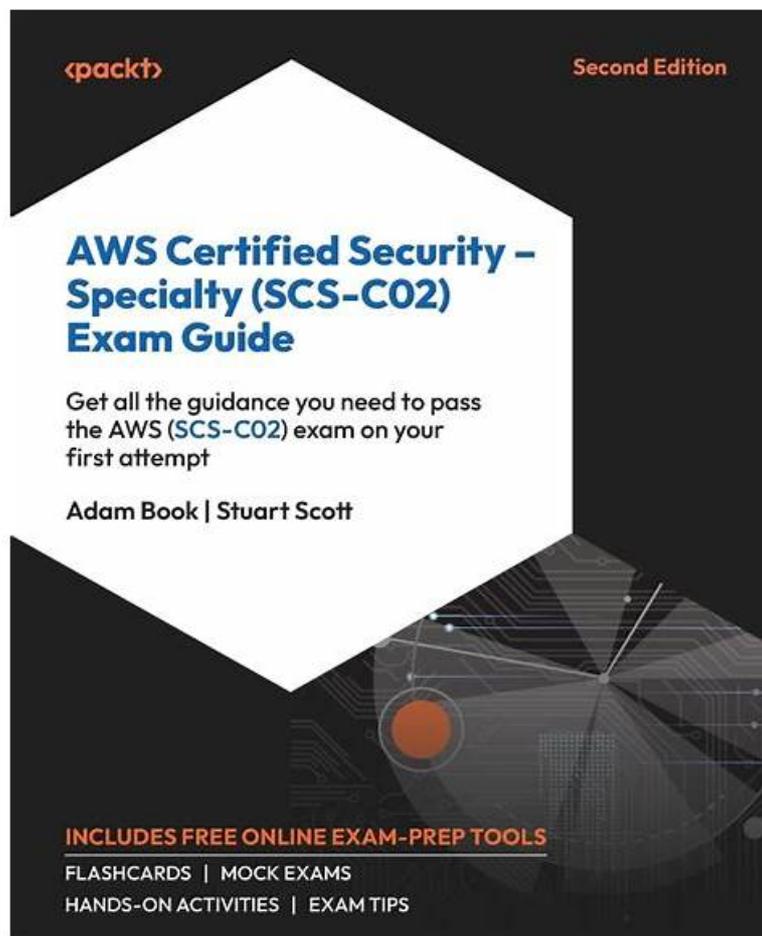


Free PDF Quiz Amazon - SCS-C02 - Perfect Reliable AWS Certified Security - Specialty Test Practice



P.S. Free 2026 Amazon SCS-C02 dumps are available on Google Drive shared by Exams4Collection: <https://drive.google.com/open?id=1K7jg0wijkmyBj3hc24mnVRxQXwNL8d5v>

We have a team of experts curating the real SCS-C02 questions and answers for the end users. We are always working on updating the latest SCS-C02 questions and providing the correct SCS-C02 answers to all of our users. We provide free updates for one year from the date of purchase. You can benefit from the updates SCS-C02 Preparation material, and you will be able to pass the SCS-C02 exam in the first attempt.

Amazon SCS-C02 Exam Syllabus Topics:

Topic	Details
Topic 1	<ul style="list-style-type: none">• Identity and Access Management: The topic equips AWS Security specialists with skills to design, implement, and troubleshoot authentication and authorization mechanisms for AWS resources. By emphasizing secure identity management practices, this area addresses foundational competencies required for effective access control, a vital aspect of the certification exam.
Topic 2	<ul style="list-style-type: none">• Security Logging and Monitoring: This topic prepares AWS Security specialists to design and implement robust monitoring and alerting systems for addressing security events. It emphasizes troubleshooting logging solutions and analyzing logs to enhance threat visibility.

Topic 3	<ul style="list-style-type: none"> • Infrastructure Security: Aspiring AWS Security specialists are trained to implement and troubleshoot security controls for edge services, networks, and compute workloads under this topic. Emphasis is placed on ensuring resilience and mitigating risks across AWS infrastructure. This section aligns closely with the exam's focus on safeguarding critical AWS services and environments.
Topic 4	<ul style="list-style-type: none"> • Threat Detection and Incident Response: In this topic, AWS Security specialists gain expertise in crafting incident response plans and detecting security threats and anomalies using AWS services. It delves into effective strategies for responding to compromised resources and workloads, ensuring readiness to manage security incidents. Mastering these concepts is critical for handling scenarios assessed in the SCS-C02 Exam.

