PassQueen
a.k.a dumpsqueen.com

Certified SOC Analyst (CSA)

ECCouncil 312-39

Version Demo

Total Demo Questions: 26

Total Premium Questions: 263

Buy Premium PDF

https://passqueen.com

support@passqueen.com

Topic Break Down

Topic	No. of Questions
Topic 1, Security Operations and Management	134
Topic 2, Security Threats, Vulnerabilities, and Attacks	48
Topic 3, Incident Response	40
Topic 4, Digital Forensics and Indicator Analysis Techniques	41
Total	263

QUESTION NO: 1

David Reynolds, a SOC analyst at a healthcare organization, is investigating suspicious login attempts flagged by the SIEM. To mitigate brute-force risk on targeted endpoints, he collaborates with IT to implement an automatic account lockout policy that temporarily disables accounts after multiple failed login attempts. Within the SOC's eradication strategy, which category of measures does this action align with?

A. Physical security measures

B. Network security measures

C. Host security measures

D. Authentication and authorization measures

ANSWER: D

Explanation:

Authentication and authorization measures is correct because an automatic account lockout policy directly controls how an identity is permitted to authenticate after repeated failed attempts. Its purpose is to limit password-guessing and credential-stuffing activity by temporarily preventing further authentication attempts once a configured threshold has been reached. In a SOC eradication or containment context, this reduces the attacker's opportunity to continue exploiting an account password through automated brute-force attempts, while generating an event that can be investigated and correlated with SIEM telemetry. The control is tied to account and credential access decisions: it evaluates failed authentication activity and changes the account's authorization to attempt sign-in for a defined period or until administrative remediation occurs. NIST identifies this capability as "Unsuccessful Logon Attempts," requiring systems to enforce actions such as account or node lockout after a defined number of invalid attempts. This makes the policy an identity and access management safeguard that supports secure authentication, incident response, and recovery of affected accounts. See NIST SP 800-53 Rev. 5 and NIST Digital Identity Guidelines.

QUESTION NO: 2

You are a Threat Hunter in an IT company's security team working to enhance threat hunting capabilities. You observed that relying solely on traditional security alerts often results in missed detections of sophisticated threats. To strengthen your approach, you decide to incorporate multiple data sources, including external threat intelligence feeds, internal security logs, network traffic data, and endpoint telemetry. To efficiently process this vast amount of data, you implement a new tool that can aggregate, normalize, and correlate threat intelligence with internal telemetry to gain a more holistic understanding of emerging threats and enhance detection accuracy. What key threat detection capability is being leveraged in this scenario?

A. Threat Reports

B. Intelligence Buy-In

C. Threat Trending

D. Data Integration

ANSWER: D

Explanation:

Data Integration is the threat detection capability being leveraged because the scenario explicitly describes collecting security-relevant information from diverse sources and making it usable together. External threat-intelligence feeds provide indicators and adversary context, while internal logs, network traffic, and endpoint telemetry provide evidence of activity within the organization. Aggregating these sources into a common platform, normalizing their formats, and correlating related events allows analysts to identify meaningful relationships that are not visible when each data source is reviewed separately.

This capability supports threat hunting by enriching internal events with intelligence about known malicious infrastructure, files, tactics, techniques, and procedures. For example, a network connection, endpoint process event, and intelligence indicator may be correlated to establish a stronger detection hypothesis. Integrated data also enables faster pivots across hosts, users, processes, and network activity, improving both detection fidelity and investigation speed. NIST describes log

management as collecting, storing, analyzing, and correlating log data to support security monitoring and incident response. The MITRE ATT&CK framework similarly helps analysts connect observed telemetry to adversary behaviors. See NIST SP 800-92: Guide to Computer Security Log Management and MITRE ATT&CK.

QUESTION NO: 3

Which of the following security technology is used to attract and trap people who attempt unauthorized or illicit utilization of the host system?

A. De-Militarized Zone (DMZ)

B. Firewall

C. Honeypot

D. Intrusion Detection System

ANSWER: C

Explanation:

Honeypot is correct because it is a deliberately deployed decoy system, service, or resource intended to appear valuable or vulnerable to an attacker. Its purpose is to lure individuals attempting unauthorized access or illicit use away from legitimate assets and into a monitored environment. Because normal users should have no legitimate reason to interact with a honeypot, any access or activity directed at it is highly suspicious and can provide security teams with high-value indicators of malicious behavior.

In SOC operations, honeypots help detect reconnaissance, exploitation attempts, credential attacks, malware activity, and attacker techniques. They can also generate forensic information such as source addresses, commands, payloads, and interaction patterns, supporting incident investigation and threat intelligence. A honeypot must be isolated and carefully monitored so that it remains a controlled detection mechanism rather than becoming a pathway to production systems. The U.S. National Institute of Standards and Technology describes honeypots as systems designed to be probed, attacked, or compromised to collect information about attackers and their methods. See NIST guidance on security monitoring and log management and Kaspersky's overview of honeypots.

