Vendor: Amazon

Exam Code:SCS-C02

Exam Name:AWS Certified Security - Specialty

Version: Demo

QUESTION 1

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

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

Which solution will meet these requirements?

A. Create an Amazon CloudWatch Logs log group. Configure the load balancers to send logs to the log group. Use the CloudWatch Logs console to search the logs. Create CloudWatch Logs filters on the logs for the required met-rics.

B. Create an Amazon S3 bucket. Configure the load balancers to send logs to the S3 bucket. Use Amazon Athena to search the logs that are in the S3 bucket. Create Amazon CloudWatch filters on the S3 log files for the re-quired metrics.

C. Create an Amazon S3 bucket. Configure the load balancers to send logs to the S3 bucket. Use Amazon Athena to search the logs that are in the S3 bucket. Create Athena queries for the required metrics. Publish the metrics to Amazon CloudWatch.

D. Create an Amazon CloudWatch Logs log group. Configure the load balancers to send logs to the log group. Use the AWS Management Console to search the logs. Create Amazon Athena queries for the required metrics. Publish the metrics to Amazon CloudWatch.

Correct Answer: C

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

QUESTION 2

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 SDK. Use each keyring individually or combine keyrings into a multi-keyring. Decrypt the data by using a keyring that has the primary key in the multi-keyring.

B. Use data key caching. Use the local cache that the AWS Encryption SDK provides with a caching cryptographic materials manager.

C. Use KMS key rotation. Use a local cache in the AWS Encryption SDK with a caching cryptographic materials manager.

D. Use keyrings with the AWS Encryption SDK. Use each keyring individually or combine keyrings into a multi-keyring. Use any of the wrapping keys in the multi-keyring to decrypt the data.

Correct Answer: B

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

QUESTION 3

A security engineer is configuring account-based access control (ABAC) to allow only specific principals to put objects into an Amazon S3 bucket. The principals already have access to Amazon S3.

The security engineer needs to configure a bucket policy that allows principals to put objects into the S3 bucket only if the value of the Team tag on the object matches the value of the Team tag that is associated with the principal. During testing, the security engineer notices that a principal can still put objects into the S3 bucket when the tag values do not match.

Which combination of factors are causing the PutObject operation to succeed when the tag values are different? (Select TWO.)

A. The principal\\'s identity-based policy grants access to put objects into the S3 bucket with no conditions.

B. The principal\\'s identity-based policy overrides the condition because the identity-based policy contains an explicit allow.

C. The S3 bucket\\'s resource policy does not deny access to put objects.

D. The S3 bucket\\'s resource policy cannot allow actions to the principal.

E. The bucket policy does not apply to principals in the same zone of trust.

Correct Answer: AC

When using ABAC, the principal\\'s identity-based policy and the S3 bucket\\'s resource policy are both evaluated to determine the effective permissions. If either policy grants access to the principal, the action is allowed. If either policy denies access to the principal, the action is denied. Therefore, to enforce the tag-based condition, both policies must deny access when the tag values do not match. In this case, the principal\\'s identity-based policy grants access to put objects into the S3 bucket with no conditions (A), which means that the policy does not check for the tag values. This policy overrides the condition in the bucket policy because an explicit allow always takes precedence over an implicit deny. The bucket policy can only allow or deny actions to the principal based on the condition, but it cannot override the identity-based policy. The S3 bucket\\'s resource policy does not deny access to put objects ? which means that it also does not check for the tag values. The bucket policy can only allow or deny actions to the principal based on the condition, but it cannot override the identity-based policy. Therefore, the combination of factors A and C are causing the PutObject operation to succeed when the tag values are different. References: Using ABAC with Amazon S3 Bucket policy examples

---

QUESTION 4

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

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

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

How will the security engineer be able to comply with these requirements?

A. Remove the existing NAT gateway. Create a new NAT gateway that only the application server subnets can use.

B. Configure the DB instanceTMs inbound network ACL to deny traffic from the security group ID of the NAT gateway.

C. Modify the route tables of the DB instance subnets to remove the default route to the NAT gateway.

D. Configure the route table of the NAT gateway to deny connections to the DB instance subnets.

Correct Answer: C

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

QUESTION 5

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

Which solution meets these criteria?

