Free Amazon SCS-C03 Practice Questions

Shared by Knight on 12-08-2026

For More Free Questions and Preparation Resources

Check the Links on Last Page

Question 1

Question Type: MultipleChoice

AWS Config cannot deliver configuration snapshots to Amazon S3.

Which TWO actions will remediate this issue?

Options:

A- Verify the S3 bucket policy allows config.amazonaws.com.
B- Verify the IAM role has s3:GetBucketAcl and s3:PutObject permissions.
C- Verify the S3 bucket can assume the IAM role.
D- Verify IAM policy allows AWS Config to write logs.
E- Modify AWS Config API permissions.

Answer:

A, B

Explanation:

AWS Config requires permissions at two levels to deliver configuration data: the AWS Config service role and the S3 bucket policy. The AWS Certified Security -- Specialty Study Guide states that the S3 bucket policy must explicitly allow the config.amazonaws.com service principal to write objects. Additionally, the IAM role used by AWS Config must allow s3:GetBucketAcl and s3:PutObject.

If either permission is missing, AWS Config cannot deliver snapshots and will log delivery errors in CloudTrail. This dual-permission model ensures least privilege while maintaining secure delivery of compliance data.

Other options reference incorrect principals or irrelevant permissions.

Referenced AWS Specialty Documents:

AWS Certified Security -- Specialty Official Study Guide

AWS Config Prerequisites

Question 2

Question Type: MultipleChoice

A company's security engineer receives an alert that indicates that an unexpected principal is accessing a company-owned Amazon Simple Queue Service (Amazon SQS) queue. All the company's accounts are within an organization in AWS Organizations. The security engineer must implement a mitigation solution that minimizes compliance violations and investment in tools outside of AWS.

What should the security engineer do to meet these requirements?

Options:

A- Create security groups and attach them to all SQS queues.
B- Modify network ACLs in all VPCs to restrict inbound traffic.
C- Create interface VPC endpoints for Amazon SQS. Restrict access using aws:SourceVpce and aws:PrincipalOrgId conditions.
D- Use a third-party cloud access security broker (CASB).

Answer:

C

Explanation:

Amazon SQS is a regional service that supports AWS PrivateLink through interface VPC endpoints. According to AWS Certified Security -- Specialty documentation, the most secure and compliant way to restrict access to AWS services is by using VPC endpoints combined with resource-based policies.

By creating interface VPC endpoints for Amazon SQS in all VPCs, traffic to SQS remains on the AWS network and does not traverse the public internet. Using the aws:SourceVpce condition in the SQS queue policy ensures that only requests originating from approved VPC endpoints can access the queue. Adding the aws:PrincipalOrgId condition further restricts access to principals that belong to the same AWS Organization.

Security groups and network ACLs do not apply to SQS because SQS is not deployed inside a VPC. Third-party CASB tools add cost and operational overhead.

Referenced AWS Specialty Documents:

AWS Certified Security -- Specialty Official Study Guide

Amazon SQS Security and VPC Endpoints

AWS Organizations Condition Keys

Question 3

Question Type: MultipleChoice

A company needs a solution to protect critical data from being permanently deleted. The data is stored in Amazon S3 buckets. The company needs to replicate the S3 objects from the company's primary AWS Region to a secondary Region to meet disaster recovery requirements. The company must also ensure that users who have administrator access cannot permanently delete the data in the secondary Region.

Which solution will meet these requirements?

Options:

A- Configure AWS Backup to perform cross-Region S3 backups. Select a backup vault in the secondary Region. Enable AWS Backup Vault Lock in governance mode for the backups in the secondary Region.
B- Implement S3 Object Lock in compliance mode in the primary Region. Configure S3 replication to replicate the objects to an S3 bucket in the secondary Region.
C- Configure S3 replication to replicate the objects to an S3 bucket in the secondary Region. Create an S3 bucket policy to deny the s3:ReplicateDelete action on the S3 bucket in the secondary Region.
D- Configure S3 replication to replicate the objects to an S3 bucket in the secondary Region. Configure S3 object versioning on the S3 bucket in the secondary Region.

