PassQueen
a.k.a dumpsqueen.com

Certified Ethical Hacker (CEH)

GAQM CEH-001

Version Demo

Total Demo Questions: 102

Total Premium Questions: 1028

Buy Premium PDF

https://passqueen.com

support@passqueen.com

Topic Break Down

Topic	No. of Questions
Topic 1, Introduction to Ethical Hacking	80
Topic 2, Footprinting and Reconnaissance	83
Topic 3, Scanning Networks	137
Topic 4, Enumeration	26
Topic 5, System Hacking	139
Topic 6, Malware Threats	55
Topic 7, Sniffing	74
Topic 8, Social Engineering	37
Topic 9, Denial of Service	40
Topic 10, Session Hijacking	32
Topic 11, Evading IDS, Firewalls, and Honeypots	97
Topic 12, Hacking Web Servers	21
Topic 13, Hacking Web Applications	44
Topic 14, SQL Injection	30
Topic 15, Hacking Wireless Networks	40
Topic 16, Hacking Mobile Platforms	1
Topic 17, Cloud Computing	1
Topic 18, Cryptography	82
Topic 19, Mix Questions	9
Total	1028

QUESTION NO: 1

Which of the following describes a component of Public Key Infrastructure (PKI) where a copy of a private key is stored to provide third-party access and to facilitate recovery operations?

A. Key registry

B. Recovery agent

C. Directory

D. Key escrow

ANSWER: D

Explanation:

Key escrow is the PKI arrangement in which a copy of a cryptographic private key is securely retained by a trusted third party, often called an escrow authority or key-recovery authority. Its purpose is to make encrypted information recoverable when the original key holder cannot provide the key, such as after key loss, employee departure, device failure, or a properly authorized legal or organizational access request. The escrowed material must be protected with strong technical and procedural safeguards because possession of a private key can permit decryption or other actions performed as that key holder. A sound key-escrow process therefore defines who may request recovery, requires authorization and auditing, and protects the stored key using secure storage and access controls. In PKI terminology, this directly matches the described storage of a private-key copy for third-party access and recovery operations. NIST discusses key recovery and key escrow as mechanisms supporting recovery of encrypted data and cryptographic keying material in its Recommendation for Key Management, Part 1. Additional PKI and key-management guidance is available from the UK National Cyber Security Centre cryptographic key-management guidance.

QUESTION NO: 2

Which of the following represent weak password? (Select 2 answers)

A. Passwords that contain letters, special characters, and numbers ExamplE. ap1$%##f@52

B. Passwords that contain only numbers ExamplE. 23698217

C. Passwords that contain only special characters ExamplE. &*#@!(%)

D. Passwords that contain letters and numbers ExamplE. meerdfget123

E. Passwords that contain only letters ExamplQWERTYKLRTY

F. Passwords that contain only special characters and numbers ExamplE. 123@$45

G. Passwords that contain only letters and special characters ExamplE. bob@&ba

H. Passwords that contain Uppercase/Lowercase from a dictionary list ExamplE. OrAnGe

ANSWER: B H

Explanation:

Passwords that contain only numbers ExamplE. 23698217 are weak because an eight-digit numeric password has a very limited search space: only 100 million possible combinations. Attackers can test such combinations efficiently in offline password-cracking scenarios, and numeric strings are often vulnerable to guessing when they resemble dates, phone fragments, or other personally meaningful values. A password restricted to digits also lacks the length and unpredictability that modern password guidance emphasizes.

Passwords that contain Uppercase/Lowercase from a dictionary list ExamplE. OrAnGe are also weak because changing capitalization does not make a common dictionary word meaningfully unpredictable. Password-cracking tools use wordlists and mangling rules that routinely test capitalization variations, including forms such as “OrAnGe.” Current guidance stresses screening chosen passwords against lists of commonly used, expected, or compromised values and prioritizing length and

resistance to guessing rather than relying on superficial character substitutions or case changes. See NIST's Digital Identity Guidelines: Authentication and Authenticator Management and CISA's Use Strong Passwords guidance.

QUESTION NO: 3

What is the correct PCAP filter to capture all TCP traffic going to or from host 192.168.0.125 on port 25?

A. tcp.src == 25 and ip.host == 192.168.0.125

B. host 192.168.0.125:25

C. port 25 and host 192.168.0.125

D. tcp.port == 25 and ip.host == 192.168.0.125

ANSWER: C

Explanation:

port 25 and host 192.168.0.125 is the correct libpcap capture-filter expression. In Berkeley Packet Filter (BPF) syntax, `port 25` matches packets whose source or destination transport-layer port is 25, and `host 192.168.0.125` matches packets whose source or destination IP address is 192.168.0.125. Combining these primitives with `and` limits the capture to traffic involving that host and SMTP port 25 in either direction. Because TCP is the protocol normally associated with SMTP on port 25, this expression captures the intended TCP communications in the usual exam context. If strict protocol qualification were required in operational use, the expression could be written as `tcp and port 25 and host 192.168.0.125`; however, the provided expression is the appropriate and valid capture-filter choice among the listed answers. Capture filters use the libpcap/tcpdump filter language, which Wireshark also uses when a filter is applied before capture begins. The tcpdump manual documents the `host` and `port` primitives and their bidirectional source-or-destination matching behavior: pcap-filter(7). Wireshark's distinction between capture and display filters is described in its Capture Filters guide.

QUESTION NO: 4

Which of the following wireless technologies can be detected by NetStumbler? (Select all that apply)

A. 802.11b

B. 802.11e

C. 802.11a

D. 802.11g

E. 802.11

ANSWER: A C D

Explanation:

NetStumbler can detect 802.11b, 802.11a, and 802.11g wireless networks when it is used with a compatible wireless adapter and driver. These are the IEEE 802.11 physical-layer variants that NetStumbler was designed to discover during wireless reconnaissance. It actively sends probe requests and records responses from access points, allowing an assessor to identify network names, channels, signal levels, access-point MAC addresses, and related network details. 802.11b and 802.11g operate in the 2.4 GHz band, while 802.11a operates in the 5 GHz band; therefore, practical detection of each technology also depends on the radio capabilities of the installed adapter. A dual-band adapter capable of the required legacy scanning functions is needed to discover networks across both bands. NetStumbler's published documentation identifies support for 802.11a, 802.11b, and 802.11g, making these the applicable selections for this question. Its historical role in wireless security assessments is as a network-discovery tool, helping ethical hackers inventory nearby Wi-Fi infrastructure before conducting authorized testing. See the NetStumbler project site and the IEEE's overview of 802.11 wireless standards.

QUESTION NO: 5

Several of your co-workers are having a discussion over the etc/passwd file. They are at odds over what types of encryption are used to secure Linux passwords.(Choose all that apply.

A. Linux passwords can be encrypted with MD5

B. Linux passwords can be encrypted with SHA

C. Linux passwords can be encrypted with DES

D. Linux passwords can be encrypted with Blowfish

E. Linux passwords are encrypted with asymmetric algrothims

ANSWER: A B C D

Explanation:

Linux password credentials are normally protected by one-way password-hashing schemes rather than reversible encryption. On contemporary systems, the password-hash field is generally stored in `/etc/shadow`; the corresponding field in `/etc/passwd` typically contains an `x` placeholder. The available hash schemes depend on the system's `crypt(3)` implementation and configuration, but Linux systems have supported traditional DES-based `crypt`, MD5-based crypt, SHA-family crypt schemes such as SHA-256 and SHA-512, and Blowfish-based bcrypt. Thus, "Linux passwords can be encrypted with MD5," "Linux passwords can be encrypted with SHA," "Linux passwords can be encrypted with DES," and "Linux passwords can be encrypted with Blowfish" identify algorithm families that may be used for Linux password hashes. In practice, administrators should choose a modern, deliberately slow password-hashing method supported by their distribution and policy; legacy DES and MD5 formats are retained largely for compatibility and are not appropriate choices for new deployments. The hash format is commonly identifiable from its prefix, which allows the authentication library to select the appropriate verification method. See the Linux crypt(3) manual page for supported hash methods and the shadow(5) manual page for password-field storage details.

QUESTION NO: 6

What is the correct command to run Netcat on a server using port 56 that spawns command shell when connected?

A. nc -port 56 -s cmd.exe

B. nc -p 56 -p -e shell.exe

C. nc -r 56 -c cmd.exe

D. nc -L 56 -t -e cmd.exe

E. nc -l -p 56 -e cmd.exe

ANSWER: E

Explanation:

`nc -l -p 56 -e cmd.exe` is the conventional Netcat command for a Windows listener that binds to port 56 and starts `cmd.exe` for a connecting client. The `-l` switch places Netcat in listening mode, meaning it waits for an inbound connection rather than initiating one. The `-p 56` switch specifies the local listening port. Finally, `-e cmd.exe` directs the traditional Netcat implementation to execute the Windows command interpreter and attach its standard input, output, and error streams to the network connection. This behavior is commonly discussed as a bind shell and should be used only in an authorized lab, assessment, or administrative recovery context because it exposes remote command execution to anyone able to connect to the listener. The documented traditional Netcat syntax identifies `-l` as listen mode, `-p` as the local port, and `-e` as execution of a program after connection. See the Netcat-traditional manual page. Modern Ncat provides similar controlled execution functionality through its documented exec options, though its command syntax differs from traditional Netcat.

QUESTION NO: 7

Because UDP is a connectionless protocol: (Select 2)

A. UDP recvfrom() and write() scanning will yield reliable results

B. It can only be used for Connect scans

C. It can only be used for SYN scans

D. There is no guarantee that the UDP packets will arrive at their destination

E. ICMP port unreachable messages may not be returned successfully

ANSWER: D E

Explanation:

“There is no guarantee that the UDP packets will arrive at their destination” is correct because UDP provides a connectionless, best-effort datagram service. It does not establish a session or include built-in delivery acknowledgments, retransmission, ordering, or recovery mechanisms. A UDP probe can therefore be lost in transit, filtered, or dropped because of congestion without the sender receiving confirmation that it reached the target. This is particularly important during UDP port scanning: lack of a response is not definitive evidence that a port is open or closed, so scanners may retransmit probes and classify many results as open|filtered.

“ICMP port unreachable messages may not be returned successfully” is also correct. A closed UDP port commonly causes the target host to send an ICMP Destination Unreachable/Port Unreachable response, but ICMP itself is not reliable and may be lost, rate-limited, filtered by a firewall, or suppressed by a host. Consequently, the absence of an ICMP error cannot be treated as proof that the UDP service is available. Reliable UDP scan interpretation requires accounting for both missing UDP probes and missing ICMP feedback. See the UDP specification (RFC 768) and Nmap’s UDP scan documentation.

QUESTION NO: 8

Which of the following are password cracking tools? (Choose three.)

A. BTCrack

B. John the Ripper

C. KerbCrack

D. Nikto

E. Cain and Abel

F. Havij

ANSWER: B C E

Explanation:

John the Ripper, **KerbCrack**, and **Cain and Abel** are password-cracking tools used in authorized password-security assessments. John the Ripper is a widely used password-auditing and recovery tool. It processes password hashes and supports dictionary, brute-force, and rule-based candidate generation to identify weak passwords. The project documentation describes its purpose and supported hash formats: Openwall John the Ripper.

KerbCrack is designed for cracking credentials associated with Kerberos authentication. In a controlled test, an assessor can use captured or obtained Kerberos material to perform offline guessing and determine whether service-account or user passwords are weak. This is especially relevant in enterprise environments using Active Directory and Kerberos.

Cain and Abel is a Windows-focused security and password-recovery suite. Its password-auditing features include collecting and recovering credentials through methods such as hash cracking and analysis of authentication traffic. Its historical tool

documentation is available from Oxid.it Cain & Abel. These tools should be used only with explicit authorization, because password auditing can expose account credentials and security weaknesses.

QUESTION NO: 9

Which of the following statements about a zone transfer correct?(Choose three.

A. A zone transfer is accomplished with the DNS

B. A zone transfer is accomplished with the nslookup service

C. A zone transfer passes all zone information that a DNS server maintains

D. A zone transfer passes all zone information that a nslookup server maintains

E. A zone transfer can be prevented by blocking all inbound TCP port 53 connections

F. Zone transfers cannot occur on the Internet

ANSWER: A C E

Explanation:

A zone transfer is a DNS function used to replicate authoritative zone data from a primary DNS server to an authorized secondary DNS server. Therefore, “A zone transfer is accomplished with the DNS” is correct: the transfer is performed through DNS zone-transfer mechanisms, principally the full-zone AXFR operation and, where supported, the incremental IXFR operation. A full AXFR transfer provides the receiving secondary server with the complete set of resource records in the zone, which supports the statement that “A zone transfer passes all zone information that a DNS server maintains.” This assumes the statement is referring to a full zone transfer rather than an incremental update.

“A zone transfer can be prevented by blocking all inbound TCP port 53 connections” is also correct from a network-control perspective. DNS zone transfers use TCP port 53, so preventing inbound TCP connections to that port prevents remote secondary servers from establishing the TCP session required to request AXFR or IXFR data. In practice, administrators normally combine firewall controls with DNS-server access-control lists, permitting transfers only to approved secondary servers rather than broadly exposing transfer capability. RFC 5936 specifies DNS zone-transfer operation and its TCP transport requirements: RFC 5936. Guidance on restricting transfers to authorized hosts is also available from ISC BIND 9 documentation.

QUESTION NO: 10

SNMP is a connectionless protocol that uses UDP instead of TCP packets (True or False)

A. true

B. false

ANSWER: A

Explanation:

The statement “true” is correct in the standard, foundational description of SNMP networking. SNMP commonly uses the User Datagram Protocol (UDP), which is connectionless: a management station can send a request without first establishing and maintaining a session-oriented connection with the managed device. SNMP requests are conventionally sent to UDP port 161, while unsolicited notifications such as traps are conventionally delivered to UDP port 162. This lightweight transport design helps SNMP perform routine monitoring and management functions with minimal protocol overhead, which is particularly useful for network devices that must expose operational information efficiently. UDP does not itself guarantee delivery, ordering, or retransmission, so SNMP managers commonly use timeouts and retries when a response is not received. The Internet Engineering Task Force’s SNMP transport mapping specifies UDP as the preferred transport mapping and documents the relevant well-known ports. Although modern SNMP implementations can support additional transport

mappings in some environments, UDP is the expected answer for this general protocol characteristic. See RFC 3417: Transport Mappings for SNMP and the IANA SNMP service-port registry.

QUESTION NO: 11

What flags are set in a X-MAS scan?(Choose all that apply.

A. SYN

B. ACK

C. FIN

D. PSH

E. RST

F. URG

ANSWER: C D F

Explanation:

A TCP X-MAS scan sets the FIN, PSH, and URG control flags in a probe packet. Its name comes from the packet appearing "lit up" with several flags enabled, analogous to a Christmas tree. The scan is a TCP flag-based reconnaissance technique designed to infer a target port's state from the response behavior defined by TCP implementations. In the usual interpretation, a closed port responds to this unexpected flag combination with a TCP reset, whereas an open port commonly gives no response. This behavior can make the scan useful when testing systems that do not filter these unusual packets, although results depend on the target operating system, TCP stack behavior, and intervening firewalls. Nmap documents the XMAS scan as a scan that turns on the FIN, PSH, and URG flags; therefore, FIN, PSH, and URG are the required selections. The packet does not use these flags to establish a normal TCP connection; instead, the distinctive combination is the defining characteristic of the X-MAS probe. See the Nmap Reference Guide: TCP NULL, FIN, and XMAS scans and the Nmap port-scanning techniques documentation for technical details.

QUESTION NO: 12

Wayne is the senior security analyst for his company. Wayne is examining some traffic logs on a server and came across some inconsistencies. Wayne finds some IP packets from a computer purporting to be on the internal network. The packets originate from 192.168.12.35 with a TTL of 15. The server replied to this computer and received a response from 192.168.12.35 with a TTL of 21. What can Wayne infer from this traffic log?

A. The initial traffic from 192.168.12.35 was being spoofed.

B. The traffic from 192.168.12.25 is from a Linux computer.

C. The TTL of 21 means that the client computer is on wireless.

D. The client computer at 192.168.12.35 is a zombie computer.

ANSWER: A

Explanation:

The initial traffic from 192.168.12.35 was being spoofed. TTL is an IP-header field intended to limit a packet's lifetime: every router that forwards an IPv4 packet decrements its TTL by at least one. Consequently, the TTL observed by a destination is influenced by both the sender's initial TTL and the number of routed hops traversed. In this scenario, the server is seeing traffic that claims to come from the same internal address but with materially inconsistent TTL behavior. A packet allegedly originating from an internal host arrives with a low remaining TTL of 15, while a subsequent response using that same claimed source address arrives with a remaining TTL of 21. In the exam's context, this inconsistency is evidence that the first packet did not actually follow the path expected for the asserted internal host and that its source address was forged. TTL-based analysis is therefore useful as a lightweight spoofing indicator when a purported source's expected network distance

is known. RFC 1812 specifies the router requirement to decrement TTL during forwarding, which is the protocol behavior underlying this inference: RFC 1812, Requirements for IPv4 Routers. The IPv4 TTL field is also defined in RFC 791, Internet Protocol.

