al 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 AM policy similar to the one that fellows Populate the ec2: ResourceTag/Team condition key with a proper team name Attach resulting policies to the corresponding IAM roles.

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

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

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

- D. Tag each IAM role with a Team tag 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"
      ],
      "Resource": "*"
    },
    {
      "Effect": "Allow",
      "Action": [
        "ec2:StopInstances",
        "ec2:TerminateInstances"
      ],
      "Resource": "*",
      "Condition": {
        "StringEquals": {
          "ec2:ResourceTag/Team": "${aws:PrincipalTag/Team}"
        }
      }
    }
  ]
}
```

- G. Tag each IAM role with the Team key, and use the team name in the tag 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": "${aws:PrincipalTag/Team}"
        }
      }
    }
  ]
}
```

**Answer: A**

#### NEW QUESTION 10

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.

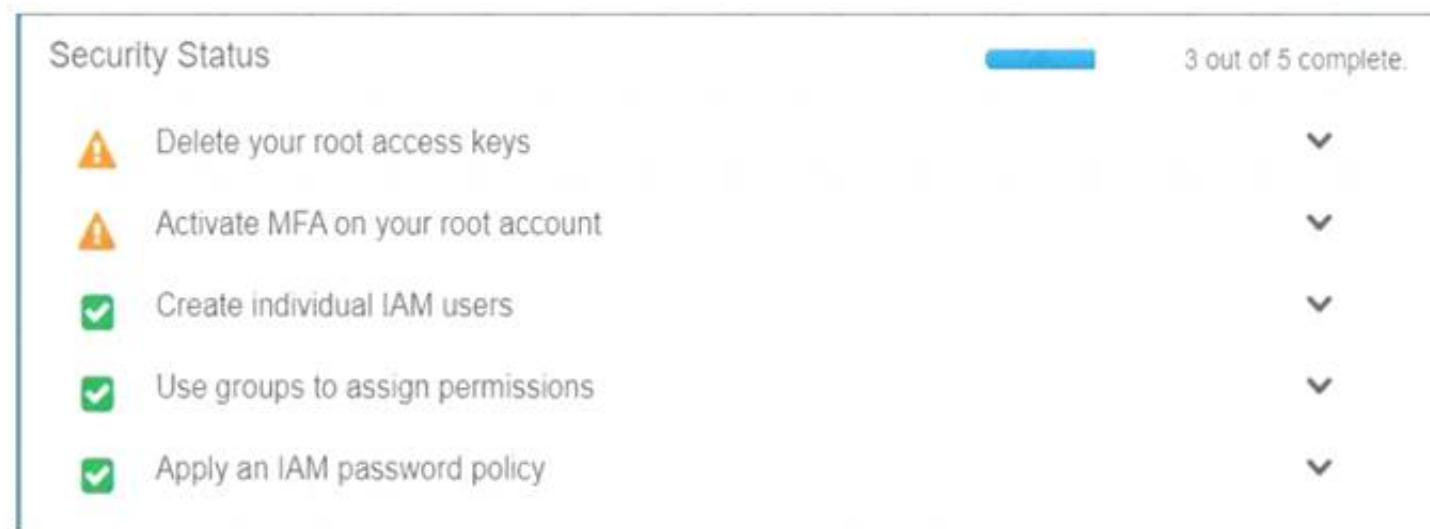
**Answer: C**

#### 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.amazonaws.com/IAM/latest/UserGuide/id\\_credentials\\_mfa.html](http://docs.IAM.amazonaws.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 13

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

Which solution will meet these requirements?

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

**Answer: A**

### NEW QUESTION 16

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

Specifically, the security team wants EventBridge to watch for the s3:PutObjectAcl, s3:DeleteBucketPolicy, and s3:PutBucketPolicy API invocation logs from CloudTrail. While developing the solution in a single account, the security team discovers that the s3:PutObjectAcl API call does not invoke an EventBridge event. However, the s3:DeleteBucketPolicy API call and the s3:PutBucketPolicy API call do invoke an event.

The security team has enabled CloudTrail for AWS management events with a basic configuration in the AWS Region in which EventBridge is being tested.

Verification of the EventBridge event pattern indicates that the pattern is set up correctly. The security team must implement a solution so that the s3:PutObjectAcl API call will invoke an EventBridge event. The solution must not generate false notifications.

Which solution will meet these requirements?

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

**Answer: D**

### Explanation:

The correct answer is D. Enable CloudTrail to monitor data events for read and write operations to S3 buckets. According to the AWS documentation<sup>1</sup>, CloudTrail data events are the resource operations performed on or within a resource. These are also known as data plane operations. Data events are often high-volume activities. For example, Amazon S3 object-level API activity (such as GetObject, DeleteObject, and PutObject) is a data event.

By default, trails do not log data events. To record CloudTrail data events, you must explicitly add the supported resources or resource types for which you want to collect activity. For more information, see Logging data events in the Amazon S3 User Guide<sup>2</sup>.

In this case, the security team wants EventBridge to watch for the s3:PutObjectAcl API invocation logs from CloudTrail. This API uses the acl subresource to set the access control list (ACL) permissions for a new or existing object in an S3 bucket<sup>3</sup>. This is a data event that affects the S3 object resource type. Therefore, the security team must enable CloudTrail to monitor data events for read and write operations to S3 buckets in order to invoke an EventBridge event for this API call.

The other options are incorrect because:

- A. Modifying the EventBridge event pattern by selecting Amazon S3 and All Events as the event type will not capture the s3:PutObjectAcl API call, because this is a data event and not a management event. Management events provide information about management operations that are performed on resources in your AWS account. These are also known as control plane operations<sup>4</sup>.
- B. Modifying the EventBridge event pattern by selecting Amazon S3 and Bucket Level Operations as the event type will not capture the s3:PutObjectAcl API call, because this is a data event that affects the S3 object resource type and not the S3 bucket resource type. Bucket level operations are management events that affect the configuration or metadata of an S3 bucket<sup>5</sup>.
- C. Enabling CloudTrail Insights to identify unusual API activity will not help the security team monitor new S3 objects or changes to any S3 bucket policy or setting that result in public access. CloudTrail Insights helps AWS users identify and respond to unusual activity associated with API calls and API error rates by continuously analyzing CloudTrail management events<sup>6</sup>. It does not analyze data events or generate EventBridge events.

References:

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

**NEW QUESTION 19**

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

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

**Answer: B**

**NEW QUESTION 20**

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

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

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

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

**Answer: C**

**Explanation:**

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

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

The other options are incorrect for the following reasons:

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

**NEW QUESTION 22**

A developer is building a serverless application hosted on AWS that uses Amazon Redshift as a data store. The application has separate modules for readwrite and read-only functionality. The modules need their own database users for compliance reasons.

Which combination of steps should a security engineer implement to grant appropriate access? (Select TWO.)

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



**Answer:** A

**Explanation:**

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

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

**NEW QUESTION 23**

A security engineer has enabled IAM Security Hub in their IAM account, and has enabled the Center for internet Security (CIS) IAM Foundations compliance standard. No evaluation results on compliance are returned in the Security Hub console after several hours. The engineer wants to ensure that Security Hub can evaluate their resources for CIS IAM Foundations compliance.