QUESTION NO: 4

A government agency needs to monitor its network for unusual data exfiltration attempts. Traditional log data is insufficient to identify traffic anomalies, so the SIEM team integrates traffic flow data to detect large transfers and unexpected spikes. The team must choose the appropriate protocol to collect IP traffic information from routers and switches. Which protocol should be used?

A. SNMP (Simple Network Management Protocol)

B. NetFlow (RFC 3954)

C. Syslog

D. IPFIX (IP Flow Information Export)

ANSWER: D

Explanation:

IPFIX (IP Flow Information Export) is the appropriate protocol because it is the IETF standard for exporting IP flow records from routers, switches, and other network devices to a flow collector or SIEM. A flow record summarizes communications using metadata such as source and destination addresses, ports, protocol, timestamps, packet counts, and byte counts. This visibility lets SOC analysts establish normal traffic baselines and identify suspiciously large outbound transfers, unusual destinations, unexpected protocols, or abrupt volume spikes that can indicate exfiltration. IPFIX is especially useful where encrypted traffic limits payload inspection, since the metadata still provides strong behavioral evidence about network

activity. Its information-element model is extensible, allowing exporters to include standard and enterprise-specific flow fields while maintaining interoperable collection and analysis. The SIEM can correlate these flow records with endpoint, DNS, identity, and proxy telemetry to assess whether a transfer is authorized or anomalous. The Internet Engineering Task Force defines IPFIX in RFC 7011 and describes the protocol's architecture for exporting flow information from metering processes to collecting processes. See IETF RFC 7011: Specification of the IP Flow Information Export (IPFIX) Protocol and the IANA IPFIX Information Elements registry.

QUESTION NO: 5

A SOC team is implementing a threat intelligence strategy to proactively defend against threats. The CISO emphasizes that collecting data is not enough; the team must allocate personnel, tools, and time to gather intelligence aligned with key concerns (fraud, phishing, nation-state threats). They must determine who will collect intelligence, which sources will be monitored, and how frequently collection occurs. What is this process called?

A. Resources

B. Tasking

C. High-level requirements

D. Prioritization

ANSWER: B

Explanation:

Tasking is correct because it converts established intelligence needs into specific, manageable collection assignments. In this scenario, the SOC has already identified the areas of concern, such as fraud, phishing, and nation-state activity. The next operational need is to direct collection activity: assigning personnel, selecting and configuring collection tools, identifying external and internal sources to monitor, and defining the cadence for that monitoring. Those actions are tasking.

Within an intelligence process, tasking ensures that analysts and technical resources are used deliberately against the organization's most relevant risks. A tasking plan may direct one analyst to monitor phishing-reporting channels each day, assign another to assess threat-actor reporting, and configure automated tools to collect relevant telemetry or intelligence-feed indicators continuously. It also establishes accountability, timing, and expected outputs, enabling the SOC to obtain intelligence that can support detection engineering, incident investigation, and risk decisions. This structured direction prevents collection from becoming an unfocused accumulation of data. The U.S. Intelligence Community describes collection as being guided by intelligence requirements and managed through direction and tasking, while CISA emphasizes operationally relevant cyber threat intelligence for defensive use. See ODNI: What Is Intelligence? and CISA: Cyber Threats and Advisories.

QUESTION NO: 6

Jannet works in a multinational corporation that operates multiple data centers, cloud environments, and on-premises systems. As a SOC analyst, she notices that security incidents are taking too long to detect and investigate. After analyzing this, she discovers that logs from firewalls, endpoint security solutions, authentication servers, and cloud applications are scattered across different systems in various formats. Her team has to manually convert logs into a readable format before investigating incidents. What approach should she implement to accept logs from heterogeneous sources with different formats, convert them into a common format, and improve incident detection and response time?

A. Log transformation

B. Log normalization

C. Log correlation

D. Log collection

ANSWER: B

Explanation:

Log normalization is the appropriate approach because it standardizes event data received from diverse technologies and vendors into a consistent schema. Firewalls, endpoint platforms, identity systems, and cloud services commonly produce records with different field names, timestamp formats, event identifiers, and message structures. Normalization parses those source-specific records and maps their contents into common fields, such as source and destination addresses, usernames, actions, event categories, and normalized timestamps. This removes the need for analysts to manually interpret or reformat each log type before beginning an investigation.

In a SOC workflow, normalized data enables SIEM detection rules, dashboards, searches, and incident-response processes to work consistently across the environment. Analysts can query the same logical fields regardless of the original source, which improves detection coverage and shortens triage and investigation time. It also provides a reliable foundation for subsequent analytics, including cross-source event correlation and alert enrichment. The NIST log-management guidance identifies the need to process logs from varied sources and make them usable for analysis, while Microsoft describes normalization as mapping disparate data into a common schema for security operations. See NIST SP 800-92: Guide to Computer Security Log Management and Microsoft Sentinel normalization documentation.