QUESTION NO: 13

Which DNS records would be useful to an ethical hacker performing authorized reconnaissance of an organization's Internet-facing services? (Choose three.)

A. MX records

B. TXT records

C. CNAME records

D. ARP cache entries

E. TCP sequence numbers

ANSWER: A B C

Explanation:

MX records are useful because they identify mail exchangers responsible for receiving email for a domain. These hosts can reveal the organization's mail infrastructure or third-party email provider. **TXT records** can expose publicly published verification and security-policy information, including SPF and DMARC policy data, domain-verification values, and occasionally operational details. **CNAME records** are useful because they identify aliases that may point to hosted applications, content-delivery networks, cloud services, or other external infrastructure.

DNS reconnaissance can help build an authorized asset inventory, but records alone do not establish that a host is vulnerable. Findings should be validated within the scope and rules of engagement. DNS record types and their semantics are documented by IANA in the DNS Parameters registry, while SPF is specified in RFC 7208.

QUESTION NO: 14

Which of the following statement correctly defines ICMP Flood Attack? (Select 2 answers)

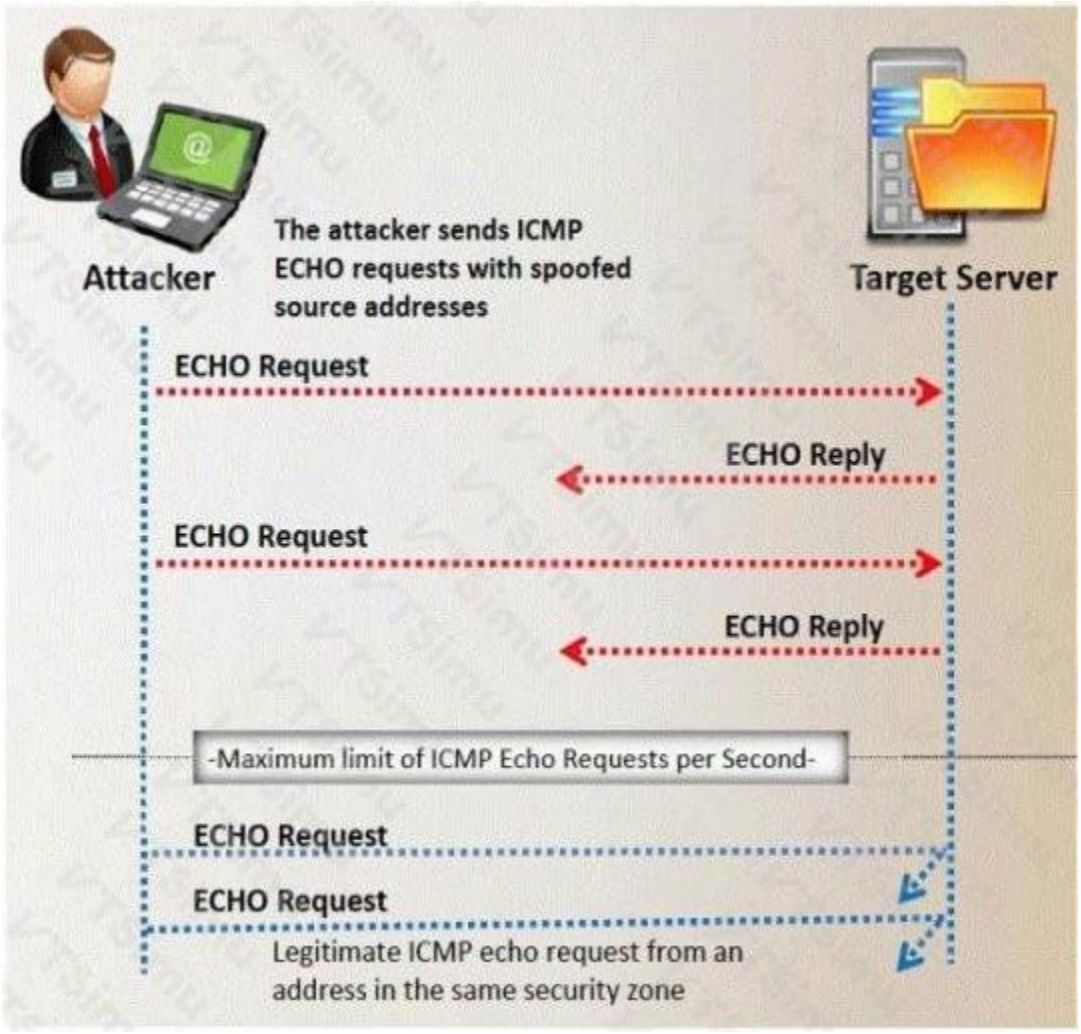

A. Bogus ECHO reply packets are flooded on the network spoofing the IP and MAC address

B. The ICMP packets signal the victim system to reply and the combination of traffic saturates the bandwidth of the victim's network

C. ECHO packets are flooded on the network saturating the bandwidth of the subnet causing denial of service

D. A DDoS ICMP flood attack occurs when the zombies send large volumes of ICMP_ECHO_REPLY packets to the victim system.

ANSWER: B D

Explanation:

An ICMP flood is a denial-of-service technique that overwhelms a target or the network path to it with a high volume of ICMP traffic. “The ICMP packets signal the victim system to reply and the combination of traffic saturates the bandwidth of the victim's network” correctly captures the amplification of traffic caused when ICMP Echo Requests prompt the victim to generate Echo Replies. At sufficient volume, both inbound requests and outbound replies consume bandwidth and processing resources, preventing legitimate communication.

“A DDoS ICMP flood attack occurs when the zombies send large volumes of ICMP_ECHO_REPLY packets to the victim system.” is also a valid description of a distributed ICMP flooding pattern: compromised hosts can coordinate to direct excessive ICMP Echo Reply traffic at a single victim. The defining characteristic is the volumetric exhaustion of available network capacity or host resources through sustained ICMP message traffic. In a distributed attack, many zombie systems make filtering and absorbing the aggregate traffic substantially more difficult. AWS describes ICMP floods as attacks using large volumes of ICMP Echo Request or Echo Reply packets to overwhelm a target: AWS: What is an ICMP flood?. Additional technical context on ICMP and Echo messaging is available in RFC 792.

QUESTION NO: 15

Which of the following are recommended controls for securing a wireless network that provides employee access? (Choose three.)

A. Use WPA3-Enterprise or WPA2-Enterprise with 802.1X authentication

B. Use strong encryption and disable WEP

C. Place guest wireless users on a separate network segment

D. Hide the SSID as the primary security control

E. Use MAC filtering as the only access-control mechanism

ANSWER: A B C

Explanation:

Use WPA3-Enterprise or WPA2-Enterprise with 802.1X authentication is correct because Enterprise Wi-Fi uses an authentication server, commonly RADIUS, to authenticate individual users or devices rather than relying on one shared wireless password. This improves accountability and allows access to be revoked for a single identity.

Use strong encryption and disable WEP is also correct. WEP has fundamental cryptographic weaknesses and should not be used to protect organizational wireless traffic. **Place guest wireless users on a separate network segment** limits access between untrusted guest devices and internal corporate systems, reducing the consequences of a compromised or malicious guest endpoint.

Hiding an SSID is not a meaningful security control because wireless clients and access points still disclose the network name during normal operation. MAC filtering can be bypassed by an attacker who observes and impersonates an authorized address, so it should not be relied on as the primary protection. See CISA guidance on Wi-Fi security and the Wi-Fi Alliance security overview.

QUESTION NO: 16

Jayden is a network administrator for her company. Jayden wants to prevent MAC spoofing on all the Cisco switches in the network. How can she accomplish this?

A. Jayden can use the commanD. ip binding set.

B. Jayden can use the commanD. no ip spoofing.

C. She should use the commanD. no dhcp spoofing.

D. She can use the commanD. ip dhcp snooping binding.

E. She should configure switch port security, for example: `switchport port-security mac-address sticky`.

ANSWER: E

Explanation:

She should configure switch port security, for example: `switchport port-security mac-address sticky`. Cisco switch port security controls which source MAC addresses are permitted on an access switch port. With sticky learning enabled, the switch learns the initially observed MAC address and converts it into a secure MAC address for that port. The port can then be limited to that authorized address, commonly by also configuring a maximum of one secure address and an appropriate violation action. If a frame arrives with an unauthorized source MAC address, the configured violation behavior can restrict, protect, or shut down the port. This makes port security the appropriate Layer 2 control for mitigating MAC-address spoofing attempts on switched access ports. In an enterprise deployment, Jayden should apply the policy consistently to user-facing access interfaces, save sticky MAC addresses as appropriate, and select a violation mode that matches operational requirements and incident-response procedures. Cisco documents that port security restricts input to an

interface by limiting and identifying MAC addresses allowed to access the port, while sticky secure MAC addressing dynamically learns secure addresses. See the Cisco port security configuration guide and the Cisco port-security overview.

QUESTION NO: 17

Which type of scan does NOT open a full TCP connection?

A. Stealth Scan

B. XMAS Scan

C. Null Scan

D. FIN Scan

ANSWER: A

Explanation:

Stealth Scan is correct because it commonly refers to a TCP SYN scan, also called a half-open scan. Rather than completing the TCP three-way handshake, the scanner sends a SYN packet and evaluates the target's reply. A SYN/ACK response indicates that the port is listening, while an RST response generally indicates that the port is closed. When a SYN/ACK is received, the scanner sends an RST instead of the final ACK required to establish a normal TCP session. Consequently, the connection is never fully opened or passed to the target application as an established session.

This approach is valuable during reconnaissance because it can identify TCP port states efficiently while avoiding the complete-connect behavior associated with a standard TCP connect scan. Nmap documents this method as SYN scan and notes that it is often called "half-open scanning" because it does not complete TCP connections. See the Nmap SYN Scan documentation and the TCP specification (RFC 793) for the TCP handshake and reset behavior.

QUESTION NO: 18

Take a look at the following attack on a Web Server using obstructed URL:

```
http://www.certifiedhacker.com/script.ext?
template=%2e%2e%2f%2e%2e%2f%2e%2e%2f%65%74%63%2f%70%61%73%73%77%64
This request is made up of:
%2e%2e%2f%2e%2e%2f%2e%2e%2f = ../../../
%65%74%63 = etc
%2f = /
%70%61%73%73%77%64 = passwd
```

How would you protect from these attacks?

A. Configure the Web Server to deny requests involving "hex encoded" characters

B. Create rules in IDS to alert on strange Unicode requests

C. Use SSL authentication on Web Servers

D. Enable Active Scripts Detection at the firewall and routers

ANSWER: B

Explanation:

Create rules in IDS to alert on strange Unicode requests is the best answer because an obstructed URL attack commonly uses alternate encodings, including malformed or overlong Unicode representations, to conceal path-traversal sequences from simple request inspection. Detection rules can normalize or identify suspicious encoded delimiters, traversal patterns, and nonstandard Unicode sequences in HTTP request targets, then generate an alert for investigation or trigger an associated blocking response when deployed inline. This provides visibility into attempted evasion even when the attacker changes the precise encoding used. Effective handling also depends on web-server patching and correct URL canonicalization, but, among the supplied responses, IDS detection focused on anomalous Unicode requests directly addresses the attack technique shown. OWASP explains that path traversal attacks can use encoded forms of directory separators and traversal strings to evade filters; see the OWASP Path Traversal guidance. The OWASP Input Validation Cheat Sheet also describes canonicalizing input before validation, an important principle for recognizing encoded malicious request paths consistently.

QUESTION NO: 19

Which of the following are variants of mandatory access control mechanisms? (Choose two.)

A. Two factor authentication

B. Acceptable use policy

C. Username / password

D. User education program

E. Sign in register

F. Rule-based access control

G. Lattice-based access control

ANSWER: F G

Explanation:

Mandatory access control is an authorization model in which a central authority defines and enforces access decisions. Individual users and resource owners cannot freely change permissions at their discretion. The two recognized variants are rule-based access control and lattice-based access control. Rule-based access control applies system-wide rules, often based on security labels, clearances, classifications, time, or other policy conditions. A rule engine or operating-system policy determines whether the requested operation is permitted. Lattice-based access control organizes security labels into an ordered structure, or lattice, and makes decisions by comparing a subject's clearance with an object's classification and category set. This is the model commonly associated with multilevel-security environments, where information flow must be controlled between sensitivity levels. Both approaches place policy enforcement under centralized administrative control and use labels or formal rules rather than allowing resource owners to delegate permissions. NIST defines mandatory access control as access control based on information sensitivity and formal authorization, and its glossary also identifies lattice-based control as a model using security labels and dominance relationships. See the NIST definition of mandatory access control and the NIST definition of lattice-based access control.

QUESTION NO: 20

Stephanie works as senior security analyst for a manufacturing company in Detroit. Stephanie manages network security throughout the organization. Her colleague Jason told her in confidence that he was able to see confidential corporate information posted on the external website http://www.jeansclothesman.com. He tries random URLs on the company's website and finds confidential information leaked over the web. Jason says this happened about a month ago. Stephanie visits the said URLs, but she finds nothing. She is very concerned about this, since someone should be held accountable if there was sensitive information posted on the website.

Where can Stephanie go to see past versions and pages of a website?

A. She should go to the web page Samspade.org to see web pages that might no longer be on the website

B. If Stephanie navigates to Search.com; she will see old versions of the company website

C. Stephanie can go to Archive.org to see past versions of the company website

D. AddressPast.com would have any web pages that are no longer hosted on the company's website

ANSWER: C

Explanation:

Stephanie can go to Archive.org to see past versions of the company website. The Internet Archive's Wayback Machine preserves historical captures of publicly accessible web pages and lets a user enter a URL to review snapshots taken on specific dates. This makes it a valuable source for investigating a report that content was available in the past but is no longer present on the live site. By selecting captures around the time Jason observed the alleged disclosure, Stephanie may be able to determine whether the relevant path, page content, directory listings, linked files, or other publicly exposed material was recorded.

In a security investigation, archived copies can help establish a timeline of external exposure and identify artifacts that should be correlated with authoritative evidence, such as web-server access logs, application deployment records, CDN logs, backups, and incident tickets. Archived material is not a complete forensic record: pages may not have been crawled, may have been excluded, or may not render exactly as they did originally. Nevertheless, it is an appropriate public resource for checking historical versions of a website and supporting preliminary validation of the reported leak. The Wayback Machine is available through web.archive.org; its collection and access information are maintained by the Internet Archive.

QUESTION NO: 21

Windows LAN Manager (LM) hashes are known to be weak. Which of the following are known weaknesses of LM? (Choose three)

A. Converts passwords to uppercase.

B. Hashes are sent in clear text over the network.

C. Makes use of only 32 bit encryption.

D. Effective length is 7 characters.

ANSWER: A B D

Explanation:

LM hashing has structural properties that substantially reduce password resistance. It converts the password to uppercase before processing, so passwords that differ only by letter case produce the same LM hash. This discards an important portion of password entropy and makes password guessing easier. LM also processes a password as two independent seven-character portions after padding or truncating it to fourteen bytes. Consequently, an attacker can crack each portion separately rather than having to search the full password as one value; this is commonly summarized as an effective seven-character cracking problem per half. In legacy LM-based authentication scenarios, the hash-derived material is exposed through network authentication exchanges rather than the plaintext password being transmitted. Capturing these exchanges can enable offline password-cracking attempts, especially given LM's uppercase conversion and independent seven-character halves. Microsoft documents that LM hashes are weak and recommends disabling their storage and use where possible, while credential-security guidance explains the risks of legacy challenge-response authentication. See Microsoft's guidance on preventing LM hash storage and MITRE ATT&CK: Security Account Manager.

QUESTION NO: 22

Attackers target HINFO record types stored on a DNS server to enumerate information. These are information records and potential source for reconnaissance. A network administrator has the option of entering host information specifically the CPU type and operating system when creating a new DNS record. An attacker can extract this type of information easily from a DNS server.

Which of the following commands extracts the HINFO record?