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

- A. Add full Amazon Inspector IAM permissions to the Security Hub service role to allow it to perform the CIS compliance evaluation
- B. Ensure that IAM Trusted Advisor Is enabled in the account and that the Security Hub service role has permissions to retrieve the Trusted Advisor security-related recommended actions
- C. Ensure that IAM Confi
- D. is enabled in the account, and that the required IAM Config rules have been created for the CIS compliance evaluation
- E. Ensure that the correct trail in IAM CloudTrail has been configured for monitoring by Security Hub and that the Security Hub service role has permissions to perform the GetObject operation on CloudTrails Amazon S3 bucket

**Answer:** C

**Explanation:**

To ensure that Security Hub can evaluate their resources for CIS AWS Foundations compliance, the security engineer should do the following:

- Ensure that AWS Config is enabled in the account. This is a service that enables continuous assessment and audit of your AWS resources for compliance.
- Ensure that the required AWS Config rules have been created for the CIS compliance evaluation. These are rules that represent your desired configuration settings for specific AWS resources or for an entire AWS account.

**NEW QUESTION 27**

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 32**

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

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

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

- A. Log in to the suspicious instance and use the netstat command to identify remote connections. Use the IP addresses from these remote connections to create deny rules in the security group of the instance. Install diagnostic tools on the instance for investigation. Update the outbound network ACL for the subnet in us-east-1b to explicitly deny all connections as the first rule during the investigation of the instance.
- B. Update the outbound network ACL for the subnet in us-east-1b to explicitly deny all connections as the first rule. Replace the security group with a new security group that allows connections only from a diagnostics security group. Update the outbound network ACL for the us-east-1b subnet to remove the deny all rule. Launch a new EC2 instance that has diagnostic tools. Assign the new security group to the new EC2 instance. Use the new EC2 instance to investigate the suspicious instance.
- C. Ensure that the Amazon Elastic Block Store (Amazon EBS) volumes that are attached to the suspicious EC2 instance will not delete upon termination. Terminate the instance. Launch a new EC2 instance in us-east-1a that has diagnostic tools. Mount the EBS volumes from the terminated instance for investigation.
- D. Create an AWS WAF web ACL that denies traffic to and from the suspicious instance. Attach the AWS WAF web ACL to the instance to mitigate the attack. Log in to the instance and install diagnostic tools to investigate the instance.

**Answer:** B

**Explanation:**

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

**NEW QUESTION 33**

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

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

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

**Answer:** BD

**Explanation:**

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

**NEW QUESTION 36**

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

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

**Answer:** C

**NEW QUESTION 40**

A company has enabled Amazon GuardDuty in all AWS Regions as part of its security monitoring strategy. In one of its VPCs, the company hosts an Amazon EC2 instance that works as an FTP server. A high number of clients from multiple locations contact the FTP server. GuardDuty identifies this activity as a brute force attack because of the high number of connections that happen every hour.

The company has flagged the finding as a false positive, but GuardDuty continues to raise the issue. A security engineer must improve the signal-to-noise ratio without compromising the company's visibility of potential anomalous behavior.

Which solution will meet these requirements?

- A. Disable the FTP rule in GuardDuty in the Region where the FTP server is deployed.
- B. Add the FTP server to a trusted IP list.
- C. Deploy the list to GuardDuty to stop receiving the notifications.
- D. Create a suppression rule in GuardDuty to filter findings by automatically archiving new findings that match the specified criteria.
- E. Create an AWS Lambda function that has the appropriate permissions to delete the finding whenever a new occurrence is reported.

**Answer:** C

**Explanation:**

"When you create an Amazon GuardDuty filter, you choose specific filter criteria, name the filter and can enable the auto-archiving of findings that the filter matches. This allows you to further tune GuardDuty to your unique environment, without degrading the ability to identify threats. With auto-archive set, all findings are still generated by GuardDuty, so you have a complete and immutable history of all suspicious activity."

**NEW QUESTION 45**

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

Which solution meets these requirements in the MOST secure way?

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

**Answer:** C

**Explanation:**

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

References: : AWS Lambda Developer Guide

#### NEW QUESTION 47

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

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

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

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

**Answer: C**

#### Explanation:

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

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

#### NEW QUESTION 48

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

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

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

**Answer: D**

#### Explanation:

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

This answer is correct because the KMS key policy is the primary way to control access to the KMS key, and it must explicitly allow the AWS account to have full access to the key. If the KMS key policy has been edited to remove this permission, then the IAM policy that grants kms:Decrypt permission to the IAM user has no effect, and the IAM user cannot decrypt the objects in the S3 bucket<sup>12</sup>.

The other options are incorrect because:

- A. The IAM policy does not need to allow the kms:DescribeKey permission, because this permission is not required for decrypting objects in S3 using SSE-KMS. The kms:DescribeKey permission allows getting information about a KMS key, such as its creation date, description, and key state<sup>3</sup>.
- B. The S3 bucket has not been changed to use the AWS managed key to encrypt objects at rest, because this would not cause an Access Denied message for the IAM user. The AWS managed key is a default KMS key that is created and managed by AWS for each AWS account and Region. The IAM user does not need any permissions on this key to use it for SSE-KMS<sup>4</sup>.
- C. An S3 bucket policy does not need to be added to allow the IAM user to access the objects, because the IAM user already has s3:List\* and s3:Get\* permissions for the S3 bucket and its objects through an IAM policy. An S3 bucket policy is an optional way to grant cross-account access or public access to an S3 bucket<sup>5</sup>.

References:

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

#### NEW QUESTION 53

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

How should the security engineer prevent unauthorized access to the EC2 instances?

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

**Answer: C**

#### Explanation:



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

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

#### NEW QUESTION 58

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

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

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

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

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

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

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

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



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

**Answer:** A

**Explanation:**

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

Options B,C and D are invalid because you have to ensure the condition of "s3:x-amz-server-side-encryption":"IAM:kms" is present

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

Submit your Feedback/Queries to our Expert

**NEW QUESTION 61**

A company is using AWS WAF to protect a customized public API service that is based on Amazon EC2 instances. The API uses an Application Load Balancer. The AWS WAF web ACL is configured with an AWS Managed Rules rule group. After a software upgrade to the API and the client application, some types of requests are no longer working and are causing application stability issues. A security engineer discovers that AWS WAF logging is not turned on for the web ACL. The security engineer needs to immediately return the application to service, resolve the issue, and ensure that logging is not turned off in the future. The security engineer turns on logging for the web ACL and specifies Amazon Cloud-Watch Logs as the destination. Which additional set of steps should the security engineer take to meet the requirements?

- A. Edit the rules in the web ACL to include rules with Count action
- B. Review the logs to determine which rule is blocking the requests
- C. Modify the IAM policy of all AWS WAF administrators so that they cannot remove the logging configuration for any AWS WAF web ACLs.
- D. Edit the rules in the web ACL to include rules with Count action
- E. Review the logs to determine which rule is blocking the requests
- F. Modify the AWS WAF resource policy so that AWS WAF administrators cannot remove the logging configuration for any AWS WAF web ACLs.
- G. Edit the rules in the web ACL to include rules with Count and Challenge action
- H. Review the logs to determine which rule is blocking the requests
- I. Modify the AWS WAF resource policy so that AWS WAF administrators cannot remove the logging configuration for any AWS WAF web ACLs.
- J. Edit the rules in the web ACL to include rules with Count and Challenge action
- K. Review the logs to determine which rule is blocking the requests
- L. Modify the IAM policy of all AWS WAF administrators so that they cannot remove the logging configuration for any AWS WAF web ACLs.