QUESTION NO: 7

A large web hosting service provider, Web4Everyone, hosts multiple major websites and platforms. You are a Level 1 SOC analyst responsible for investigating web server logs for potential malicious activity. Recently, your team detected multiple failed login attempts and unusual traffic patterns targeting the companyâ€™s web application. To efficiently analyze the logs and identify key details such as remote host, username, timestamp, requested resource, HTTP status code, and user-agent, you need a structured log format that ensures quick and accurate parsing. Which standardized log format will you choose for this scenario?

A. JSON Format

B. Common Log Format (CLF)

C. Tab-Separated Format

D. Extended Log Format (ELF)

ANSWER: D

Explanation:

Extended Log Format (ELF) is the appropriate choice because it provides a standardized, field-oriented web logging structure that can record the complete set of investigation details required in this scenario. Its extensible field directives allow an administrator to include client address or remote host, authenticated user name, date and time, requested URI/resource, HTTP status, and the User-Agent request header. This gives a SOC analyst consistent, parseable records for investigating password-spraying attempts, identifying repeated requests from an originating host, correlating activity by account and time, and recognizing scanners, bots, or suspicious automation through User-Agent values.

The format was defined by the W3C specifically to support configurable HTTP log fields while retaining a predictable syntax. That balance is valuable in a hosting environment, where analysts may need to normalize logs from many applications and quickly ingest them into a SIEM or log-analysis pipeline. Including the User-Agent alongside request and response metadata adds useful client context without sacrificing structured parsing. The W3C specification documents fields such as `c-ip`, `cs-username`, `date`, `time`, `cs-uri-stem`, `sc-status`, and `cs(User-Agent)`. See the W3C Extended Log File Format specification and Microsoft IIS extended logging field documentation.

QUESTION NO: 8

Wesley is an incident handler in a company named Maddison Tech. One day, he was learning techniques for eradicating the insecure deserialization attacks.

What among the following should Wesley avoid from considering?

A. Deserialization of trusted data must cross a trust boundary

B. Understand the security permissions given to serialization and deserialization

C. Allow serialization for security-sensitive classes

D. Validate untrusted input, which is to be serialized to ensure that serialized data contain only trusted classes

ANSWER: C

Explanation:

Allow serialization for security-sensitive classes is the practice Wesley should avoid. Serialization converts an object's state into a transportable or storable format, and deserialization reconstructs that object. When classes embody privileged operations, sensitive configuration, authorization state, credentials, or security-critical business logic, making them serializable expands the attack surface. An attacker who can influence serialized content may attempt to manipulate object state or exploit deserialization behavior before normal application validation and authorization controls take effect.

A secure design minimizes the set of serializable types and excludes classes whose state or behavior could affect security decisions. Where object persistence or interchange is required, applications should use constrained data-transfer objects containing only necessary primitive data, apply strict type allowlists and integrity protections, and validate data before it is used. This follows the broader secure-development principle of reducing attack surface and avoiding unsafe deserialization of attacker-controlled data. OWASP identifies deserialization of untrusted data as a source of remote code execution, privilege escalation, and other serious impacts; see the OWASP Deserialization Cheat Sheet and the OWASP guidance on deserialization of untrusted data.

QUESTION NO: 9

Which encoding replaces unusual ASCII characters with "%" followed by the characterâ€™s two-digit ASCII code expressed in hexadecimal?

A. Unicode Encoding

B. UTF Encoding

C. Base64 Encoding

D. URL Encoding

ANSWER: D

Explanation:

URL Encoding is correct. Also called percent-encoding, it represents characters that are not permitted or have reserved meanings in a URI by replacing each relevant byte with a percent sign followed by two hexadecimal digits. For example, an ASCII space has hexadecimal value `20`, so it is commonly represented as `%20` in a URI. This encoding lets browsers, web servers, APIs, and security tools transmit URI components reliably without ambiguity caused by delimiters, whitespace, or non-ASCII data. In formal URI syntax, percent-encoding operates on bytes: non-ASCII characters are first encoded into bytes, typically using UTF-8, and each byte is then written as `%` plus its hexadecimal value. Thus, the question's description of a percent sign followed by a two-digit hexadecimal ASCII value identifies URL Encoding, even though the more precise standards term is percent-encoding. Understanding this format is particularly useful in SOC analysis because encoded URLs, request parameters, and log entries may conceal significant paths, query values, or payload content that must be decoded during triage. RFC 3986 defines the percent-encoding mechanism for URIs: RFC 3986, Section 2.1. Additional practical guidance is available in MDN's Percent-encoding reference.

QUESTION NO: 10

A mid-sized financial institution's SOC is overwhelmed by thousands of daily alerts, many based on Indicators of Compromise (IoCs) such as suspicious IPs, hashes, and domains. These alerts lack context about whether they truly pose a threat. Analysts waste time on low-priority incidents while severe threats may be missed. The team lacks tools and intelligence to correlate IoCs with real-world threats, making prioritization difficult and causing alert fatigue. Which poses the greatest challenge in this environment?