A
```
c:> nslookup
> Set type=hinfo
> certhack-srv
Server: dns.certifiedhacker.com
Address: 10.0.0.4
sales.certifiedhacker.com      CPU = Intel Quad Chip OS=Linux 2.8
dns.certifiedhacker.com        Internet address = 10.0.0.56
```

B
```
c:> nslookup
> Set dns=hinfo
> certhack-srv
Server: dns.certifiedhacker.com
IP: 10.0.0.4
sales.certifiedhacker.com      CPU = Intel Quad Chip OS=Linux 2.8
dns.certifiedhacker.com        Internet address = 10.0.0.56
```

C
```
c:> nslookup
> Set record=hinfo
> certhack-srv
host: dns.certifiedhacker.com
Address: 10.0.0.4
sales.certifiedhacker.com      CPU = Intel Quad Chip OS=Linux 2.8
dns.certifiedhacker.com        Internet address = 10.0.0.56
```

D
```
c:> nslookup
> Configure type=hinfo
> certhack-srv
Host: dns.certifiedhacker.com
IP: 10.0.0.4
sales.certifiedhacker.com      CPU = Intel Quad Chip OS=Linux 2.8
dns.certifiedhacker.com Internet address = 10.0.0.56
```

A. Option A

B. Option B

C. Option C

D. Option D

ANSWER: A

Explanation:

The correct command is the one that explicitly requests the `HINFO` DNS resource-record type for the target host, such as `nslookup -type=HINFO target.example` (or the equivalent `dig target.example HINFO`). DNS clients do not infer HINFO data from a normal hostname lookup; the query must specify the desired record type. When an authoritative DNS server has published an HINFO record and permits the query, its response contains the host-information fields defined for this record: a character string identifying the CPU type and a character string identifying the operating system. This can be useful in authorized inventory work, but it may also expose system-profile information during reconnaissance, which is why administrators generally avoid publishing unnecessary HINFO data on public zones. The DNS protocol defines HINFO as a standard resource record with CPU and OS character-string fields in RFC 1035. DNS query tools such as `nslookup` support selecting the requested record type through the `-type` parameter. See RFC 1035, section 3.3.2 for the HINFO record format and Microsoft nslookup documentation for query-type usage.

QUESTION NO: 23

Name two software tools used for OS guessing? (Choose two.

A. Nmap

B. Snadboy

C. Queso

D. UserInfo

E. NetBus

ANSWER: A C

Explanation:

Nmap and **Queso** are tools associated with remote operating-system identification, often called OS fingerprinting or OS guessing. Nmap performs active OS detection by sending carefully selected TCP/IP probes to a target and evaluating characteristics of the responses, such as TCP sequence behavior, IP and TCP header values, and responses to uncommon flag combinations. It compares the observed behavior with its maintained fingerprint database and reports the most likely operating system, device type, and related confidence information. This makes it a standard reconnaissance tool when used with proper authorization.

Queso is a classic TCP/IP stack fingerprinting utility. Like Nmap's OS-detection capability, it identifies implementation-specific network-stack behavior that can distinguish operating systems and versions. Its historical role in security training and older penetration-testing material is specifically as an OS-guessing/fingerprinting program. Both tools illustrate the core fingerprinting principle: network protocol standards allow implementation choices, and those subtle choices can produce a recognizable signature. For current Nmap OS-detection usage and requirements, consult the Nmap OS Detection documentation. Background on TCP/IP fingerprinting is also available in the Nmap Network Scanning guide.

QUESTION NO: 24

Which of the following are appropriate controls for protecting password hashes stored by an application? (Choose three.)

A. Use a modern password-hashing function such as Argon2id, bcrypt, PBKDF2, or scrypt

B. Use a unique random salt for each password

C. Configure an appropriate work factor

D. Store passwords using reversible encryption for easier recovery

E. Use the same static salt for every user account

ANSWER: A B C

Explanation:

Using a modern password-hashing function, assigning a unique salt to each password, and applying an appropriate work factor are correct controls. Passwords should not be stored using reversible encryption or fast general-purpose hashes alone. Purpose-built password-hashing functions such as Argon2id, bcrypt, PBKDF2, or scrypt are designed to make large-scale offline guessing more expensive.

A unique random salt prevents identical passwords from producing identical stored values and reduces the effectiveness of precomputed rainbow tables. A configurable work factor increases the computing cost of each password-verification attempt and should be selected based on the organization’s environment and periodically reviewed as hardware capabilities change. OWASP provides password-storage recommendations in its Password Storage Cheat Sheet. NIST also addresses password-verifier storage in SP 800-63B.

QUESTION NO: 25

Which of the following network attacks relies on sending an abnormally large packet size that exceeds TCP/IP specifications?

A. Ping of death

B. SYN flooding

C. TCP hijacking

D. Smurf attack

ANSWER: A

Explanation:

Ping of death is correct because it exploits the maximum size limit for an IPv4 packet. The IPv4 Total Length field is 16 bits, so a fully reassembled IP datagram cannot exceed 65,535 bytes, including its header. In a ping-of-death attack, an attacker sends malformed, fragmented ICMP echo-request traffic which, after the target attempts IP-fragment reassembly, exceeds that permitted size. Historically, vulnerable TCP/IP implementations could mishandle this oversized reassembly operation, causing memory corruption, crashes, freezes, or denial of service. Although current operating systems and network stacks have long been patched against the original vulnerability, the attack remains an important example of a denial-of-service technique based on malformed oversized IP packets. RFC 791 specifies the IPv4 datagram format and defines Total Length as a 16-bit field, establishing the relevant maximum datagram size. The CERT Coordination Center also documents the historical vulnerability involving oversized IP packets and fragmentation handling. See RFC 791: Internet Protocol and CERT Vulnerability Note VU#638652.

QUESTION NO: 26

Which of the following type of scanning utilizes automated process of proactively identifying vulnerabilities of the computing systems present on a network?

A. Port Scanning

B. Single Scanning

C. External Scanning

D. Vulnerability Scanning

ANSWER: D

Explanation:

Vulnerability Scanning is the automated, systematic process of identifying known security weaknesses in hosts, operating systems, applications, services, and network devices. A vulnerability scanner assesses discovered assets against a database of known vulnerabilities, insecure configurations, missing patches, and exposed services, then reports findings that can be prioritized and remediated. This makes it a proactive security practice: organizations run scans regularly to identify and address weaknesses before they can be exploited. In an ethical-hacking workflow, vulnerability scanning commonly follows host and service discovery and provides a more detailed assessment of potential security exposures across the environment. Scan results should be validated where appropriate, because factors such as incomplete credentialed access, environmental context, and detection limitations can affect findings. The National Institute of Standards and Technology describes vulnerability monitoring and scanning as activities for identifying vulnerabilities and configuration weaknesses, while the Cybersecurity and Infrastructure Security Agency recommends continuous vulnerability management to reduce organizational risk. See NIST SP 800-115: Technical Guide to Information Security Testing and Assessment and CISA: Vulnerabilities and Exploits.

QUESTION NO: 27

Doug is conducting a port scan of a target network. He knows that his client target network has a web server and that there is a mail server also which is up and running. Doug has been sweeping the network but has not been able to elicit any response from the remote target. Which of the following could be the most likely cause behind this lack of response? Select 4.

A. UDP is filtered by a gateway

B. The packet TTL value is too low and cannot reach the target

C. The host might be down

D. The destination network might be down

E. The TCP windows size does not match

F. ICMP is filtered by a gateway

ANSWER: A B C F

Explanation:

“UDP is filtered by a gateway,” “The packet TTL value is too low and cannot reach the target,” “The host might be down,” and “ICMP is filtered by a gateway” are valid causes of an apparently silent target during reconnaissance. A security gateway can silently discard UDP probes rather than return an application response or an ICMP unreachable message. Similarly, filtering ICMP can suppress diagnostic responses that a scanner may use for host discovery, path assessment, or interpreting unreachable conditions. This can make reachable systems appear unresponsive even when services are operating behind the filtering device. A packet’s Time To Live is decremented at every routed hop; if it expires before reaching the target, the probe cannot arrive and therefore cannot produce a target response. This behavior is defined in RFC 1812. Finally, a specific scanned host can be unavailable, powered off, disconnected, or otherwise unable to respond, even where other systems such as web and mail servers are known to exist on the broader target network. Nmap also documents that firewalls and packet filtering commonly cause probes to be dropped or produce ambiguous scan results: Nmap Network Scanning Reference Guide.

QUESTION NO: 28

An Nmap scan shows the following open ports, and nmap also reports that the OS guessing results to match too many signatures hence it cannot reliably be identified:

21 ftp

23 telnet

80 http

443 https

What does this suggest?

A. This is a Windows Domain Controller

B. The host is not firewalled

C. The host is not a Linux or Solaris system

D. The host is not properly patched

E. The scan alone is insufficient to reliably determine the host's OS, firewall configuration, or patch status.

ANSWER: E

Explanation:

The scan alone is insufficient to reliably determine the host's OS, firewall configuration, or patch status. The listed results establish only that FTP, Telnet, HTTP, and HTTPS services are reachable on their respective ports from the scanning system at the time of the scan. Nmap's statement that too many OS signatures match means its TCP/IP fingerprinting did not produce enough distinctive evidence for a reliable operating-system identification. OS detection is probabilistic: factors such as network devices, packet normalization, virtualized networking, service configurations, and partial filtering can make multiple fingerprints appear plausible. Open service ports also do not provide reliable evidence of a specific platform, whether a firewall exists, or the host's patch level. Those conclusions require further evidence. A careful next step is service and version detection, such as Nmap's `-sV` scan, followed by validation of discovered versions and authorized vulnerability

assessment. Nmap also supports more detailed OS-detection controls, but results should be treated as estimates rather than definitive identification when the fingerprint is ambiguous. See the Nmap OS Detection documentation and the Nmap Version Detection documentation.

QUESTION NO: 29

SNMP is a protocol used to query hosts, servers, and devices about performance or health status data. This protocol has long been used by hackers to gather great amount of information about remote hosts.

Which of the following features makes this possible? (Choose two)

A. It used TCP as the underlying protocol.

B. It uses community string that is transmitted in clear text.

C. It is susceptible to sniffing.

D. It is used by all network devices on the market.

ANSWER: B C

Explanation:

“It uses community string that is transmitted in clear text.” is correct because SNMP versions 1 and 2c use a community string as a shared credential, and it is commonly sent without encryption. If an attacker captures SNMP traffic or guesses a weak/default community string, the credential can permit access to Management Information Base data. That data may reveal valuable reconnaissance details, including interface information, routing-related data, system descriptions, software details, and operational status.

“It is susceptible to sniffing.” is also correct because the unprotected transmission of SNMPv1/v2c community strings and management data allows a party able to observe relevant network traffic to capture that information. This is especially significant where SNMP requests traverse shared, exposed, or insufficiently segmented network paths. Secure Network Management Protocol version 3 addresses these risks by supporting authentication and privacy (encryption), so organizations should use SNMPv3 where possible, restrict SNMP management access with network controls, and replace default community strings. See the SNMP Architecture RFC 3411 and the CISA guidance on SNMP vulnerabilities.

QUESTION NO: 30

The SYN flood attack sends TCP connections requests faster than a machine can process them.

Attacker creates a random source address for each packet

SYN flag set in each packet is a request to open a new connection to the server from the spoofed IP address

Victim responds to spoofed IP address, then waits for confirmation that never arrives (timeout wait is about 3 minutes)

Victim ' s connection table fills up waiting for replies and ignores new connections

Legitimate users are ignored and will not be able to access the server

How do you protect your network against SYN Flood attacks?

A. SYN cookies. Instead of allocating a record, send a SYN-ACK with a carefully constructed sequence number generated as a hash of the clients IP address, port number, and other information. When the client responds with a normal ACK, that special sequence number will be included, which the server then verifies. Thus, the server first allocates memory on the third packet of the handshake, not the first.

B. RST cookies - The server sends a wrong SYN/ACK back to the client. The client should then generate a RST packet telling the server that something is wrong. At this point, the server knows the client is valid and will now accept incoming connections from that client normally

C. Check the incoming packet ' s IP address with the SPAM database on the Internet and enable the filter using ACLs at the Firewall

D. Stack Tweaking. TCP stacks can be tweaked in order to reduce the effect of SYN floods. Reduce the timeout before a stack frees up the memory allocated for a connection

E. Micro Blocks. Instead of allocating a complete connection, simply allocate a micro record of 16-bytes for the incoming SYN object

ANSWER: A B D E

Explanation:

SYN cookies are an established defense because they defer allocation of per-connection state until the endpoint returns a valid final ACK. The server encodes sufficient validation information into the SYN-ACK sequence number, allowing it to validate the completed handshake without retaining a normal half-open connection record for every received SYN. This directly protects the backlog from exhaustion during a spoofed-source flood. RST cookies provide another validation-oriented approach: a deliberately invalid SYN-ACK causes a genuine host that receives it to send a reset, enabling the defender to distinguish a reachable source before allowing normal connection handling. Stack Tweaking is also appropriate because reducing the half-open connection timeout and tuning backlog-related TCP parameters limits how long incomplete handshakes consume resources. Micro Blocks are a resource-conservation technique: retaining only minimal state for an initial SYN substantially reduces memory pressure compared with allocating a full connection-control structure. These controls address the defining condition of a SYN flood—large numbers of incomplete TCP handshakes—by validating peers before committing normal state, reducing retained state, or reclaiming it faster. RFC 4987 describes SYN flooding and common mitigations, including SYN cookies and backlog/resource-management measures: RFC 4987. Linux also documents SYN-cookie behavior and its role when the SYN backlog is overwhelmed: Linux IP sysctl documentation.

QUESTION NO: 31

Which one of the following network attacks takes advantages of weaknesses in the fragment reassembly functionality of the TCP/IP protocol stack?

A. Teardrop

B. Smurf

C. Ping of Death

D. SYN flood

E. SNMP Attack

ANSWER: A

Explanation:

Teardrop is correct because it exploits flaws in how vulnerable TCP/IP implementations reassemble fragmented IP packets. In normal IP fragmentation, each fragment contains an offset indicating its position within the original datagram, and the receiving host uses those offsets to reconstruct the packet. A teardrop attack sends deliberately malformed fragments, commonly with overlapping or otherwise inconsistent offset and length values. Older or unpatched operating systems could mishandle these invalid fragment sequences during reassembly, potentially causing instability, a system crash, or a reboot. The attack therefore specifically targets the fragment-reassembly functionality in the IP layer of the TCP/IP stack rather than an application-level service. This behavior is documented in MITRE's CAPEC attack pattern for teardrop-style packet fragmentation attacks, which describes the use of overlapping IP fragments to exploit reassembly weaknesses. Modern systems generally include stronger validation and are less susceptible, but the attack remains an important historical example of denial-of-service risks arising from improper network-protocol parsing. See MITRE CAPEC-23: File and Directory Information Obfuscation and RFC 791: Internet Protocol for IP fragmentation fields and processing context.

QUESTION NO: 32

Which of the following is a preventive control?

A. Smart card authentication

B. Security policy

C. Audit trail

D. Continuity of operations plan

ANSWER: A

Explanation:

Smart card authentication is a preventive control because it acts before access is granted, verifying that a person possesses an authorized credential and, commonly, knows the associated PIN. A smart card contains cryptographic credentials that can be used to strongly authenticate a user to a workstation, application, or protected system. By requiring successful authentication before allowing entry, it prevents unauthorized users from reaching resources in the first place. This is the essential purpose of a preventive control: reducing the likelihood of a security incident through an enforced safeguard at the point of attempted access.

Smart-card-based authentication is especially valuable because the credential is a physical token and can support certificate-based cryptography, making it more resistant to credential replay and password-only compromise when properly implemented. NIST identifies identification and authentication as a control family and describes multifactor authentication requirements for system access. Its digital identity guidance also addresses the use of cryptographic authenticators, including hardware-based authenticators. See NIST SP 800-53, Revision 5 and NIST SP 800-63B: Digital Identity Guidelines.

QUESTION NO: 33

What are the default passwords used by SNMP? (Choose two.)

A. Password

B. SA

C. Private

D. Administrator

E. Public

F. Blank

ANSWER: C E

Explanation:

“Private” and “Public” are the conventional default SNMP community strings associated with SNMPv1 and SNMPv2c deployments. More precisely, these values are community strings rather than passwords: they are transmitted with requests and act as a shared, weak form of access control. “Public” is traditionally configured for read-only access, allowing a management station to retrieve device information such as interface status, system details, and counters. “Private” is traditionally associated with read-write access, permitting authorized management functions where that capability is enabled. These well-known defaults are significant during an ethical security assessment because devices that retain them may expose management information or accept unauthorized configuration changes. A tester should identify SNMP services only within the approved scope and verify whether default community strings remain enabled, then report the finding so the organization can replace them with strong unique values or, preferably, use SNMPv3. SNMPv3 provides authentication and privacy features that address the fundamental weaknesses of plaintext community strings. Cisco’s SNMP guidance describes community strings and their access-control role: Cisco: SNMP Community Strings. The IETF SNMP architecture specification also distinguishes the community-based model from SNMPv3 security mechanisms: RFC 3411.