**Answer:** A

**Explanation:**

This answer is correct because it meets the requirements of returning the application to service, resolving the issue, and ensuring that logging is not turned off in the future. By editing the rules in the web ACL to include rules with Count actions, the security engineer can test the effect of each rule without blocking or allowing requests. By reviewing the logs, the security engineer can identify which rule is causing the problem and modify or delete it accordingly. By modifying the IAM policy of all AWS WAF administrators, the security engineer can restrict their permissions to prevent them from removing the logging configuration for any AWS WAF web ACLs.

**NEW QUESTION 65**

A developer is building a serverless application hosted on AWS Lambda that uses Amazon Redshift in a data store. The application has separate modules for read/write and read-only functionality. The modules need their own database users for compliance reasons.

Which combination of steps should a security engineer implement to grant appropriate access? (Select TWO )

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

**Answer:** CD

**Explanation:**

To grant appropriate access to the application modules, the security engineer should do the following:

- Configure an IAM policy for each module. Specify the ARN of an Amazon Redshift database user that allows the GetClusterCredentials API call. This allows the application modules to use temporary credentials to access the database with the permissions of the specified user.
- Create local database users for each module. This allows the security engineer to create separate users for read/write and read-only functionality, and to assign them different privileges on the database tables.

**NEW QUESTION 68**

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

Please select:

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

**Answer:** D

**Explanation:**

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

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

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

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

**NEW QUESTION 72**

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/how-use-aws-systems-manager-to-view-vulnerability-id>

**NEW QUESTION 73**

An international company has established a new business entity in South Korea. The company also has established a new AWS account to contain the workload for the South Korean region. The company has set up the workload in the new account in the ap-northeast-2 Region. The workload consists of three Auto Scaling groups of Amazon EC2 instances. All workloads that operate in this Region must keep system logs and application logs for 7 years.

A security engineer must implement a solution to ensure that no logging data is lost for each instance during scaling activities. The solution also must keep the logs for only the required period of 7 years.

Which combination of steps should the security engineer take to meet these requirements? (Choose three.)

- A. Ensure that the Amazon CloudWatch agent is installed on all the EC2 instances that the Auto Scaling groups launch
- B. Generate a CloudWatch agent configuration file to forward the required logs to Amazon CloudWatch Logs.
- C. Set the log retention for desired log groups to 7 years.
- D. Attach an IAM role to the launch configuration or launch template that the Auto Scaling groups use. Configure the role to provide the necessary permissions to forward logs to Amazon CloudWatch Logs.
- E. Attach an IAM role to the launch configuration or launch template that the Auto Scaling groups use. Configure the role to provide the necessary permissions to forward logs to Amazon S3.
- F. Ensure that a log forwarding application is installed on all the EC2 instances that the Auto Scaling groups launch
- G. Configure the log forwarding application to periodically bundle the logs and forward the logs to Amazon S3.
- H. Configure an Amazon S3 Lifecycle policy on the target S3 bucket to expire objects after 7 years.

**Answer:** ABC

**Explanation:**

The correct combination of steps that the security engineer should take to meet these requirements are A. Ensure that the Amazon CloudWatch agent is installed on all the EC2 instances that the Auto Scaling groups launch. Generate a CloudWatch agent configuration file to forward the required logs to Amazon CloudWatch Logs., B. Set the log retention for desired log groups to 7 years., and C. Attach an IAM role to the launch configuration or launch template that the Auto Scaling groups use. Configure the role to provide the necessary permissions to forward logs to Amazon CloudWatch Logs.

\* A. This answer is correct because it meets the requirement of ensuring that no logging data is lost for each instance during scaling activities. By installing the CloudWatch agent on all the EC2 instances, the security engineer can collect and send system logs and application logs to CloudWatch Logs, which is a service that stores and monitors log data. By generating a CloudWatch agent configuration file, the security engineer can specify which logs to forward and how often.

\* B. This answer is correct because it meets the requirement of keeping the logs for only the required period of 7 years. By setting the log retention for desired log groups, the security engineer can control how long CloudWatch Logs retains log events before deleting them. The security engineer can choose a predefined retention period of 7 years, or use a custom value.

\* C. This answer is correct because it meets the requirement of providing the necessary permissions to forward logs to CloudWatch Logs. By attaching an IAM role to the launch configuration or launch template that the Auto Scaling groups use, the security engineer can grant permissions to the EC2 instances that are launched by the Auto Scaling groups. By configuring the role to provide the necessary permissions, such as cloudwatch:PutLogEvents and cloudwatch:CreateLogStream, the security engineer can allow the EC2 instances to send log data to CloudWatch Logs.

**NEW QUESTION 78**

A company is deploying an Amazon EC2-based application. The application will include a custom health-checking component that produces health status data in JSON format. A Security Engineer must

implement a secure solution to monitor application availability in near-real time by analyzing the health status data.

Which approach should the Security Engineer use?

- A. Use Amazon CloudWatch monitoring to capture Amazon EC2 and networking metrics Visualizemetrics 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 Redshif
- C. (hen 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 availability<sup>1</sup>. The health status data in JSON format can be sent to CloudWatch as custom metrics<sup>2</sup>, and then displayed in CloudWatch dashboards<sup>3</sup>. The other options are either inefficient or insecure for monitoring application availability in near-real time.

**NEW QUESTION 80**

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 m an regions when it opened the account Which of the following will allow (he Security Engineer 10 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 lo 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 tor 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 SQL<sup>1</sup>. 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 range<sup>2</sup>. The other options are not effective ways to review the use of the exposed access key.

**NEW QUESTION 82**

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 85**

An organization must establish the ability to delete an IAM KMS Customer Master Key (CMK) within a 24- hour timeframe to keep it from being used for encrypt or decrypt operations Which of tne following actions will address this requirement?

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

**Answer:** C

**Explanation:**

the schedule key deletion function within KMS allows you to specify a waiting period before deleting a customer master key (CMK)<sup>4</sup>. The minimum waiting period is 7 days and the maximum is 30 days<sup>5</sup>. This function prevents the CMK from being used for encryption or decryption operations during the waiting period<sup>4</sup>. The other options are either invalid or ineffective for deleting a CMK within a 24-hour timeframe.

**NEW QUESTION 87**

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 92**

An ecommerce company is developing new architecture for an application release. The company needs to implement TLS for incoming traffic to the application. Traffic for the application will originate from the internet TLS does not have to be implemented in an end-to-end configuration because the company is concerned about impacts on performance. The incoming traffic types will be HTTP and HTTPS The application uses ports 80 and 443. What should a security engineer do to meet these requirements?