A. Malware-centric and CTI are not equivalent

B. Information overload

C. Budget and enterprise skill

D. Distinguishing IoC from CTI

ANSWER: D

Explanation:

Distinguishing IoC from CTI is the greatest challenge because the SOC is receiving raw technical observables without the contextual intelligence needed to determine their operational significance. An indicator such as an IP address, file hash, or domain can be associated with malicious activity, but it does not by itself establish who is using it, what adversary behavior it supports, whether it is current, how reliable the source is, or whether it is relevant to the institution's assets and risk profile. Cyber threat intelligence adds that decision-making context by connecting indicators to adversaries, campaigns, tactics, techniques, procedures, targeting, and confidence assessments. This enrichment enables analysts to assess likely impact and urgency rather than treating every indicator-driven alert as equally important. In the described environment, the inability to make that distinction directly causes poor prioritization, excessive analyst effort, and the risk that genuinely dangerous activity will be lost among low-value alerts. The distinction is central to turning high-volume detection data into actionable, risk-informed SOC decisions. See the NIST discussion of cyber threat information in NIST SP 800-150 and MITRE's overview of using threat intelligence to understand adversary behavior in MITRE ATT&CK Resources.

QUESTION NO: 11

In which of the following incident handling and response stages, the root cause of the incident must be found from the forensic results?

A. Evidence Gathering

B. Evidence Handling

C. Eradication

D. Systems Recovery

ANSWER: C

Explanation:

Eradication is the incident-response stage in which the responder uses the forensic findings gathered during the investigation to identify the underlying cause of the compromise and eliminate it from the environment. The goal is not merely to remove visible malicious files or stop an active alert; it is to establish how the incident occurred, such as through an exploited vulnerability, stolen credentials, a malicious persistence mechanism, or a configuration weakness, and then remove the attacker's foothold and remediate the condition that enabled it.

This root-cause work is essential before restoring normal operations because recovery should occur only after the organization has confidence that the malicious artifacts, compromised accounts, persistence paths, and exploited weaknesses have been addressed. NIST describes eradication as removing incident components and identifying and mitigating the vulnerabilities exploited by the attacker, which aligns directly with determining the cause from forensic results. See NIST SP 800-61 Rev. 2, Computer Security Incident Handling Guide. EC-Council's incident-handling curriculum likewise treats eradication as the phase where the incident's cause is removed so the organization can proceed safely toward recovery. See the EC-Council Certified Incident Handler program overview.

QUESTION NO: 12

Identify the password cracking attempt involving a precomputed dictionary of plaintext passwords and their corresponding hash values to crack the password.

A. Dictionary Attack

B. Rainbow Table Attack

C. Bruteforce Attack

D. Syllable Attack

ANSWER: B

Explanation:

Rainbow Table Attack is correct because rainbow tables are precomputed data structures used to reverse password hashes efficiently. An attacker generates hashes for a large set of candidate plaintext passwords in advance, then stores the results in a space-efficient chain-based format. When the attacker obtains a target password hash, the table can be searched to identify a plaintext password that produces the same hash. This moves much of the computational cost from the time of the attack to an earlier precomputation phase, making lookup attacks against unsalted, fast password hashes substantially faster than calculating every candidate hash during the attack.

The key characteristic in the question is the use of a precomputed collection associating candidate plaintexts with hash results. Rainbow tables are specifically designed for that purpose. Unique, randomly generated salts prevent a single precomputed table from applying broadly because the salt changes the effective input to the hash for each stored password. Modern password storage should therefore use unique salts together with deliberately slow, password-specific hashing algorithms such as Argon2, bcrypt, PBKDF2, or scrypt. See the NIST definition of rainbow tables and OWASP's Password Storage Cheat Sheet.

QUESTION NO: 13

Where will you find the reputation IP database, if you want to monitor traffic from known bad IP reputation using OSSIM SIEM?

A. /etc/ossim/reputation

B. /etc/ossim/siem/server/reputation/data

C. /etc/siem/ossim/server/reputation.data

D. /etc/ossim/server/reputation.data

ANSWER: D

Explanation:

The OSSIM reputation IP database is stored at `/etc/ossim/server/reputation.data`. OSSIM uses this local reputation-data file as part of its reputation framework to associate observed source or destination IP addresses with known malicious, suspicious, or otherwise negatively rated IP intelligence. When events or network traffic are processed, the SIEM can compare IP values against this reputation data and generate relevant correlation results, alarms, or context for analysts.

This location is specifically within the OSSIM server configuration hierarchy because reputation lookups are handled by the server-side SIEM and correlation components. The `reputation.data` filename denotes the reputation database/data file itself, rather than a generic configuration directory. In operational use, analysts should ensure that reputation feeds and updates are correctly maintained so detections reflect current threat intelligence and avoid relying on stale IP classifications. OSSIM/USM documentation describes reputation intelligence as a mechanism for enriching and correlating monitored activity with known threat information. See the OSSIM source repository and the USM Appliance documentation for product architecture and threat-intelligence context.