QUESTION NO: 34

Bob reads an article about how insecure wireless networks can be. He gets approval from his management to implement a policy of not allowing any wireless devices on the network. What other steps does Bob have to take in order to successfully implement this? (Select 2 answer.)

A. Train users in the new policy.

B. Disable all wireless protocols at the firewall.

C. Disable SNMP on the network so that wireless devices cannot be configured.

D. Continuously survey the area for wireless devices.

ANSWER: A D

Explanation:

“Train users in the new policy.” is essential because a prohibition on wireless devices depends on employees, contractors, and other authorized users understanding what is forbidden, why the restriction exists, and how to comply. Training should cover organization-issued and personally owned wireless equipment, such as Wi-Fi access points, wireless routers, hotspots, and other radio-capable devices, along with the reporting process for suspected unauthorized equipment. Clear communication makes the policy actionable and supports consistent enforcement.

“Continuously survey the area for wireless devices.” is also necessary because policy alone cannot identify rogue or accidentally deployed wireless equipment. Regular or continuous wireless monitoring helps detect unauthorized access points and clients, locate their physical source, validate that the wireless prohibition remains effective, and enable prompt incident handling. NIST wireless-security guidance emphasizes maintaining an inventory, monitoring for rogue devices, and using detection capabilities as part of wireless security management. These administrative and technical measures work together: user training reduces the likelihood of unauthorized deployment, while ongoing surveys provide verification and detection when violations occur. See NIST SP 800-153, Guidelines for Securing Wireless Local Area Networks and NIST SP 800-97, Establishing Wireless Robust Security Networks.

QUESTION NO: 35

During a forensic investigation, which practices help preserve the integrity and admissibility of collected digital evidence? (Choose three.)

A. Document the chain of custody

B. Create a forensic image and verify it with cryptographic hashes

C. Use write blockers when acquiring storage media

D. Perform all analysis directly on the original disk

E. Delete temporary files before imaging the system

ANSWER: A B C

Explanation:

Document the chain of custody is correct because investigators must record who collected, handled, transferred, stored, and examined evidence, along with relevant dates and times. This provides accountability and helps demonstrate that the evidence was controlled throughout the investigation.

Create a forensic image and verify it with cryptographic hashes is also correct. A bit-for-bit forensic image preserves the original media for analysis, while hash values can demonstrate whether the image or original evidence has changed. **Use write blockers when acquiring storage media** helps prevent the forensic workstation from writing to the evidence drive and altering its contents during acquisition.

Investigators should avoid conducting analysis directly on original media when a verified forensic copy can be used. NIST describes evidence preservation, acquisition, hashing, and chain-of-custody practices in NIST SP 800-86.

QUESTION NO: 36

Ann would like to perform a reliable scan against a remote target. She is not concerned about being stealth at this point.

Which of the following type of scans would be the most accurate and reliable option?

A. A half-scan

B. A UDP scan

C. A TCP Connect scan

D. A FIN scan

ANSWER: C

Explanation:

A TCP Connect scan is the most accurate and reliable choice when stealth is not a concern because it asks the operating system to establish a normal, full TCP connection to each target port. A successful connection completes the standard three-way TCP handshake: SYN, SYN/ACK, and ACK. This provides clear confirmation that the target is reachable, the TCP service is listening on that port, and a connection can be established from the scanning host. After the connection result is recorded, the scanner closes it, typically with a reset or a normal close sequence.

Because it uses the host operating system's normal networking interface rather than requiring raw-packet privileges, a TCP Connect scan is broadly compatible and dependable in environments where lower-level packet crafting may not be available. It is also useful for remote testing because its outcome reflects whether a real client connection can be made across the network path, including intervening filtering and routing conditions. Nmap documents this technique as the TCP connect scan (`-sT`) and describes its use of the operating system's `connect` system call. The TCP handshake behavior is defined in the TCP specification. See Nmap: TCP Connect Scan and RFC 793: Transmission Control Protocol.

QUESTION NO: 37

Bob, an Administrator at XYZ was furious when he discovered that his buddy Trent, has launched a session hijack attack against his network, and sniffed on his communication, including administrative tasks suck as configuring routers, firewalls, IDS, via Telnet.

Bob, being an unhappy administrator, seeks your help to assist him in ensuring that attackers such as Trent will not be able to launch a session hijack in XYZ.

Based on the above scenario, please choose which would be your corrective measurement actions. (Choose two)

A. Use encrypted protocols, like those found in the OpenSSH suite.

B. Implement FAT32 filesystem for faster indexing and improved performance.

C. Configure the appropriate spoof rules on gateways (internal and external).

D. Monitor for CRP caches, by using IDS products.

ANSWER: A C

Explanation:

Use encrypted protocols, like those found in the OpenSSH suite, is a key corrective measure because Telnet transmits credentials, commands, and session traffic in cleartext. SSH replaces Telnet with an encrypted, integrity-protected channel and authenticates the remote server, substantially reducing an attacker's ability to sniff administrative credentials or inject

traffic into an observed administrative session. Administrative access to routers, firewalls, and intrusion-detection systems should therefore use SSH and other secure management protocols rather than plaintext remote-access services.

Configure the appropriate spoof rules on gateways (internal and external) is also an important network-layer control. Anti-spoofing rules validate that packets arriving at an interface have source addresses that are plausible for that interface and network direction. Applying ingress and egress filtering at organizational boundaries helps prevent forged-source traffic from entering or leaving the network, reducing opportunities for attacks that rely on IP spoofing to interfere with or impersonate session traffic. These measures work together: encrypted management sessions protect the content and authenticity of administrator communications, while gateway anti-spoofing controls limit malicious forged traffic at the network perimeter. See RFC 4251: The Secure Shell Protocol Architecture and RFC 2827: Network Ingress Filtering.

QUESTION NO: 38

You want to hide a secret.txt document inside c:\windows\system32\tcpip.dll kernel library using ADS streams. How will you accomplish this?

A. copy secret.txt c:\windows\system32\tcpip.dll kernel>secret.txt

B. copy secret.txt c:\windows\system32\tcpip.dll:secret.txt

C. copy secret.txt c:\windows\system32\tcpip.dll |secret.txt

D. copy secret.txt >< c:\windows\system32\tcpip.dll kernel secret.txt

ANSWER: B

Explanation:

`copy secret.txt c:\windows\system32\tcpip.dll:secret.txt` is correct because it uses the NTFS alternate data stream naming convention: a normal host-file path followed by a colon and a stream name. In this case, `c:\windows\system32\tcpip.dll` is the host file and `secret.txt` is the named stream attached to that file. NTFS supports multiple streams associated with one file; the ordinary file contents are held in the unnamed default stream, while a named stream is addressed with the `filename:streamname` form. The Windows `copy` command can therefore copy the contents of the source file into the named stream when the target is specified in this format. The operation requires an NTFS volume and permissions sufficient to write to the target location; a protected system directory normally requires elevated privileges. Microsoft documents the colon-based syntax for accessing named data streams and notes that applications can create, read, and write them through standard file operations. See Microsoft: File Streams and Microsoft: copy command.

QUESTION NO: 39

The intrusion detection system at a software development company suddenly generates multiple alerts regarding attacks against the company's external webserver, VPN concentrator, and DNS servers. What should the security team do to determine which alerts to check first?

A. Investigate based on the maintenance schedule of the affected systems.

B. Investigate based on the service level agreements of the systems.

C. Investigate based on the potential effect of the incident.

D. Investigate based on the order that the alerts arrived in.

ANSWER: C

Explanation:

Investigate based on the potential effect of the incident. Alert triage should prioritize events according to their likely business and security impact rather than operational timing or alert arrival sequence. In this situation, the team should rapidly assess which suspected attack could most seriously affect confidentiality, integrity, availability, public-facing services, critical business processes, or access to internal systems. For example, a credible compromise that could provide remote access to

the corporate network, interrupt name resolution, expose sensitive customer or source-code data, or take down a key externally available service warrants prompt investigation because its consequences may be severe and widespread.

This is an impact-based prioritization decision: analysts use available alert context, asset criticality, exposure, threat indicators, and plausible consequences to establish the order of response. That approach helps ensure limited incident-response resources are directed first to the events where delay could cause the greatest harm. NIST incident-handling guidance describes prioritization using factors including functional impact, information impact, and recoverability, and recommends documenting criteria so incidents are handled consistently. See NIST SP 800-61 Rev. 2, Computer Security Incident Handling Guide and CISA Incident Response resources.

QUESTION NO: 40

Which tool would be used to collect wireless packet data?

A. NetStumbler

B. John the Ripper

C. Nessus

D. Netcat

E. Kismet

ANSWER: E

Explanation:

Kismet is the appropriate tool for collecting wireless packet data because it is designed for wireless network detection, monitoring, and packet capture. When used with a compatible wireless adapter in monitor mode, Kismet passively receives 802.11 management, control, and data frames rather than associating with an access point. It can record the observed traffic and associated metadata, including detected access points, clients, channels, encryption details, and packet logs, for later examination in packet-analysis software. This passive collection capability is especially useful during an authorized wireless security assessment because it enables the tester to survey wireless activity and preserve evidence without actively transmitting probes to the target network. Kismet supports multiple capture sources and logging formats, including packet capture output, making it suitable for gathering the raw wireless traffic required for subsequent analysis. The Kismet documentation describes it as a wireless network and device detector, sniffer, and intrusion-detection framework, and its capture-source documentation explains how it receives wireless traffic from supported interfaces. See the Kismet introduction and Kismet Linux Wi-Fi capture-source documentation.

QUESTION NO: 41

As a securing consultant, what are some of the things you would recommend to a company to ensure DNS security? Select the best answers.

A. Use the same machines for DNS and other applications

B. Harden DNS servers

C. Use split-horizon operation for DNS servers

D. Restrict Zone transfers

E. Have subnet diversity between DNS servers

ANSWER: B C D E

Explanation:

Hardening DNS servers, using split-horizon DNS, restricting zone transfers, and maintaining subnet diversity between DNS servers are complementary DNS-security controls. DNS server hardening reduces the exposed attack surface by keeping

software current, disabling unnecessary services, applying secure configuration baselines, and limiting administrative access. CISA emphasizes the importance of protecting DNS infrastructure and adopting resilient, securely configured DNS services in its CISA DNS infrastructure guidance.

Split-horizon operation presents different DNS views to internal and external clients. This supports separation of private internal names and addresses from the public namespace, reducing unnecessary disclosure of internal infrastructure. Restricting zone transfers is essential because authoritative zone data can reveal hosts and service records; transfers should be permitted only to explicitly authorized secondary servers, typically with authenticated transfer mechanisms. The DNS Zone Transfer Protocol specification (RFC 5936) describes the zone-transfer mechanism that administrators should tightly control.

Finally, placing authoritative or recursive DNS servers on diverse subnets improves resilience and availability. Network-path diversity helps prevent a single subnet outage, routing issue, or localized denial-of-service event from disabling all DNS resolution capability. Together, these measures protect DNS confidentiality, integrity, and availability.

QUESTION NO: 42

Which DNS resource record can indicate how long any "DNS poisoning" could last?

A. MX

B. SOA

C. NS

D. TIMEOUT

ANSWER: B

Explanation:

SOA is correct because the Start of Authority record provides zone-level timing information relevant to DNS caching behavior. In particular, the SOA RDATA includes a MINIMUM field. In modern DNS use, this field is the negative-cache TTL: it tells recursive resolvers how long they may cache an authoritative negative response, such as a response indicating that a queried name does not exist. A poisoned DNS cache remains useful to an attacker only while the resolver continues serving the cached false information; therefore, DNS TTL values are central to estimating the potential duration of an effective poisoning incident. The SOA record also identifies the authoritative zone and contains administrative timing parameters, making it the DNS resource record associated with these zone caching and propagation controls. RFC 2308 specifies that the TTL for negative caching is derived from the SOA record, using the lower of the SOA RR TTL and the SOA MINIMUM field. This makes examining the SOA record an appropriate way to assess the relevant cache lifetime for negative DNS data. See RFC 2308: Negative Caching of DNS Queries and RFC 1035: Domain Names—Implementation and Specification.

QUESTION NO: 43

In the context of Trojans, what is the definition of a Wrapper?

A. An encryption tool to protect the Trojan

B. A tool used to bind the Trojan with a legitimate file

C. A tool used to calculate bandwidth and CPU cycles wasted by the Trojan

D. A tool used to encapsulate packets within a new header and footer

ANSWER: B

Explanation:

A tool used to bind the Trojan with a legitimate file is correct. In Trojan terminology, a wrapper (also called a binder) combines a malicious payload with an apparently harmless or desirable program into a single executable. The goal is to make the target more likely to run the file: when the user launches what appears to be a legitimate application, the wrapper

can start the legitimate program while also executing the concealed Trojan payload. This is a social-engineering and execution-concealment technique, rather than a network-packet encapsulation mechanism. Understanding wrappers is useful in ethical hacking because it illustrates why users and defenders must not rely solely on a file name, icon, or expected application behavior when deciding whether software is safe. Defensive controls include downloading software only from trusted sources, applying endpoint protection and application-control policies, and examining suspicious files in an isolated analysis environment. MITRE ATT&CK describes user execution as a technique in which adversaries rely on a user to run a malicious file or program, a behavior commonly associated with Trojan delivery methods: MITRE ATT&CK: User Execution. CISA likewise recommends using trusted sources and exercising caution with unexpected files and downloads: CISA: Recognize and Report Phishing.

QUESTION NO: 44

Which of the following controls help protect an organization against DHCP spoofing attacks? (Choose three.)

A. Enable DHCP snooping on user VLANs

B. Mark only authorized DHCP server uplinks as trusted

C. Use Dynamic ARP Inspection with DHCP snooping bindings

D. Configure all switch access ports as trusted

E. Disable all endpoint MAC addresses in the CAM table

ANSWER: A B C

Explanation:

DHCP snooping, marking legitimate DHCP-server uplinks as trusted, and using Dynamic ARP Inspection with DHCP snooping bindings are appropriate controls. DHCP snooping examines DHCP messages on untrusted access ports and blocks unauthorized DHCP server responses, such as DHCPOFFER and DHCPACK packets sent by a rogue endpoint. The switch port leading to the authorized DHCP server or relay must be configured as trusted so legitimate DHCP responses are permitted.

DHCP snooping also creates an IP-to-MAC-to-port binding database. Dynamic ARP Inspection can use this database to validate ARP packets, which helps prevent an attacker from using false ARP mappings after attempting to influence address assignment. Cisco documents DHCP snooping and the trusted-port model in its DHCP snooping and Dynamic ARP Inspection guidance.

QUESTION NO: 45

You want to carry out session hijacking on a remote server. The server and the client are communicating via TCP after a successful TCP three way handshake. The server has just received packet #120 from the client. The client has a receive window of 200 and the server has a receive window of 250.

Within what range of sequence numbers should a packet, sent by the client fall in order to be accepted by the server?

A. 200-250

B. 121-371

C. 120-321

D. 121-231

E. 120-370

F. 121-370

ANSWER: F

Explanation:

121-370 is the correct range when the question is interpreted using the normal simplified TCP receive-window calculation for individual sequence-number positions. After receiving sequence number 120, the server's next expected sequence number (RCV.NXT) is 121. The relevant advertised window is the server's receive window, which is 250 bytes; the client's own receive window governs traffic travelling in the opposite direction and does not determine what the server will accept from the client. TCP therefore accepts sequence numbers beginning at 121 and continuing for 250 positions. Because the upper edge of a TCP receive window is exclusive, the valid sequence numbers are 121 through 370 inclusive: 121 + 250 = 371, with 371 being the first sequence number outside this window. This follows TCP's sequence-space acceptability rule, which expresses an open upper bound as RCV.NXT ≤ SEG.SEQ < RCV.NXT + RCV.WND. See the TCP standard's receive-window and segment-acceptability requirements in RFC 9293, Section 3.10.7.4 and the original TCP sequence-space definition in RFC 793.

