DUMPSARENA

Understanding Cisco Cybersecurity Operations Fundamentals (CBROPS)

Cisco 200-201

Version Demo

Total Demo Questions: 55

Total Premium Questions: 553

Buy Premium PDF

https://dumpsarena.co

sales@dumpsarena.co

sales@dumpsarena.co
dumpsarena.co

Topic Break Down

Topic	No. of Questions
Topic 1, Security Concepts	203
Topic 2, Security Monitoring	98
Topic 3, Host-Based Analysis	78
Topic 4, Network Intrusion Analysis	130
Topic 5, Security Policies and Procedures	43
Topic 6, Mix Questions	1
Total	553

QUESTION NO: 1

An engineer must compare NIST vs ISO frameworks The engineer deeded to compare as readable documentation and also to watch a comparison video review. Using Windows 10 OS. the engineer started a browser and searched for a NIST document and then opened a new tab in the same browser and searched for an ISO document for comparison

The engineer tried to watch the video, but there 'was an audio problem with OS so the engineer had to troubleshoot it At first the engineer started CMD and looked fee a driver path then locked for a corresponding registry in the registry editor The engineer enabled "Audiosrv" in task manager and put it on auto start and the problem was solved Which two components of the OS did the engineer touch? (Choose two)

A. permissions

B. PowerShell logs

C. service

D. MBR

E. process and thread

ANSWER: C E

Explanation:

The correct components are **service** and **process and thread**. **Audiosrv** is the service name for Windows Audio. Windows services are long-running operating-system components that can start automatically and provide functionality to applications and users without requiring a user to launch them manually. Configuring Windows Audio to start automatically directly involves service management and ensures audio capabilities are available after the computer starts.

The troubleshooting also involves **process and thread** management because Task Manager exposes the running activity associated with applications, background processes, and services. A Windows service executes through one or more processes and threads; inspecting or enabling the service through Task Manager therefore interacts with this OS execution-management component. The browser's ability to play the comparison video depends on the active Windows Audio service and its associated running execution context. Microsoft identifies Windows Audio as the service that manages audio for Windows-based programs, and its service configuration includes an automatic startup type. See Microsoft documentation on Windows services and Microsoft documentation on the Windows Audio service.

QUESTION NO: 2

Which type of data is used to detect anomalies in the network?

A. statistical data

B. alert data

C. transaction data

D. metadata

ANSWER: A

Explanation:

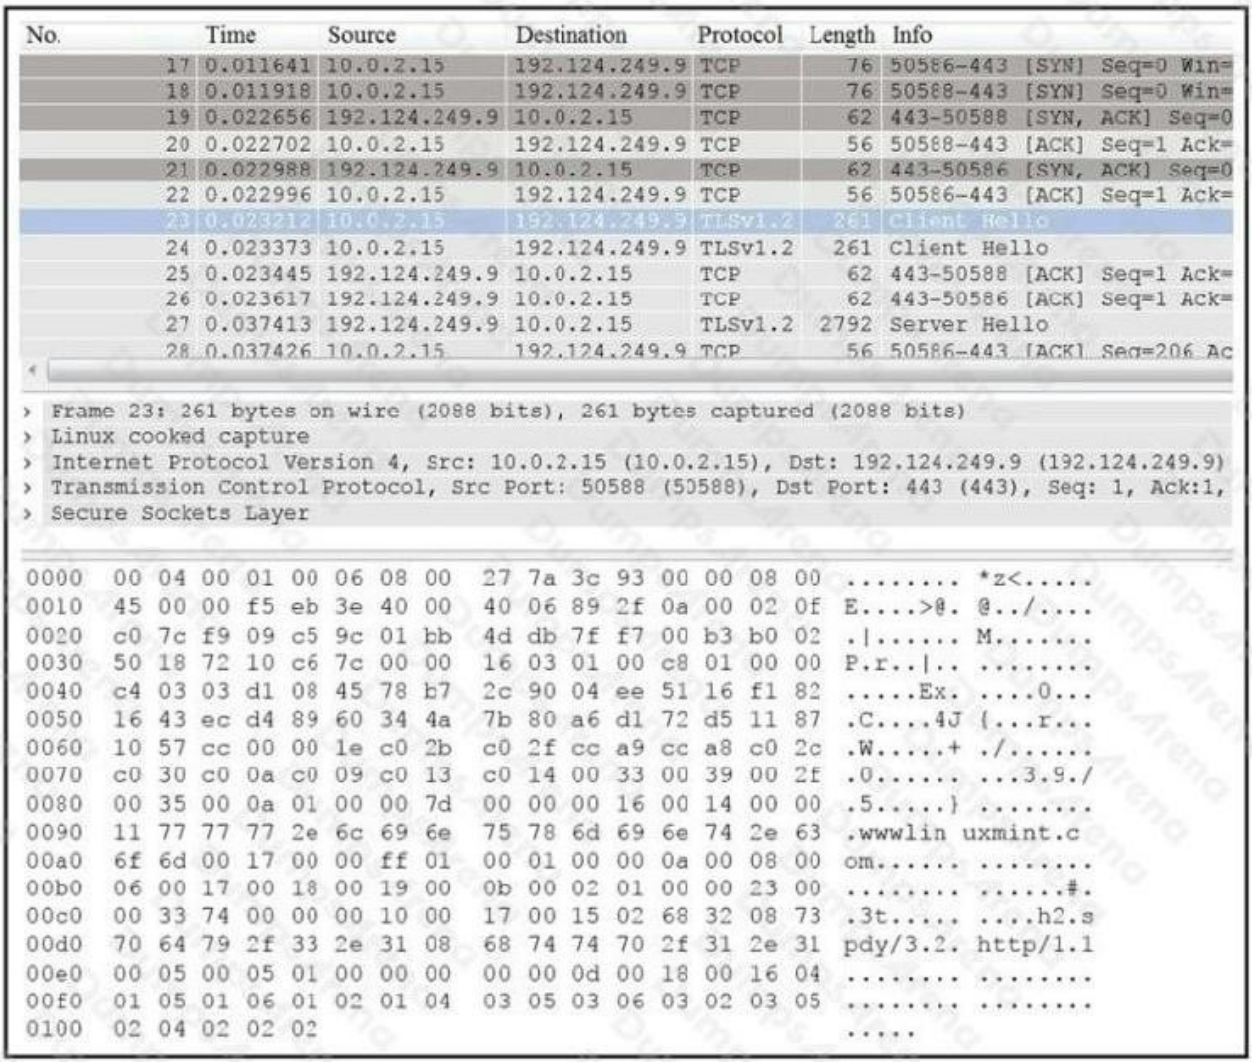
statistical data is used for network anomaly detection because it establishes a measurable profile, or baseline, of expected network behavior over time. Security operations tools collect and analyze values such as normal traffic volume, connection rates, protocol use, packet sizes, bandwidth consumption, login frequency, and communication patterns between hosts. From these observations, the tools can calculate distributions, averages, ranges, and other thresholds that represent ordinary activity for a particular network, device, user, or time period.

When newly observed activity differs significantly from that established statistical baseline, it can be identified as anomalous and investigated. For example, an unusual increase in outbound traffic, a sudden change in the number of connections from a host, or network activity at an atypical time can indicate a potential security event. This approach is especially valuable for

detecting previously unknown threats because it focuses on deviations in behavior rather than relying solely on a known malicious signature. Cisco's CyberOps training covers the use of network telemetry and analysis to recognize suspicious activity, while NIST describes anomaly detection as identifying observations that depart from expected behavior. See Cisco CyberOps Associate training and NIST: Anomaly Detection.

QUESTION NO: 3 - (DRAG DROP)

DRAG DROP



Refer to the exhibit. Drag and drop the element name from the left onto the correct piece of the PCAP file on the right.

Select and Place:

source address	10.0.2.15
destination address	50588
source port	443
destination port	192.124.249.9
Network Protocol	Transmission Control Protocol
Transport Protocol	Internet Protocol v4
Application Protocol	Transport Layer Security v1.2

ANSWER:

source address	source address
destination address	source port
source port	destination port
destination port	destination address
Network Protocol	Transport Protocol
Transport Protocol	Network Protocol
Application Protocol	Application Protocol

Explanation:

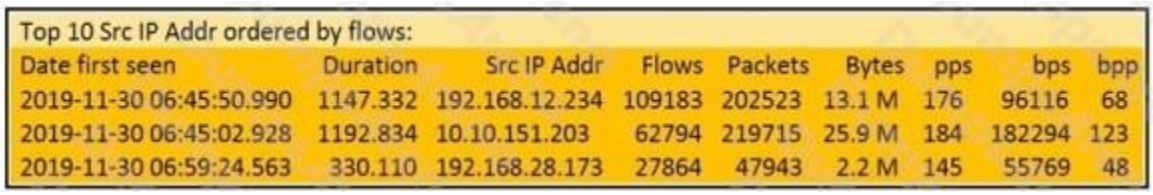
A packet capture presents communications in layers, and each displayed value belongs to a specific header or protocol layer. The IPv4 header supplies the endpoint addresses: **10.0.2.15** is the source address and **192.124.249.9** is the destination address. Accordingly, the correct network protocol identification is **Internet Protocol v4**. Cisco's overview of IP explains that IP provides addressing for packets moving between network endpoints: Cisco: What Is Internet Protocol?.

Inside the IP payload, the TCP header identifies the communication endpoints at the transport layer. The client-side ephemeral source port is **50588**, while the service destination port is **443**. The transport protocol is therefore **Transmission Control Protocol**. TCP is the connection-oriented transport protocol that uses source and destination port fields to identify the sending process and receiving service; Cisco describes TCP/IP protocol responsibilities in its TCP/IP overview.

Finally, the protected application session is identified as **Transport Layer Security v1.2**. TLS operates above TCP and protects application data through encryption and authentication. The protocol's versioned specification is published by the IETF in RFC 5246. Thus, the complete correct relationship follows the packet's layered structure: IPv4 supplies addresses, TCP supplies ports, and TLS identifies the secured application protocol.

QUESTION NO: 4

Refer to the exhibit.



Top 10 Src IP Addr ordered by flows:

Date first seen	Duration	Src IP Addr	Flows	Packets	Bytes	pps	bps	bpp
2019-11-30 06:45:50.990	1147.332	192.168.12.234	109183	202523	13.1 M	176	96116	68
2019-11-30 06:45:02.928	1192.834	10.10.151.203	62794	219715	25.9 M	184	182294	123
2019-11-30 06:59:24.563	330.110	192.168.28.173	27864	47943	2.2 M	145	55769	48

What information is depicted?

A. IIS data

B. NetFlow data

C. network discovery event

D. IPS event data

ANSWER: B

Explanation:

NetFlow data is depicted. NetFlow is flow telemetry generated by network devices to summarize communications observed on the network. Rather than retaining every packet's full contents, a NetFlow record identifies a conversation using characteristics such as source and destination IP addresses, source and destination ports, protocol, interfaces, timestamps, packet counts, and byte counts. This makes NetFlow especially useful to a security operations team for establishing normal traffic baselines, investigating connections between hosts, identifying unusual service use, and understanding the volume and direction of network activity. In a monitoring or analytics interface, NetFlow-derived information is commonly presented as traffic flows, conversations, top talkers, protocol or port summaries, and source/destination relationships. Cisco describes NetFlow as a technology that collects IP traffic information as it enters or exits an interface, enabling network traffic monitoring and analysis. Cisco's flow-monitoring documentation also explains that Flexible NetFlow gathers and exports flow records containing selected traffic attributes and counters. See Cisco's NetFlow overview and Flexible NetFlow configuration guide.

