

SCS-C02^{Q&As}

AWS Certified Security - Specialty

Pass Amazon SCS-C02 Exam with 100% Guarantee

Free Download Real Questions & Answers **PDF** and **VCE** file from:

<https://www.pass2lead.com/scs-c02.html>

100% Passing Guarantee
100% Money Back Assurance

Following Questions and Answers are all new published by Amazon
Official Exam Center

-  **Instant Download** After Purchase
-  **100% Money Back** Guarantee
-  **365 Days** Free Update
-  **800,000+** Satisfied Customers



QUESTION 1

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 mode. Place objects in the S3 buckets.
- B. Use S3 Glacier Vault Lock to attach a Vault Lock policy to new S3 buckets. Wait 24 hours to complete the Vault Lock process. Place objects in the S3 buckets.
- C. Create new S3 buckets with S3 Object Lock enabled in governance mode. Place objects in the S3 buckets.
- D. Create new S3 buckets with S3 Object Lock enabled in governance mode. Add a legal hold to the S3 buckets. Place objects in the S3 buckets.

Correct Answer: A

QUESTION 2

An international company wants to combine AWS Security Hub findings across all the company's AWS Regions and from multiple accounts. In addition, the company

wants to create a centralized custom dashboard to correlate these findings with operational data for deeper analysis and insights. The company needs an analytics tool to search and visualize Security Hub findings.

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

- A. Designate an AWS account as a delegated administrator for Security Hub. Publish events to Amazon CloudWatch from the delegated administrator account, all member accounts, and required Regions that are enabled for Security Hub findings.
- B. Designate an AWS account in an organization in AWS Organizations as a delegated administrator for Security Hub. Publish events to Amazon EventBridge from the delegated administrator account, all member accounts, and required Regions that are enabled for Security Hub findings.
- C. In each Region, create an Amazon EventBridge rule to deliver findings to an Amazon Kinesis data stream. Configure the Kinesis data streams to output the logs to a single Amazon S3 bucket.
- D. In each Region, create an Amazon EventBridge rule to deliver findings to an Amazon Kinesis Data Firehose delivery stream. Configure the Kinesis Data Firehose delivery streams to deliver the logs to a single Amazon S3 bucket.
- E. Use AWS Glue DataBrew to crawl the Amazon S3 bucket and build the schema. Use AWS Glue Data Catalog to query the data and create views to flatten nested attributes. Build Amazon QuickSight dashboards by using Amazon Athena.
- F. Partition the Amazon S3 data. Use AWS Glue to crawl the S3 bucket and build the schema. Use Amazon Athena to query the data and create views to flatten nested attributes. Build Amazon QuickSight dashboards that use the Athena views.

Correct Answer: BDF

The correct answer is B, D, and F. Designate an AWS account in an organization in AWS Organizations as a delegated administrator for Security Hub. Publish events to Amazon EventBridge from the delegated administrator account, all member accounts, and required Regions that are enabled for Security Hub findings. In each Region, create an Amazon EventBridge rule to deliver findings to an Amazon Kinesis Data Firehose delivery stream. Configure the Kinesis Data Firehose delivery streams to deliver the logs to a single Amazon S3 bucket. Partition the Amazon S3 data. Use AWS Glue to crawl the S3 bucket and build the schema. Use Amazon Athena to query the data and create views to flatten nested attributes. Build Amazon QuickSight dashboards that use the Athena views. According to the AWS documentation, AWS Security Hub is a service that provides you with a comprehensive view of your security state across your AWS accounts, and helps you check your environment against security standards and best practices. You can use Security Hub to aggregate security findings from various sources, such as AWS services, partner products, or your own applications. To use Security Hub 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 Security Hub as a service principal for AWS Organizations, which lets you designate a delegated administrator account for Security Hub. The delegated administrator account can enable Security Hub automatically in all existing and future accounts in your organization, and can view and manage findings from all accounts. According to the AWS documentation, Amazon EventBridge is a serverless event bus that makes it easy to connect applications using data from your own applications, integrated software as a service (SaaS) applications, and AWS services. You can use EventBridge to create rules that match events from various sources and route them to targets for processing. To use EventBridge with Security Hub findings, you need to enable Security Hub as an event source in EventBridge. This will allow you to publish events from Security Hub to EventBridge in the same Region. You can then create EventBridge rules that match Security Hub findings based on criteria such as severity, type, or resource. You can also specify targets for your rules, such as Lambda functions, SNS topics, or Kinesis Data Firehose delivery streams. According to the AWS documentation, Amazon Kinesis Data Firehose is a fully managed service that delivers real-time streaming data to destinations such as Amazon S3, Amazon Redshift, Amazon Elasticsearch Service (Amazon ES), and Splunk. You can use Kinesis Data Firehose to transform and enrich your data before delivering it to your destination. To use Kinesis Data Firehose with Security Hub findings, you need to create a Kinesis Data Firehose delivery stream in each Region where you have enabled Security Hub. You can then configure the delivery stream to receive events from EventBridge as a source, and deliver the logs to a single S3 bucket as a destination. You can also enable data transformation or compression on the delivery stream if needed. 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 Security Hub findings, you need to create an S3 bucket that will store the logs from Kinesis Data Firehose delivery streams. You can then partition the data in the bucket by using prefixes such as account ID or Region. This will improve the performance and cost-effectiveness of querying the data. According to the AWS documentation, AWS Glue is a fully managed extract, transform, and load (ETL) service that makes it easy to prepare and load your data for analytics. You can use Glue to crawl your data sources, identify data formats, and suggest schemas and transformations. You can also use Glue Data Catalog as a central metadata repository for your data assets. To use Glue with Security Hub findings, you need to create a Glue crawler that will crawl the S3 bucket and build the schema for the data. The crawler will create tables in the Glue Data Catalog that you can query using standard SQL. According to the AWS documentation, Amazon Athena is an interactive query service that makes it easy to analyze data in Amazon S3 using standard SQL. Athena is serverless, so there is no infrastructure to manage, and you pay only for the queries that you run. You can use Athena with Glue Data Catalog as a metadata store for your tables. To use Athena with Security Hub findings, you need to create views in Athena that will flatten nested attributes in the data. For example, you can create views that extract fields such as account ID, Region, resource type, resource ID, finding type, finding title, and finding description from the JSON data. You can then query the views using SQL and join them with other tables if needed. According to the AWS documentation, Amazon QuickSight is a fast, cloud-powered business intelligence service that makes it easy to deliver insights to everyone in your organization. You can use QuickSight to create and publish interactive dashboards that include machine learning insights. You can also use QuickSight to connect to various data sources, such as Athena, S3, or RDS. To use QuickSight with Security Hub findings, you need to create QuickSight dashboards that use the Athena views as data sources. You can then visualize and analyze the findings using charts, graphs, maps, or tables. You can also apply filters, calculations, or aggregations to the data. You can then share the dashboards with your users or embed them in your applications.