QUESTION NO: 46

Kevin is an IT security analyst working for Emerson Time Makers, a watch manufacturing company in Miami. Kevin and his girlfriend Katy recently broke up after a big fight. Kevin believes that she was seeing another person. Kevin, who has an online email account that he uses for most of his mail, knows that Katy has an account with that same company. Kevin logs into his email account online and gets the following URL after successfully logged in: http://www.youremailhere.com/mail.asp?mailbox=Kevin&Smith=121%22 Kevin changes the URL to: http://www.youremailhere.com/mail.asp?mailbox=Katy&Sanchez=121%22 Kevin is trying to access her email account to see if he can find out any information. What is Kevin attempting here to gain access to Katy's mailbox?

A. This type of attempt is called URL obfuscation when someone manually changes a URL to try and gain unauthorized access

B. By changing the mailbox's name in the URL, Kevin is attempting directory transversal

C. Kevin is trying to utilize query string manipulation to gain access to her email account

D. He is attempting a path-string attack to gain access to her mailbox

ANSWER: C

Explanation:

Kevin is trying to utilize query string manipulation to gain access to her email account. A query string is the portion of a URL after the question mark that supplies named parameters to a web application. In this case, the URL contains a `mailbox` parameter, and Kevin changes its value from his own identifier, `Kevin`, to `Katy`. This is an attempt to make the server return a resource belonging to a different user by tampering with client-controlled input. If the application relies on that parameter without verifying that the authenticated session is authorized to access the requested mailbox, it has an insecure direct object reference, commonly categorized under broken access control. Proper security requires server-side, object-level authorization for every mailbox request; being logged in must not by itself permit a user to select another user's mailbox through a modified URL. OWASP describes this risk and the need to validate authorization for each requested object in its Broken Access Control guidance. OWASP's Authorization Cheat Sheet likewise recommends enforcing authorization checks server-side and applying them consistently to every request.

QUESTION NO: 47

What ports should be blocked on the firewall to prevent NetBIOS traffic from not coming through the firewall if your network is comprised of Windows NT, 2000, and XP?(Choose all that apply.

A. 110

B. 135

C. 139

D. 161

E. 445

F. 1024

G. 137 and 138

ANSWER: C E G

Explanation:

139 and **137 and 138** are the core NetBIOS-related firewall ports. NetBIOS Name Service uses UDP port 137 for name registration and resolution, NetBIOS Datagram Service uses UDP port 138 for connectionless datagram traffic, and NetBIOS Session Service uses TCP port 139 for session-oriented communications, including legacy Windows file and printer sharing. Blocking these ports at a network boundary prevents those NetBIOS services from being exposed or traversing the firewall.

445 should also be blocked when the objective is to prevent legacy Windows hosts from exposing file-sharing traffic externally. Windows 2000 and later systems, including Windows XP, can use SMB directly over TCP 445 rather than encapsulating SMB in a NetBIOS session on TCP 139. Consequently, filtering TCP 445 alongside the NetBIOS ports closes the direct-hosted SMB path that could otherwise provide equivalent network file-sharing access. Microsoft's Windows service port guidance identifies UDP 137, UDP 138, TCP 139, and TCP 445 as the relevant NetBIOS/SMB service ports. See Microsoft's service and network port requirements and the Microsoft SMB overview.

QUESTION NO: 48

If an e-commerce site was put into a live environment and the programmers failed to remove the secret entry point that was used during the application development, what is this secret entry point known as?

A. SDLC process

B. Honey pot

C. SQL injection

D. Trap door

ANSWER: D

Explanation:

Trap door is correct. A trap door, also commonly called a backdoor, is a hidden or undocumented mechanism that bypasses normal authentication, authorization, or application workflow controls. Developers may create such an entry point temporarily while building or testing an application to simplify administration, troubleshooting, or access to restricted functions. If it remains when the e-commerce application is deployed to production, an attacker who discovers it could gain unauthorized access or execute functions outside the intended security model. This creates a serious security exposure because the entry point is not subject to the same visible, reviewed access controls that protect ordinary users. Secure software-development practice requires teams to identify and remove development-only accounts, test interfaces, debugging functions, hard-coded credentials, and undocumented access paths before release. Production deployment should also include code review, security testing, and configuration validation to ensure hidden functionality has not been left enabled. OWASP describes backdoors as malicious or unauthorized functionality that can provide unintended access, and its testing guidance emphasizes reviewing applications for such hidden paths. See the OWASP Backdoors reference and the OWASP Web Security Testing Guide.

QUESTION NO: 49

What makes web application vulnerabilities so aggravating? (Choose two)

A. They can be launched through an authorized port.

B. A firewall will not stop them.

C. They exist only on the Linux platform.

D. They are detectable by most leading antivirus software.

ANSWER: A B

Explanation:

Web application vulnerabilities are especially difficult to manage because malicious requests can be delivered through an authorized port. Web applications normally accept client traffic over standard web ports such as HTTP port 80 and HTTPS port 443. Those ports must commonly remain available for legitimate users, so an attacker can place malicious input inside apparently valid web requests instead of needing an unusual network service or blocked destination port. OWASP describes how web-application attacks exploit weaknesses in application logic and the handling of untrusted input, often within ordinary HTTP requests.

A firewall will not stop them when the firewall is a conventional network firewall that primarily permits or denies traffic according to addresses, ports, and protocols. If it allows normal web traffic to reach the application, it does not inherently determine whether submitted parameters, cookies, headers, uploads, or API payloads are malicious. Effective mitigation therefore requires secure application design, input validation, authentication and authorization controls, patching, logging, and, where appropriate, application-layer protections such as a web application firewall. See the OWASP Top 10 and OWASP's Web Application Firewall guidance.

QUESTION NO: 50

Which of the following LM hashes represent a password of less than 8 characters? (Select 2)

A. BA810DBA98995F1817306D272A9441BB

B. 44EFCE164AB921C0AAD3B435B51404EE

C. 0182BD0BD4444BF836077A718CCDF409

D. CEC52EB9C8E3455DC2265B23734E0DAC

E. B757BF5C0D87772FAAD3B435B51404EE

F. E52CAC67419A9A224A3B108F3FA6CB6D

ANSWER: B E

Explanation:

LM hashing processes a password as two independent seven-character portions. Before hashing, the password is converted to uppercase, padded or truncated to fourteen characters, and split into two seven-byte blocks. For a password shorter than eight characters, the second block contains only null padding. The LM hash for that all-null second block is the well-known constant `AAD3B435B51404EE`. Consequently, an LM hash ending in this value indicates that no characters occupied the second seven-character portion, meaning the original password contained seven or fewer characters. Both `44EFCE164AB921C0AAD3B435B51404EE` and `B757BF5C0D87772FAAD3B435B51404EE` have this required second half and therefore represent passwords of fewer than eight characters. The first half remains password-specific because it is derived from the first seven-character block. This recognizable null-half value is commonly used in credential-auditing and forensic work to identify short passwords stored with legacy LM hashing. Microsoft describes the LM one-way-function construction, including the fourteen-byte input and two seven-byte sections, in its NTLM protocol specification. Additional technical background on LM hashes is available from Openwall.

QUESTION NO: 51

One of your junior administrator is concerned with Windows LM hashes and password cracking. In your discussion with them, which of the following are true statements that you would point out?

Select the best answers.

A. John the Ripper can be used to crack a variety of passwords, but one limitation is that the output doesn't show if the password is upper or lower case.

B. BY using NTLMV1, you have implemented an effective countermeasure to password cracking.

C. SYSKEY is an effective countermeasure.

D. If a Windows LM password is 7 characters or less, the hash will be passed with the following characters, in HEX-0011223344556677889.

E. Enforcing Windows complex passwords is an effective countermeasure.

ANSWER: A C E

Explanation:

“John the Ripper can be used to crack a variety of passwords, but one limitation is that the output doesn't show if the password is upper or lower case.” is correct in the specific context of LM hashes. Before generating an LM hash, Windows converts the password to uppercase. Consequently, recovering an LM password reveals its uppercase form and cannot establish the user’s original capitalization. John the Ripper supports LM hash formats and is commonly used to demonstrate this legacy weakness in authorized password-auditing exercises.

“SYSKEY is an effective countermeasure.” is correct in the historical Windows/legacy-exam context: SYSKEY added protection for the Security Accounts Manager data containing password-hash material, making some offline attacks more difficult. Current environments should instead use supported modern OS protections and avoid retaining LM hashes, because Microsoft has retired SYSKEY in current Windows releases.

“Enforcing Windows complex passwords is an effective countermeasure.” is also correct. Complexity requirements increase the password search space and help resist guessing and cracking attacks; strong password length and uniqueness are especially important. Microsoft documents both the importance of preventing LM-hash storage and the Windows complexity policy at Prevent Windows from storing an LM hash and Password must meet complexity requirements.

QUESTION NO: 52

Which of the following Registry location does a Trojan add entries to make it persistent on Windows 7? (Select 2 answers)

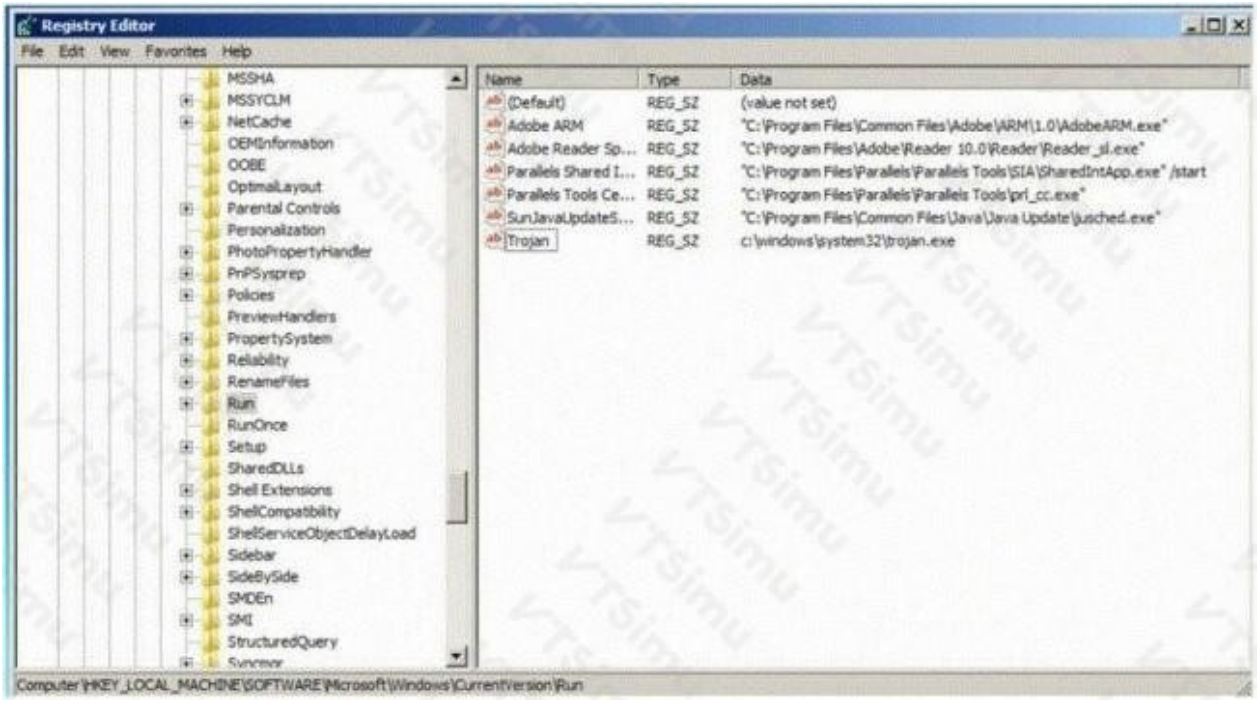

A. HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Windows\CurrentVersion\Run

B. HKEY_CURRENT_USER\Software\Microsoft\Windows\System32\CurrentVersion\Run

C. HKEY_CURRENT_USER\Software\Microsoft\Windows\CurrentVersion\Run

ANSWER: A C

Explanation:

`HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Windows\CurrentVersion\Run` and `HKEY_CURRENT_USER\Software\Microsoft\Windows\CurrentVersion\Run` are standard Windows Run registry keys that can be used to achieve persistence. Windows processes the commands or executable paths stored as values under these keys when a user logs on. The machine-wide location applies to users of the computer and normally requires administrative privileges to modify, whereas the current-user location applies only to the currently logged-on user and can generally be modified without elevation. A Trojan can therefore place a value pointing to its executable in either location so that its payload is launched again after logon, allowing it to survive a reboot or a user sign-out/sign-in cycle. This technique is commonly known as Registry Run Keys / Startup Folder persistence and is documented by MITRE ATT&CK as sub-technique T1547.001. Microsoft also identifies registry-based startup locations, including the Run keys, as locations used to start applications automatically at user logon. See MITRE ATT&CK: Registry Run Keys / Startup Folder and Microsoft: Run and RunOnce Registry Keys.

QUESTION NO: 53

The network administrator at Spears Technology, Inc has configured the default gateway Cisco router's access-list as below:

```
Current configuration : 1206 bytes
!
version 12.3
!
hostname Victim
!
enable secret 5 $1$h2iz$DHYpcqURFOAPD2aDuA.YXO
!
interface Ethernet0/0
ip address dhcp
ip nat outside
half-duplex
!
interface Ethernet0/1
ip address 192.168.1.1 255.255.255.0
ip nat inside
half-duplex
!
router rip
network 192.168.1.0
!
ip nat inside source list 102 interface Ethernet0/0 overload
no ip http server
ip classless
!
access-list 1 permit 192.168.1.0 0.0.0.255
access-list 102 permit ip any any
!
snmp-server community public RO

snmp-server community private RW 1
snmp-server enable traps tty
!
line con 0
logging synchronous
login
line aux 0
line vty 0 4
password secret
login
!!
end
```

You are hired to conduct security testing on their network. You successfully brute-force the SNMP community string using a SNMP crack tool. The access-list configured at the router prevents you from establishing a successful connection. You want to retrieve the Cisco configuration from the router. How would you proceed?

A. Use the Cisco's TFTP default password to connect and download the configuration file

B. Run a network sniffer and capture the returned traffic with the configuration file from the router

C. Run Generic Routing Encapsulation (GRE) tunneling protocol from your computer to the router masking your IP address

D. Send a customized SNMP set request with a spoofed source IP address in the range - 192.168.1.0

ANSWER: B D

Explanation:

The appropriate technique is to send a customized SNMP set request with a spoofed source IP address in the permitted 192.168.1.0 range, then run a network sniffer and capture the returned traffic with the configuration file from the router. A Cisco SNMP access list limits which management-station source addresses may use the configured community string. If a tester can generate a packet whose apparent source matches an allowed address, the router may process the request despite the tester's actual address being outside that range. In a controlled assessment, a write-capable SNMP request can be used with Cisco configuration-copy management objects to instruct the device to transfer its running or startup configuration through the network. Because the traffic containing that transfer traverses a segment observable by the tester, packet capture can recover the configuration data or otherwise validate that the copy operation occurred. This scenario illustrates why source-address ACLs alone are not sufficient protection for SNMP write access: spoofing and traffic observation can undermine a community-string-based design. Cisco documents configuration copying through SNMP in its SNMP configuration-copy guidance; the underlying SNMP management framework is specified in RFC 3411.

QUESTION NO: 54

How do you defend against MAC attacks on a switch?

A. Disable SPAN port on the switch

B. Enable SNMP Trap on the switch

C. Configure IP security on the switch

D. Enable Port Security on the switch

ANSWER: D

Explanation:

Enable Port Security on the switch is the appropriate defense against common Layer 2 MAC-based attacks, particularly MAC-address flooding of a switch's forwarding table and unauthorized devices connecting to an access port. Port security lets an administrator restrict how many source MAC addresses may be learned on a port and optionally specify the permitted addresses, either statically or through secure dynamic learning. If an attacker sends frames using many fabricated source MAC addresses, the configured maximum is exceeded and the switch can protect the port by restricting traffic or shutting it down, depending on the violation mode. This prevents uncontrolled MAC-table learning and limits access to approved endpoints. In a properly designed access-layer configuration, port security should be applied to user-facing access ports, with an appropriate maximum-MAC limit, a deliberate violation action, and operational monitoring/recovery procedures. Cisco documents the available secure MAC-address limits and violation handling behavior in its Port Security Configuration Guide. The broader role of access-port controls in mitigating unauthorized endpoint connections is also covered in the CISA secure network access guidance.

QUESTION NO: 55

What ICMP message types are used by the ping command?