QUESTION NO: 5

Which security technology allows only a set of pre-approved applications to run on a system?

A. application-level blacklisting

B. host-based IPS

C. application-level whitelisting

D. antivirus

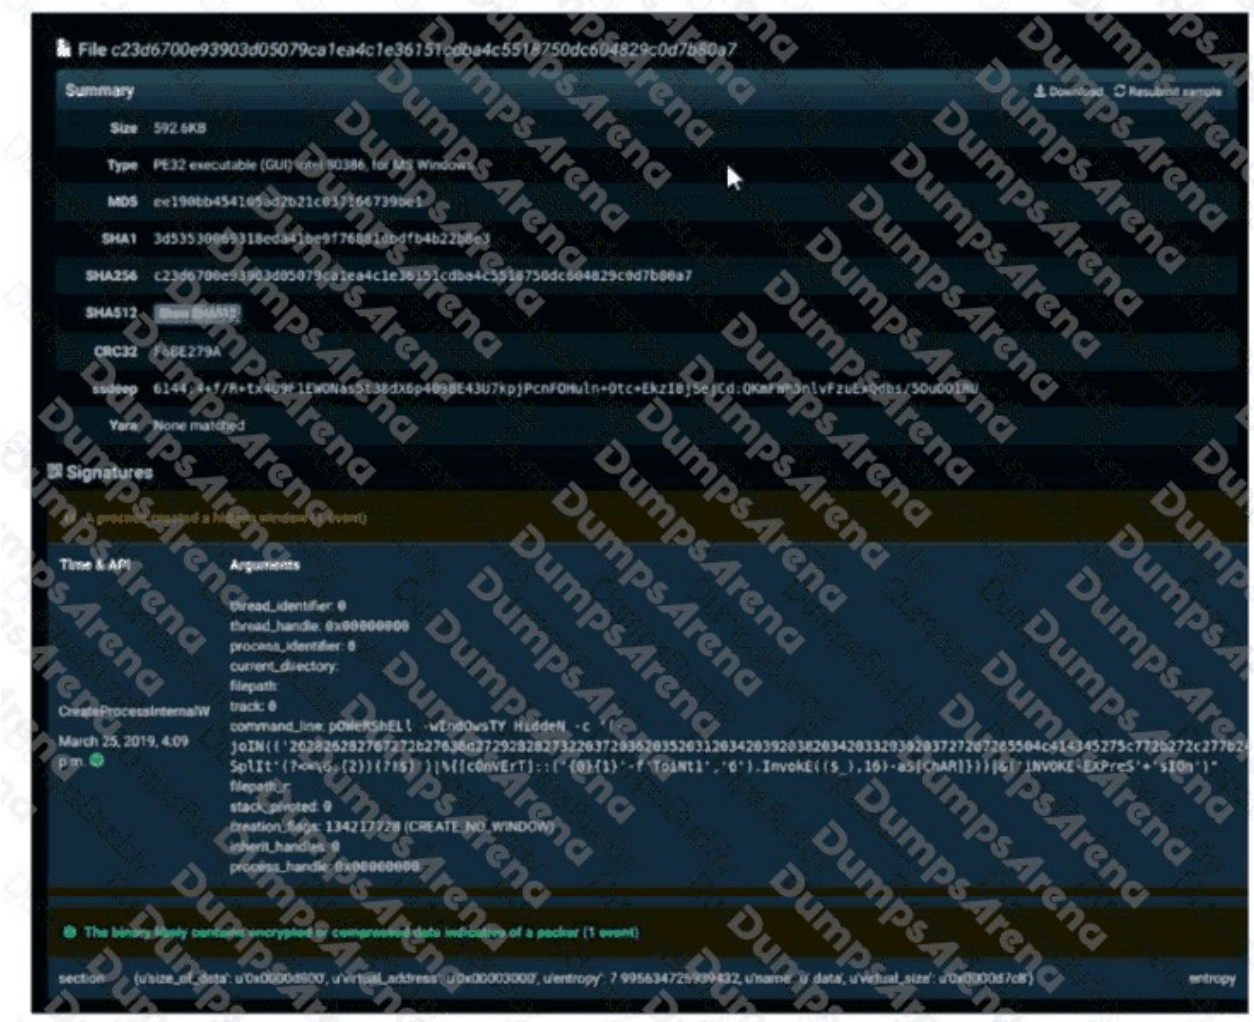
ANSWER: C

Explanation:

Application-level whitelisting is correct because it enforces an allow-by-exception execution policy: software is permitted to run only when it appears on an explicitly approved list. The approval criteria can include an application's filename, path, publisher certificate, cryptographic hash, or a managed policy rule. This approach reduces the attack surface on endpoints by preventing unauthorized executables, including many forms of malware and unapproved tools, from launching even if they reach the device. In an enterprise environment, the approved set is normally defined and maintained centrally, allowing administrators to ensure that users and systems run only known, authorized software required for their roles. Cisco's endpoint-security guidance identifies application whitelisting as a control that permits only authorized applications to execute. The National Institute of Standards and Technology also describes application allowlisting as a technique for restricting

execution to authorized software. See Cisco endpoint security resources and NIST SP 800-167, Guide to Application Whitelisting.

QUESTION NO: 6



Refer to the exhibit. A SOC engineer is analyzing Cuckoo Sandbox report for a file that has been identified as suspicious by the endpoint security system. What is the state of the file?

A. The file was identified as PE32 executable with a high level of entropy to bypass AV via encryption.

B. The file was detected as an executable binary file, but no suspicious activity was detected and it is false positive.

C. The file was detected as executable and was marked by the SSDeep hashing algorithm as suspicious.

D. The file identified as an executable binary for Microsoft Word with macros creating hidden process via PowerShell.

ANSWER: A

Explanation:

The file was identified as PE32 executable with a high level of entropy to bypass AV via encryption. Cuckoo Sandbox performs static analysis before and alongside behavioral execution, including identifying the submitted file's format and calculating entropy values. A PE32 classification identifies a 32-bit Windows Portable Executable, which is a native executable format commonly used by Windows malware as well as legitimate software. In malware triage, an unusually high entropy measurement is a meaningful indicator because encrypted, compressed, or packed content has a more random byte distribution than ordinary executable sections. Attackers commonly use packing or encryption to conceal code and strings from signature-based antivirus inspection, then unpack or decrypt the payload during execution. Therefore, the reported combination of a PE32 executable classification and high entropy supports treating the sample as suspicious and potentially

protected through encryption or packing rather than as a plainly readable binary. The SOC engineer should preserve the file hash and sandbox artifacts, then correlate the static result with the behavioral processes, network activity, and indicators recorded in the report. Cuckoo documents the static metadata and analysis results it produces in its analysis results documentation. Cisco also describes sandboxing as a way to analyze suspicious files and observe malicious behavior in its Advanced Malware Protection overview.

QUESTION NO: 7

An engineer discovered a breach, identified the threat's entry point, and removed access. The engineer was able to identify the host, the IP address of the threat actor, and the application the threat actor targeted. What is the next step the engineer should take according to the NIST SP 800-61 Incident handling guide?

A. Recover from the threat.

B. Analyze the threat.

C. Identify lessons learned from the threat.

D. Reduce the probability of similar threats.

ANSWER: A

Explanation:

Recover from the threat. is the appropriate next step. The engineer has already performed core detection and analysis activities by confirming the breach, identifying the affected host and targeted application, and determining the threat actor's source address. Removing the actor's access is a containment action that stops or limits ongoing unauthorized activity. Once the incident is contained and the immediate threat has been removed, the incident-handling process proceeds to recovery.

In NIST SP 800-61 Revision 2, recovery is part of the containment, eradication, and recovery phase. It focuses on safely restoring affected systems and services to normal operation, validating that the environment is operating correctly, and monitoring for signs that the threat may recur. Recovery can include restoring systems from known-good backups, correcting compromised configurations, applying needed patches, rotating affected credentials, and increasing monitoring of the affected host and application. The organization should return services in a controlled manner and confirm that the remediation remains effective before declaring normal operations fully restored. NIST's incident-handling lifecycle is described in NIST SP 800-61 Revision 2, particularly its containment, eradication, and recovery guidance. NIST also provides current incident-response resources through its Cybersecurity Framework.

QUESTION NO: 8

Why is encryption challenging to security monitoring?

A. Encryption analysis is used by attackers to monitor VPN tunnels.

B. Encryption is used by threat actors as a method of evasion and obfuscation.

C. Encryption introduces additional processing requirements by the CPU.

D. Encryption introduces larger packet sizes to analyze and store.

ANSWER: B

Explanation:

Encryption is used by threat actors as a method of evasion and obfuscation. Security monitoring traditionally relies heavily on inspecting packet payloads, URLs, commands, file transfers, and application-layer indicators. When an attacker uses encrypted channels such as TLS, SSH, or a VPN tunnel, those contents are protected from passive network inspection. This can conceal command-and-control traffic, malware downloads, credential theft, and data exfiltration within traffic that may resemble legitimate encrypted business activity.

Encryption itself remains essential for protecting confidentiality and integrity; the monitoring challenge is the loss of direct payload visibility when decryption is unavailable, inappropriate, or operationally impractical. Defenders therefore supplement inspection with metadata and behavioral analysis, including flow records, certificate details, TLS fingerprints, traffic timing, packet sizes, destination reputation, and unusual communication patterns. Cisco's Encrypted Traffic Analytics approach is based on deriving security-relevant intelligence from encrypted traffic without decrypting the payload, illustrating why adversaries can use encryption as a concealment mechanism while defenders must apply additional detection methods. See Cisco's Encrypted Traffic Analytics white paper and Cisco Secure Network Analytics.

QUESTION NO: 9

What is rule-based detection when compared to statistical detection?

A. proof of a user's identity

B. proof of a user's action

C. likelihood of user's action

D. falsification of a user's identity

ANSWER: B

Explanation:

proof of a user's action is correct because rule-based detection applies defined, deterministic conditions to observed events. A security analyst or detection platform creates rules that specify the evidence required for an alert, such as a particular protocol sequence, a known malicious indicator, an unauthorized command, or a disallowed authentication event. When telemetry satisfies the rule's conditions, the system can assert that the defined action or pattern occurred in the collected data. This makes rule-based detection especially useful for identifying known behaviors and enforcing clearly stated security policies, because the alert can be traced directly to the matched rule and its supporting event records.

In Cisco security operations, intrusion and event-detection rules are built around identifiable traffic characteristics and conditions. The rule match supplies concrete evidence that the monitored activity met the configured detection criteria, enabling analysts to validate and investigate the action using logs, packet data, and other telemetry. Cisco documents how intrusion rules inspect traffic and generate events when their specified conditions are met: Cisco Secure Firewall Management Center—Intrusion Rules. Cisco's security monitoring guidance also emphasizes collecting and correlating event evidence for incident analysis: Cisco Secure Network Analytics.

QUESTION NO: 10

Which security technology guarantees the integrity and authenticity of all messages transferred to and from a web application?