- A. Create a public Application Load Balance
- B. Create two listeners one listener on port 80 and one listener on port 443. Create one target grou
- C. Create a rule to forward traffic from port 80 to the listener on port 443 Provision a public TLS certificate in AWS Certificate Manager (ACM). Attach the certificate to the listener on port 443.
- D. Create a public Application Load Balance
- E. Create two listeners one listener on port 80 and one listener on port 443. Create one target grou
- F. Create a rule to forward traffic from port 80 to the listener on port 443 Provision a public TLS certificate in AWS Certificate Manager (ACM). Attach the certificate to the listener on port 80.
- G. Create a public Network Load Balance
- H. Create two listeners one listener on port 80 and one listener on port 443. Create one target grou
- I. Create a rule to forward traffic from port 80 to the listener on port 443. Set the protocol for the listener on port 443 to TLS.
- J. Create a public Network Load Balance
- K. Create a listener on port 443. Create one target grou
- L. Create a rule to forward traffic from port 443 to the target grou
- M. Set the protocol for the listener on port 443 to TLS.

**Answer: A**

**Explanation:**

An Application Load Balancer (ALB) is a type of load balancer that operates at the application layer (layer 7) of the OSI model. It can distribute incoming traffic based on the content of the request, such as the host header, path, or query parameters. An ALB can also terminate TLS connections and decrypt requests from clients before sending them to the targets.

To implement TLS for incoming traffic to the application, the following steps are required:

- Create a public ALB in a public subnet and register the EC2 instances as targets in a target group.
- Create two listeners for the ALB, one on port 80 for HTTP traffic and one on port 443 for HTTPS traffic.
- Create a rule for the listener on port 80 to redirect HTTP requests to HTTPS using the same host, path, and query parameters.
- Provision a public TLS certificate in AWS Certificate Manager (ACM) for the domain name of the application. ACM is a service that lets you easily provision, manage, and deploy public and private SSL/TLS certificates for use with AWS services and your internal connected resources.
- Attach the certificate to the listener on port 443 and configure the security policy to negotiate secure connections between clients and the ALB.
- Configure the security groups for the ALB and the EC2 instances to allow inbound traffic on ports 80 and 443 from the internet and outbound traffic on any port to the EC2 instances.

This solution will meet the requirements of implementing TLS for incoming traffic without impacting performance or requiring end-to-end encryption. The ALB will handle the TLS termination and decryption, while forwarding unencrypted requests to the EC2 instances.

Verified References:

- <https://docs.aws.amazon.com/elasticloadbalancing/latest/application/introduction.html>
- <https://docs.aws.amazon.com/elasticloadbalancing/latest/application/create-https-listener.html>
- <https://docs.aws.amazon.com/acm/latest/userguide/acm-overview.html>

**NEW QUESTION 97**

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 repositorie
- B. Install Amazon Inspector Agent from the ECS container instances' user dat



- 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 101

Your development team is using access keys to develop an application that has access to S3 and DynamoDB. A new security policy has outlined that the credentials should not be older than 2 months, and should be rotated. How can you achieve this?  
Please select:

- A. Use the application to rotate the keys in every 2 months via the SDK
- B. Use a script to query the creation date of the key
- C. If older than 2 months, create new access key and update all applications to use it inactivate the old key and delete it.
- D. Delete the user associated with the keys after every 2 month
- E. Then recreate the user again.
- F. Delete the IAM Role associated with the keys after every 2 month
- G. Then recreate the IAM Role again.

**Answer: B**

#### Explanation:

One can use the CLI command list-access-keys to get the access keys. This command also returns the "CreateDate" of the keys. If the CreateDate is older than 2 months, then the keys can be deleted.

The Returns list-access-keys CLI command returns information about the access key IDs associated with the specified IAM user. If there are none, the action returns an empty list

Option A is incorrect because you might as use a script for such maintenance activities Option C is incorrect because you would not rotate the users themselves

Option D is incorrect because you don't use IAM roles for such a purpose For more information on the CLI command, please refer to the below Link:

<http://docs.IAM.amazon.com/cli/latest/reference/iam/list-access-keys.html>

The correct answer is: Use a script to query the creation date of the keys. If older than 2 months, create new access key and update all applications to use it inactivate the old key and delete it.

Submit your Feedback/Queries to our Experts

#### NEW QUESTION 104

A website currently runs on Amazon EC2, with mostly static content on the site. Recently the site was subjected to a DDoS attack a security engineer was (asked) was redesigning the edge security to help

Mitigate this risk in the future.

What are some ways the engineer could achieve this (Select THREE)?

- A. Use IAM X-Ray to inspect the traffic going to the EC2 instances.
- B. Move the static content to Amazon S3, and front this with an Amazon CloudFront distribution.
- C. Change the security group configuration to block the source of the attack traffic
- D. Use IAM WAF security rules to inspect the inbound traffic.
- E. Use Amazon Inspector assessment templates to inspect the inbound traffic.
- F. Use Amazon Route 53 to distribute traffic.

**Answer: BDF**

#### Explanation:

To redesign the edge security to help mitigate the DDoS attack risk in the future, the engineer could do the following:

- Move the static content to Amazon S3, and front this with an Amazon CloudFront distribution. This allows the engineer to use a global content delivery network that can cache static content at edge locations and reduce the load on the origin servers.
- Use AWS WAF security rules to inspect the inbound traffic. This allows the engineer to use web application firewall rules that can filter malicious requests based on IP addresses, headers, body, or URI strings, and block them before they reach the web servers.
- Use Amazon Route 53 to distribute traffic. This allows the engineer to use a scalable and highly available DNS service that can route traffic based on different policies, such as latency, geolocation, or health checks.

#### NEW QUESTION 108

A Security Engineer is working with a Product team building a web application on AWS. The application uses Amazon S3 to host the static content, Amazon API Gateway to provide RESTful services; and Amazon DynamoDB as the backend data store. The users already exist in a directory that is exposed through a SAML identity provider.

Which combination of the following actions should the Engineer take to enable users to be authenticated into the web application and call APIs? (Choose three.)

- A. Create a custom authorization service using AWS Lambda.
- B. Configure a SAML identity provider in Amazon Cognito to map attributes to the Amazon Cognito user pool attributes.
- C. Configure the SAML identity provider to add the Amazon Cognito user pool as a relying party.
- D. Configure an Amazon Cognito identity pool to integrate with social login providers.
- E. Update DynamoDB to store the user email addresses and passwords.
- F. Update API Gateway to use a COGNITO\_USER\_POOLS authorizer.

**Answer: BCF**

#### Explanation:

The combination of the following actions should the Engineer take to enable users to be authenticated into the web application and call APIs are:

- B. Configure a SAML identity provider in Amazon Cognito to map attributes to the Amazon Cognito user pool attributes. This is a necessary step to federate the existing users from the SAML identity provider to the Amazon Cognito user pool, which will be used for authentication and authorization1.
- C. Configure the SAML identity provider to add the Amazon Cognito user pool as a relying party. This is a necessary step to establish a trust relationship between the SAML identity provider and the Amazon Cognito user pool, which will allow the users to sign in using their existing credentials2.
- F. Update API Gateway to use a COGNITO\_USER\_POOLS authorizer. This is a necessary step to enable API Gateway to use the Amazon Cognito user pool as an authorizer for the RESTful services, which will validate the identity or access tokens that are issued by Amazon Cognito when a user signs in successfully3. The other options are incorrect because:
- A. Creating a custom authorization service using AWS Lambda is not a necessary step, because Amazon Cognito user pools can provide built-in authorization features, such as scopes and groups, that can be used to control access to API resources4.
- D. Configuring an Amazon Cognito identity pool to integrate with social login providers is not a necessary step, because the users already exist in a directory that is exposed through a SAML identity provider, and there is no requirement to support social login providers5.
- E. Updating DynamoDB to store the user email addresses and passwords is not a necessary step, because the user credentials are already stored in the SAML identity provider, and there is no need to duplicate them in DynamoDB6.