A. Timestamp request (13) and timestamp reply (14)

B. Echo request (8) and Echo reply (0)

C. Echo request (0) and Echo reply (1)

D. Ping request (1) and Ping reply (2)

ANSWER: B

Explanation:

The correct answer is **Echo request (8) and Echo reply (0)**. For IPv4, the traditional `ping` utility tests IP-layer reachability by transmitting ICMP Echo Request messages and listening for matching ICMP Echo Reply messages from the destination host. ICMP Type 8 identifies an Echo Request, while ICMP Type 0 identifies an Echo Reply. The messages include an identifier and sequence number, allowing the originating system to associate replies with requests and calculate packet loss and round-trip time. This simple request-and-response exchange is why ping is widely used for initial connectivity diagnostics between hosts. ICMP is defined as part of the Internet Protocol suite; its control and diagnostic messages are carried directly inside IP packets rather than using TCP or UDP ports. The authoritative IPv4 ICMP specification defines the Echo and Echo Reply type values in its message-format sections. See RFC 792: Internet Control Message Protocol and the IANA ICMP Parameters registry for the registered ICMP type assignments.

QUESTION NO: 56

Exhibit:

Based on the following extract from the log of a compromised machine, what is the hacker really trying to steal?

A. har.txt

B. SAM file

C. wwwroot

D. Repair file

ANSWER: B

Explanation:

SAM file is correct because the Security Accounts Manager database is a high-value Windows credential store. The log activity described in the supplied explanation shows that the attacker is obtaining `har.txt` as an intermediate or disguised file, but that file contains a copied SAM hive. The attacker's actual objective is therefore the account and password-verifier data held in the SAM database, rather than the temporary filename used to stage or transfer it.

Windows uses the SAM database to maintain local user and group account information. When obtained together with the relevant system boot-key material, SAM data can be used in offline credential attacks, such as extracting password hashes for cracking or pass-the-hash-related assessment activity. This makes copying the SAM hive a common post-compromise credential-access technique. Microsoft documents that the SAM registry hive is stored under `%SystemRoot%\System32\config\SAM` and is protected while Windows is running; attackers may use backups, shadow copies, or other methods to obtain a copy. See Microsoft's Reg save documentation and the MITRE ATT&CK entry for Security Account Manager credential dumping.

QUESTION NO: 57

SSL has been seen as the solution to a lot of common security problems. Administrator will often time make use of SSL to encrypt communications from points A to point B. Why do you think this could be a bad idea if there is an Intrusion Detection System deployed to monitor the traffic between point A and B?

A. SSL is redundant if you already have IDS's in place

B. SSL will trigger rules at regular interval and force the administrator to turn them off

C. SSL will slow down the IDS while it is breaking the encryption to see the packet content

D. SSL will blind the content of the packet and Intrusion Detection Systems will not be able to detect them

ANSWER: D

Explanation:

SSL/TLS encryption protects the confidentiality and integrity of application data while it travels between the two points. When a network-based intrusion detection system is positioned only on the path between those endpoints, it can generally observe packet headers, addresses, ports, TLS handshake metadata, traffic volume, and timing, but it cannot inspect the encrypted application payload. As a result, signatures or detections that depend on seeing commands, URLs, exploit strings, malware content, or protocol fields inside the protected session cannot operate on that hidden content. This is why “SSL will blind the content of the packet and Intrusion Detection Systems will not be able to detect them” is the best answer: the key limitation is loss of payload visibility at the network monitoring point, not that encryption itself is unnecessary. Organizations can restore selected inspection capabilities by terminating or decrypting TLS at a properly controlled inspection proxy, or by using endpoint-based telemetry where data is available before encryption or after decryption. NIST notes that encrypted traffic limits an IDPS’s ability to inspect network content and discusses deployment considerations for encrypted communications. See NIST SP 800-94, Guide to Intrusion Detection and Prevention Systems and the TLS 1.3 specification (RFC 8446).

QUESTION NO: 58

Which of the following LM hashes represents a password of less than 8 characters?

A. 0182BD0BD4444BF836077A718CCDF409

B. 44EFCE164AB921CAAAD3B435B51404EE

C. BA810DBA98995F1817306D272A9441BB

D. CEC52EB9C8E3455DC2265B23734E0DAC

E. B757BF5C0D87772FAAD3B435B51404EE

F. E52CAC67419A9A224A3B108F3FA6CB6D

ANSWER: B E

Explanation:

LM hashing divides the uppercase password into two independent seven-character blocks and produces a 16-byte, 32-hexadecimal-character result. When a password contains fewer than eight characters, its second seven-character block is empty. The LM hash of that empty block is the well-known constant `AAD3B435B51404EE`, so a valid LM hash ending in that value indicates that the password did not extend into the second block. Therefore, `44EFCE164AB921CAAAD3B435B51404EE` and `B757BF5C0D87772FAAD3B435B51404EE` represent passwords of fewer than eight characters. The first value was normalized to valid hexadecimal notation by removing the stray non-hexadecimal character in the supplied text; the resulting value has the required 32 hexadecimal characters. This recognition pattern is useful in authorized password-auditing and incident-response work because the LM construction exposes whether the second password block was unused. Microsoft documents that LM hashes are a legacy password-hash format and explains the security implications of retaining compatibility with it; see Microsoft's password technical overview. The LM algorithm’s seven-character split and empty-half behavior are also documented in Openwall John the Ripper documentation.

QUESTION NO: 59

Snort is an open source Intrusion Detection system. However, it can also be used for a few other purposes as well.

Which of the choices below indicate the other features offered by Snort?

A. IDS, Packet Logger, Sniffer

B. IDS, Firewall, Sniffer

C. IDS, Sniffer, Proxy

D. IDS, Sniffer, content inspector

ANSWER: A

Explanation:

IDS, Packet Logger, Sniffer is correct because Snort is designed to operate in three principal modes: packet sniffer mode, packet logger mode, and network intrusion detection mode. In sniffer mode, it reads packets from a network interface and displays packet headers and/or payload data for immediate observation. In packet logger mode, it captures traffic and writes it to logs for later review, which supports packet-level investigation and forensic analysis. In intrusion detection mode, Snort applies its configured rules to observed traffic and generates alerts when traffic matches defined signatures or detection logic. These modes are fundamental to Snort's traditional architecture and explain why it is useful not only for real-time detection but also for traffic visibility and evidence collection. The Snort documentation describes command-line operation and configuration for packet capture, logging, and inspection, while the Snort FAQ identifies it as an open-source network intrusion detection and prevention system. See the official Snort documentation and the Snort FAQ for further details.

QUESTION NO: 60

Keystroke logging is the action of tracking (or logging) the keys struck on a keyboard, typically in a covert manner so that the person using the keyboard is unaware that their actions are being monitored.

How will you defend against hardware keyloggers when using public computers and Internet Kiosks? (Select 4 answers)

A. Alternate between typing the login credentials and typing characters somewhere else in the focus window

B. Type a wrong password first, later type the correct password on the login page defeating the keylogger recording

C. Type a password beginning with the last letter and then using the mouse to move the cursor for each subsequent letter.

D. The next key typed replaces selected text portion. E.g. if the password is "secret", one could type "s", then some dummy keys "asdfsd".Then these dummies could be selected with mouse, and next character from the password "e" is typed, which replaces the dummies"asdfsd"

E. The next key typed replaces selected text portion. E.g. if the password is "secret", one could type "s", then some dummy keys "asdfsd".Then these dummies could be selected with mouse, and next character from the password "e" is typed, which replaces the dummies"asdfsd"

ANSWER: A C D E

Explanation:

Alternating between credential characters and unrelated characters entered elsewhere in the active window can obscure which keystrokes form the final credential, particularly when an attacker has only the raw keyboard-event stream and cannot reliably associate every event with its destination. Entering a password from its last character while using the mouse to reposition the insertion point similarly separates the physical typing sequence from the final character order displayed in the password field. A hardware logger records key events, whereas cursor placement is performed with the mouse; this makes reconstruction less straightforward.

The selected-text replacement technique also adds deliberate noise. Dummy characters can be typed, selected with the mouse, and replaced by the next required password character, leaving a final value that differs from the complete sequence of keys captured by the logger. Both occurrences of this technique describe the same valid keystroke-obfuscation countermeasure and are therefore correctly selected as presented. These are legacy mitigation techniques rather than guarantees: public systems may be compromised in additional ways, so users should avoid entering sensitive credentials on untrusted kiosks whenever possible. CISA describes keylogging as a method used to capture keyboard input in its keylogger guidance; NIST also emphasizes using phishing-resistant authentication methods where feasible in SP 800-63B.

QUESTION NO: 61

What sequence of packets is sent during the initial TCP three-way handshake?

A. SYN, SYN-ACK, ACK

B. SYN, URG, ACK

C. SYN, ACK, SYN-ACK

D. FIN, FIN-ACK, ACK

ANSWER: A

Explanation:

SYN, SYN-ACK, ACK is the sequence used to establish a TCP connection. The initiating host first sends a TCP segment with the SYN flag set, proposing an initial sequence number and requesting that a connection be opened. The receiving host responds with a segment that has both SYN and ACK set. Its SYN supplies the receiver's own initial sequence number, while its ACK acknowledges the initiator's SYN and sequence number. Finally, the initiating host sends an ACK segment acknowledging the receiver's SYN. At that point, both endpoints have confirmed bidirectional reachability, synchronized their initial sequence numbers, and can begin exchanging application data over the established connection. This exchange is called a three-way handshake because it requires these three directed messages to establish the reliable, stateful TCP session. Packet captures commonly show these control flags at the beginning of a normal TCP connection, making the pattern important for network analysis and ethical security testing. The TCP state-transition and synchronization behavior is specified in RFC 9293, Transmission Control Protocol. For a practical protocol overview, see the Wireshark TCP analysis documentation.

QUESTION NO: 62

NSLookup is a good tool to use to gain additional information about a target network. What does the following command accomplish?

nslookup

> server

> set type =any

> ls -d

A. Enables DNS spoofing

B. Loads bogus entries into the DNS table

C. Verifies zone security

D. Performs a zone transfer

E. Resets the DNS cache

ANSWER: D

Explanation:

Performs a zone transfer is correct. In interactive NSLookup, `server <ipaddress>` selects the DNS server that will receive subsequent queries. The `ls` command requests a listing of DNS information for the specified domain; the `-d` modifier requests all records in that domain's zone. This is commonly described as attempting a DNS zone transfer, specifically an AXFR-style full-zone transfer. A successful transfer can reveal the zone's resource records, such as host names, IP addresses, mail exchangers, name servers, and other DNS data, which is why it is useful during authorized network reconnaissance and security assessments. The `set type=any` setting asks for records broadly, although the key operation in the displayed sequence is `ls -d target.com`. Whether the command succeeds depends on the authoritative

DNS server's zone-transfer policy: properly configured servers restrict transfers to authorized secondary DNS servers or other explicitly permitted clients. Microsoft documents `ls` as listing domain information in NSLookup, and RFC 5936 defines the DNS AXFR mechanism used for full zone transfers. See Microsoft NSLookup ls documentation and RFC 5936: DNS Zone Transfer Protocol.

QUESTION NO: 63

You are trying to package a RAT Trojan so that Anti-Virus software will not detect it. Which of the listed technique will NOT be effective in evading Anti-Virus scanner?

A. Convert the Trojan.exe file extension to Trojan.txt disguising as text file

B. Break the Trojan into multiple smaller files and zip the individual pieces

C. Change the content of the Trojan using hex editor and modify the checksum

D. Encrypt the Trojan using multiple hashing algorithms like MD5 and SHA-1

ANSWER: D

Explanation:

"Encrypt the Trojan using multiple hashing algorithms like MD5 and SHA-1" is the correct answer because MD5 and SHA-1 are cryptographic hash functions, not encryption algorithms. A hash function accepts input data and produces a fixed-length digest; it does not produce a reversibly protected version of the original executable that can later be restored and run. Consequently, applying a hash to a program would yield a digest value rather than an operational Trojan package. Hashing therefore cannot serve as an encryption-based packaging mechanism or change an executable into a runnable, concealed form. NIST defines secure hash algorithms as functions that generate message digests for integrity-related cryptographic uses, rather than reversible confidentiality mechanisms. See NIST FIPS 180-4, Secure Hash Standard. NIST also distinguishes encryption, which protects data confidentiality through reversible cryptographic transformation with a key, from hashing, which is designed to be one-way; see the NIST privacy and cybersecurity framework FAQs. Additionally, MD5 and SHA-1 are legacy hash algorithms with known collision weaknesses, but that fact does not make them encryption methods. The fundamental issue is that hashing cannot encrypt or preserve a runnable program for later execution.

QUESTION NO: 64

Melissa is a virus that attacks Microsoft Windows platforms.

To which category does this virus belong?

A. Polymorphic

B. Boot Sector infector

C. System

D. Macro

ANSWER: D

Explanation:

Macro is correct because the Melissa malware was a Microsoft Word macro virus. It was distributed in a Word document, commonly identified as *list.doc*, that contained malicious Visual Basic for Applications (VBA) macro code. When a recipient opened the document and enabled macros, the code executed within Word, modified the user's Normal template, and used Microsoft Outlook to mail copies of the infected document to contacts. This behavior illustrates the defining characteristic of a macro virus: malicious instructions are embedded in the macro language of an office document or template and execute through the associated productivity application rather than as a standalone Windows executable.

Melissa became a major incident in 1999 because its automated Outlook mail propagation created substantial email volume and operational disruption across organizations. Its reliance on Word document macros and VBA code is the basis for classifying it as a macro virus. The CERT Coordination Center's historical advisory documents the Melissa incident and its Word/Outlook propagation mechanism: CERT Advisory CA-1999-04. Microsoft's guidance on VBA macros provides relevant background on the Office macro execution environment: Microsoft Office security and compliance documentation.

QUESTION NO: 65

Vulnerability scanners are automated tools that are used to identify vulnerabilities and misconfigurations of hosts. They also provide information regarding mitigating discovered vulnerabilities.

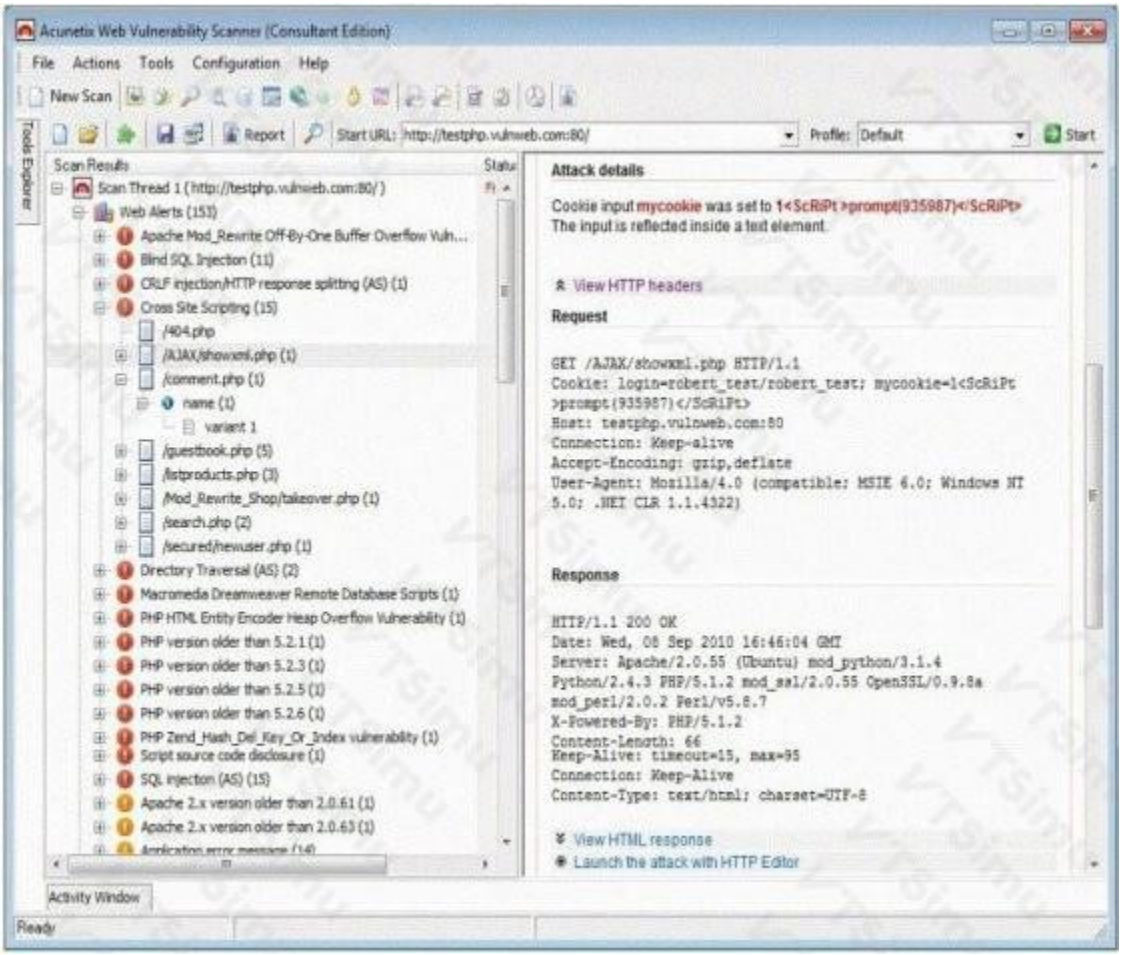

Which of the following statements is incorrect?