A. Hypertext Transfer Protocol

B. SSL Certificate

C. Tunneling

D. VPN

E. Transport Layer Security (TLS)

ANSWER: E

Explanation:

Transport Layer Security (TLS) is the technology that protects communications between a web application and its clients. When a web application is accessed through HTTPS, TLS establishes a protected session before application data is exchanged. During the TLS handshake, the server proves its identity using a certificate and the associated private key; this enables the client to authenticate the intended server, subject to proper certificate validation. TLS then applies authenticated

encryption to records carrying application data. Authenticated encryption both encrypts the data and detects any modification, insertion, deletion, or replay-related violation of the protected record stream, providing integrity for messages exchanged within the session. TLS can also authenticate the client when mutual TLS is configured, using a client certificate. A certificate is an important component of server identity validation, but TLS is the protocol that negotiates cryptographic keys and applies integrity protection to the transmitted messages. RFC 8446 specifies that TLS is designed to provide a secure channel with authentication, confidentiality, and integrity protection; it also defines the authenticated-encryption record protection used by TLS 1.3. See the TLS 1.3 specification (RFC 8446) and Cisco's overview of SSL/TLS.

QUESTION NO: 11

An engineer needs to have visibility on TCP bandwidth usage, response time, and latency, combined with deep packet inspection to identify unknown software by its network traffic flow. Which two features of Cisco Application Visibility and Control should the engineer use to accomplish this goal? (Choose two.)

A. management and reporting

B. traffic filtering

C. adaptive AVC

D. metrics collection and exporting

E. application recognition

ANSWER: D E

Explanation:

metrics collection and exporting is required to obtain the operational data in the scenario. Cisco AVC uses performance and flow telemetry to collect application and TCP-flow statistics, including bandwidth utilization and performance measures such as response time and network latency. Exporting that telemetry enables a monitoring or analytics platform to retain, correlate, and report on the measurements across devices and flows. This makes the feature central to gaining the requested visibility rather than merely enforcing a policy.

application recognition supplies the deep packet inspection capability. Cisco’s Network-Based Application Recognition technology identifies and classifies applications from traffic characteristics and packet contents, even when port-based identification is insufficient. The resulting application identity can then be associated with the collected flow and performance data, allowing the engineer to identify the software responsible for observed bandwidth consumption or latency. Together, application recognition establishes what the traffic is, while metrics collection and exporting provides the measurable TCP usage and performance evidence needed for analysis. Cisco describes AVC as combining application awareness with monitoring and control capabilities, and documents NBAR as the application-classification foundation: Cisco NBAR2 Overview and Cisco Application Visibility and Control.

QUESTION NO: 12

Which two fields are typically included in a NetFlow record? (Choose two.)

A. packet count

B. byte count

C. file hash

D. username

E. email message subject

ANSWER: A B

Explanation:

packet count and **byte count** are included in NetFlow records to describe the volume of traffic observed for a flow. Flow telemetry summarizes communications rather than retaining each packet payload. Analysts use packet and byte counts with timestamps, addresses, ports, and protocol values to identify high-volume transfers, establish traffic baselines, and investigate anomalous activity. A username and file hash are not standard NetFlow flow-record fields. See Cisco NetFlow and RFC 3954: Cisco NetFlow Services Export Version 9.

QUESTION NO: 13

According to the NIST SP 800-86. which two types of data are considered volatile? (Choose two.)

A. swap files

B. temporary files

C. login sessions

D. dump files

E. free space

F. active network connections

ANSWER: C F

Explanation:

NIST SP 800-86 characterizes volatile data as information that is likely to be lost when a system is powered off or that exists only while the operating system is actively running. **login sessions** are volatile because their state is maintained by the active operating system and associated processes; that information generally disappears when the host is shut down or restarted. **active network connections** are also volatile because connection state, such as current TCP sessions and their associated in-memory data structures, exists only while the system and network stack are operating. During incident response and forensic collection, this evidence should be acquired early—before shutdown, reboot, or other actions that could alter memory-resident state. NIST specifically identifies volatile operating-system data as including items such as network connections, logged-on users, running processes, and open files. This supports collecting login-session and active-connection information as part of live-response evidence gathering. See NIST SP 800-86: Guide to Integrating Forensic Techniques into Incident Response and NIST's Cybersecurity Framework resources for incident-response and evidence-handling context.

QUESTION NO: 14

Which difficulty occurs when log messages are compared from two devices separated by a Layer 3 device that performs Network Address Translation?

A. IP addresses in the log messages match

B. Timestamps of the log messages are different.

C. Log messages contain incorrect information

D. IP addresses in the log messages do not match

ANSWER: D

Explanation:

IP addresses in the log messages do not match because Network Address Translation changes packet address information as traffic crosses the translating Layer 3 device. A device on the inside of the NAT boundary can log the original, private source address, while a device on the outside logs the translated public address for the same connection. For example, an internal endpoint may generate a log entry using its RFC 1918 address, but an external firewall, server, or sensor records the NAT-assigned address instead. This creates a correlation challenge for security operations: analysts must use NAT translation records, firewall connection logs, and matching details such as ports, protocols, and times to associate the events

with the same session. Cisco describes NAT as translating private addresses to public addresses for traffic leaving a network, which explains why observations on opposite sides of the translation boundary use different IP values. Maintaining accurate NAT logs and synchronized time sources supports reliable event investigation and incident response. See Cisco's Network Address Translation overview and the Cisco NAT command reference.

QUESTION NO: 15

How does an attacker observe network traffic exchanged between two users?

A. port scanning

B. man-in-the-middle

C. command injection

D. denial of service

ANSWER: B

Explanation:

Man-in-the-middle is correct because this attack places the attacker logically between two communicating parties, enabling the attacker to intercept and observe traffic as it passes between them. The attacker may achieve this position through techniques such as ARP spoofing on a local network, rogue wireless access points, or manipulation of name-resolution traffic. Once positioned in the communication path, the attacker can capture unencrypted data and, depending on the circumstances, may also alter or inject traffic. Encryption protocols such as TLS help protect the confidentiality and integrity of application data, but users and systems must still validate certificates and avoid untrusted networks or rogue access points to reduce exposure to interception. Cisco describes man-in-the-middle attacks as attacks in which an adversary intercepts communications between two parties, and NIST identifies interception as a threat to information confidentiality. See Cisco: Common Cyberattacks and NIST: Man-in-the-Middle Attack.

QUESTION NO: 16

What are the two differences between stateful and deep packet inspection? (Choose two)

A. Stateful inspection is capable of TCP state tracking, and deep packet filtering checks only TCP source and destination ports

B. Deep packet inspection is capable of malware blocking, and stateful inspection is not

C. Deep packet inspection operates on Layer 3 and 4. and stateful inspection operates on Layer 3 of the OSI model

D. Deep packet inspection is capable of TCP state monitoring only, and stateful inspection can inspect TCP and UDP.

E. Stateful inspection is capable of packet data inspections, and deep packet inspection is not

F. Stateful inspection tracks connection state, whereas deep packet inspection examines application-layer payload content.

ANSWER: B F

Explanation:

Deep packet inspection is capable of malware blocking, and stateful inspection is not, because DPI examines packet contents beyond basic header fields and connection state. This payload-aware analysis enables a next-generation firewall to identify applications, recognize threat signatures, and enforce security controls such as intrusion prevention and malware protection. Cisco describes next-generation firewall capabilities as including application visibility and control, advanced threat protection, and intrusion prevention, all of which depend on inspection beyond traditional stateful filtering.

Stateful inspection tracks the state and context of network connections, such as TCP session establishment and teardown, and makes allow/deny decisions primarily using protocol behavior and Layer 3/Layer 4 information. Deep packet inspection extends this approach by examining the application-layer payload within permitted traffic. Therefore, the key distinction is

that stateful inspection focuses on whether traffic belongs to a valid, established session, whereas deep packet inspection can evaluate what the traffic actually contains and apply content- or threat-based policy. This distinction is foundational to Cisco Secure Firewall's next-generation security services. See Cisco's firewall overview and Cisco Secure Firewall documentation.

QUESTION NO: 17

An organization is cooperating with several third-party companies. Data exchange is on an unsecured channel using port 80 Internal employees use the FTP service to upload and download sensitive data An engineer must ensure confidentiality while preserving the integrity of the communication. Which technology must the engineer implement in this scenario'?

A. X 509 certificates

B. RADIUS server

C. CA server

D. web application firewall

E. TLS with X.509 certificates (HTTPS and FTPS)

ANSWER: E

Explanation:

TLS with X.509 certificates, deployed as HTTPS for web traffic and FTPS for FTP traffic, is the appropriate solution because it protects data while it traverses an untrusted network. TLS establishes an authenticated secure session between communicating endpoints and uses symmetric encryption for the application data. This provides confidentiality, preventing an observer on the channel from reading sensitive transfers. TLS also applies authenticated integrity protection to transmitted records, allowing the recipient to detect modification, insertion, or replay attempts within the protected session.

X.509 certificates are an important part of this design: they bind a public key to a verified identity and allow a client or partner system to authenticate the server, and optionally the client, during the TLS handshake. For the port-80 web service, the organization should make the service available through HTTPS rather than cleartext HTTP. For employee file transfers, it should configure explicit or implicit FTPS so that FTP authentication and data channels are protected by TLS. The TLS protocol requirements are defined in RFC 8446, and the application of TLS to FTP is specified in RFC 4217.

QUESTION NO: 18

Which two elements are used by the defense-in-depth strategy? (Choose two)

A. packet segmentation

B. least privilege principle

C. single unified security solution

D. distributed database management system

E. firewalls

ANSWER: B E

Explanation:

Defense in depth applies multiple, complementary safeguards so that the failure or bypass of one control does not leave an organization without protection. **Firewalls** are a core technical layer in this approach because they enforce network-security policy by controlling traffic between network zones and can limit unauthorized connections before they reach internal assets. They provide an important boundary control alongside other preventive, detective, and response capabilities.

The **least privilege principle** is also fundamental to layered security. It requires that users, processes, and service accounts receive only the access and permissions necessary to perform their authorized tasks. Restricting privileges reduces the potential impact of compromised credentials, malware execution, and accidental administrative actions. When least privilege is combined with network controls such as firewalls, an attacker who overcomes one layer still faces additional restrictions on access and movement. NIST identifies least privilege as a core access-control principle in control AC-6, and Cisco describes firewalls as an essential capability for enforcing security policy and protecting network environments. See NIST SP 800-53 Rev. 5, AC-6 Least Privilege and Cisco: What Is a Firewall?.

QUESTION NO: 19

Which two components reduce the attack surface on an endpoint? (Choose two.)

A. secure boot

B. load balancing

C. increased audit log levels

D. restricting USB ports

E. full packet captures at the endpoint

ANSWER: A D

Explanation:

Secure boot and restricting USB ports reduce endpoint attack surface by limiting opportunities for unauthorized code execution and unwanted device interaction. Secure Boot is a UEFI security feature that verifies cryptographic signatures in the boot chain, allowing the device to load trusted boot software and helping protect against bootkits or other pre-operating-system malware. This establishes a trusted foundation before the operating system and endpoint protections start. Restricting USB ports reduces the physical-device attack surface by preventing or tightly controlling removable storage and peripheral devices. USB media can introduce malicious files, emulate keyboards or network adapters, or enable unauthorized data transfer; enforcing port controls or device allowlists limits those paths. Together, these controls address distinct endpoint exposure points: the system startup process and externally connected hardware. Cisco endpoint-security guidance emphasizes hardening hosts by reducing unnecessary access and controlling removable media, while UEFI defines Secure Boot as a mechanism for verifying authorized boot components. See Cisco's endpoint-security overview at Cisco Secure Endpoint and the UEFI specification's Secure Boot and Driver Signing section.

QUESTION NO: 20

Which two protocols are used for DDoS amplification attacks? (Choose two.)