References:

1: Using Tokens with User Pools 2: Adding SAML Identity Providers to a User Pool 3: Control Access to a REST API Using Amazon Cognito User Pools as Authorizer 4: API Authorization with Resource Servers and OAuth 2.0 Scopes 5: Using Identity Pools (Federated Identities) 6: Amazon DynamoDB

#### NEW QUESTION 110

A security engineer is trying to use Amazon EC2 Image Builder to create an image of an EC2 instance. The security engineer has configured the pipeline to send logs to an Amazon S3 bucket. When the security engineer runs the pipeline, the build fails with the following error: "AccessDenied: Access Denied status code: 403".

The security engineer must resolve the error by implementing a solution that complies with best practices for least privilege access.

Which combination of steps will meet these requirements? (Choose two.)

- A. Ensure that the following policies are attached to the IAM role that the security engineer is using: EC2InstanceProfileForImageBuilder, EC2InstanceProfileForImageBuilderECRContainerBuilds, and AmazonSSMManagedInstanceCore.
- B. Ensure that the following policies are attached to the instance profile for the EC2 instance: EC2InstanceProfileForImageBuilder, EC2InstanceProfileForImageBuilderECRContainerBuilds, and AmazonSSMManagedInstanceCore.
- C. Ensure that the AWSImageBuilderFullAccess policy is attached to the instance profile for the EC2 instance.
- D. Ensure that the security engineer's IAM role has the s3:PutObject permission for the S3 bucket.
- E. Ensure that the instance profile for the EC2 instance has the s3:PutObject permission for the S3 bucket.

**Answer: BE**

#### Explanation:

The most likely cause of the error is that the instance profile for the EC2 instance does not have the s3:PutObject permission for the S3 bucket. This permission is needed to upload logs to the bucket. Therefore, the security engineer should ensure that the instance profile has this permission.

One possible solution is to attach the AWSImageBuilderFullAccess policy to the instance profile for the EC2 instance. This policy grants full access to Image Builder resources and related AWS services, including the s3:PutObject permission for any bucket with "imagebuilder" in its name. However, this policy may grant more permissions than necessary, which violates the principle of least privilege.

Another possible solution is to create a custom policy that only grants the s3:PutObject permission for the specific S3 bucket that is used for logging. This policy can be attached to the instance profile along with the other policies that are required for Image Builder functionality: EC2InstanceProfileForImageBuilder, EC2InstanceProfileForImageBuilderECRContainerBuilds, and AmazonSSMManagedInstanceCore. This solution follows the principle of least privilege more closely than the previous one.

- Ensure that the following policies are attached to the instance profile for the EC2 instance: EC2InstanceProfileForImageBuilder, EC2InstanceProfileForImageBuilderECRContainerBuilds, and AmazonSSMManagedInstanceCore.

- Ensure that the instance profile for the EC2 instance has the s3:PutObject permission for the S3 bucket.

This can be done by either attaching the AWSImageBuilderFullAccess policy or creating a custom policy with this permission.

1: Using managed policies for EC2 Image Builder - EC2 Image Builder 2: PutObject - Amazon Simple Storage Service 3: AWSImageBuilderFullAccess - AWS Managed Policy

#### NEW QUESTION 111

A company has an AWS account that hosts a production application. The company receives an email notification that Amazon GuardDuty has detected an Impact:IAMUser/AnomalousBehavior finding in the account. A security engineer needs to run the investigation playbook for this security incident and must collect and analyze the information without affecting the application.

Which solution will meet these requirements MOST quickly?

- A. Log in to the AWS account by using read-only credential
- B. Review the GuardDuty finding for details about the IAM credentials that were use
- C. Use the IAM console to add a DenyAll policy to the IAM principal.
- D. Log in to the AWS account by using read-only credential
- E. Review the GuardDuty finding to determine which API calls initiated the findin
- F. Use Amazon Detective to review the API calls in context.
- G. Log in to the AWS account by using administrator credential
- H. Review the GuardDuty finding for details about the IAM credentials that were use
- I. Use the IAM console to add a DenyAll policy to the IAM principal.
- J. Log in to the AWS account by using read-only credential
- K. Review the GuardDuty finding to determinewhich API calls initiated the findin
- L. Use AWS CloudTrail Insights and AWS CloudTrail Lake to review the API calls in context.

**Answer: B**

#### Explanation:

This answer is correct because logging in with read-only credentials minimizes the risk of accidental or malicious changes to the AWS account. Reviewing the GuardDuty finding can help identify which API calls initiated the finding and which IAM principal was involved. Using Amazon Detective can help analyze and visualize the API calls in context, such as which resources were affected, which IP addresses were used, and how the activity deviated from normal patterns. Amazon Detective can also help identify related findings from other sources, such as AWS Config or AWS Audit Manager.

### NEW QUESTION 112

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

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

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

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

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

**Answer: D**

#### Explanation:

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

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

The other options are incorrect because:

- A. Adding a resource policy that allows each member of the group to access Amazon SES is not a solution, because resource policies are not supported by Amazon SES<sup>1</sup>. Resource policies are policies that are attached to AWS resources, such as S3 buckets or SNS topics, to control access to those resources<sup>2</sup>. Amazon SES does not have any resources that can have resource policies attached to them.
- B. Adding a resource policy that allows "Principal": {"AWS": "arn:aws:iam::account-number:group/Dev"} is not a solution, because resource policies do not support IAM groups as principals<sup>3</sup>. Principals are entities that can perform actions on AWS resources, such as IAM users, roles, or AWS accounts<sup>4</sup>. IAM groups are not principals, but collections of IAM users that share the same permissions<sup>5</sup>.
- C. Removing the AWS Control Tower control (guardrail) that restricts access to Amazon SES is not a solution, because AWS Control Tower does not have any guardrails that restrict access to Amazon SES<sup>6</sup>. Guardrails are high-level rules that govern the overall behavior of an organization's accounts<sup>7</sup>. AWS Control Tower provides a set of predefined guardrails that cover security, compliance, and operations domains<sup>8</sup>.

References:

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

### NEW QUESTION 113

A company uses infrastructure as code (IaC) to create AWS infrastructure. The company writes the code as AWS CloudFormation templates to deploy the infrastructure. The company has an existing CI/CD pipeline that the company can use to deploy these templates.

After a recent security audit, the company decides to adopt a policy-as-code approach to improve the company's security posture on AWS. The company must prevent the deployment of any infrastructure that would violate a security policy, such as an unencrypted Amazon Elastic Block Store (Amazon EBS) volume.

Which solution will meet these requirements?

- A. Turn on AWS Trusted Advisor
- B. Configure security notifications as webhooks in the preferences section of the CI/CD pipeline.
- C. Turn on AWS Config
- D. Use the prebuilt rules or customized rule
- E. Subscribe the CI/CD pipeline to an Amazon Simple Notification Service (Amazon SNS) topic that receives notifications from AWS Config.
- F. Create rule sets in AWS CloudFormation Guard
- G. Run validation checks for CloudFormation templates as a phase of the CI/CD process.