A. Vulnerability scanners attempt to identify vulnerabilities in the hosts scanned.

B. Vulnerability scanners can help identify out-of-date software versions, missing patches, or system upgrades

C. They can validate compliance with or deviations from the organization's security policy

D. Vulnerability scanners can identify weakness and automatically fix and patch the vulnerabilities without user intervention

ANSWER: D

Explanation:

"Vulnerability scanners can identify weakness and automatically fix and patch the vulnerabilities without user intervention " is the incorrect statement. Vulnerability scanning is principally an assessment and discovery activity: the scanner probes systems, correlates detected versions and configurations with known vulnerability information, and reports findings that require remediation. Remediation commonly involves reviewing the finding, confirming its applicability and severity, selecting an appropriate update or configuration change, testing it where necessary, approving the change, deploying it, and rescanning to verify that the issue has been addressed. This process is important because patches can affect application

compatibility, availability, and operational controls. A scanner may integrate with separate patch-management, orchestration, or configuration-management products, but that integration does not make vulnerability scanning itself an automatic, unattended patching capability. NIST describes vulnerability scanning as identifying vulnerabilities and validating the effectiveness of remediation, while treating remediation as a distinct activity within technical security testing and assessment. See NIST SP 800-115, Technical Guide to Information Security Testing and Assessment and CISA's Known Exploited Vulnerabilities Catalog for authoritative guidance on identifying vulnerabilities and prioritizing remediation.

QUESTION NO: 66

Which of the following buffer overflow exploits are related to Microsoft IIS web server? (Choose three)

A. Internet Printing Protocol (IPP) buffer overflow

B. Code Red Worm

C. Indexing services ISAPI extension buffer overflow

D. NeXT buffer overflow

ANSWER: A B C

Explanation:

The correct selections are Internet Printing Protocol (IPP) buffer overflow, Code Red Worm, and Indexing services ISAPI extension buffer overflow. These are all historically associated with vulnerabilities in Microsoft Internet Information Services and its installed components or script mappings. Microsoft documented the Internet Printing Protocol issue as a buffer-overrun vulnerability affecting the IIS printing capability; successful exploitation could permit code execution in the security context of the IIS service. The Code Red worm notoriously exploited a buffer overflow in the IIS Index Server's Indexing Service ISAPI extension, accessed through the vulnerable `idq.dll` mapping. This flaw enabled unauthenticated remote attackers to execute arbitrary code and allowed the worm to propagate rapidly across vulnerable servers. Microsoft also identified a distinct buffer-overrun issue in the Indexing Service ISAPI extension that could allow an attacker to run code on an affected IIS server. These examples demonstrate why unnecessary IIS extensions and services should be removed or disabled, security updates applied promptly, and exposed web-server functionality minimized. Microsoft's IIS security bulletins for the printing vulnerability and Index Server vulnerability provide the primary technical context: MS01-023 and MS01-033.

QUESTION NO: 67

Your computer is infected by E-mail tracking and spying Trojan. This Trojan infects the computer with a single file - emos.sys

Which step would you perform to detect this type of Trojan?

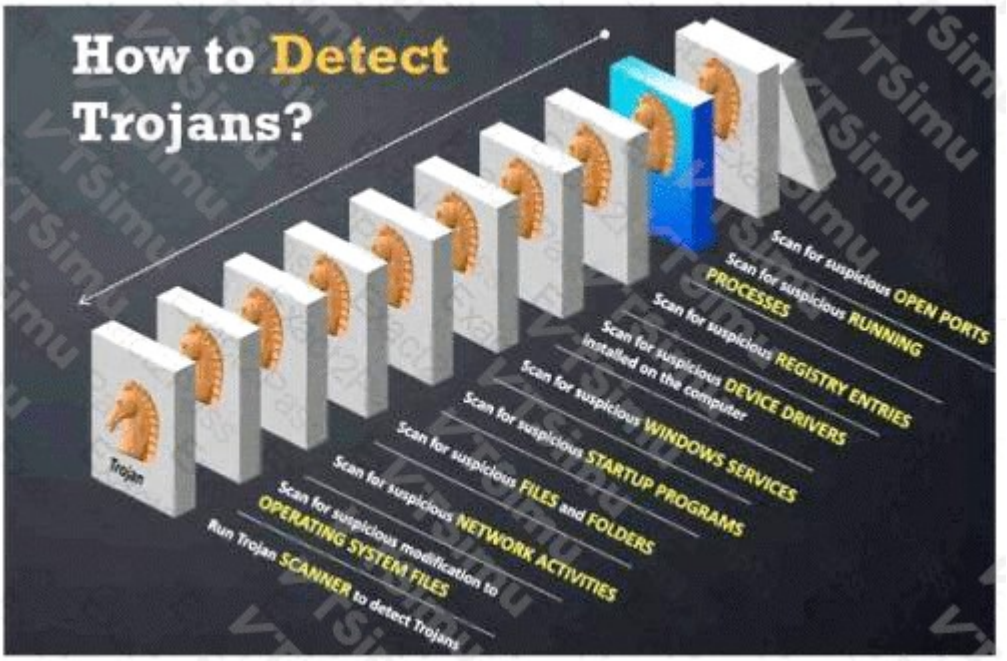

A. Scan for suspicious startup programs using msconfig

B. Scan for suspicious network activities using Wireshark

C. Scan for suspicious device drivers in c:\windows\system32\drivers

D. Scan for suspicious open ports using netstat

ANSWER: C

Explanation:

Scan for suspicious device drivers in c:\windows\system32\drivers is correct because the identified malicious file, `emos.sys`, uses the `.sys` extension associated with Windows kernel-mode driver files. Windows commonly stores installed driver binaries in `%SystemRoot%\System32\drivers`, so examining that directory for an unexpected driver with the stated filename is the most direct detection step. An investigator should verify the file's full path, creation time, digital signature, publisher, and cryptographic hash, then correlate those findings with installed driver-service configuration. Microsoft documents that driver packages include driver binaries and that Windows maintains driver-related installation information; suspicious or unsigned driver binaries warrant careful incident-response handling rather than execution or deletion without evidence preservation. Sysinternals Autoruns can also enumerate configured drivers and services, helping establish whether a suspicious `.sys` file is configured to load during system startup. These checks directly connect the known indicator, `emos.sys`, to the Windows location and persistence mechanism most relevant to a driver-based Trojan. See Microsoft's driver package documentation and Microsoft Sysinternals Autoruns.

QUESTION NO: 68

Which of the following resources does NMAP need to be used as a basic vulnerability scanner covering several vectors like SMB, HTTP and FTP?

A. Metasploit scripting engine

B. Nessus scripting engine

C. NMAP scripting engine

D. SAINT scripting engine

ANSWER: C

Explanation:

NMAP needs the **NMAP scripting engine**, more commonly called the Nmap Scripting Engine (NSE), to function as a basic vulnerability scanner across services such as SMB, HTTP, and FTP. NSE extends Nmap beyond host discovery and port/service enumeration by allowing scripts to interact with discovered services, identify insecure configurations, collect service-specific information, and check for known vulnerability conditions. Scripts are organized into categories, including *vuln*, which is intended for scripts that check services for known vulnerabilities. For example, after identifying an SMB, web, or FTP service during a scan, an assessor can select relevant NSE scripts to perform targeted checks against that service. Nmap's official documentation describes NSE as a system for writing and sharing scripts that automate a wide variety of networking tasks, and it documents script categories and usage through the `--script` option. This makes NSE the Nmap-native resource that provides the requested multi-vector basic vulnerability-scanning capability. See the official Nmap Scripting Engine documentation and the NSE vulnerability-script category reference.

QUESTION NO: 69

Where should a security tester be looking for information that could be used by an attacker against an organization? (Select all that apply)

A. CHAT rooms

B. WHOIS database

C. News groups

D. Web sites

E. Search engines

F. Organization's own web site

ANSWER: A B C D E F

Explanation:

All listed sources are relevant to open-source intelligence gathering and should be reviewed during an authorized security assessment. CHAT rooms, News groups, and Web sites can reveal employee discussions, technical questions, shared files, project references, email-address formats, software versions, or information about business relationships. Search engines can expose indexed documents, cached content, error messages, public cloud resources, and unintentionally published data. An Organization's own web site is especially valuable because it commonly identifies personnel, locations, technologies, contact details, partners, and organizational structure.

The WHOIS database can provide domain-registration and related infrastructure information, subject to the privacy redaction practices of the applicable registry and registrar. This intelligence helps a tester understand the organization's public-facing footprint and identify information that could assist attacker reconnaissance. The work must remain within written authorization and applicable law, but examining publicly available information is a core part of assessing exposure before active testing begins. CISA describes open-source intelligence as information collected from publicly available sources, and NIST's guidance recognizes reconnaissance as an important stage in assessing systems and identifying relevant attack surface. See CISA Open Source Intelligence and NIST SP 800-115.

QUESTION NO: 70

The SNMP Read-Only Community String is like a password. The string is sent along with each SNMP Get-Request and allows (or denies) access to a device. Most network vendors ship their equipment with a default password of "public". This is the so-called "default public community string". How would you keep intruders from getting sensitive information regarding the network devices using SNMP? (Select 2 answers)

A. Enable SNMPv3 which encrypts username/password authentication

B. Use your company name as the public community string replacing the default 'public'

C. Enable IP filtering to limit access to SNMP device

D. The default configuration provided by device vendors is highly secure and you don't need to change anything

ANSWER: A C

Explanation:

Enable SNMPv3 which encrypts username/password authentication is correct because SNMPv3 replaces the shared, plaintext community-string model used by earlier SNMP versions with the User-based Security Model. It supports authenticated management messages and, when privacy is configured, encrypts SNMP protocol data. Using SNMPv3 with authentication and privacy protects management queries and responses from disclosure or tampering by parties observing network traffic. Administrators should use strong, unique SNMPv3 credentials and select an authentication-and-privacy security level rather than relying on a default community string.

Enable IP filtering to limit access to SNMP device is also correct because access-control lists or firewall rules can restrict UDP SNMP traffic to explicitly authorized management stations. This materially reduces the systems that can send SNMP requests to a device and limits exposure of device inventory, interface, routing, and configuration-related information. Filtering should be applied consistently on the device management plane and, where appropriate, at network boundaries, allowing only the necessary SNMP version, source addresses, and ports. NIST recommends using SNMPv3 security features and restricting management access through controls such as filtering. See RFC 3414: User-based Security Model for SNMPv3 and NIST SP 800-123: Guide to General Server Security.

QUESTION NO: 71

Anonymizer sites access the Internet on your behalf, protecting your personal information from disclosure. An anonymizer protects all of your computer's identifying information while it surfs for you, enabling you to remain at least one step removed from the sites you visit.

You can visit Web sites without allowing anyone to gather information on sites visited by you. Services that provide anonymity disable pop-up windows and cookies, and conceal visitor's IP address.

These services typically use a proxy server to process each HTTP request. When the user requests a Web page by clicking a hyperlink or typing a URL into their browser, the service retrieves and displays the information using its own server. The remote server (where the requested Web page resides) receives information on the anonymous Web surfing service in place of your information.

In which situations would you want to use anonymizer? (Select 3 answers)

A. Increase your Web browsing bandwidth speed by using Anonymizer

B. To protect your privacy and Identity on the Internet

C. To bypass blocking applications that would prevent access to Web sites or parts of sites that you want to visit.

D. Post negative entries in blogs without revealing your IP identity

ANSWER: B C D

Explanation:

“To protect your privacy and Identity on the Internet” is a primary use of an anonymizing proxy. The proxy submits the web request on the user’s behalf, so the destination website ordinarily sees the proxy service’s network address rather than the user’s public IP address. This provides a degree of separation between the user and the visited service, reducing direct exposure of identifying network information.

“To bypass blocking applications that would prevent access to Web sites or parts of sites that you want to visit.” is also a common technical use. Because the browser connects to an intermediary that retrieves the requested content, a proxy can sometimes provide access where direct requests are blocked. Such use must remain consistent with organizational policy, law, and authorization requirements.

“Post negative entries in blogs without revealing your IP identity” reflects the same IP-address masking capability. A blog may receive the anonymizer’s address instead of the poster’s address, which can support anonymous expression. However, anonymity is not absolute: account logins, browser fingerprinting, tracking scripts, provider logs, and content itself can still

identify a person. The Electronic Frontier Foundation discusses practical limitations and protections for online anonymity in its Surveillance Self-Defense guide, while the HTTP proxy model is defined in RFC 9110.

QUESTION NO: 72

What is the command used to create a binary log file using tcpdump?

A. tcpdump -w ./log

B. tcpdump -r log

C. tcpdump -vde logtcpdump -vde ? log

D. tcpdump -l /var/log/

ANSWER: A

Explanation:

`tcpdump -w ./log` is correct because the `-w` option instructs tcpdump to write captured packets to a file rather than displaying packet summaries on standard output. The resulting `./log` file is written in tcpdump's packet-capture (pcap) save-file format, which is a binary format containing packet records and capture metadata. The filename does not need a particular extension, although extensions such as `.pcap` or `.cap` are commonly used to make the file type clear. For example, a practical capture command could be `tcpdump -i eth0 -w capture.pcap`, which records traffic seen on `eth0` for later inspection. A binary capture file is especially useful in ethical-hacking and forensic workflows because it preserves packet-level evidence that can subsequently be opened, filtered, and analyzed with tools such as tcpdump or Wireshark. The tcpdump manual specifies that `-w` writes raw packets to a file for later analysis, rather than parsing and printing them while capturing. See the official tcpdump manual and Wireshark's capture file format documentation.

QUESTION NO: 73

Assuring two systems that are using IPSec to protect traffic over the internet, what type of general attack could compromise the data?

A. Spoof Attack

B. Smurf Attack

C. Man inthe Middle Attack

D. Trojan Horse Attack

E. Back Orifice Attack

ANSWER: D E

Explanation:

"Trojan Horse Attack" and "Back Orifice Attack

" are correct because IPsec protects packets while they traverse the network, but its security boundary is normally the communicating endpoints. A Trojan horse can execute on an endpoint and collect, alter, or exfiltrate information before the host submits it to IPsec for encryption. It can also access information after IPsec has received and decrypted it. Thus, the traffic may remain cryptographically protected in transit while the underlying data is nevertheless compromised on the system where it is created or consumed.

Back Orifice is a remote-administration backdoor/Trojan that enables unauthorized control of a compromised Windows host. An attacker with this level of endpoint access can interact with files, processes, and user data in cleartext at the host itself, placing the sensitive information outside the protection provided by the IPsec tunnel. This illustrates a core operational security principle: encrypted transport must be paired with endpoint hardening, malware prevention, patching, and monitoring. IPsec's architecture and security services are defined for protection of IP traffic, while endpoint compromise

requires host-level controls. See RFC 4301: Security Architecture for the Internet Protocol and MITRE ATT&CK: Back Orifice.

QUESTION NO: 74

During a penetration test, a tester finds that the web application being analyzed is vulnerable to Cross Site Scripting (XSS). Which of the following conditions must be met to exploit this vulnerability?

A. The web application does not have the secure flag set.

B. The session cookies do not have the HttpOnly flag set.

C. The victim user should not have an endpoint security solution.

D. The victim's browser must have ActiveX technology enabled.

E. A victim's browser must load a page in which the injected script is returned and executed.

ANSWER: E

Explanation:

A victim's browser must load a page in which the injected script is returned and executed. XSS is exploitable when attacker-controlled content reaches a browser in an executable web context—for example, as executable JavaScript, an event-handler attribute, or unsafe HTML that results in script execution—and the target page is rendered by the victim's browser. The key security impact comes from the fact that the browser executes the injected code with the origin and privileges of the vulnerable application. This can allow the script to interact with page content, make same-origin requests as the victim, alter what the victim sees, or perform actions available to that authenticated user. In practice, reflected XSS commonly requires the victim to follow a crafted link or submit a crafted request, while stored XSS requires the victim to visit a page that displays the persisted payload. The essential exploitation event in either case is loading and executing the injected content in the vulnerable application's origin. OWASP describes XSS as the injection of malicious scripts into otherwise trusted websites and provides prevention and context-specific guidance at OWASP Cross Site Scripting. Additional technical guidance on injection contexts and defenses is available in the OWASP XSS Prevention Cheat Sheet.