A. ICMPv6

B. DNS

C. NTP

D. TCP

E. HTTP

ANSWER: B C

Explanation:

DNS and NTP are commonly abused in reflected amplification DDoS attacks because both can operate over UDP and can produce responses substantially larger than a request. An attacker sends a small request to publicly reachable servers while spoofing the victim's source IP address. The servers then send their much larger replies to the victim, combining traffic from many reflectors and consuming the victim's bandwidth or network-device resources. With DNS, requests designed to return large records or DNSSEC-related data can create a significant response-to-request size ratio. With NTP, older or

misconfigured servers that respond to the `monlist` query have historically enabled especially large replies. The attacker does not need to establish a session with the victim; the key conditions are source-address spoofing, UDP-based request handling, and services that return amplified responses. Cisco identifies DNS and NTP among the UDP services leveraged for reflection and amplification attacks, while CISA emphasizes reducing exposure by preventing IP source-address spoofing through network ingress filtering. See Cisco's DDoS mitigation guidance and the CISA NTP amplification alert.

QUESTION NO: 21

A security analyst received a ticket about suspicious traffic from one of the workstations. During the investigation, the analyst discovered that the workstation was communicating with an external IP. The analyst could not investigate further and escalated the case to a T2 security analyst. What are the two data visibility challenges that the security analyst should identify? (Choose two.)

A. A default user agent is present in the headers.

B. Traffic is not encrypted.

C. HTTP requests and responses are sent in plaintext.

D. POST requests have a " Microsoft-IIS/7.5 " server header.

E. Encrypted data is being transmitted.

ANSWER: C E

Explanation:

HTTP requests and responses are sent in plaintext. is a data-visibility concern because unencrypted HTTP exposes the full request and response contents to any party able to observe the traffic path. This can include URLs, form values, session identifiers, downloaded content, and data submitted in POST bodies. From a security-operations perspective, the analyst must recognize that plaintext application traffic can disclose sensitive information and may enable unauthorized observation or collection of data. Cisco identifies the protection of data in transit as an important security objective, including the use of secure protocols where appropriate.

Encrypted data is being transmitted. is also a visibility challenge because encryption protects confidentiality but limits direct inspection of the application payload by network-based monitoring tools. A security analyst may still use metadata such as source and destination addresses, ports, TLS handshake details, certificate information, flow volume, timing, and reputation data, but may need approved decryption, endpoint telemetry, or additional investigation by a higher-tier analyst to determine what the encrypted session contains. These two conditions illustrate the operational balance between protecting communications and maintaining sufficient telemetry for threat detection and investigation. See Cisco's guidance on encrypted-traffic visibility at Cisco Encrypted Traffic Analytics and Cisco's overview of cybersecurity operations at Cisco Security Operations.

QUESTION NO: 22

Which are two denial-of-service attacks? (Choose two.)

A. TCP connections

B. ping of death

C. man-in-the-middle

D. code-red

E. UDP flooding

ANSWER: B E

Explanation:

ping of death and **UDP flooding** are denial-of-service attacks because each targets the availability of a host, network service, or network path. A ping-of-death attack uses improperly formed or oversized ICMP echo traffic, historically exploiting weaknesses in how affected systems reassembled fragmented packets. Processing the malformed traffic could cause a vulnerable system to become unstable, crash, or stop responding, thereby denying legitimate users access to the service. Cisco identifies malformed-packet attacks, including Ping of Death, as a denial-of-service threat type.

A UDP flood sends a high volume of UDP datagrams to a target, often across many destination ports. The target and its network infrastructure must receive, inspect, and potentially generate responses for this traffic, consuming bandwidth, CPU, memory, and other packet-processing resources. Once those resources are exhausted, legitimate traffic can be delayed or dropped and services become unavailable. This makes UDP flooding a common volumetric or resource-exhaustion denial-of-service technique. These attacks are designed around disrupting availability, one of the core security objectives. See Cisco's overview of denial-of-service attack mitigation and CISA guidance on denial-of-service attacks.

QUESTION NO: 23

A SOC analyst detected connections to known C&C and port scanning activity to main HR database servers from one of the HR endpoints via Cisco StealthWatch. What are the two next steps of the SOC team according to the NISTSP800-61 incident handling process? (Choose two)

A. Isolate affected endpoints and take disk images for analysis

B. Provide security awareness training to HR managers and employees

C. Block connection to this C&C server on the perimeter next-generation firewall

D. Update antivirus signature databases on affected endpoints to block connections to C&C

E. Detect the attack vector and analyze C&C connections

ANSWER: A C

Explanation:

"Isolate affected endpoints and take disk images for analysis" and "Block connection to this C&C server on the perimeter next-generation firewall" are the appropriate immediate actions under the containment portion of the NIST incident-handling lifecycle. The observed command-and-control communication and internal port scanning indicate that the HR endpoint may be compromised and may be actively receiving attacker instructions or attempting lateral movement toward sensitive database systems. Isolating the affected endpoint limits additional scanning, data access, malware propagation, and attacker control. Obtaining a disk image preserves volatile investigative value in the endpoint's stored artifacts, such as malware, persistence mechanisms, logs, and configuration files, before remediation changes the system state.

Blocking the known command-and-control destination at the perimeter firewall is a complementary containment measure. It disrupts the adversary's external communication channel and helps protect other hosts if they attempt to contact the same infrastructure. NIST describes containment as selecting and implementing strategies that limit an incident's scope and damage while preserving needed evidence. These actions can proceed in parallel as part of a coordinated response, with the SOC documenting actions and maintaining appropriate evidence handling. See NIST's Computer Security Incident Handling Guide and its updated incident response guidance.

QUESTION NO: 24

Which incidence response step includes identifying all hosts affected by an attack?

A. detection and analysis

B. post-incident activity

C. preparation

D. containment, eradication, and recovery

ANSWER: A

Explanation:

Detection and analysis is the incident-response step that establishes the scope and nature of a suspected security incident. After an alert or report is received, responders validate whether an incident has occurred, collect and correlate evidence from logs, endpoints, network telemetry, and security tools, and determine which systems, accounts, services, and data are involved. Identifying every host affected by an attack is a core scoping activity: the team must understand the incident's reach before it can reliably select containment actions and measure recovery. This investigation may include identifying the initially compromised system, discovering lateral movement, finding systems with matching indicators of compromise, and determining whether related credentials or applications are affected. NIST's incident-handling guidance describes detection and analysis as the phase for examining incident indicators, evaluating their impact, and documenting findings needed to characterize the incident. A complete inventory of affected hosts gives responders the evidence needed to make informed, proportionate response decisions and to maintain an accurate incident record. See NIST SP 800-61 Rev. 2 and the associated incident handling guide (PDF).

QUESTION NO: 25

Refer to the exhibit.

File name	CVE-2009-4324 PDF 2009-11-30 note200911.pdf
File size	400918 bytes
File type	PDF document, version 1.6
CRC32	11638A9B
MD5	61baabd6fc12e01ff73ceacc07c84f9a
SHA1	0805d0ae62f5358b9a3f4c1868d552f5c3561b17
SHA256	27cced58a0fcbb0bbe3894f74d3014611039fefdf3bd2b0ba7ad85b18194c
SHA512	5a43bc7eef279b209e2590432cc3e2eb480d0f78004e265f00b98b4afdc9a
Ssdeep	1536:p0AAH2KthGBjcdBj8VETeePxsT65ZZ3pdx/ves/SQR/875+:prahGV6B
PEiD	None matched
Yara	• embedded_pe (Contains an embedded PE32 file) • embedded_win_api (A non-Windows executable contains win32 API • vmdetect (Possibly employs anti-virtualization techniques)
VirusTotal	Permalink VirusTotal Scan Date: 2013-12-27 06:51:52 Detection Rate: 32/46 (collapse)

An engineer is analyzing this Cuckoo Sandbox report for a PDF file that has been downloaded from an email. What is the state of this file?

A. The file has an embedded executable and was matched by PEiD threat signatures for further analysis.

B. The file has an embedded non-Windows executable but no suspicious features are identified.

C. The file has an embedded Windows 32 executable and the Yara field lists suspicious features for further analysis.

D. The file was matched by PEiD threat signatures but no suspicious features are identified since the signature list is up to date.

ANSWER: C

Explanation:

The file has an embedded Windows 32 executable and the Yara field lists suspicious features for further analysis. This conclusion comes from the static-analysis details shown in the Cuckoo report: the PDF contains an embedded portable executable identified as a Windows 32-bit binary. A PDF can act as a delivery container for embedded content, so identifying a PE within it is a significant indicator requiring investigation. The report's YARA results also identify features or patterns associated with suspicious content. YARA rules are used to classify files by matching defined byte strings, structural characteristics, and other indicators, making these matches useful triage evidence rather than a final verdict by themselves. In this situation, the embedded executable should be extracted or analyzed within the sandbox workflow, and its behavioral findings, network activity, dropped files, and process activity should be correlated with the static indicators. Cuckoo documents static analysis as a report component that identifies file properties and signatures, while YARA describes rules as pattern-based detection logic for identifying and classifying malware-related artifacts. See the Cuckoo report documentation and the YARA rule-writing documentation.

QUESTION NO: 26

Refer to the exhibit A penetration tester runs the Nmap scan against the company server to uncover possible vulnerabilities and exploit them Which two elements can the penetration tester identity from the scan results? (Choose two.)

A. UIDs and group identifiers

B. number of concurrent connections the server can handle

C. running services and applications

D. server uptime and internal clock

E. server purpose and functionality

ANSWER: C E

Explanation:

"running services and applications" is correct because Nmap identifies open ports and can probe them to determine the service listening on each port. When service and version detection is enabled, Nmap sends protocol-specific probes and reports details such as the application protocol, product, and often the product version. This information lets a penetration tester identify exposed attack surface and correlate discovered software with relevant vulnerability information.

"server purpose and functionality" is also correct because the combination of discovered services commonly reveals the server's likely operational role. For example, web-related services suggest web hosting or application delivery, file-sharing services indicate file-serving functionality, and database listeners indicate a database role. Nmap's operating-system and service-detection capabilities are specifically intended to gather this kind of host and network-service intelligence during reconnaissance. The scan does not need direct authenticated access to establish these externally observable characteristics; it derives them from reachable ports and responses to service probes. Cisco security operations personnel use this reconnaissance data to understand an asset's exposed services, validate whether those services are expected, and prioritize investigation of potentially vulnerable software. See the Nmap service and version detection documentation and the Nmap operating-system detection documentation.

QUESTION NO: 27

How does statistical detection differ from rule-based detection?

A. Statistical detection involves the evaluation of events, and rule-based detection requires an evaluated set of events to function.

B. Statistical detection defines legitimate data over time, and rule-based detection works on a predefined set of rules

C. Rule-based detection involves the evaluation of events, and statistical detection requires an evaluated set of events to function Rule-based detection defines

D. legitimate data over a period of time, and statistical detection works on a predefined set of rules

ANSWER: B

Explanation:

Statistical detection defines legitimate data over time, and rule-based detection works on a predefined set of rules. Statistical detection establishes a baseline of expected behavior by observing data across a period of time. This baseline can include such characteristics as usual connection volumes, protocol use, login timing, data-transfer patterns, or host-to-host communications. Once the system has sufficient observations, it can identify behavior that deviates materially from the learned norm and flag it as anomalous. This approach is valuable for detecting unusual or previously unseen activity because the detection is based on behavioral variance rather than solely on a known indicator.