- H. Create rule sets as SCP
- I. Integrate the SCPs as a part of validation control in a phase of the CI/CD process.

**Answer: C**

**Explanation:**

The correct answer is C. Create rule sets in AWS CloudFormation Guard. Run validation checks for CloudFormation templates as a phase of the CI/CD process. This answer is correct because AWS CloudFormation Guard is a tool that helps you implement policy-as-code for your CloudFormation templates. You can use Guard to write rules that define your security policies, such as requiring encryption for EBS volumes, and then validate your templates against those rules before deploying them. You can integrate Guard into your CI/CD pipeline as a step that runs the validation checks and prevents the deployment of any non-compliant templates<sup>12</sup>.

The other options are incorrect because:

- A. Turning on AWS Trusted Advisor and configuring security notifications as webhooks in the preferences section of the CI/CD pipeline is not a solution, because AWS Trusted Advisor is not a policy-as-code tool, but a service that provides recommendations to help you follow AWS best practices. Trusted Advisor does not allow you to define your own security policies or validate your CloudFormation templates against them<sup>3</sup>.
- B. Turning on AWS Config and using the prebuilt or customized rules is not a solution, because AWS Config is not a policy-as-code tool, but a service that monitors and records the configuration changes of your AWS resources. AWS Config does not allow you to validate your CloudFormation templates before deploying them, but only evaluates the compliance of your resources after they are created<sup>4</sup>.
- D. Creating rule sets as SCPs and integrating them as a part of validation control in a phase of the CI/CD process is not a solution, because SCPs are not policy-as-code tools, but policies that you can use to manage permissions in your AWS Organizations. SCPs do not allow you to validate your CloudFormation templates, but only restrict the actions that users and roles can perform in your accounts<sup>5</sup>.

References:

1: What is AWS CloudFormation Guard? 2: Introducing AWS CloudFormation Guard 2.0 3: AWS Trusted Advisor 4: What Is AWS Config? 5: Service control policies - AWS Organizations

**NEW QUESTION 115**

There are currently multiple applications hosted in a VPC. During monitoring it has been noticed that multiple port scans are coming in from a specific IP Address block. The internal security team has requested that all offending IP Addresses be denied for the next 24 hours. Which of the following is the best method to quickly and temporarily deny access from the specified IP Address's.

Please select:

- A. Create an AD policy to modify the Windows Firewall settings on all hosts in the VPC to deny access from the IP Address block.
- B. Modify the Network ACLs associated with all public subnets in the VPC to deny access from the IP Address block.
- C. Add a rule to all of the VPC Security Groups to deny access from the IP Address block.
- D. Modify the Windows Firewall settings on all AMI'S that your organization uses in that VPC to deny access from the IP address block.

**Answer: B**

**Explanation:**

NACL acts as a firewall at the subnet level of the VPC and we can deny the offending IP address block at the subnet level using NACL rules to block the incoming traffic to the VPC instances. Since NACL rules are applied as per the Rule numbers make sure that this rule number should take precedence over other rule numbers if there are any such rules that will allow traffic from these IP ranges. The lowest rule number has more precedence over a rule that has a higher number. The IAM Documentation mentions the following as a best practices for IAM users

For extra security, enable multi-factor authentication (MFA) for privileged IAM users (users who are allowed access to sensitive resources or APIs). With MFA, users have a device that generates a unique authentication code (a one-time password, or OTP). Users must provide both their normal credentials (like their user name and password) and the OTP. The MFA device can either be a special piece of hardware, or it can be a virtual device (for example, it can run in an app on a smartphone).

Options C is invalid because these options are not available Option D is invalid because there is not root access for users

For more information on IAM best practices, please visit the below URL: <https://docs.IAM.amazon.com/IAM/latest/UserGuide/best-practices.html>

The correct answer is: Modify the Network ACLs associated with all public subnets in the VPC to deny access from the IP Address block.

omit your Feedback/Queries to our Experts

**NEW QUESTION 119**

A company stores images for a website in an Amazon S3 bucket. The company is using Amazon CloudFront to serve the images to end users. The company recently discovered that the images are being accessed from countries where the company does not have a distribution license.

Which actions should the company take to secure the images to limit 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:**

To secure the images to limit their distribution, the company should take the following actions:

- Update the S3 bucket policy to restrict access to a CloudFront origin access identity (OAI). This allows the company to use a special CloudFront user that can access objects in their S3 bucket, and prevent anyone else from accessing them directly.
- Add a CloudFront geo restriction deny list of countries where the company lacks a license. This allows the company to use a feature that controls access to their content based on the geographic location of their viewers, and block requests from countries where they do not have a distribution license.

**NEW QUESTION 120**

While securing the connection between a company's VPC and its on-premises data center, a Security Engineer sent a ping command from an on-premises host (IP address 203.0.113.12) to an Amazon EC2 instance (IP address 172.31.16.139).

The ping command did not return a response. The flow log in the VPC showed the following:

```
2 123456789010 eni-1235b8ca 203.0.113.12 172.31.16.139 0 0 1 4 336 1432917027 1432917142 ACCEPT OK
2 123456789010 eni-1235b8ca 172.31.16.139 203.0.113.12 0 0 1 4 336 1432917094 1432917142 REJECT OK
```



What action should be performed to allow the ping to work?

- A. In the security group of the EC2 instance, allow inbound ICMP traffic.
- B. In the security group of the EC2 instance, allow outbound ICMP traffic.
- C. In the VPC's NACL, allow inbound ICMP traffic.
- D. In the VPC's NACL, allow outbound ICMP traffic.

**Answer:** D

#### NEW QUESTION 121

A company finds that one of its Amazon EC2 instances suddenly has a high CPU usage. The company does not know whether the EC2 instance is compromised or whether the operating system is performing background cleanup.

Which combination of steps should a security engineer take before investigating the issue? (Select THREE.)

- A. Disable termination protection for the EC2 instance if termination protection has not been disabled.
- B. Enable termination protection for the EC2 instance if termination protection has not been enabled.
- C. Take snapshots of the Amazon Elastic Block Store (Amazon EBS) data volumes that are attached to the EC2 instance.
- D. Remove all snapshots of the Amazon Elastic Block Store (Amazon EBS) data volumes that are attached to the EC2 instance.
- E. Capture the EC2 instance metadata, and then tag the EC2 instance as under quarantine.
- F. Immediately remove any entries in the EC2 instance metadata that contain sensitive information.

**Answer:** BCE

#### Explanation:

[https://d1.awsstatic.com/WWPS/pdf/aws\\_security\\_incident\\_response.pdf](https://d1.awsstatic.com/WWPS/pdf/aws_security_incident_response.pdf)

#### NEW QUESTION 123

A company uses Amazon RDS for MySQL as a database engine for its applications. A recent security audit revealed an RDS instance that is not compliant with company policy for encrypting data at rest. A security engineer at the company needs to ensure that all existing RDS databases are encrypted using server-side encryption and that any future deviations from the policy are detected.

Which combination of steps should the security engineer take to accomplish this? (Select TWO.)