>> **Reliable SCS-C02 Test Practice** <<

SCS-C02 Key Concepts, Valid Braindumps SCS-C02 Ebook

With our SCS-C02 training braindumps, you must feel respected. We believe that every individual has his or her own will, and we will not force you to make any decision. What we can do is to make our SCS-C02 learning prep perfect as much as possible, and let our SCS-C02 practice quiz conquer you with your own charm. And there are three versions of the SCS-C02 exam questions: the PDF, Software and APP online which you can choose as you like.

Amazon AWS Certified Security - Specialty Sample Questions (Q25-Q30):

NEW QUESTION # 25

A security engineer has noticed an unusually high amount of traffic coming from a single IP address. This was discovered by analyzing the Application Load Balancer's access logs.

How can the security engineer limit the number of requests from a specific IP address without blocking the IP address?

- A. Use AWS Shield to limit the originating traffic hit rate.
- B. Add a rule to the Application Load Balancer to route the traffic originating from the IP address in question and show a static webpage.
- C. Implement the GeoLocation feature in Amazon Route 53.
- **D. Implement a rate-based rule with AWS WAF.**

Answer: D

Explanation:

To mitigate traffic volume from a specific IP address without entirely blocking it, AWS WAF's rate-based rules are the appropriate solution. AWS WAF (Web Application Firewall) provides rate-based rules that allow a user to count and limit the rate of requests from individual IP addresses.

A rate-based rule tracks the number of requests that each originating IP makes in a rolling five-minute period.

If the number of requests exceeds a specified threshold, WAF applies an action such as block or count.

This makes AWS WAF an ideal tool to throttle traffic rather than block it, which directly meets the use case described.

Reference from AWS Certified Security - Specialty Official Guide:

This capability is part of AWS WAF's standard feature set, explicitly covered under the topics of Logging and Monitoring and Mitigating DDoS and Abnormal Behavior. Rate-based rules are discussed as a method for limiting the number of incoming requests based on request patterns without denying access outright.

NEW QUESTION # 26

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?

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

Answer: B

Explanation:

The correct answer is C. The GuardDuty integration with Security Hub was never activated in the AWS account where the finding was generated.

According to the AWS documentation¹, GuardDuty findings are automatically sent to Security Hub only if the GuardDuty integration with Security Hub is enabled in the same account and Region. This means that the security tooling account, which is the delegated administrator for both GuardDuty and Security Hub, must enable the GuardDuty integration with Security Hub in each member account and Region where GuardDuty is enabled. Otherwise, the findings from GuardDuty will not be visible in Security Hub.

The other options are incorrect because:

* VPC flow logs are not required for GuardDuty to generate DNS findings. GuardDuty uses VPC DNS logs, which are automatically enabled for all VPCs, to detect malicious or unauthorized DNS activity.

* The DHCP option configured for a custom OpenDNS resolver does not affect GuardDuty's ability to generate DNS findings.

GuardDuty uses its own threat intelligence sources to identify malicious domains, regardless of the DNS resolver used by the EC2 instance.

* Cross-Region aggregation in Security Hub is not relevant for this scenario, because the company operates out of a single AWS Region. Cross-Region aggregation allows Security Hub to aggregate findings from multiple Regions into a single Region.

References:

1: Managing GuardDuty accounts with AWS Organizations : Amazon GuardDuty Findings : How Amazon GuardDuty Works : Cross-Region aggregation in AWS Security Hub

NEW QUESTION # 27

A company is migrating its Amazon EC2 based applications to use Instance Metadata Service Version 2 (IMDSv2). A security engineer needs to determine whether any of the EC2 instances are still using Instance Metadata Service Version 1 (IMDSv1). What should the security engineer do to confirm that the IMDSv1 endpoint is no longer being used?

- A. Create a security group that blocks access to HTTP for the IMDSv1 endpoint. Attach the security group to all EC2 instances.
- B. Configure user data scripts for all EC2 instances to send logging information to AWS CloudTrail when IMDSv1 is used. Create a metric filter and an Amazon CloudWatch dashboard. Track the metric in the dashboard.
- **C. Create an Amazon CloudWatch dashboard. Verify that the EC2:MetadataNoToken metric is zero across all EC2 instances. Monitor the dashboard.**
- D. Configure logging on the Amazon CloudWatch agent for IMDSv1 as part of EC2 instance startup. Create a metric filter and a CloudWatch dashboard. Track the metric in the dashboard.