Rule-based detection, in contrast, applies explicitly configured logic to incoming events. Those rules can match known malicious indicators, event attributes, correlation conditions, thresholds, or policy violations. An alert is generated when the observed event data satisfies the defined rule conditions. Cisco's CyberOps Associate certification objectives include understanding security monitoring concepts and event analysis, where both rule-driven alerts and behavioral/anomaly-oriented analysis are relevant: Cisco CyberOps Associate Certification. Cisco also describes behavioral analytics as a means to identify anomalous activity in network traffic: Cisco Secure Network Analytics.

QUESTION NO: 28

What specific type of analysis is assigning values to the scenario to see expected outcomes?

A. deterministic

B. exploratory

C. probabilistic

D. descriptive

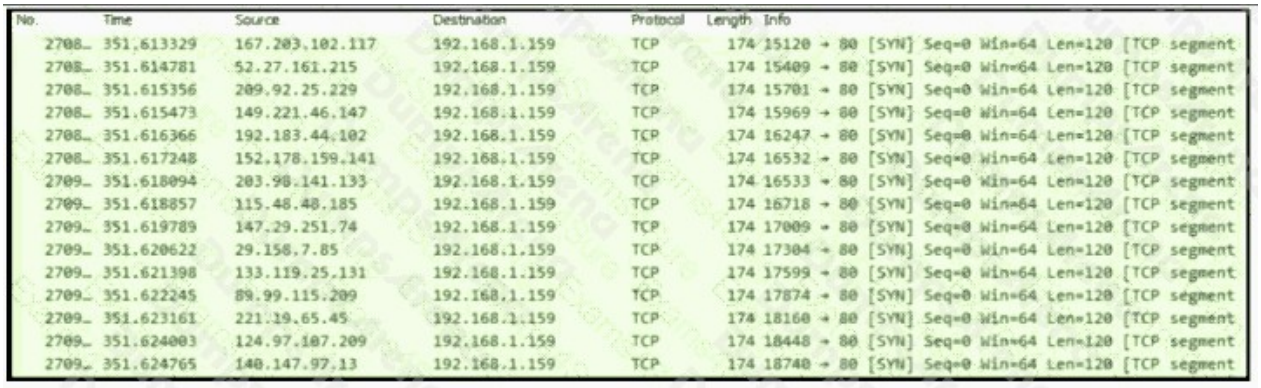
ANSWER: A

Explanation:

Deterministic is correct because deterministic analysis uses defined, fixed input values and a known model or set of rules to calculate an expected outcome. Given the same scenario, assumptions, and assigned values, the analysis will produce the same result every time. This makes it useful when an analyst needs to model a specific security scenario with explicit assumptions, such as the anticipated operational impact of a control failure, the projected cost of an incident under stated conditions, or the expected result of a policy change. The analyst assigns values to relevant variables, processes those inputs through the model, and obtains a single predictable outcome for that set of conditions. This approach supports repeatable evaluation and clear documentation of the assumptions underlying a conclusion. In cybersecurity operations, such structured modeling can help teams communicate expected effects and make decisions using consistent criteria. The National Institute of Standards and Technology defines deterministic in terms of outcomes being fixed by specified conditions, and its risk-assessment guidance emphasizes documenting assumptions, threat sources, impacts, and analysis inputs. See the NIST deterministic glossary entry and NIST SP 800-30 Rev. 1.

QUESTION NO: 29

Refer to exhibit.



No.	Time	Source	Destination	Protocol	Length	Info
2708…	351.613329	167.203.102.117	192.168.1.159	TCP	174	15120 → 80 [SYN] Seq=0 Win=64 Len=120 [TCP segment
2708…	351.614781	52.27.161.215	192.168.1.159	TCP	174	15409 → 80 [SYN] Seq=0 Win=64 Len=120 [TCP segment
2708…	351.615356	209.92.25.229	192.168.1.159	TCP	174	15701 → 80 [SYN] Seq=0 Win=64 Len=120 [TCP segment
2708…	351.615473	149.221.46.147	192.168.1.159	TCP	174	15969 → 80 [SYN] Seq=0 Win=64 Len=120 [TCP segment
2708…	351.616366	192.183.44.102	192.168.1.159	TCP	174	16247 → 80 [SYN] Seq=0 Win=64 Len=120 [TCP segment
2708…	351.617248	152.178.159.141	192.168.1.159	TCP	174	16532 → 80 [SYN] Seq=0 Win=64 Len=120 [TCP segment
2709…	351.618094	203.98.141.133	192.168.1.159	TCP	174	16533 → 80 [SYN] Seq=0 Win=64 Len=120 [TCP segment
2709…	351.618857	115.48.48.185	192.168.1.159	TCP	174	16718 → 80 [SYN] Seq=0 Win=64 Len=120 [TCP segment
2709…	351.619789	147.29.251.74	192.168.1.159	TCP	174	17009 → 80 [SYN] Seq=0 Win=64 Len=120 [TCP segment
2709…	351.620622	29.158.7.85	192.168.1.159	TCP	174	17304 → 80 [SYN] Seq=0 Win=64 Len=120 [TCP segment
2709…	351.621398	133.119.25.131	192.168.1.159	TCP	174	17599 → 80 [SYN] Seq=0 Win=64 Len=120 [TCP segment
2709…	351.622245	89.99.115.209	192.168.1.159	TCP	174	17874 → 80 [SYN] Seq=0 Win=64 Len=120 [TCP segment
2709…	351.623161	221.19.65.45	192.168.1.159	TCP	174	18160 → 80 [SYN] Seq=0 Win=64 Len=120 [TCP segment
2709…	351.624003	124.97.187.209	192.168.1.159	TCP	174	18448 → 80 [SYN] Seq=0 Win=64 Len=120 [TCP segment
2709…	351.624765	140.147.97.13	192.168.1.159	TCP	174	18740 → 80 [SYN] Seq=0 Win=64 Len=120 [TCP segment

An engineer is Investigating an Intrusion and Is analyzing the pcap file. Which two key elements must an engineer consider? (Choose two.)

A. Variable "info" field and unchanging sequence number

B. High volume of SYN packets with very little variance in time

C. identical length of 120 and window size (64)

D. SYN packets acknowledged from several source IP addresses

E. same source IP address with a destination port 80

ANSWER: B D

Explanation:

High volume of SYN packets with very little variance in time is a key indicator because a TCP SYN flood attempts to consume connection-state resources by rapidly sending connection requests. In normal TCP operation, SYN packets are ordinarily followed by corresponding SYN-ACK and ACK exchanges that establish connections. A dense concentration of initial SYN packets over a short interval, particularly when the capture shows little evidence of completed handshakes, is therefore important evidence of an availability attack.

SYN packets acknowledged from several source IP addresses is also significant because activity involving many apparent sources can indicate a distributed attack. A DDoS campaign commonly uses numerous compromised systems to generate traffic toward one target, making source diversity an important investigative characteristic. The analyst should correlate the packet timestamps, source addresses, destination, TCP flags, and handshake completion behavior to determine the scale and likely nature of the event. TCP SYN flooding is specifically documented as a denial-of-service technique in RFC 4987. Cisco also describes DDoS attacks as attempts to overwhelm a target using traffic from multiple distributed sources in its DDoS attack overview.

QUESTION NO: 30

An organization that develops high-end technology is going through an internal audit The organization uses two databases The main database stores patent information and a secondary database stores employee names and contact information A compliance team is asked to analyze the infrastructure and identify protected data Which two types of protected data should be identified? (Choose two)

A. Personally Identifiable Information (PII)

B. Payment Card Industry (PCI)

C. Protected Health Information (PHI)

D. Intellectual Property (IP)

E. Sarbanes-Oxley (SOX)

ANSWER: A D

Explanation:

Personally Identifiable Information (PII) applies to the employee database because employee names and contact details are information that can identify, contact, or distinguish a specific individual. A compliance review must identify this data so that appropriate controls can be applied throughout its lifecycle, including access restrictions, secure storage, monitoring, retention, and secure disposal. NIST defines PII as information that can be used to distinguish or trace an individual's identity, either by itself or when combined with other information. See the NIST definition of PII.

Intellectual Property (IP) applies to the patent-information database. Patent records concern inventions and associated technical knowledge that have substantial business value and are protected under intellectual-property law. This makes them sensitive organizational assets that should be classified and protected against unauthorized disclosure, alteration, or loss. During the audit, identifying patent information as IP supports appropriate data classification and controls, such as need-to-know access, encryption where appropriate, and logging of access to proprietary records. The U.S. Patent and Trademark Office describes patents as grants providing rights relating to inventions; see the USPTO patent basics.

QUESTION NO: 31

What should a security analyst consider when comparing inline traffic interrogation with traffic tapping to determine which approach to use in the network?

A. Tapping interrogation replicates signals to a separate port for analyzing traffic

B. Tapping interrogations detect and block malicious traffic

C. Inline interrogation enables viewing a copy of traffic to ensure traffic is in compliance with security policies

D. Inline interrogation detects malicious traffic but does not block the traffic

ANSWER: A

Explanation:

Tapping interrogation replicates signals to a separate port for analyzing traffic. A network tap is a passive monitoring mechanism placed on a link so that a security tool, such as an intrusion detection sensor or packet analyzer, can receive a copy of the packets without being in the forwarding path. This architecture lets the analyst observe network behavior, inspect packet contents, and collect evidence while the original traffic continues between the communicating devices. The copy is delivered to a monitoring interface or other designated output for analysis, which is the key characteristic to consider when choosing a tap-based deployment.

This distinction matters operationally: tapping is appropriate when visibility and passive analysis are required and the organization does not want the inspection device to become part of the production traffic path. Cisco describes network traffic monitoring as collecting traffic through mechanisms such as SPAN and taps, which provide packets to monitoring tools for visibility. See Cisco's traffic mirroring documentation and its overview of network taps.

QUESTION NO: 32

Which two actions are appropriate during the containment phase of an incident response? (Choose two.)

A. isolating an affected endpoint

B. blocking a malicious domain

C. rebuilding the affected system

D. conducting a lessons-learned meeting

E. retiring the incident case

ANSWER: A B

Explanation:

Isolating an affected endpoint and **blocking a malicious domain** are containment actions because they limit the spread and impact of an incident while preserving the opportunity for further investigation. Endpoint isolation can prevent lateral movement and command-and-control communication. Blocking a known malicious domain at DNS, proxy, or firewall controls can stop additional connections to attacker infrastructure. Rebuilding systems and conducting lessons learned are generally performed later during recovery and post-incident activity. See NIST SP 800-61 Rev. 2 and CISA Incident Response.

QUESTION NO: 33

Which two elements of the incident response process are stated in NIST SP 800-61 r2? (Choose two.)

A. detection and analysis

B. post-incident activity

C. vulnerability scoring

D. vulnerability management

E. risk assessment

ANSWER: A B

Explanation:

NIST SP 800-61 Revision 2, *Computer Security Incident Handling Guide*, organizes the incident-response life cycle into four phases: preparation; detection and analysis; containment, eradication, and recovery; and post-incident activity. Therefore, **detection and analysis** and **post-incident activity** are explicitly named elements of the NIST incident-response process. Detection and analysis covers identifying potential security events, determining whether an incident has occurred, assessing its scope and impact, and documenting the findings so that responders can make informed decisions. This phase turns monitoring alerts and reported anomalies into a validated understanding of the incident. Post-incident activity occurs after response and recovery actions and focuses on lessons learned, preserving and reviewing incident information, identifying process improvements, and implementing changes that strengthen future detection, response, and resilience. NIST emphasizes that these activities support continuous improvement of an organization's incident-handling capability rather than treating an incident as complete immediately after systems are restored. The complete lifecycle diagram and descriptions appear in section 3.2 of the guide. See NIST SP 800-61 Revision 2 and the associated NIST publication PDF.