Answer:

B

Explanation:

The requirement is twofold:cross-Region replicationfor disaster recovery andimmutabilityso that even administrators cannot permanently delete data in the secondary Region. The S3-native feature designed to prevent deletion (including version deletion) for a defined retention period isS3 Object Lock. When Object Lock is configured incompliance mode,no user, including the root user and administrators, can remove Object Lock protections or permanently delete protected object versions before retention expires. This meets the ''admins cannot permanently delete'' requirement far better than policy-based controls.

To replicate the data for DR, the company can configureS3 replicationfrom the primary Region bucket to a bucket in the secondary Region. With Object Lock in place (and the destination bucket appropriately configured to support Object Lock and versioning), replicated objects can be retained immutably, providing a strongly protected copy for recovery.

Option D (versioning alone) prevents accidental overwrites but does not stop an admin from deleting all versions. Option C only blocks replication deletes; it does not prevent an admin from

deleting objects directly in the secondary bucket. Option A uses AWS Backup vault lock ingovernancemode, which still allows privileged users with special permissions to override retention; it also creates backups rather than true object-level replication. Therefore, Object Lock in compliance mode combined with replication is the best match.

Question 4

Question Type: MultipleChoice

A security engineer is designing a solution that will provide end-to-end encryption between clients and Docker containers running in Amazon Elastic Container Service (Amazon ECS). This solution must also handle volatile traffic patterns.

Which solution would have the MOST scalability and LOWEST latency?

Options:

A- Configure a Network Load Balancer to terminate the TLS traffic and then re-encrypt the traffic to the containers.
B- Configure an Application Load Balancer to terminate the TLS traffic and then re-encrypt the traffic to the containers.
C- Configure a Network Load Balancer with a TCP listener to pass through TLS traffic to the containers.
D- Configure Amazon Route 53 to use multivalue answer routing to send traffic to the containers.

Answer:

C

Explanation:

Network Load Balancers operate at Layer 4 and are optimized for extreme performance, ultra-low latency, and handling sudden traffic spikes. According to AWS Certified Security -- Specialty documentation, using a TCP listener on an NLB allows TLS traffic to pass through directly to backend containers without termination, preserving true end-to-end encryption.

This approach eliminates the overhead of decrypting and re-encrypting traffic at the load balancer, reducing latency and maximizing throughput. NLBs scale automatically to handle volatile traffic patterns and millions of requests per second.

Application Load Balancers operate at Layer 7 and introduce additional latency due to TLS termination and HTTP processing. Route 53 multivalue routing does not provide load balancing

at the transport layer and does not ensure encryption handling.

AWS recommends NLB TCP pass-through for high-performance, end-to-end encrypted container workloads.

Referenced AWS Specialty Documents:

AWS Certified Security -- Specialty Official Study Guide

Elastic Load Balancing Architecture

Network Load Balancer Performance Characteristics

Question 5

Question Type: MultipleChoice

A security engineer recently rotated the host keys for an Amazon EC2 instance. The security engineer is trying to access the EC2 instance by using the EC2 Instance Connect feature. However, the security engineer receives an error for failed host key validation. Before the rotation of the host keys, EC2 Instance Connect worked correctly with this EC2 instance.

What should the security engineer do to resolve this error?

Options:

A- Import the key material into AWS Key Management Service (AWS KMS).
B- Manually upload the new host key to the AWS trusted host keys database.
C- Ensure that the AmazonSSMManagedInstanceCore policy is attached to the EC2 instance profile.
D- Create a new SSH key pair for the EC2 instance.