QUESTION NO: 14

A SIEM alert is triggered due to unusual network traffic involving NetBIOS. The system log shows: "The TCP/IP NetBIOS Helper service entered the running state." Concurrently, Windows Security Event ID 4624 ("An account was successfully

logged on") appears for multiple machines within a short time frame. The logon type is 3 (Network logon). Which of the following security incidents is the SIEM detecting?

A. An attacker performing lateral movement within the network

B. A user connecting to shared files from multiple workstations

C. A network administrator conducting routine maintenance

D. A malware infection spreading via SMB protocol

ANSWER: A

Explanation:

An attacker performing lateral movement within the network is the best classification for this correlated alert pattern. Windows Security Event ID 4624 with logon type 3 records a successful network logon, meaning credentials were used to access a system remotely over the network. A concentrated burst of these successful logons across multiple endpoints is a high-value indicator of host-to-host activity, particularly when it is anomalous for the account, source host, timing, or number of destination systems. NetBIOS-related traffic and the TCP/IP NetBIOS Helper service can support Windows name-resolution and SMB-based connectivity, which are commonly present when attackers enumerate systems, access administrative shares, authenticate remotely, or execute actions on additional hosts. In a SOC investigation, the analyst should correlate the event records' account name, source network address, workstation name, destination hosts, and process or share-access telemetry to establish the path of movement. This behavior aligns with the MITRE ATT&CK Lateral Movement tactic, where adversaries use valid accounts or remote services to expand access after an initial foothold. Microsoft documents logon type 3 as a network logon in its Event 4624 guidance; relevant adversary behavior is also described in MITRE ATT&CK Lateral Movement.

QUESTION NO: 15

The Security Operations Center (SOC) team at Rapid Response Group, a leading cybersecurity firm, is facing challenges in managing security incidents efficiently. With an increasing volume of alerts and security events being generated daily in their Microsoft Sentinel environment, the team is struggling to respond to threats quickly and consistently. To enhance their incident response capabilities, they aim to automate routine security tasks, such as log collection, alert triaging, remediation steps, and notifications to stakeholders. By implementing automated workflows, they seek to reduce response times, eliminate manual intervention for repetitive actions, and ensure a standardized approach to handling security threats across the organization. Which component of Microsoft Sentinel should they utilize to create these automated workflows for incident response?

A. Community

B. Playbooks

C. Workspace

D. Analytics

ANSWER: B

Explanation:

Playbooks are the Microsoft Sentinel component designed to automate and orchestrate incident-response workflows. A playbook is built on Azure Logic Apps and can be triggered manually by an analyst or automatically through an automation rule when an alert or incident is created. This allows a SOC to perform repeatable actions consistently at scale, including enriching incidents with threat-intelligence or asset data, collecting relevant context, creating service-management tickets, sending stakeholder notifications, requesting approvals, and initiating supported containment or remediation actions.

For the described environment, playbooks directly reduce the operational burden caused by high alert volumes. Rather than requiring an analyst to manually execute the same enrichment, triage, communication, and response steps for every relevant incident, the SOC can define a standardized workflow once and apply it reliably. This improves response consistency and helps reduce time to acknowledge and respond, while retaining analyst oversight for actions that require judgment. Microsoft documents Sentinel playbooks as collections of procedures that can run in response to incidents, alerts, or entities and that

use Azure Logic Apps connectors to integrate with security and business systems. See Microsoft Sentinel playbooks documentation and Microsoft Sentinel automation rules documentation.

QUESTION NO: 16

A Security Operations Center (SOC) analyst receives a high-priority alert indicating unusual user activity. An employee account is attempting to access company resources from a different country and outside of their normal working hours. This behavior raises concerns about potential account compromise or unauthorized access. To automate the initial response and quickly restrict access while further investigating the incident, which SOAR playbook would be relevant to adapt and implement?

A. Alert Enrichment SOAR Playbook

B. Deprovisioning Users SOAR Playbook

C. Malware Containment SOAR Playbook

D. Phishing Investigations SOAR Playbook

ANSWER: B

Explanation:

Deprovisioning Users SOAR Playbook is the appropriate playbook to adapt because the alert indicates a possible identity compromise and the stated immediate objective is to restrict the user's access while the SOC investigates. A deprovisioning workflow can automate containment actions such as disabling the account, revoking active sessions and authentication tokens, blocking sign-in, requiring a password reset, and notifying the identity or incident-response team. These steps reduce the opportunity for an attacker using valid credentials to move laterally, access sensitive resources, or maintain an active session.