QUESTION NO: 34

What are two denial-of-service (DoS) attacks? (Choose two)

A. port scan

B. SYN flood

C. man-in-the-middle

D. phishing

E. teardrop

ANSWER: B E

Explanation:

SYN flood and **teardrop** are denial-of-service attacks because each seeks to make a host or service unavailable by exhausting resources or disrupting protocol processing. A SYN flood abuses the TCP three-way handshake: the attacker sends a large number of SYN requests, often using spoofed source addresses, and does not complete the connections. The target retains many half-open connection states until its backlog or available resources are consumed, preventing legitimate users from establishing sessions. A teardrop attack uses malformed, overlapping, or otherwise problematic IP fragment offsets. Historically, vulnerable systems could mishandle fragment reassembly, leading to instability, resource exhaustion, or a crash. Both attacks therefore target availability, which is the security objective affected by DoS activity. Cisco security

training treats flooding and malformed-packet attacks as common mechanisms used to deny legitimate access to network services. For a technical description of SYN-flood behavior and mitigations, see Cisco: TCP SYN Flooding Attacks and Common Mitigations. Additional background on denial-of-service attacks and their availability impact is available from CISA: Denial-of-Service Attacks.

QUESTION NO: 35

An engineer received an alert affecting the degraded performance of a critical server. Analysis showed a heavy CPU and memory load. What is the next step the engineer should take to investigate this resource usage?

A. Run "ps -d" to decrease the priority state of high load processes to avoid resource exhaustion.

B. Run "ps -u" to find out who executed additional processes that caused a high load on a server.

C. Run "ps -ef" to understand which processes are taking a high amount of resources.

D. Run "ps -m" to capture the existing state of daemons and map required processes to find the gap.

E. Run "top" to identify the processes consuming high CPU or memory resources.

ANSWER: E

Explanation:

Run `top` to identify the processes consuming high CPU or memory resources. This is the appropriate immediate investigation step because `top` provides a continuously refreshed, system-wide view of process activity alongside overall CPU and memory statistics. The engineer can sort the process display by CPU consumption or resident memory use, identify the relevant process ID, account, command, and resource percentage, and then determine whether the load is expected or requires containment and further incident investigation. This is especially useful during an active degradation event because resource-intensive processes can change rapidly; a static process listing alone may not reveal which processes are currently responsible for the contention. The `top` interactive display includes fields such as CPU usage, memory usage, process state, and command name, giving the engineer actionable evidence for the next operational decision. After identifying the process, the engineer should validate its legitimacy and business purpose before taking action such as stopping it, adjusting its limits, or escalating the issue. The Linux `top` documentation describes its process-monitoring and sorting capabilities in the top(1) manual page. Cisco's incident-handling guidance similarly emphasizes collecting and analyzing evidence before containment actions: Cisco Incident Response Services.

QUESTION NO: 36

What are the two characteristics of the full packet captures? (Choose two.)

A. Identifying network loops and collision domains.

B. Troubleshooting the cause of security and performance issues.

C. Reassembling fragmented traffic from raw data.

D. Detecting common hardware faults and identify faulty assets.

E. Providing a historical record of a network transaction.

ANSWER: C E

Explanation:

Full packet capture records the actual packets observed on an interface, including packet headers and, when capture policy permits, packet payloads. This level of detail preserves the information needed for forensic and protocol analysis after the event has occurred. As a result, it provides a historical record of a network transaction: an analyst can review the packets belonging to a connection, inspect protocol fields, correlate request and response activity, and establish what traffic was

present at a particular time. Cisco describes packet capture as a way to collect packets for later analysis, and packet-analysis tools can reconstruct conversations from the captured packet data.

Because a full capture retains the individual raw packets rather than only flow metadata or summary statistics, analysis software can also reassemble fragmented traffic from the captured data. IP fragmentation and segmented application exchanges may span multiple packets; preserving the packet sequence, addressing information, protocol headers, and available payload enables a decoder to reconstruct the higher-level traffic for inspection. This makes full packet capture particularly valuable when analysts need evidentiary detail and protocol-level visibility. See Cisco's Embedded Packet Capture Overview and Wireshark's packet reassembly documentation.

QUESTION NO: 37

A cyberattacker notices a security flaw in a software that a company is using They decide to tailor a specific worm to exploit this flaw and extract saved passwords from the software To which category of the Cyber Kill Cham model does this event belong?

A. reconnaissance

B. delivery

C. weaponization

D. exploitation

ANSWER: C

Explanation:

weaponization is correct because the attacker is preparing a tailored malicious capability for a known target weakness. In the Cyber Kill Chain model, weaponization is the stage in which an adversary combines an exploit with a payload to create a usable attack tool or artifact. Here, the attacker has identified a flaw and then customizes a worm specifically to take advantage of it. The worm's credential-stealing functionality is the payload, while the code that abuses the software flaw is the exploit component. Together, these form a weapon designed for later use against the organization's environment. This stage is especially significant to defenders because customized malware may be designed to evade standard security controls and to achieve a particular objective, such as collecting saved passwords. The event describes creation and tailoring of the malicious tool, rather than its transmission to a target or execution on a target system. Cisco's overview of the model describes weaponization as preparing the malicious payload and exploit for use in an intrusion: Cisco: What Is the Cyber Kill Chain?. The original model is further documented by Lockheed Martin in its Intelligence-Driven Computer Network Defense paper.

QUESTION NO: 38

How is NetFlow different from traffic mirroring?

A. NetFlow collects metadata and traffic mirroring clones data.

B. Traffic mirroring impacts switch performance and NetFlow does not.

C. Traffic mirroring costs less to operate than NetFlow.

D. NetFlow generates more data than traffic mirroring.

ANSWER: A

Explanation:

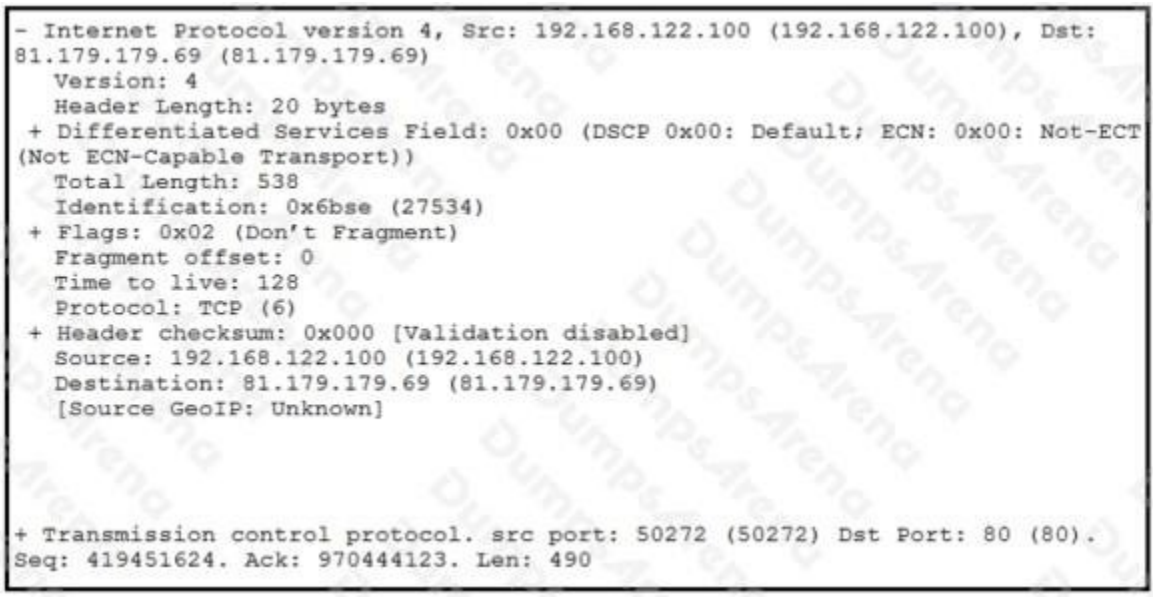
NetFlow collects metadata and traffic mirroring clones data. NetFlow is a flow-monitoring technology: network devices summarize observed communications into flow records. These records commonly include details such as source and destination addresses, protocol, ports, packet and byte counts, timestamps, and TCP flags. Because it exports summarized

information about conversations rather than the packets themselves, NetFlow provides scalable visibility for traffic analysis, capacity planning, anomaly detection, and security monitoring.

Traffic mirroring, often implemented with a switched port analyzer (SPAN) session, makes copies of packets seen on a source interface, VLAN, or other defined traffic source and forwards those copies to a monitoring destination. This permits a packet analyzer, IDS, or other security tool to inspect packet contents and protocol details present in the copied traffic. The core distinction is therefore the type of telemetry produced: NetFlow exports structured flow metadata, while traffic mirroring delivers copies of actual packet data for inspection. Cisco describes NetFlow as exporting flow information to collectors and SPAN as forwarding copies of network traffic to an analyzer. See Cisco NetFlow overview and Cisco SPAN configuration guidance.

QUESTION NO: 39

Refer to the exhibit.



What should be interpreted from this packet capture?

A. 81.179.179.69 is sending a packet from port 80 to port 50272 of IP address 192.168.122.100 using UDP protocol.

B. 192.168.122.100 is sending a packet from port 50272 to port 80 of IP address 81.179.179.69 using TCP protocol.

C. 192.168.122.100 is sending a packet from port 80 to port 50272 of IP address 81.179.179.69 using UDP protocol.

D. 81.179.179.69 is sending a packet from port 50272 to port 80 of IP address 192.168.122.100 using TCP UDP protocol.

ANSWER: B

Explanation:

192.168.122.100 is sending a packet from port 50272 to port 80 of IP address 81.179.179.69 using TCP protocol. Packet-capture summaries conventionally present the source endpoint first and the destination endpoint second. In the displayed flow, the endpoint before the direction indicator is 192.168.122.100 with source port 50272, while the endpoint after it is 81.179.179.69 with destination port 80. The protocol field identifies the packet as TCP. Port 50272 is an ephemeral client-side port selected for this connection, and port 80 is the well-known service port commonly associated with HTTP. Therefore, this packet represents traffic from the internal host toward the remote web service endpoint. Interpreting the five-tuple—source IP address, source port, destination IP address, destination port, and transport protocol—is essential during security investigations because it establishes traffic direction and identifies the service being contacted. TCP provides a connection-oriented transport service, with sequence and acknowledgment information used to track reliable data exchange. This interpretation is consistent with TCP's endpoint and port model described in RFC 793 and Wireshark's guidance on analyzing TCP conversations in the Wireshark User's Guide.

QUESTION NO: 40

Which type of access control depends on the job function of the user?

A. discretionary access control

B. nondiscretionary access control

C. role-based access control

D. rule-based access control

ANSWER: C

Explanation:

Role-based access control is correct because it assigns permissions according to a user's organizational role or job function, rather than granting access individually to each user. For example, people assigned a network administrator role can receive permissions needed to manage network devices, while users in an accounting role receive access to relevant financial systems and records. This approach lets an organization define a consistent set of privileges for each role and then grant those privileges when users are assigned to that role. It supports least privilege by making permissions align with the responsibilities required for a position, and it simplifies administration when employees join, change jobs, or leave the organization. Cisco's cybersecurity training identifies role-based access control as an access-control approach in which rights and permissions are assigned based on job functions and responsibilities. The National Institute of Standards and Technology similarly defines RBAC as controlling system access based on roles and associated operations, with users assigned to appropriate roles. See Cisco's access control overview and NIST's Role Based Access Control project.