A. A customer managed CMK that uses customer provided key material

B. A customer managed CMK that uses AWS provided key material

C. An AWS managed CMK

D. Operation system-native encryption that uses GnuPG

Correct Answer: A

https://awscli.amazonaws.com/v2/documentation/api/latest/reference/kms/import-key-material.html

aws kms import-key-material \

--key-id 1234abcd-12ab-34cd-56ef-1234567890ab \

--encrypted-key-material fileb://EncryptedKeyMaterial.bin \ --import-token fileb://ImportToken.bin \

--expiration-model KEY_MATERIAL_EXPIRES \

--valid-to 2021-09-21T19:00:00Z

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

A customer managed CMK is a KMS key that you create, own, and manage in your AWS account. You have full control over the key configuration, permissions, rotation, and deletion. You can use a customer managed CMK to encrypt and

decrypt data in AWS services that are integrated with AWS KMS, such as Amazon EBS1. A customer managed CMK can use either AWS provided key material or customer provided key material. AWS provided key material is generated by

AWS KMS and never leaves the service unencrypted. Customer provided key material is generated outside of AWS KMS and imported into a customer managed CMK. You can specify an expiration date for the imported key material,

after

which the CMK becomes unusable until you reimport new key material2.

To meet the criteria of automatically expiring the key material in 90 days, you need to use customer provided key material and set the expiration date accordingly. This way, you can ensure that the data encrypted with the CMK will not be

accessible after 90 days unless you reimport new key material and re-encrypt the data. The other options are incorrect for the following reasons:

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

C. An AWS managed CMK is a KMS key that is created, owned, and managed by an AWS service on your behalf. You have limited control over the key configuration, permissions, rotation, and deletion. You cannot use an AWS managed CMK to encrypt data in other AWS services or applications. You also cannot set an expiration date for the key material of an AWS managed CMK4. D. Operation system-native encryption that uses GnuPG is not a solution that uses AWS KMS. GnuPG is a command line tool that implements the OpenPGP standard for encrypting and signing data. It does not integrate with Amazon EBS or other AWS services. It also does not provide a way to automatically expire the key material used for encryption5. References:

1: Customer Managed Keys -AWS Key Management Service

2: [Importing Key Material in AWS Key Management Service (AWS KMS) -AWS Key Management Service]

3: [Rotating Customer Master Keys -AWS Key Management Service]

4: [AWS Managed Keys -AWS Key Management Service]

5: The GNU Privacy Guard

QUESTION 6

A company has a group of Amazon EC2 instances in a single private subnet of a VPC with no internet gateway attached. A security engineer has installed the Amazon CloudWatch agent on all instances in that subnet to capture logs from a specific application. To ensure that the logs flow securely, the company\\'s networking team has created VPC endpoints for CloudWatch monitoring and CloudWatch logs. The networking team has attached the endpoints to the VPC. The application is generating logs. However, when the security engineer queries CloudWatch, the logs do not appear.

Which combination of steps should the security engineer take to troubleshoot this issue? (Choose three.)

A. Ensure that the EC2 instance profile that is attached to the EC2 instances has permissions to create log streams and write logs.

B. Create a metric filter on the logs so that they can be viewed in the AWS Management Console.

C. Check the CloudWatch agent configuration file on each EC2 instance to make sure that the CloudWatch agent is collecting the proper log files.

D. Check the VPC endpoint policies of both VPC endpoints to ensure that the EC2 instances have permissions to use them.

E. Create a NAT gateway in the subnet so that the EC2 instances can communicate with CloudWatch.

F. Ensure that the security groups allow all the EC2 instances to communicate with each other to aggregate logs before sending.

Correct Answer: ACD

The possible steps to troubleshoot this issue are:

A. Ensure that the EC2 instance profile that is attached to the EC2 instances has permissions to create log streams and write logs. This is a necessary step because the CloudWatch agent uses the credentials from the instance profile to

communicate with CloudWatch1.

C. Check the CloudWatch agent configuration file on each EC2 instance to make sure that the CloudWatch agent is collecting the proper log files. This is a necessary step because the CloudWatch agent needs to know which log files to

monitor and send to CloudWatch2.

D. Check the VPC endpoint policies of both VPC endpoints to ensure that the EC2 instances have permissions to use them. This is a necessary step because the VPC endpoint policies control which principals can access the AWS services

through the endpoints3.

The other options are incorrect because:

B. Creating a metric filter on the logs is not a troubleshooting step, but a way to extract metric data from the logs. Metric filters do not affect the visibility of the logs in the AWS Management Console. E. Creating a NAT gateway in the subnet is

not a solution, because the EC2 instances do not need internet access to communicate with CloudWatch through the VPC endpoints. A NAT gateway would also incur additional costs. F. Ensuring that the security groups allow all the EC2

instances to communicate with each other is not a necessary step, because the CloudWatch agent does not require log aggregation before sending. Each EC2 instance can send its own logs independently to CloudWatch.

References:

1: IAM Roles for Amazon EC2

2: CloudWatch Agent Configuration File: Logs Section

3: Using Amazon VPC Endpoints : Metric Filters : NAT Gateways : CloudWatch Agent Reference: Log Aggregation

---

QUESTION 7

A company needs to retain tog data archives for several years to be compliant with regulations. The tog data is no longer used but It must be retained

What Is the MOST secure and cost-effective solution to meet these requirements?