The unusual country and out-of-hours pattern are identity-risk indicators that should trigger rapid, reversible containment rather than waiting for a full manual investigation. The workflow can also collect relevant identity-provider logs, create an incident ticket, preserve evidence, and route an approval step for exceptions when legitimate travel is possible. After the user's identity and activity are verified, access can be restored under appropriate controls such as multifactor authentication and conditional-access policies. This approach follows incident-response guidance to contain a suspected compromise promptly while preserving the ability to investigate it. See CISA incident response resources and Microsoft Entra ID identity risk guidance.

QUESTION NO: 17

Which of the following technique involves scanning the headers of IP packets leaving a network to make sure that the unauthorized or malicious traffic never leaves the internal network?

A. Egress Filtering

B. Throttling

C. Rate Limiting

D. Ingress Filtering

ANSWER: A

Explanation:

Egress Filtering is correct because it examines and controls outbound network traffic as it leaves an internal network. Security devices such as perimeter firewalls, routers, and intrusion-prevention systems can inspect packet-header fields—including source and destination IP addresses, ports, protocols, and TCP flags—and apply rules that prevent prohibited, spoofed, or malicious traffic from exiting. This helps contain compromised hosts by blocking communications with unauthorized external systems, command-and-control infrastructure, or destinations not permitted by organizational policy. It

also supports anti-spoofing controls: an organization can reject outbound packets whose source addresses do not belong to its assigned internal address ranges. The term “egress” specifically denotes traffic moving out of a network boundary, so filtering packets leaving the internal environment directly matches the activity in the question. RFC 2827 describes network ingress filtering as a method for limiting source-address spoofing and notes that filtering should also be applied at appropriate network boundaries; the same anti-spoofing principle is commonly implemented for outbound traffic as egress filtering. NIST guidance likewise identifies egress filtering as a firewall capability for restricting outbound communications according to policy. See RFC 2827: Network Ingress Filtering and NIST SP 800-41 Rev. 1: Guidelines on Firewalls and Firewall Policy.

QUESTION NO: 18

Harley is working as a SOC analyst with Powell Tech. Powell Inc. is using Internet Information Service (IIS) version 7.0 to host their website.

Where will Harley find the web server logs, if he wants to investigate them for any anomalies?

A. %SystemDrive%\inetpub\logs\LogFiles\W3SVCN

B. SystemDrive%\LogFiles\inetpub\logs\W3SVCN

C. %SystemDrive%\LogFiles\logs\W3SVCN

D. SystemDrive%\ inetpub\LogFiles\logs\W3SVCN

ANSWER: A

Explanation:

The default IIS web-server log location is `%SystemDrive%\inetpub\logs\LogFiles\W3SVCN`, where `%SystemDrive%` normally resolves to `C:` and `N` represents the numeric IIS site identifier. For example, logs for the Default Web Site are commonly stored in a directory such as `C:\inetpub\logs\LogFiles\W3SVC1`. IIS creates a separate `W3SVC<site ID>` folder for each website so that HTTP request activity can be investigated on a per-site basis. These log files commonly contain valuable SOC investigation evidence, including client IP addresses, request methods and URIs, HTTP status codes, user agents, referrers, and timestamps. Harley can use the site-specific folder to identify anomalous requests, scanning activity, suspicious status-code patterns, or unusual access to web resources. The log directory can be changed by an administrator at the server or individual-site level, but this path is the standard default for IIS logging. Microsoft documents the default log-file directory as `%SystemDrive%\inetpub\logs\LogFiles`; the site identifier is appended beneath that location. See Microsoft IIS logFile configuration documentation and Microsoft guidance for configuring IIS logging.

QUESTION NO: 19

Charline is working as an L2 SOC Analyst. One day, an L1 SOC Analyst escalated an incident to her for furtherinvestigation and confirmation. Charline, after a thorough investigation, confirmed the incident and assigned it with an initial priority.

What would be her next action according to the SOC workflow?

A. She should immediately escalate this issue to the management

B. She should immediately contact the network administrator to solve the problem

C. She should communicate this incident to the media immediately

D. She should formally raise a ticket and forward it to the IRT

ANSWER: D

Explanation:

She should formally raise a ticket and forward it to the IRT. Once an L2 analyst has validated that an alert is a genuine incident and has assigned an initial priority, the SOC workflow requires the event to be formally documented and routed to

the Incident Response Team (IRT) for coordinated handling. The ticket creates the authoritative incident record, preserving investigation findings, affected assets, indicators, timestamps, severity or priority, and the actions already performed. This documentation enables accountable ownership, consistent communications, auditability, and appropriate tracking through containment, eradication, recovery, and closure. Forwarding the documented incident to the IRT also ensures that the personnel with incident-response authority can begin or coordinate the required response actions according to the organization's incident-response plan. Initial prioritization is particularly important because it gives the receiving response team the context needed to allocate resources and act within the applicable service-level objectives. This flow aligns with NIST guidance that incident handling should include recording and reporting incidents as part of a managed response process. See NIST SP 800-61 Rev. 2, Computer Security Incident Handling Guide and CISA Incident Response guidance.