Answer:

B

Explanation:

EC2 Instance Connect can performserver/host authenticity checksby validating the instance's SSHhost keyagainst atrusted host keyssource. When you rotate the instance's host keys, the host presents anewfingerprint. If the trusted host keys source still contains theoldhost key, connections that enforce host key verification will fail with ahost key validationerror. The fix is to update the trusted host key record so the new host key fingerprint is recognized as valid. Therefore, the correct action is toupload the new host keyto the trusted host keys database used

for EC2 Instance Connect host key verification.

Option A is unrelated: AWS KMS does not store or manage SSH host keys for EC2 Instance Connect validation. Option C is for AWS Systems Manager managed instances and has no effect on SSH host key validation. Option D rotatesuser/client authentication keys(SSH key pair used to log in) but does not resolve a failure that occurs specifically because theserver host keychanged and is no longer trusted. Updating the trusted host keys database restores the expected trust chain and allows Instance Connect to work again with the rotated host keys.

Question 6

Question Type: MultipleChoice

A security engineer is troubleshooting an AWS Lambda function that is namedMyLambdaFunction. The function is encountering an error when the function attempts to read the objects in an Amazon S3 bucket that is namedDOC-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?

Options:

A- Remove the Condition element. Change the Principal element to the following:{ 'AWS': 'arn:aws:lambda:::function:MyLambdaFunction' }

B- Change the Action element to the following:['s3:GetObject*', 's3:GetBucket*']
C- Change the Resource element to'arn:aws:s3:::DOC-EXAMPLE-BUCKET/*'.
D- 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 policy currently grants s3:GetObject but targets thebucket ARN(arn:aws:s3:::DOC-EXAMPLE-BUCKET). For Amazon S3, object-level actions such asGetObjectmust referenceobject ARNs, not the bucket ARN. The correct resource pattern is the bucket ARNwith /*appended (for example, arn:aws:s3:::DOC-EXAMPLE-BUCKET/*) so the permission applies to objects within the bucket. Without this, S3 evaluates the request against a resource that does not match the requested object, resulting in an access denial even though the action appears correct.

The other options do not address the root cause. Expanding actions (Option B) is unnecessary and overly permissive, and it still would not fix the incorrect resource ARN for object reads. Changing principals or removing conditions (Option A) is not required just to allow reads---Lambda typically accesses S3 using the function'sexecution role, and bucket policies are commonly used for cross-account or service-based access control, but the immediate failure here is the mismatch between s3:GetObject and the bucket-only resource. Option D is invalid because it inverts principal/service usage and sets an incorrect resource type for S3 authorization.

Question 7

Question Type: MultipleChoice

A company has hundreds of AWS accounts in an organization in AWS Organizations. The company operates out of a single AWS Region. The company has a dedicated security tooling AWS account in the organization. The security tooling account is configured as the organization's delegated administrator for Amazon GuardDuty and AWS Security Hub. The company has configured the environment to automatically enable GuardDuty and Security Hub for existing AWS accounts and new AWS accounts.

The company is performing control tests on specific GuardDuty findings to make sure that the company's security team can detect and respond to security events. The security team launched an Amazon EC2 instance and attempted to run DNS requests against a test domain,example.com, to generate a DNS finding. However, the GuardDuty finding was never created in the Security Hub delegated administrator account.

Why was the finding not created in the Security Hub delegated administrator account?

Options:

A- VPC flow logs were not turned on for the VPC where the EC2 instance was launched.
B- The VPC where the EC2 instance was launched had the DHCP option configured for a custom OpenDNS resolver.
C- The GuardDuty integration with Security Hub was never activated in the AWS account where the finding was generated.
D- Cross-Region aggregation in Security Hub was not configured.

Answer:

B

Explanation:

GuardDuty's DNS-related detections depend on GuardDuty being able to observeDNS query behaviorthrough AWS-provided DNS resolution paths in the VPC. If a VPC is configured to use acustom DNS resolvervia DHCP options (for example, an OpenDNS resolver) instead of the AmazonProvidedDNS resolver, DNS queries may bypass the visibility path GuardDuty relies on for DNS analysis and pattern detection. In that case, the test traffic (queries to example.com) might not be evaluated by GuardDuty's DNS finding logic, so no DNS finding is generated---and therefore nothing is forwarded into Security Hub.