QUESTION 3

A company is evaluating its security posture. In the past, the company has observed issues with specific hosts and host header combinations that affected

the company's business. The company has configured AWS WAF web ACLs as an initial step to mitigate these issues.

The company must create a log analysis solution for the AWS WAF web ACLs to monitor problematic activity. The company wants to process all the AWS WAF logs in a central location. The company must have the ability to filter out requests based on specific hosts.

A security engineer starts to enable access logging for the AWS WAF web ACLs.

What should the security engineer do next to meet these requirements with the MOST operational efficiency?

A. Specify Amazon Redshift as the destination for the access logs. Deploy the Amazon Athena Redshift connector. Use Athena to query the data from Amazon Redshift and to filter the logs by host.

B. Specify Amazon CloudWatch as the destination for the access logs. Use Amazon CloudWatch Logs Insights to design a query to filter the logs by host.

C. Specify Amazon CloudWatch as the destination for the access logs. Export the CloudWatch logs to an Amazon S3 bucket. Use Amazon Athena to query the logs and to filter the logs by host.

D. Specify Amazon CloudWatch as the destination for the access logs. Use Amazon Redshift Spectrum to query the logs and to filter the logs by host.

Correct Answer: C

Specify Amazon CloudWatch as the destination for the access logs. Export the CloudWatch logs to an Amazon S3 bucket. Use Amazon Athena to query the logs and to filter the logs by host.

According to the AWS documentation, AWS WAF offers logging for the traffic that your web ACLs analyze. The logs include information such as the time that AWS WAF received the request from your protected AWS resource, detailed information about the request, and the action setting for the rule that the request matched. You can send your logs to an Amazon CloudWatch Logs log group, an Amazon Simple Storage Service (Amazon S3) bucket, or an Amazon Kinesis Data Firehose. To create a log analysis solution for the AWS WAF web ACLs, you can use Amazon Athena, which is an interactive query service that makes it easy to analyze data in Amazon S3 using standard SQL. You can use Athena to query and filter the AWS WAF logs by host or any other criteria. Athena is serverless, so there is no infrastructure to manage, and you pay only for the queries that you run.