QUESTION NO: 41

Refer to the exhibit.

An attacker Infiltrated an organization's network and ran a scan to advance with the lateral movement technique. Which two elements from the scan assist the attacker? (Choose two.)

A. function and service the server is providing

B. CPU and vendor version of the asset

C. running services and ports

D. security identifiers of logged-in accounts

E. latency and MS information to calculate delays for a command injection

ANSWER: A C

Explanation:

"function and service the server is providing" assists lateral movement because service discovery lets an attacker infer a discovered host's purpose and its likely role in the environment. For example, Microsoft RPC, NetBIOS, and SMB-related services can indicate a Windows system that offers remote management, name-sharing, or file-sharing capabilities. That context helps the attacker prioritize hosts that may expose useful administrative paths, shared resources, or credentials and select an appropriate next technique for accessing another system.

"running services and ports" is also directly useful because a scan identifies reachable listening services and the network ports through which they can be contacted. Open ports define the attack surface available from the attacker's current position, while detected service information supports targeting protocols commonly used to move within enterprise networks. Cisco describes lateral movement as adversaries moving through a network after gaining an initial foothold, often by leveraging remote services and valid accounts. MITRE similarly documents remote services as a technique used to access internal systems through exposed management and file-sharing protocols. This scan information therefore provides the host and service reconnaissance needed to identify viable internal movement paths. See MITRE ATT&CK: Lateral Movement and MITRE ATT&CK: Remote Services.

QUESTION NO: 42



Category	Started On	Completed On	Duration	Cuckoo Version
FILE	2014-02-23 21:52:16	2014-02-23 21:52:34	18 seconds	1.0

File Details

File name	[illegible]
File size	414720 bytes
File type	PE32 executable (GUI) Intel 80386, for MS Windows
CRC32	8048E2EA
MD5	[illegible]
SHA1	[illegible]
SHA256	f4855d1b10f7ab1a2e6099016437f72c5f98579d69f88b6312cc24400f483177
SHA512	9756e0af8981bc9296a3879fe02d0e102c5557ba99a004230ca4f1dfd03592cf497c123d2a6a05596b07432188aaef42976e80d9da742c09002758e721db2595
Ssdeep	6144:Eu2UY7eiLnfnB7pRiBI+5zLqiZ49XCUgNqGyCYUE/1rWDepfYXt+o6YUPL:Eu2UY7eandid=5VGCUgMDCk/1r7EE
PEiD	None matched
Yara	• shellcode (Matched shellcode byte patterns)
VirusTotal	Permalink VirusTotal Scan Date: 2014-01-12 23:43:56 Detection Rate: 26/47 (collapse)

Refer to the exhibit. An employee received an email from an unknown sender with an attachment and reported it as a phishing attempt. An engineer uploaded the file to Cuckoo for further analysis. What should an engineer interpret from the provided Cuckoo report?

A. Win32.polip.a.exe is an executable file and should be flagged as malicious.

B. The file is clean and does not represent a risk.

C. Cuckoo cleaned the malicious file and prepared it for usage.

D. MD5 of the file was not identified as malicious.

ANSWER: D

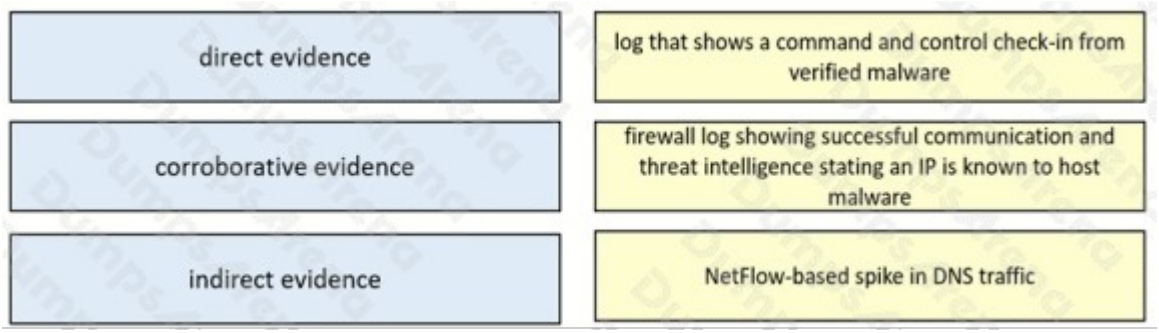
Explanation:

MD5 of the file was not identified as malicious. This interpretation follows from the report's hash-reputation result: the submitted file's MD5 value does not have a malicious identification in the referenced reputation lookup. An MD5 hash is a file fingerprint, and sandbox reports commonly use it to correlate a submitted sample with previously known files and threat-intelligence results. When no malicious reputation is returned for that hash, the correct conclusion is specifically that the hash was not identified as malicious at the time of the lookup.

This result is a reputation observation, rather than a guarantee about the attachment's safety. A file can be new, modified, targeted, or otherwise absent from public intelligence sources while still requiring behavioral analysis and normal incident-handling precautions. Cuckoo reports present file metadata, hashes, behavioral results, and integrated intelligence information so analysts can combine static and dynamic evidence during triage. Cisco security operations practice likewise treats file reputation as one useful signal within a broader investigation, not as a standalone proof of file safety. See the Cuckoo Sandbox report-results documentation and Cisco's Advanced Malware Protection overview.

QUESTION NO: 43 - (DRAG DROP)

Drag and drop the type of evidence from the left onto the description of that evidence on the right.



ANSWER:



Explanation:

Evidence classification depends on how conclusively an artifact establishes the investigated activity. A log that records a command-and-control check-in from malware that has already been verified provides direct evidence. The event is an explicit record of the known malicious program communicating with its command-and-control infrastructure, rather than merely an unusual condition that requires a further inference.

A firewall log showing successful communication with an IP address that threat intelligence identifies as hosting malware is corroborative evidence. It independently supports the investigation by connecting the observed host activity to intelligence associated with malicious infrastructure. This type of evidence is especially useful for adding confidence and context to an established or suspected finding. Cisco's incident-handling guidance emphasizes collecting and analyzing relevant artifacts and contextual information during investigations; see the Cisco Incident Response Services overview.

A NetFlow-based spike in DNS traffic is indirect evidence. NetFlow supplies metadata about communications, such as traffic volumes and flows, and an unusual DNS increase can identify behavior worth investigating. The spike can indicate activity patterns consistent with compromise or misuse, but its meaning depends on further analysis of the involved endpoints, domains, queries, and related telemetry. This is consistent with the role of flow data in security monitoring described in Cisco's network telemetry guidance.

Therefore, place direct evidence with the verified-malware command-and-control check-in, corroborative evidence with the firewall communication to the threat-intelligence malware IP, and indirect evidence with the NetFlow DNS-traffic spike.

QUESTION NO: 44

Refer to the exhibit.

What information is depicted?

A. IIS data

B. NetFlow data

C. network discovery event

D. IPS event data

ANSWER: B

Explanation:

NetFlow data is correct because NetFlow is flow telemetry that summarizes communications observed by network devices. Rather than storing every packet payload, a NetFlow record identifies characteristics of a conversation, commonly including source and destination IP addresses, source and destination ports, IP protocol, ingress or egress interfaces, timestamps, packet counts, byte counts, and TCP flags. This metadata gives security operations personnel a compact, scalable view of who communicated with whom, over which service, for how long, and at what volume. Those properties make it particularly useful for traffic baselining, identifying unexpected internal or external connections, investigating command-and-control patterns, and finding unusual data-transfer volumes. Cisco describes NetFlow as technology for collecting IP network traffic information and exporting flow records to a collector for analysis. Cisco security analytics products also use flow telemetry to provide visibility into network activity and help detect threats. The format and fields normally shown in a flow-oriented exhibit therefore represent NetFlow information. See Cisco's NetFlow overview and Cisco Secure Network Analytics overview.

QUESTION NO: 45

What is a collection of compromised machines that attackers use to carry out a DDoS attack?

A. subnet

B. botnet

C. VLAN

D. command and control

ANSWER: B

Explanation:

The correct answer is **botnet**. A botnet is a network of computers, servers, Internet of Things devices, or other endpoints that have been compromised with malware and can be remotely directed by an attacker. Each compromised endpoint, often called a bot or zombie, can generate traffic or requests toward a chosen target. By coordinating large numbers of these devices, an attacker can create the distributed traffic volume characteristic of a distributed denial-of-service attack. The distribution is important: traffic originates from many distinct systems and networks, which can overwhelm the target's bandwidth, processing capacity, or supporting services and can make filtering more challenging. Botnets are generally managed through an attacker-controlled infrastructure that sends instructions to the compromised devices, enabling the attacker to activate the bots for a campaign without directly operating every device. Cisco describes botnets as networks of infected devices controlled by a threat actor and notes their use in DDoS activity. For additional context, see Cisco: What Is a Botnet? and CISA: Malware and Botnets.

QUESTION NO: 46

Which two data sources can an analyst use to investigate a suspected password-spraying attack against enterprise user accounts? (Choose two.)

A. identity-provider authentication logs

B. VPN authentication logs

C. printer toner logs

D. switch power-supply logs

E. disk defragmentation reports

ANSWER: A B

Explanation:

Identity-provider authentication logs and **VPN authentication logs** are correct because password spraying involves repeated authentication attempts against many accounts using one or a small number of passwords. Identity-provider logs can show failed sign-in attempts, source IP addresses, targeted users, device information, conditional-access results, and time patterns. VPN authentication logs provide similar evidence for remote-access services, including repeated failures from the same source against multiple usernames.

Correlating these sources helps an analyst distinguish password spraying from a single user forgetting a password. Typical signs include a source address attempting one password against many accounts, widespread failures in a short period, and later successful authentication to one of the targeted accounts. MITRE ATT&CK documents password spraying as a sub-technique of brute force at MITRE ATT&CK: Password Spraying. Cisco also describes the role of identity controls in its Zero Trust overview.

QUESTION NO: 47

An engineer is working with the compliance teams to identify the data passing through the network. During analysis, the engineer informs the compliance team that external penmeter data flows contain records, writings, and artwork Internal segregated network flows contain the customer choices by gender, addresses, and product preferences by age. The engineer must identify protected data. Which two types of data must be identified'? (Choose two.)

A. SOX

B. PII

C. PHI

D. PCI

E. copyright

ANSWER: B E

Explanation:

PII must be identified because the internal network flow includes customer addresses and demographic or preference attributes. An address can identify or help identify a particular individual, and the associated gender, age-related preferences, and purchasing information increase the sensitivity of that customer profile. Identifying these flows enables the organization to apply appropriate data-classification, access-control, monitoring, retention, and privacy-handling requirements. NIST defines personally identifiable information as information that can distinguish or trace an individual's identity, either alone or when combined with other linked or linkable information. See NIST's PII definition.

copyright must also be identified because records, writings, and artwork are categories of original works that can be protected by copyright. Such content may be owned by the organization, a customer, or another rights holder, so recognizing it in network flows supports controls against unauthorized copying, distribution, disclosure, or use. Copyright protection applies to original works of authorship fixed in a tangible medium, including literary, written, and artistic works. The U.S. Copyright Office describes these protected categories at Copyright Basics FAQ. Therefore, the relevant protected-data classifications for the described flows are PII and copyright.

QUESTION NO: 48

An engineer received an alert affecting the degraded performance of a critical server. Analysis showed a heavy CPU and memory load. What is the next step the engineer should take to investigate this resource usage?