Answer: C

Explanation:

<https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/configuring-instance-metadata-service.html> IMDSv2 uses token-backed sessions, while IMDSv1 does not. The MetadataNoToken CloudWatch metric tracks the number of calls to the instance metadata service that are using IMDSv1. By tracking this metric to zero, you can determine if and when all of your software has been upgraded to use IMDSv2.

NEW QUESTION # 28

A security engineer recently rotated all IAM access keys in an AWS account. The security engineer then configured AWS Config and enabled the following AWS Config managed rules; mfa-enabled-for-iam-console-access, iam-user-mfa-enabled, access-key-rotated, and iam-user-unused-credentials-check.

The security engineer notices that all resources are displaying as noncompliant after the IAM GenerateCredentialReport API operation is invoked.

What could be the reason for the noncompliant status?

- A. The AWS Config rules have a MaximumExecutionFrequency value of 24 hours.
- B. The security engineer does not have the GetCredentialReport permission.
- C. The IAM credential report was generated within the past 4 hours.
- D. The security engineer does not have the GenerateCredentialReport permission.

Answer: A

Explanation:

The correct answer is D. The AWS Config rules have a MaximumExecutionFrequency value of 24 hours.

According to the AWS documentation¹, the MaximumExecutionFrequency parameter specifies the maximum frequency with which AWS Config runs evaluations for a rule. For AWS Config managed rules, this value can be one of the following:

One_Hour

Three_Hours

Six_Hours

Twelve_Hours

TwentyFour_Hours

If the rule is triggered by configuration changes, it will still run evaluations when AWS Config delivers the configuration snapshot.

However, if the rule is triggered periodically, it will not run evaluations more often than the specified frequency.

In this case, the security engineer enabled four AWS Config managed rules that are triggered periodically. Therefore, these rules will only run evaluations every 24 hours, regardless of when the IAM credential report is generated. This means that the resources will display as noncompliant until the next evaluation cycle, which could take up to 24 hours after the IAM access keys are rotated.

The other options are incorrect because:

A) The IAM credential report can be generated at any time, but it will not affect the compliance status of the resources until the next evaluation cycle of the AWS Config rules.

B) The security engineer was able to invoke the IAM GenerateCredentialReport API operation, which means they have the GenerateCredentialReport permission. This permission is required to generate a credential report that lists all IAM users in an AWS account and their credential status².

C) The security engineer does not need the GetCredentialReport permission to enable or evaluate AWS Config rules. This permission is required to retrieve a credential report that was previously generated by using the GenerateCredentialReport operation².

Reference:

1: AWS::Config::ConfigRule - AWS CloudFormation 2: IAM: Generate and retrieve IAM credential reports

NEW QUESTION # 29

A company has an application that runs on Amazon EC2 instances behind an Application Load Balancer (ALB). The instances are in an Amazon EC2 Auto Scaling group and are attached to Amazon Elastic Block Store (Amazon EBS) volumes.

A security engineer needs to preserve all forensic evidence from one of the instances.

Which order of steps should the security engineer use to meet this requirement?

- A. Take an EBS volume snapshot of the instance and store the snapshot in an Amazon S3 bucket. Take a memory snapshot of the instance and store the snapshot in an S3 bucket. Detach the instance from the Auto Scaling group. Deregister the instance from the ALB. Stop the instance.
- B. Detach the instance from the Auto Scaling group. Deregister the instance from the ALB. Take an EBS volume snapshot of the instance and store the snapshot in an Amazon S3 bucket. Take a memory snapshot of the instance and store the snapshot in an S3 bucket. Stop the instance.
- C. Detach the instance from the Auto Scaling group. Deregister the instance from the ALB. Stop the instance. Take a memory snapshot of the instance and store the snapshot in an Amazon S3 bucket. Take an EBS volume snapshot of the instance and store the snapshot in an S3 bucket.
- D. Take a memory snapshot of the instance and store the snapshot in an Amazon S3 bucket. Stop the instance. Take an EBS volume snapshot of the instance and store the snapshot in an S3 bucket. Detach the instance from the Auto Scaling group. Deregister the instance from the ALB.

Answer: D

Explanation:

The correct answer is B because it preserves the forensic evidence from the instance in the correct order. The first step is to take a memory snapshot of the instance and store it in an S3 bucket, as memory data is volatile and can be lost when the instance is stopped. The second step is to stop the instance, which will prevent any further changes to the EBS volume. The third step is to take an EBS volume snapshot of the instance and store it in an S3 bucket, which will capture the disk state of the instance. The last two steps are to detach the instance from the Auto Scaling group and deregister it from the ALB, which will isolate the instance from the