To use Athena with AWS WAF logs, you need to export the CloudWatch logs to an S3 bucket. You can do this by creating a subscription filter that sends your log events to a Kinesis Data Firehose delivery stream, which then delivers the data to an S3 bucket. Alternatively, you can use AWS DMS to migrate your CloudWatch logs to S3. After you have exported your CloudWatch logs to S3, you can create a table in Athena that points to your S3 bucket and use the AWS service log format that matches your log schema. For example, if you are using JSON format for your AWS WAF logs, you can use the AWSJSONSerDe serde. Then you can run SQL queries on your Athena table and filter the results by host or any other field in your log data. Therefore, this solution meets the requirements of creating a log analysis solution for the AWS WAF web ACLs with the most operational efficiency. This solution does not require setting up any additional infrastructure or services, and it leverages the existing capabilities of CloudWatch, S3, and Athena. The other options are incorrect because:

A. Specifying Amazon Redshift as the destination for the access logs is not possible, because AWS WAF does not support sending logs directly to Redshift. You would need to use an intermediate service such as Kinesis Data Firehose or AWS DMS to load the data from CloudWatch or S3 to Redshift. Deploying the Amazon Athena Redshift connector is not necessary, because you can query Redshift data directly from Athena without using a connector. This solution would also incur additional costs and operational overhead of managing a Redshift cluster.

B. Specifying Amazon CloudWatch as the destination for the access logs is possible, but using Amazon CloudWatch Logs Insights to design a query to filter the logs by host is not efficient or scalable. CloudWatch Logs Insights is a feature that enables you to interactively search and analyze your log data in CloudWatch Logs. However, CloudWatch Logs Insights has some limitations, such as a maximum query duration of 20 minutes, a maximum of 20 log groups per query, and a maximum retention period of 24 months. These limitations may affect your ability to perform complex and long-running analysis on your AWS WAF logs. D. Specifying Amazon CloudWatch as the destination for the access logs is possible, but using Amazon Redshift Spectrum to query the logs and filter them by host is not efficient or cost-effective. Redshift Spectrum is a feature of Amazon Redshift that enables you to run queries against exabytes of data in S3 without loading or transforming any data. However, Redshift Spectrum requires a Redshift cluster to process the queries, which adds additional costs and operational overhead. Redshift Spectrum also charges you based on the number of bytes scanned by each query, which can be expensive if you have large volumes of log data.

References:

- 1: Logging AWS WAF web ACL traffic -Amazon Web Services
 - 2: What Is Amazon Athena? -Amazon Athena
 - 3: Streaming CloudWatch Logs Data to Amazon S3 -Amazon CloudWatch Logs
 - 4: Migrate data from CloudWatch Logs using AWS Database Migration Service -AWS Database Migration Service
 - 5: Querying AWS service logs -Amazon Athena
 - 6: Querying data from Amazon Redshift -Amazon Athena
 - 7: Analyzing log data with CloudWatch Logs Insights -Amazon CloudWatch Logs
 - 8: CloudWatch Logs Insights quotas -Amazon CloudWatch
 - 9: Querying external data using Amazon Redshift Spectrum -Amazon Redshift
 - 10: Amazon Redshift Spectrum pricing -Amazon Redshift
-

QUESTION 4

A company hosts an application on Amazon EC2 that is subject to specific rules for regulatory compliance. One rule states that traffic to and from the workload must be inspected for network-level attacks. This involves inspecting the whole packet.

To comply with this regulatory rule, a security engineer must install intrusion detection software on a c5n.4xlarge EC2 instance. The engineer must then configure the software to monitor traffic to and from the application instances. What should the security engineer do next?

- A. Place the network interface in promiscuous mode to capture the traffic.
- B. Configure VPC Flow Logs to send traffic to the monitoring EC2 instance using a Network Load Balancer.
- C. Configure VPC traffic mirroring to send traffic to the monitoring EC2 instance using a Network Load Balancer.
- D. Use Amazon Inspector to detect network-level attacks and trigger an IAM Lambda function to send the suspicious packets to the EC2 instance.

Correct Answer: D

QUESTION 5

A web application gives users the ability to log in verify their membership's validity and browse artifacts that are stored in an Amazon S3 bucket. When a user attempts to download an object, the application must verify the permission to access the object and allow the user to download the object from a custom domain name such as example.com. What is the MOST secure way for a security engineer to implement this functionality?

- A. Configure read-only access to the object by using a bucket ACL. Remove the access after a set time has elapsed.
- B. Implement an IAM policy to give the user read access to the S3 bucket.
- C. Create an S3 presigned URL Provide the S3 presigned URL to the user through the application.
- D. Create an Amazon CloudFront signed URL. Provide the CloudFront signed URL to the user through the application.

Correct Answer: D

For this scenario you would need to set up static website hosting because a custom domain name is listed as a requirement. "Amazon S3 website endpoints do not support HTTPS or access points. If you want to use HTTPS, you can use Amazon CloudFront to serve a static website hosted on Amazon S3." This is not secure.

<https://docs.aws.amazon.com/AmazonS3/latest/userguide/website-hosting-custom-domain-walkthrough.html>

CloudFront signed URLs allow much more fine-grained control as well as HTTPS access with custom domain names:

<https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/private-content-signed-urls.html>

[Latest SCS-C02 Dumps](#)

[SCS-C02 VCE Dumps](#)

[SCS-C02 Study Guide](#)