QUESTION NO: 20

Secuzin Corp. is a large enterprise performing millions of financial transactions daily, making it critical to analyze security logs efficiently, detect suspicious activities, and respond to incidents in real time. Its SOC is responsible for managing security logs from various network devices, including firewalls, intrusion detection systems (IDS), authentication servers, and cloud services. To fulfill compliance and regulatory requirements that mandate long-term archival of logs, you need to provide a log storage solution that is scalable to handle increasing log volumes, provides encryption for data security, and is seamlessly accessible. Which storage solution should you choose to meet these long-term log storage requirements?

A. Distributed storage system

B. Hybrid storage system

C. Local storage

D. Cloud storage

ANSWER: D

Explanation:

Cloud storage is the appropriate solution because it is designed to provide elastic capacity for data whose volume grows over time, such as security-event and audit logs generated by enterprise infrastructure. A SOC can retain large log datasets without procuring and maintaining additional physical storage whenever transaction volume, telemetry sources, or retention requirements increase. Major cloud object-storage services provide high durability and availability, encryption in transit and at rest, identity-based access controls, and centralized auditing. These capabilities help protect the confidentiality and integrity of regulated log data while allowing authorized SOC analysts and auditors to retrieve it when needed.

Cloud storage also supports lifecycle policies that automatically transition logs from frequently accessed tiers to lower-cost archival tiers as they age. This allows an organization to balance immediate investigative access with multi-year compliance retention economically. Retention controls and immutable or write-once-read-many configurations can further preserve log records against alteration or deletion, which is important when records may serve as audit evidence or support incident investigations. For example, Amazon S3 documents lifecycle management and retention controls through S3 Object Lock, while Microsoft documents immutable storage for Azure Blob Storage. See Amazon S3 Object Lock documentation and Microsoft Azure immutable blob storage documentation.

QUESTION NO: 21

Which of the following command is used to enable logging in iptables?

A. $ iptables -B INPUT -j LOG

B. $ iptables -A OUTPUT -j LOG

C. $ iptables -A INPUT -j LOG

D. $ iptables -B OUTPUT -j LOG

ANSWER: B

Explanation:

`iptables -A OUTPUT -j LOG` enables logging for packets that traverse the `OUTPUT` chain by appending a rule whose target is `LOG`. In iptables syntax, `-A` means “append”: the rule is added to the specified chain. The `OUTPUT` chain handles locally generated traffic leaving the host, and the `LOG` target sends a record of each matching packet to the kernel logging facility. Depending on the system’s logging configuration, those records can then be viewed through the system journal or written to log files such as the kernel log. This type of rule is useful to a SOC analyst when establishing visibility into outbound connections, investigating suspicious egress activity, or validating firewall policy behavior. The LOG target records packet metadata and then processing continues to subsequent rules, so logging is commonly placed with appropriate match conditions and rate limits in production environments to avoid excessive log volume. The iptables manual documents both the append operation and chain processing, while the Netfilter documentation describes the LOG target and its logging behavior. See the iptables manual page and iptables extensions documentation.

QUESTION NO: 22

If the SIEM generates the following four alerts at the same time:

I. Firewall blocking traffic from getting into the network alerts

II. SQL injection attempt alerts

III. Data deletion attempt alertsIV. Brute-force attempt alerts

Which alert should be given least priority as per effective alert triaging?

A. III

B. IV

C. II

D. I

E. Firewall blocking traffic from getting into the network alerts
II. SQL injection attempt alerts
III. Data deletion attempt alertsIV. Brute-force attempt alerts
Which alert should be given least priority as per effective alert triaging?

ANSWER: D

Explanation:

Firewall blocking traffic from getting into the network alerts should receive the least priority in this scenario because the security control has already prevented the suspicious traffic from entering the environment. Effective alert triage considers both the potential impact of an event and whether the threat is ongoing or has been successfully contained. A firewall block is evidence that a preventive control operated as intended; although the event should be retained for correlation, trend analysis, and possible tuning, it generally requires less immediate analyst action than alerts that may indicate active attempts to compromise an application, gain unauthorized access, or destroy data.

SOC analysts should still validate that the block was made by the expected policy, identify repeated source activity, and correlate it with other events. Repeated blocked traffic can reveal reconnaissance or a broader campaign, but the individual alert has a lower immediate urgency when there is no indication that traffic bypassed the control. This prioritization approach helps analysts focus limited response capacity on alerts with a greater likelihood of active compromise or business impact. Guidance on alert prioritization and incident analysis is available from CISA and NIST's Cybersecurity Framework.

QUESTION NO: 23

Which of the following tool is used to recover from web application incident?

A. CrowdStrike FalconTM Orchestrator

B. Symantec Secure Web Gateway

C. Smoothwall SWG

D. Proxy Workbench

ANSWER: A

Explanation:

CrowdStrike FalconTM Orchestrator is the best answer because security orchestration is designed to coordinate and automate incident-response actions, including containment, remediation, and recovery workflows. Following a web-application security incident, an orchestration capability can run predefined playbooks that collect alerts and context, create and assign cases, isolate or contain affected resources where appropriate, initiate remediation tasks, and document recovery actions. These coordinated workflows reduce manual response time and help analysts execute a consistent recovery process after an incident has been identified and investigated.