A. Run "ps -d" to decrease the priority state of high load processes to avoid resource exhaustion.

B. Run "ps -u" to find out who executed additional processes that caused a high load on a server.

C. Run "ps -ef" to understand which processes are taking a high amount of resources.

D. Run "ps -m" to capture the existing state of daemons and map required processes to find the gap.

ANSWER: C

Explanation:

Run "ps -ef" to understand which processes are taking a high amount of resources. This is the appropriate next investigative action because the engineer needs to identify the running workloads behind the server-level CPU and memory alert before deciding on containment or remediation. The `ps` utility reports process information, while `-e` selects every process and `-f` requests a full-format listing. The resulting process inventory includes useful investigation fields such as the user identity, process ID, parent process ID, CPU-time-related information, process start time, and complete command line. These details allow the engineer to associate suspect activity with a specific executable and account, understand its parent-child relationship, and determine whether a legitimate service, scheduled job, or unexpected command is responsible for the load. Process IDs and command lines gathered from this output also provide the context needed to perform focused follow-up monitoring or inspect the relevant application and service logs. In an incident workflow, identifying the process responsible for abnormal resource consumption is essential evidence and should precede changes that could disrupt a critical server. The Linux `ps` manual documents the process-selection and full-format behavior at man7.org ps(1). Additional operational guidance for monitoring process activity is available in the Red Hat Performance Tuning Guide.

QUESTION NO: 49

What are the two differences between vulnerability and exploit? (Choose two.)

A. Known vulnerabilities are assigned special CVE numbers, and exploits are using process to take advantage of vulnerabilities.

B. Vulnerabilities can be found in hardware and software, and exploits can be used only for software-based vulnerabilities.

C. Zero-day exploit can be used to take advantage of a vulnerability until the vulnerable software or hardware is patched.

D. Vulnerabilities are usually populated in the dark web, and exploit tools and methods can be found in the public web.

E. Zero-day exploit can be used for taking advantage of a known vulnerability, and cyber-attack can be performed on company assets.

ANSWER: A C

Explanation:

“Known vulnerabilities are assigned special CVE numbers, and exploits are using process to take advantage of vulnerabilities.” is correct because a vulnerability is a weakness in a system, application, configuration, or process that can potentially be abused. Publicly disclosed vulnerabilities are commonly tracked with Common Vulnerabilities and Exposures identifiers, which provide a standardized reference for security teams, vendors, and vulnerability-management tools. An exploit is the code, technique, procedure, or method used to trigger or leverage such a weakness to produce an unintended result, such as unauthorized access, execution of code, or disclosure of data. NIST defines a vulnerability as a weakness that could be exploited by a threat source. See the NIST vulnerability glossary entry.

“Zero-day exploit can be used to take advantage of a vulnerability until the vulnerable software or hardware is patched.” is also correct. A zero-day exploit targets a vulnerability for which defenders have had no effective remediation time, often because the flaw is newly discovered or lacks an available vendor patch. Attackers can use the exploit while the weakness remains present; applying a vendor-provided patch or otherwise mitigating the vulnerable condition closes the exposure targeted by that exploit. CISA describes zero-day vulnerabilities and the urgency of remediation in its guidance on understanding zero-day vulnerabilities.

QUESTION NO: 50

An engineer must compare NIST vs ISO frameworks The engineer deeded to compare as readable documentation and also to watch a comparison video review. Using Windows 10 OS. the engineer started a browser and searched for a NIST document and then opened a new tab in the same browser and searched for an ISO document for comparison

The engineer tried to watch the video, but there 'was an audio problem with OS so the engineer had to troubleshoot it At first the engineer started CMD and looked fee a driver path then locked for a corresponding registry in the registry editor The

engineer enabled "Audiosrv" in task manager and put it on auto start and the problem was solved Which two components of the OS did the engineer touch? (Choose two)

A. permissions

B. PowerShell logs

C. service

D. MBR

E. process and thread

F. registry

ANSWER: C F

Explanation:

The engineer interacted with the **registry** and a **service**. The Windows Registry is the hierarchical configuration database used by Windows and applications to store system-wide and user-specific settings. Opening Registry Editor and locating a corresponding registry entry is direct interaction with this operating-system component. Registry values can contain configuration information relevant to hardware, drivers, applications, and service behavior. Microsoft describes the registry as a central database for configuration data; see Windows Registry.

The engineer also worked with a Windows service by enabling **Audiosrv**, the Windows Audio service, and configuring it to start automatically. A service is a long-running executable managed by the Service Control Manager that can run without a user actively interacting with it. Configuring a service startup type to Automatic allows Windows to start the service during system startup, ensuring that dependent audio functionality is available. Microsoft's service architecture documentation explains how services are configured and managed: About Services.

Therefore, the scenario explicitly identifies registry access and Windows Audio service management as the two operating-system components involved in resolving the audio issue.

QUESTION NO: 51

Which two benefits are provided by network segmentation? (Choose two.)

A. limiting lateral movement

B. reducing the scope of a security incident

C. eliminating the need for endpoint protection

D. encrypting all traffic automatically

E. removing the need for access-control policies

ANSWER: A B

Explanation:

Limiting lateral movement and **reducing the scope of a security incident** are correct. Network segmentation divides an environment into smaller security zones and controls traffic between those zones with mechanisms such as firewalls, access control lists, virtual LANs, and security group policies. If an attacker compromises a host in one segment, segmentation can prevent or restrict direct access to systems in other segments, limiting the attacker's ability to move laterally.

Segmentation also reduces incident scope because affected systems can be isolated without necessarily disrupting unrelated business services. For example, a compromised user subnet can be separated from server, payment-card, or industrial-control segments while the incident is investigated. Cisco recommends segmentation as part of a zero-trust approach to limit access and contain threats. See Cisco: What Is Network Segmentation? and Cisco: What Is Zero Trust?.

QUESTION NO: 52

Which data type is necessary to get information about source/destination ports?

A. statistical data

B. session data

C. connectivity data

D. alert data

ANSWER: B

Explanation:

session data is correct because it records the attributes of an observed communication session, or network flow, between endpoints. Source and destination port numbers are core session attributes: together with source and destination IP addresses and the transport protocol, they identify the conversation and the specific service involved. For example, a session record can show that a client used an ephemeral source port to connect to a server's destination port for HTTPS, DNS, or another application service. This context helps a security analyst establish who communicated with whom, over which protocol and service, when the communication occurred, and how much traffic was exchanged. Cisco flow-monitoring technologies use flow keys that include transport-layer source and destination ports, allowing analysts to distinguish separate conversations even when the same hosts communicate concurrently. This metadata is especially valuable for incident triage, threat hunting, traffic baselining, and correlating network activity with firewall, endpoint, and intrusion-detection telemetry. Cisco's Flexible NetFlow overview describes flow records and the fields used to characterize traffic, including transport source and destination port information: Cisco Flexible NetFlow Overview. Cisco's NetFlow documentation also explains the use of flow data for visibility into network communications: Cisco NetFlow.

QUESTION NO: 53

Refer to the exhibit.

```
Nmap scan report for 192.168.233.129
Host is up (0.00086s latency).
Not shown: 65530 filtered ports
PORT      STATE SERVICE
135/tcp   open  msrpc
139/tcp   open  netbios-ssn
445/tcp   open  microsoft-ds
49154/tcp open  unknown
49157/tcp open  unknown
MAC Address: 00:0C:29:BA:EF:BA (VMware)

Nmap done: 1 IP address (1 host up) scanned in 124.12 seconds
```

An attacker Infiltrated an organization ' s network and ran a scan to advance with the lateral movement technique. Which two elements from the scan assist the attacker? (Choose two.)

A. function and service the server is providing

B. CPU and vendor version of the asset

C. running services and ports

D. security identifiers of logged-in accounts

E. latency and MS information to calculate delays for a command injection

ANSWER: A C

Explanation:

“function and service the server is providing” and “running services and ports” are the scan elements that directly help an attacker plan lateral movement. A port-and-service scan identifies reachable network services and commonly provides service names, protocol details, and sometimes product or version information. This lets an attacker recognize a host’s likely role and select a suitable next action, such as attempting authenticated access to a file-sharing service, connecting to a remote-management service, or investigating a Windows RPC endpoint. In an internal environment, services exposed through ports associated with SMB, NetBIOS, and Microsoft RPC are particularly useful reconnaissance findings because they reveal potential paths to other systems, shared resources, and remote administration capabilities.

Understanding the function a server is providing also helps prioritize targets. For example, a system exposing enterprise file-sharing or directory-related services may be more valuable for credential reuse, access to shared data, or pivoting than a host with no relevant reachable services. Nmap’s service and version detection is designed to identify the application listening on discovered ports, which is exactly the type of evidence needed to map available internal attack surfaces. Cisco’s security guidance similarly treats service discovery as foundational network reconnaissance for identifying exposed services and reducing attack surface. See the Nmap service and version detection documentation and Cisco’s network security overview.

QUESTION NO: 54

Refer to the exhibit.

```
root@:~# cat access-logs/access_130603.txt | grep '192.168.1.91' | cut -d "\"" -f 2 |
uniq -c
   1 GET /portal.php?mode=addevent&date=2018-05-01 HTTP/1.1
   1 GET /blog/?attachment_id=2910 HTTP/1.1
   1 GET /blog/?attachment_id=2998&feed=rss2 HTTP/1.1
   1 GET /blog/?attachment_id=3156 HTTP/1.1
```

What is depicted in the exhibit?

A. Windows Event logs

B. Apache logs

C. IIS logs

D. UNIX-based syslog

ANSWER: B

Explanation:

Apache logs is correct. Apache HTTP Server commonly records web requests in the Common Log Format or Combined Log Format. These entries are structured as one line per HTTP transaction and typically contain the client IP address, an identity placeholder, an authenticated user placeholder, a timestamp in square brackets, the HTTP request line in quotation marks, the response status code, and the size of the response. Combined-format records can also include the referring URL and the user-agent string. This consistent request-oriented structure lets security analysts reconstruct client activity, identify requested resources, review HTTP methods, locate error responses, and investigate suspicious web traffic such as scanning, traversal attempts, or repeated requests from a source address. Apache documents that its access-log format is configurable, with the common and combined formats being standard examples used for web-server logging and analysis. The Apache HTTP Server logging documentation describes access logs and their fields at Apache HTTP Server Log Files, while the precise format directives and the predefined common and combined formats are documented at mod_log_config.

QUESTION NO: 55

Which two characteristics are associated with the TCP three-way handshake? (Choose two.)

A. A client sends a SYN packet to initiate a connection.

B. A server responds with a SYN-ACK packet.

C. A client sends a FIN packet before sending application data.

D. A server sends an RST packet to complete the handshake.

E. UDP sequence numbers are synchronized.

ANSWER: A B

Explanation:

A client sends a SYN packet to initiate a connection and **a server responds with a SYN-ACK packet** are correct. TCP uses the three-way handshake to establish connection state and synchronize sequence numbers before application data is exchanged. The initiating client sends a segment with the SYN flag set. If the destination service is reachable and accepting connections, the server replies with both the SYN and ACK flags set. The client then sends a final ACK, completing the handshake.

Security operations teams examine these flags in packet captures, firewall logs, and intrusion alerts to distinguish normal connection establishment from behaviors such as SYN flooding, scanning, and failed connection attempts. TCP connection establishment and flag behavior are specified in RFC 9293. Cisco provides additional security context in TCP SYN Flooding Attacks and Common Mitigations.