Option A is incorrect because VPC flow logs are not a prerequisite for GuardDuty to produce DNS findings; GuardDuty uses native telemetry sources and does not require customer-managed flow logs to be enabled. Option D is irrelevant because the company operates in a single Region, and cross-Region aggregation would not be required for the delegated administrator to see findings from the same Region. Option C is less likely given the setup explicitly states GuardDuty and Security Hub are automatically enabled across accounts; also, even if enabled, the specific lack of a DNS finding points to DNS visibility/configuration rather than a downstream integration toggle.

Question 8

Question Type: MultipleChoice

A security engineer configured VPC Flow Logs to publish to Amazon CloudWatch Logs. After 10 minutes, no logs appear. The issue is isolated to the IAM role associated with VPC Flow Logs.

What could be the reason?

Options:

A- logs:GetLogEvents is missing.
B- The engineer cannot assume the role.
C- The vpc-flow-logs.amazonaws.com principal cannot assume the role.
D- The role cannot tag the log stream.

Answer:

C

Explanation:

VPC Flow Logs require an IAM role that CloudWatch Logs can use to publish flow log records. AWS documentation and AWS Certified Security -- Specialty materials explain that the VPC Flow Logs service must be able to assume the IAM role through its trust policy. The trust relationship must include the service principal vpc-flow-logs.amazonaws.com. If the trust policy does not allow this principal to assume the role, flow logs cannot be delivered and no records will appear in the CloudWatch Logs log group even when traffic exists. logs:GetLogEvents is not required for delivery; it is used for reading logs. The security engineer's ability to assume the role is not relevant because the service, not the engineer, assumes it. Tagging permissions are not required for basic log delivery. Therefore, the most likely cause is an incorrect trust policy that prevents the VPC Flow Logs service principal from assuming the role.

Referenced AWS Specialty Documents:

AWS Certified Security -- Specialty Official Study Guide

Amazon VPC Flow Logs IAM Role Requirements

IAM Trust Policies for AWS Services

Question 9

Question Type: MultipleChoice

A company's security engineer receives an abuse notification from AWS indicating that malware is being hosted from the company's AWS account. The security engineer discovers that an IAM user created a new Amazon S3 bucket without authorization.

Which combination of steps should the security engineer take to MINIMIZE the consequences of this compromise? (Select THREE.)

Options:

A- Encrypt all AWS CloudTrail logs.
B- Turn on Amazon GuardDuty.
C- Change the password for all IAM users.
D- Rotate or delete all AWS access keys.
E- Take snapshots of all Amazon Elastic Block Store (Amazon EBS) volumes.
F- Delete any resources that are unrecognized or unauthorized.

Answer:

B, D, F

Explanation:

AWS incident response guidance emphasizes immediate containment, credential invalidation, and removal of malicious resources. According to the AWS Certified Security -- Specialty documentation, compromised credentials must be rotated or deleted immediately to prevent further unauthorized actions. Rotating or deleting access keys directly mitigates ongoing abuse.

Deleting unrecognized or unauthorized resources, such as the malicious S3 bucket, removes the active threat and limits further damage. Enabling Amazon GuardDuty provides continuous monitoring and helps identify additional compromised resources or malicious behavior that may not yet be visible.

Changing passwords for all IAM users is disruptive and unnecessary if compromise scope is limited. Encrypting CloudTrail logs does not reduce active impact. Taking EBS snapshots is primarily for forensic investigation, not immediate consequence minimization.

AWS best practices recommend GuardDuty activation, credential rotation, and removal of malicious resources as first-response actions.

Referenced AWS Specialty Documents:

AWS Certified Security -- Specialty Official Study Guide

AWS Incident Response Best Practices

Amazon GuardDuty Threat Detection