In CrowdStrike's Falcon platform, orchestration and workflow automation support response processes by connecting security operations activities and automating repetitive actions. This makes an orchestration tool appropriate for the recovery stage of incident handling, where the organization needs to restore normal operations safely and verify that response actions have been completed. NIST's incident-handling guidance likewise treats recovery as a formal phase involving restoration and validation after containment and eradication. See CrowdStrike Falcon Fusion SOAR and NIST SP 800-61 Rev. 2, Computer Security Incident Handling Guide.

QUESTION NO: 24

A SOC analyst monitoring authentication logs detects a sudden and significant spike in failed login attempts targeting multiple critical servers during non-business hours. These repeated authentication failures are abnormal compared to typical login activity. All attempts originate from a single external IP address, indicating a targeted attack rather than random scanning. Some login attempts use legitimate employee usernames, suggesting credential stuffing using previously compromised credentials or an ongoing brute-force attempt. Given this suspicious activity and its potential to escalate into unauthorized access, what is the appropriate next step in the threat-hunting process to assess the situation further?

A. Rapid response

B. Continuous improvement

C. Establish a baseline

D. Investigate and analyze

ANSWER: D

Explanation:

Investigate and analyze is the appropriate next step because the SOC analyst has already identified an anomalous authentication pattern and formed a credible hypothesis involving credential stuffing or brute-force activity. The priority is now to validate the hypothesis, determine the activity's full scope, and establish whether it resulted in successful access or follow-on malicious behavior. This includes pivoting across authentication, endpoint, and network telemetry to identify successful logons from the source IP address; determine every targeted account, server, and authentication protocol; and assess whether successful sessions were followed by privilege changes, unusual token use, remote-service activity, or suspicious processes. The analyst should also enrich the IP address with threat-intelligence and geolocation context and check whether the affected usernames appear in prior credential-exposure records. This evidence-based analysis establishes the severity and impact of the event and provides the facts needed to make proportionate containment decisions, such as blocking the source, resetting affected credentials, or enforcing stronger authentication controls. NIST's incident-handling guidance describes analysis as the activity used to validate and understand detected events before selecting response actions. MITRE ATT&CK likewise documents password spraying and credential stuffing-related authentication abuse as techniques requiring correlation of failed and successful authentication evidence. See NIST SP 800-61 Rev. 2 and MITRE ATT&CK: Brute Force.

QUESTION NO: 25

A healthcare organization's SIEM detects unusual HTTP requests targeting its patient portal. The requests originate from a foreign IP address and occur during non-business hours. The methods used are primarily TRACE and OPTIONS, which are rarely seen in normal web traffic. The SIEM correlates these with increased reconnaissance activity on other servers within the same subnet. What is the primary security concern with TRACE and OPTIONS requests?

A. They expose information about server-supported methods and request headers

B. They can be used to upload malicious payloads directly to the server

C. They make Distributed Denial of Service (DDoS) attacks easier

D. They allow attackers to bypass authentication controls

ANSWER: A

Explanation:

They expose information about server-supported methods and request headers. This is the primary concern because OPTIONS is designed to communicate the HTTP methods and other communication options a target resource supports. An attacker can use this response during reconnaissance to identify the application's attack surface, including whether potentially sensitive or unnecessary methods have been enabled. TRACE instructs the server to return the received request as seen by the server, which can disclose the handling of request headers and intermediary behavior. Although modern browser and server protections mitigate many historical cross-site tracing scenarios, externally initiated TRACE requests remain a useful signal of probing and configuration discovery in SOC monitoring. In this scenario, the unusual timing, foreign source, rare HTTP methods, and correlated subnet reconnaissance make server and application mapping the most relevant interpretation. Security teams should validate whether TRACE is required, disable it where there is no business need, limit allowed methods to those the portal actually uses, and investigate the source activity and related requests. HTTP defines OPTIONS and TRACE semantics in RFC 9110, Section 9.2 and Section 9.3.

QUESTION NO: 26

Which threat intelligence standard is designed to represent and exchange cyber threat intelligence in a structured format?

A. STIX

B. Syslog

C. NetFlow

D. CVSS

ANSWER: A

Explanation:

STIX is correct. Structured Threat Information Expression (STIX) is a standardized language and data model for representing cyber threat intelligence. It can describe intelligence objects such as indicators, malware, threat actors, attack patterns, relationships, observed data, campaigns, and courses of action in a consistent format.

STIX enables threat-intelligence producers and consumers to share information in a form that can be understood by analysts and processed by security platforms. STIX data is commonly transported using TAXII, which is a separate protocol designed to exchange CTI collections and objects. The OASIS specification is available at STIX Version 2.1.