A. Archive the data to Amazon S3 and apply a restrictive bucket policy to deny the s3 DeleteOotect API

B. Archive the data to Amazon S3 Glacier and apply a Vault Lock policy

C. Archive the data to Amazon S3 and replicate it to a second bucket in a second IAM Region Choose the S3 Standard-Infrequent Access (S3 Standard-1A) storage class and apply a restrictive bucket policy to deny the s3 DeleteObject API

D. Migrate the log data to a 16 T8 Amazon Elastic Block Store (Amazon EBS) volume Create a snapshot of the EBS volume

Correct Answer: B

To securely and cost-effectively retain log data archives for several years, the company should do the following:

Archive the data to Amazon S3 Glacier and apply a Vault Lock policy. This allows the company to use a low-cost storage class that is designed for long-term archival of data that is rarely accessed. It also allows the company to enforce

compliance controls on their S3 Glacier vault by locking a vault access policy that cannot be changed.

QUESTION 8

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

Which solution will meet this requirement?

A. Use Macie to detect an active DDoS event. Create Amazon CloudWatch alarms that respond to Macie findings.

B. Use Amazon Inspector to review resources and to invoke Amazon CloudWatch alarms for any resources that are vulnerable to DDoS attacks.

C. Create an Amazon CloudWatch alarm that monitors Firewall Manager metrics for an active DDoS event.

D. Create an Amazon CloudWatch alarm that monitors Shield Advanced metrics for an active DDoS event.

Correct Answer: D

This answer is correct because AWS Shield Advanced is a service that provides comprehensive protection against DDoS attacks of any size or duration. It also provides metrics and reports on the DDoS attack vectors, duration, and size. You

can create an Amazon CloudWatch alarm that monitors Shield Advanced metrics such as DDoSAttackBitsPerSecond, DDoSAttackPacketsPerSecond, and DDoSAttackRequestsPerSecond to receive alerts if a DDoS attack occurs against

your account.

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

QUESTION 9

A company has a set of EC2 Instances hosted in IAM. The EC2 Instances have EBS volumes which is used to store critical information. There is a business continuity requirement to ensure high availability for the EBS volumes. How can you achieve this?

A. Use lifecycle policies for the EBS volumes

B. Use EBS Snapshots

C. Use EBS volume replication

D. Use EBS volume encryption

Correct Answer: B

Data stored in Amazon EBS volumes is redundantly stored in multiple physical locations as part of normal operation of those services and at no additional charge. However, Amazon EBS replication is stored within the same availability zone, not across multiple zones; therefore, it is highly recommended that you conduct regular snapshots to Amazon S3 for long-term data durability Option A is invalid because there is no lifecycle policy for EBS volumes Option C is invalid because there is no EBS volume replication Option D is invalid because EBS volume encryption will not ensure business continuity For information on security for Compute Resources, please visit the below URL: https://d1.awsstatic.com/whitepapers/Security/Security_Compute_Services_Whitepaper.pdf

---

QUESTION 10

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. Update the IAM principals in the role trust policy as required.

B. Deploy a VPC endpoint for Secrets Manager. Create and attach an endpoint policy that specifies the IAM principals that are allowed to access the secret. Update the list of IAM principals as required.

C. Use a tag-based approach by attaching a resource policy to the secret. Apply tags to the secret and the IAM principals. Use the aws:PrincipalTag and aws:ResourceTag IAM condition keys to control access.

D. Use a deny-by-default approach by using IAM policies to deny access to the secret explicitly. Attach the policies to an IAM group. Add all IAM principals to the IAM group. Remove principals from the group when they need access. Add the principals to the group again when access is no longer allowed.

Correct Answer: C

QUESTION 11

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

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

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

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

A. Create a pre sign-up AWS Lambda trigger. Associate the Amazon Cognito function with the Amazon Cognito user pool.

B. Use a geographic match rule statement to configure an AWS WAF web ACL. Associate the web ACL with the Amazon Cognito user pool.

C. Configure an app client for the application\\'s Amazon Cognito user pool. Use the app client ID to validate the requests in the hosted Ul.

D. Update the application\\'s Amazon Cognito user pool to configure a geographic restriction setting.

E. Use Amazon Cognito to configure a social identity provider (IdP) to validate the requests on the hosted Ul.

Correct Answer: B

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

---

QUESTION 12

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.

Correct Answer: D

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

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

The other options are incorrect because:

A. The IAM policy does not need to allow the kms:DescribeKey permission, because this permission is not required for decrypting objects in S3 using SSE-KMS. The kms:DescribeKey permission allows getting information about a KMS key, such as its creation date, description, and key state.

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.

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.

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