- A. Create an IAM Config rule to detect the creation of unencrypted RDS database
- B. Create an Amazon EventBridge (Amazon CloudWatch Events) rule to trigger on the IAM Config rules compliance state change and use Amazon Simple Notification Service (Amazon SNS) to notify the security operations team.
- C. Use IAM System Manager State Manager to detect RDS database encryption configuration drift
- D. Create an Amazon EventBridge (Amazon CloudWatch Events) rule to track state changes and use Amazon Simple Notification Service (Amazon SNS) to notify the security operations team.
- E. Create a read replica for the existing unencrypted RDS database and enable replica encryption in the process
- F. Once the replica becomes active, promote it into a standalone database instance and terminate the unencrypted database instance.
- G. Take a snapshot of the unencrypted RDS database
- H. Copy the snapshot and enable snapshot encryption in the process
- I. Restore the database instance from the newly created encrypted snapshot
- J. Terminate the unencrypted database instance.
- K. Enable encryption for the identified unencrypted RDS instance by changing the configurations of the existing database

**Answer:** AD

#### NEW QUESTION 128

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 130

A developer has created an AWS Lambda function in a company's development account. The Lambda function requires the use of an AWS Key Management Service (AWS KMS) customer managed key that exists in a security account that the company's security team controls. The developer obtains the ARN of the KMS key from a previous Lambda function in the development account. The previous Lambda function had been working properly with the KMS key. When the developer uses the ARN and tests the new Lambda function an error message states that access is denied to the KMS key in the security account. The developer tests the previous Lambda function that uses the same KMS key and discovers that the previous Lambda function still can encrypt data as expected. A security engineer must resolve the problem so that the new Lambda function in the development account can use the KMS key from the security account. Which combination of steps should the security engineer take to meet these requirements? (Select TWO.)

- A. In the security account configure an IAM role for the new Lambda function
- B. Attach an IAM policy that allows access to the KMS key in the security account.
- C. In the development account configure an IAM role for the new Lambda function
- D. Attach a key policy that allows access to the KMS key in the security account.
- E. In the development account configure an IAM role for the new Lambda function
- F. Attach an IAM policy that allows access to the KMS key in the security account.
- G. Configure a key policy for the KMS key in the security account to allow access to the IAM role of the new Lambda function in the security account.
- H. Configure a key policy for the KMS key in the security account to allow access to the IAM role of the new Lambda function in the development account.

**Answer: CE**

### Explanation:

To allow cross-account access to a KMS key, the key policy of the KMS key must grant permission to the external account or principal, and the IAM policy of the external account or principal must delegate the key policy permission. In this case, the new Lambda function in the development account needs to use the KMS key in the security account, so the key policy of the KMS key must allow access to the IAM role of the new Lambda function in the development account (option E), and the IAM role of the new Lambda function in the development account must have an IAM policy that allows access to the KMS key in the security account (option C). Option A is incorrect because it creates an IAM role for the new Lambda function in the security account, not in the development account. Option B is incorrect because it attaches a key policy to an IAM role, which is not valid. Option D is incorrect because it allows access to the IAM role of the new Lambda function in the security account, not in the development account. Verified References:

➤ <https://docs.aws.amazon.com/autoscaling/ec2/userguide/key-policy-requirements-EBS-encryption.html>

### NEW QUESTION 133

A System Administrator is unable to start an Amazon EC2 instance in the eu-west-1 Region using an IAM role. The same System Administrator is able to start an EC2 instance in the eu-west-2 and eu-west-3 Regions. The IAMSystemAdministrator access policy attached to the System Administrator IAM role allows unconditional access to all IAM services and resources within the account. Which configuration caused this issue?

- A) An SCP is attached to the account with the following permission statement:

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "All",
      "Action": "*",
      "Resource": "*"
    },
    {
      "Effect": "Deny",
      "NotAction": [
        "iam:*",
        "organizations:*",
        "route53:*",
        "budgets:*",
        "waf:*",
        "cloudfront:*",
        "globalaccelerator:*",
        "importexport:*",
        "support:*"
      ],
      "Resource": "*",
      "Condition": {
        "StringEquals": {
          "aws:RequestedRegion": [
            "eu-west-*"
          ]
        }
      }
    }
  ]
}
```

- B)  
A permission boundary policy is attached to the System Administrator role with the following permission statement:

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "ec2:*"
      ],
      "Resource": "*"
    },
    {
      "Effect": "Deny",
      "NotAction": [
        "iam:*",
        "organizations:*",
        "route53:*",
        "budgets:*",
        "waf:*",
        "cloudfront:*",
        "globalaccelerator:*",
        "importexport:*",
        "support:*",
        "ec2:*"
      ],
      "Resource": "*",
      "Condition": {
        "StringEquals": {
          "aws:RequestedRegion": [
            "eu-west-1"
          ]
        }
      }
    }
  ]
}
```

C)  
 A permission boundary is attached to the System Administrator role with the following permission statement:

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "ec2:*"
      ],
      "Resource": "*"
    },
    {
      "Version": "2012-10-17",
      "Statement": [
        {
          "Effect": "Allow",
          "Action": "ec2:*",
          "Resource": "*",
          "Condition": {
            "StringEquals": {
              "aws:RequestedRegion": [
                "eu-west-1"
              ]
            }
          }
        }
      ]
    }
  ]
}
```

D)  
 An SCP is attached to the account with the following statement:



```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": "*",
      "Resource": "*"
    },
    {
      "Effect": "Deny",
      "NotAction": [
        "iam:*",
        "organizations:*",
        "route53:*",
        "budgets:*",
        "waf:*",
        "cloudfront:*",
        "globalaccelerators:*",
        "importexport:*",
        "support:*",
        "ec2:*"
      ],
      "Resource": "*",
      "Condition": {
        "StringEquals": {
          "aws:RequestedRegion": "eu-west-1"
        }
      }
    }
  ]
}
```

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

**Answer: B**

#### NEW QUESTION 134

Auditors for a health care company have mandated that all data volumes be encrypted at rest Infrastructure is deployed mainly via IAM CloudFormation however third-party frameworks and manual deployment are required on some legacy systems  
What is the BEST way to monitor, on a recurring basis, whether all EBS volumes are encrypted?

- A. On a recurring basis, update an IAM user policies to require that EC2 instances are created with an encrypted volume
- B. Configure an IAM Config rule to run on a recurring basis for volume encryption
- C. Set up Amazon Inspector rules for volume encryption to run on a recurring schedule
- D. Use CloudWatch Logs to determine whether instances were created with an encrypted volume

**Answer: B**

#### Explanation:

To support answer B, use the reference <https://d1.IAMstatic.com/whitepapers/IAM-security-whitepaper.pdf> "For example, IAM Config provides a managed IAM Config Rules to ensure that encryption is turned on for all EBS volumes in your account."

#### NEW QUESTION 136

An Application team has requested a new IAM KMS master key for use with Amazon S3, but the organizational security policy requires separate master keys for different IAM services to limit blast radius.  
How can an IAM KMS customer master key (CMK) be constrained to work with only Amazon S3?

- A. Configure the CMK key policy to allow only the Amazon S3 service to use the kms Encrypt action
- B. Configure the CMK key policy to allow IAM KMS actions only when the kms:ViaService condition matches the Amazon S3 service name.
- C. Configure the IAM user's policy to allow KMS to pass a role to Amazon S3
- D. Configure the IAM user's policy to allow only Amazon S3 operations when they are combined with the CMK

**Answer: B**

#### Explanation:

the kms:ViaService condition key can be used to restrict a CMK to work with only a specific AWS service<sup>6</sup>. By configuring the CMK key policy to allow KMS actions only when the kms:ViaService condition matches the Amazon S3 service name, you can ensure that only Amazon S3 can use the CMK<sup>7</sup>. The other options are either incorrect or insufficient for constraining a CMK to work with only Amazon S3.

#### NEW QUESTION 141

A company wants to establish separate IAM Key Management Service (IAM KMS) keys to use for different IAM services. The company's security engineer created the following key policy to allow the infrastructure deployment team to create encrypted Amazon Elastic Block Store (Amazon EBS) volumes by assuming the InfrastructureDeployment IAM role:

```
{
  "Version": "2012-10-17",
  "Id": "key-policy-eps",
  "Statement": [
    {
      "Sid": "Enable IAM User Permissions",
      "Effect": "Allow",
      "Principal": {
        "AWS": "arn:aws:iam::123456789012:root"
      },
      "Action": "kms:*",
      "Resource": "*"
    },
    {
      "Sid": "Allow use of the key",
      "Effect": "Allow",
      "Principal": {
        "AWS": "arn:aws:iam::123456789012:role/aws-reserved/sso.amazonaws.com/InfrastructureDeployment"
      },
      "Action": [
        "kms:Encrypt",
        "kms:Decrypt",
        "kms:ReEncrypt*",
        "kms:GenerateDataKey*",
        "kms:DescribeKey",
        "kms:CreateGrant",
        "kms:ListGrants",
        "kms:RevokeGrant"
      ],
      "Resource": "*",
      "Condition": {
        "StringEquals": {
          "kms:ViaService": "ec2.us-west-2.amazonaws.com"
        }
      }
    }
  ]
}
```

The security engineer recently discovered that IAM roles other than the InfrastructureDeployment role used this key (or other services). Which change to the policy should the security engineer make to resolve these issues?

- A. In the statement block that contains the Sid "Allow use of the key", under the "Condition" block, change StringEquals to StringLike.
- B. In the policy document, remove the statement block that contains the Sid "Enable IAM User Permissions". Add key management policies to the KMS policy.
- C. In the statement block that contains the Sid "Allow use of the Key", under the "Condition" block, change the Kms:ViaService value to ec2.us-east-1.amazonaws.com.
- D. In the policy document, add a new statement block that grants the kms:Disable permission to the security engineer's IAM role.

**Answer: C**

**Explanation:**

To resolve the issues, the security engineer should make the following change to the policy:

➤ In the statement block that contains the Sid "Allow use of the key", under the "Condition" block, change the Kms:ViaService value to ec2.us-east-1.amazonaws.com. This allows the security engineer to restrict the use of the key to only EC2 service in the us-east-1 region, and prevent other services from using the key.

**NEW QUESTION 145**

A security engineer is creating an AWS Lambda function. The Lambda function needs to use a role that is named LambdaAuditRole to assume a role that is named AcmeAuditFactoryRole in a different AWS account.

When the code is processed, the following error message appears: "An error occurred (AccessDenied) when calling the AssumeRole operation."

Which combination of steps should the security engineer take to resolve this error? (Select TWO.)

- A. Ensure that LambdaAuditRole has the sts:AssumeRole permission for AcmeAuditFactoryRole.
- B. Ensure that LambdaAuditRole has the AWSLambdaBasicExecutionRole managed policy attached.
- C. Ensure that the trust policy for AcmeAuditFactoryRole allows the sts:AssumeRole action from LambdaAuditRole.
- D. Ensure that the trust policy for LambdaAuditRole allows the sts:AssumeRole action from the lambda.amazonaws.com service.
- E. Ensure that the sts:AssumeRole API call is being issued to the us-east-1 Region endpoint.

**Answer: AC**

**NEW QUESTION 149**

A Security Engineer is asked to update an AWS CloudTrail log file prefix for an existing trail. When attempting to save the change in the CloudTrail console, the Security Engineer receives the following error message: "There is a problem with the bucket policy." What will enable the Security Engineer to save the change?

- A. Create a new trail with the updated log file prefix, and then delete the original trail.
- B. Update the existing bucket policy in the Amazon S3 console with the new log file prefix, and then update the log file prefix in the CloudTrail console.
- C. Update the existing bucket policy in the Amazon S3 console to allow the Security Engineer's Principal to perform PutBucketPolicy, and then update the log file prefix in the CloudTrail console.
- D. Update the existing bucket policy in the Amazon S3 console with the new log file prefix, and then update the log file prefix in the CloudTrail console.
- E. Update the existing bucket policy in the Amazon S3 console to allow the Security Engineer's Principal to perform GetBucketPolicy, and then update the log file prefix in the CloudTrail console.

**Answer: C**

**Explanation:**

The correct answer is C. Update the existing bucket policy in the Amazon S3 console with the new log file prefix, and then update the log file prefix in the CloudTrail console.

According to the AWS documentation<sup>1</sup>, a bucket policy is a resource-based policy that you can use to grant access permissions to your Amazon S3 bucket and the objects in it. Only the bucket owner can associate a policy with a bucket. The permissions attached to the bucket apply to all of the objects in the bucket that are owned by the bucket owner.

When you create a trail in CloudTrail, you can specify an existing S3 bucket or create a new one to store your log files. CloudTrail automatically creates a bucket policy for your S3 bucket that grants CloudTrail write-only access to deliver log files to your bucket. The bucket policy also grants read-only access to AWS services that you can use to view and analyze your log data, such as Amazon Athena, Amazon CloudWatch Logs, and Amazon QuickSight.

If you want to update the log file prefix for an existing trail, you must also update the existing bucket policy in the S3 console with the new log file prefix. The log file prefix is part of the resource ARN that identifies the objects in your bucket that CloudTrail can access. If you don't update the bucket policy with the new log file prefix, CloudTrail will not be able to deliver log files to your bucket, and you will receive an error message when you try to save the change in the CloudTrail console.

The other options are incorrect because:

- A. Creating a new trail with the updated log file prefix, and then deleting the original trail is not necessary and may cause data loss or inconsistency. You can simply update the existing trail and its associated bucket policy with the new log file prefix.
- B. Updating the existing bucket policy in the S3 console to allow the Security Engineer's Principal to perform PutBucketPolicy is not relevant to this issue. The PutBucketPolicy action allows you to create or replace a policy on a bucket, but it does not affect CloudTrail's ability to deliver log files to your bucket. You still need to update the existing bucket policy with the new log file prefix.
- D. Updating the existing bucket policy in the S3 console to allow the Security Engineer's Principal to perform GetBucketPolicy is not relevant to this issue. The GetBucketPolicy action allows you to retrieve a policy on a bucket, but it does not affect CloudTrail's ability to deliver log files to your bucket. You still need to update the existing bucket policy with the new log file prefix.

References:

1: Using bucket policies - Amazon Simple Storage Service

**NEW QUESTION 150**

.....



## Thank You for Trying Our Product

### We offer two products:

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

2nd - Questions and Answers in PDF Format

### SCS-C02 Practice Exam Features:

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

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