QUESTION NO: 75

Which of the following is a symmetric cryptographic standard?

A. DSA

B. PKI

C. RSA

D. 3DES

ANSWER: D

Explanation:

3DES is a symmetric-key encryption standard. Symmetric cryptography uses the same secret key for both encryption and decryption, so authorized parties must securely share and protect that key. Triple Data Encryption Standard applies the Data Encryption Standard cipher operation three times to each data block, traditionally with either two or three key values, to provide greater security than single DES. It is therefore classified as a symmetric block cipher. The National Institute of Standards and Technology describes Triple DES as a block-cipher algorithm and specifies its use with cryptographic keys in its guidance on key management. Although 3DES has been widely superseded by the more secure and efficient Advanced Encryption Standard, its design and classification remain those of a symmetric cryptographic standard. For modern systems, NIST has deprecated 3DES encryption because of its limited 64-bit block size and recommends transitioning to stronger alternatives such as AES. See NIST SP 800-67 Rev. 2: Recommendation for the Triple Data Encryption Algorithm and NIST SP 800-131A Rev. 2: Transitioning the Use of Cryptographic Algorithms and Key Lengths.

QUESTION NO: 76

How do you defend against ARP Spoofing? Select three.

A. Use ARPWALL system and block ARP spoofing attacks

B. Tune IDS Sensors to look for large amount of ARP traffic on local subnets

C. Use private VLANS

D. Place static ARP entries on servers, workstation and routers

ANSWER: A C D

Explanation:

“Use ARPWALL system and block ARP spoofing attacks,” “Use private VLANS,” and “Place static ARP entries on servers, workstation and routers

” are appropriate defensive measures against ARP spoofing. ARPWALL is specifically intended to detect and prevent malicious or inconsistent ARP mappings, helping protect hosts from forged ARP replies that could redirect local traffic through an attacker. Private VLANs provide Layer-2 isolation between hosts in the same VLAN, which reduces opportunities for one endpoint to directly impersonate another endpoint’s MAC/IP association or intercept its traffic. Static ARP entries bind a known IP address to a known MAC address on critical systems such as servers, routers, and administrative workstations. Because those bindings are not dynamically overwritten by unsolicited ARP replies, static mappings can prevent an attacker’s forged ARP response from replacing the legitimate association. In enterprise networks, these measures are commonly complemented by switch-based Dynamic ARP Inspection, which validates ARP packets against trusted address-binding information. Cisco describes Dynamic ARP Inspection and related Layer-2 protections in its Dynamic ARP Inspection guidance; the ARPWALL project is documented at SourceForge.

QUESTION NO: 77

Which of the following Trojans would be considered 'Botnet Command Control Center'?

A. YouKill DOOM

B. Damen Rock

C. Poison Ivy

D. Matten Kit

ANSWER: C

Explanation:

Poison Ivy is correct because it is a remote-access Trojan designed to give an operator persistent remote control over compromised hosts. In a botnet context, this type of malware supports command-and-control activity: infected machines establish communication with infrastructure controlled by the attacker, allowing the operator to issue commands, collect data, manage victims, and execute additional actions remotely. Poison Ivy has historically been used in targeted intrusions and is recognized as malware capable of communicating with a command-and-control server to receive instructions. A command-and-control center is not merely the infected endpoint; it is the attacker-controlled operational capability through which compromised systems are directed. Poison Ivy’s remote-administration architecture and operator tooling make it the fitting choice for a Trojan associated with botnet command-and-control functions. MITRE ATT&CK documents Poison Ivy as software used by threat actors and describes its command-and-control-related behavior: MITRE ATT&CK: PoisonIvy. Additional technical background on Poison Ivy and its use for remote access is available from Trend Micro Threat Encyclopedia.

QUESTION NO: 78

Ursula is a college student at a University in Amsterdam. Ursula originally went to college to study engineering but later changed to marine biology after spending a month at sea with her friends. These friends frequently go out to sea to follow and harass fishing fleets that illegally fish in foreign waters. Ursula eventually wants to put companies practicing illegal fishing out of business. Ursula decides to hack into the parent company's computers and destroy critical data knowing fully well that, if caught, she probably would be sent to jail for a very long time. What would Ursula be considered?

A. Ursula would be considered a gray hat since she is performing an act against illegal activities.

B. She would be considered a suicide hacker.

C. She would be called a cracker.

D. Ursula would be considered a black hat.

ANSWER: B

Explanation:

She would be considered a suicide hacker. This term describes an attacker who is willing to accept severe personal consequences, including a lengthy prison sentence, in pursuit of a strongly held cause or objective. The central fact pattern is not simply that Ursula plans an unauthorized intrusion or that she wants to damage a company's data; it is that she proceeds while fully recognizing the substantial likelihood and seriousness of punishment if she is identified. Her environmental and anti-illegal-fishing motivation supplies the ideological purpose commonly associated with this classification, while her willingness to sacrifice her freedom distinguishes the scenario. Destroying critical data on a parent company's systems would be unauthorized, destructive computer activity and can expose an attacker to serious criminal liability regardless of the cause asserted. The U.S. Department of Justice explains that intentionally accessing protected computers without authorization and causing damage is addressed under the Computer Fraud and Abuse Act: DOJ Computer Fraud guidance. For broader context on the harms and consequences of malicious cyber activity, see CISA Cyber Threats and Advisories.

QUESTION NO: 79

John runs a Web server, IDS and firewall on his network. Recently his Web server has been under constant hacking attacks. He looks up the IDS log files and sees no intrusion attempts but the Web server constantly locks up and needs rebooting due to various brute force and buffer overflow attacks but still the IDS alerts no intrusion whatsoever. John becomes suspicious and views the Firewall logs and he notices huge SSL connections constantly hitting his Web server. Hackers have been using the encrypted HTTPS protocol to send exploits to the Web server and that was the reason the IDS did not detect the intrusions. How would John protect his network from these types of attacks?

A. Install a proxy server and terminate SSL at the proxy

B. Enable the IDS to filter encrypted HTTPS traffic

C. Install a hardware SSL "accelerator" and terminate SSL at this layer

D. Enable the Firewall to filter encrypted HTTPS traffic

ANSWER: A C

Explanation:

Installing a proxy server and terminating SSL at the proxy is correct because the proxy acts as the HTTPS endpoint: it decrypts inbound TLS traffic before forwarding the request to the web application. An IDS or intrusion-prevention control positioned on the decrypted side of that proxy can then inspect HTTP requests and responses for attack signatures, anomalous request patterns, brute-force activity, and exploit payloads that would otherwise remain hidden inside encrypted traffic. The proxy can also enforce application-layer controls such as request-size limits, rate limiting, and web-application firewall policies.

Installing a hardware SSL "accelerator" and terminating SSL at this layer provides the same essential security capability while offloading computationally expensive TLS handshakes and cryptographic operations from the web server. Once the accelerator terminates TLS, security tooling can inspect plaintext traffic on the protected internal path, provided the architecture places the inspection control after decryption and the re-encryption path is appropriately protected. This design

preserves encrypted client access while restoring visibility into the application-layer content that must be examined for web attacks. See Cloudflare's SSL termination proxy overview and F5's SSL offloading description.

QUESTION NO: 80

A network admin contacts you. He is concerned that ARP spoofing or poisoning might occur on his network. What are some things he can do to prevent it?

Select the best answers.

A. Use port security on his switches.

B. Use a tool like ARPwatch to monitor for strange ARP activity.

C. Use a firewall between all LAN segments.

D. If you have a small network, use static ARP entries.

E. Use only static IP addresses on all PC's.

ANSWER: A B D

Explanation:

Use port security on his switches. is a useful layer of defense because it restricts which device MAC addresses may connect through an access port. Limiting unauthorized endpoint access reduces the opportunity for an attacker to join a switched LAN and send forged ARP replies. **Use a tool like ARPwatch to monitor for strange ARP activity.** provides essential detection: ARPwatch records IP-to-MAC associations and can alert administrators when an address unexpectedly changes its hardware address. This enables rapid investigation and containment of an attempted poisoning attack. **If you have a small network, use static ARP entries.** directly prevents hosts from dynamically accepting a fraudulent MAC-address mapping for the specified IP address, since the mapping is fixed rather than learned through ARP traffic. Static mappings are most practical for a limited number of important systems, such as gateways, servers, and administrative workstations. These measures are complementary: switch controls reduce exposure, ARP monitoring identifies suspicious changes, and static mappings protect selected critical communications. ARP behavior is defined in RFC 826; Cisco also documents switch-based ARP protection mechanisms in its Dynamic ARP Inspection guidance.

QUESTION NO: 81

Take a look at the following attack on a Web Server using obstructed URL:

http://www.example.com/script.ext?template%2e%2e%2e%2e%2e%2f%2e%2f%65%74%63%2f%70%61%73%73%77%64

The request is made up oF.

□ %2e%2e%2f%2e%2e%2f%2e%2f% = ../../../

□ %65%74%63 = etc

□ %2f = /

□ %70%61%73%73%77%64 = passwd

How would you protect information systems from these attacks?

A. Configure Web Server to deny requests involving Unicode characters.

B. Create rules in IDS to alert on strange Unicode requests.

C. Use SSL authentication on Web Servers.

D. Enable Active Scripts Detection at the firewall and routers.

ANSWER: B

Explanation:

Create rules in IDS to alert on strange Unicode requests is the best answer because the request attempts to conceal a directory-traversal payload through URL encoding. Once decoded and normalized, the encoded separators and dot characters can represent traversal sequences intended to move outside the web application's permitted directory and access sensitive operating-system files such as `/etc/passwd`. An intrusion detection system can inspect HTTP requests, apply URL decoding or normalization, and alert when it identifies suspicious traversal patterns, encoded path separators, repeated dot sequences, or requests for sensitive file paths. This provides defenders with visibility into attempted exploitation, supporting rapid investigation, blocking actions when integrated with prevention controls, and tuning based on observed attack traffic. Directory traversal is a well-established attack pattern in which attacker-controlled path input is manipulated to access files or directories outside the intended web root. OWASP describes the attack and emphasizes the importance of handling encoded traversal input safely: OWASP Path Traversal. Detection signatures should account for encoded payload forms because content inspection rules can match relevant request payload patterns after suitable normalization, as covered in the Snort content rule-option documentation.

QUESTION NO: 82

Which of the following is true of the wireless Service Set ID (SSID)? (Select all that apply.)

A. Identifies the wireless network

B. Acts as a password for network access

C. Should be left at the factory default setting

D. Not broadcasting the SSID defeats NetStumbler and other wireless discovery tools

E. Can be learned from wireless management traffic even when it is not broadcast in beacon frames

ANSWER: A E

Explanation:

An SSID identifies a wireless LAN and lets client devices distinguish the intended Wi-Fi network from other nearby networks. It is the human-readable network name configured on an access point or wireless router and selected by a user or device when joining that network. The SSID itself is not an authentication secret; access control is provided by mechanisms such as WPA2/WPA3 authentication and encryption, typically using a passphrase or enterprise credentials.

It is also correct that an SSID may be transmitted or otherwise learned through normal 802.11 management traffic, including beacon and probe-related exchanges. Consequently, suppressing the SSID from routine beacon advertisements does not provide meaningful protection against wireless discovery or an attacker monitoring traffic. Wireless-security best practice is to use WPA3 where feasible, or WPA2 with strong authentication and a strong passphrase, rather than treating the network name as a security control. Cisco's overview of SSIDs describes their network-identification role, and the Wi-Fi Alliance explains the authentication and protection provided by WPA3: Cisco: Understanding SSIDs; Wi-Fi Alliance: Wi-Fi Security.

QUESTION NO: 83

An attacker has been successfully modifying the purchase price of items purchased on the company's web site. The security administrators verify the web server and Oracle database have not been compromised directly. They have also verified the Intrusion Detection System (IDS) logs and found no attacks that could have caused this. What is the mostly likely way the attacker has been able to modify the purchase price?

A. By using SQL injection

B. By changing hidden form values

C. By using cross site scripting

D. By utilizing a buffer overflow attack

ANSWER: B

Explanation:

By changing hidden form values is correct because hidden fields are delivered to, and reside in, the client's browser. "Hidden" affects only how an HTML input is displayed; it does not provide integrity protection or make the value trustworthy. An attacker can inspect a purchase form with browser developer tools or intercept the request with a web proxy, alter a client-supplied price field, and submit the modified request through the application's normal purchase endpoint. This can occur without directly compromising either the web server or Oracle database and may not generate a recognizable intrusion-detection signature, since the resulting request can be syntactically valid application traffic.

Secure e-commerce applications must treat every client-provided value, including hidden inputs, as untrusted. The server should determine the authoritative item price from server-side product data using a product identifier, calculate totals on the server, and enforce authorization and business rules before completing the transaction. If state must be maintained across requests, it should be stored server-side or protected with appropriate integrity controls, but price calculation should still remain server-controlled. OWASP identifies parameter tampering as manipulation of parameters exchanged between client and server and recommends server-side validation and enforcement. See the OWASP Web Parameter Tampering guidance and the OWASP Input Validation Cheat Sheet.

QUESTION NO: 84

Which are true statements concerning the BugBear and Pretty Park worms?

Select the best answers.

A. Both programs use email to do their work.

B. Pretty Park propagates via network shares and email

C. BugBear propagates via network shares and email

D. Pretty Park tries to connect to an IRC server to send your personal passwords.

E. Pretty Park can terminate anti-virus applications that might be running to bypass them.

ANSWER: A C D

Explanation:

"Both programs use email to do their work" is correct because both worms use email-related functionality as part of their operation and propagation. Pretty Park uses the Windows Messaging API to distribute itself to contacts, while BugBear spreads through malicious email messages, typically relying on recipients to open its attachment. "BugBear propagates via network shares and email" is also correct: BugBear combines email propagation with the ability to copy itself across accessible network shares, increasing its reach within a connected environment. "Pretty Park tries to connect to an IRC server to send your personal passwords" accurately reflects its backdoor-like behavior. Pretty Park attempts to contact predefined Internet Relay Chat infrastructure, through which it can report information obtained from an infected system, including credentials or password-related information. These behaviors illustrate why email-borne worms can pose risks beyond initial infection: they can spread internally, communicate externally, and expose sensitive user data. Historical malware descriptions for these families are available from F-Secure's Pretty Park description and F-Secure's BugBear description.

QUESTION NO: 85

What techniques would you use to evade IDS during a Port Scan? (Select 4 answers)

A. Use fragmented IP packets

B. Spoof your IP address when launching attacks and sniff responses from the server

C. Overload the IDS with Junk traffic to mask your scan

D. Use source routing (if possible)

E. Connect to proxy servers or compromised Trojaned machines to launch attacks

ANSWER: A B D E

Explanation:

For an authorized security assessment, *Use fragmented IP packets* is a classic IDS-evasion technique because a scanner can split probe data across IP fragments, requiring the monitoring device to reassemble traffic accurately before applying port-scan detection logic. *Spoof your IP address when launching attacks and sniff responses from the server* can conceal the apparent origin of probes when return traffic can be observed through an appropriate network position or testing arrangement. *Use source routing (if possible)* leverages IPv4 source-route options to influence the path a packet takes, potentially avoiding a sensor located only on the normal route; it depends on network devices allowing those options. *Connect to proxy servers or compromised Trojaned machines to launch attacks* places the scan origin at an intermediary system, so network monitoring sees the intermediary rather than the assessor's direct address. These are established examples of traffic-manipulation, path-selection, and indirection techniques discussed in IDS/firewall-evasion material. Modern IDS/IPS platforms commonly normalize fragments, reject spoofed or source-routed traffic, and correlate distributed scans, so these methods are best understood for defensive testing and detection validation rather than as reliable bypasses. See the Nmap Reference Guide on firewall and IDS evasion and RFC 791 for IPv4 fragmentation and source-routing mechanisms.

QUESTION NO: 86

The traditional traceroute sends out ICMP ECHO packets with a TTL of one, and increments the TTL until the destination has been reached. By printing the gateways that generate ICMP time exceeded messages along the way, it is able to determine the path packets take to reach the destination.

The problem is that with the widespread use of firewalls on the Internet today, many of the packets that traceroute sends out end up being filtered, making it impossible to completely trace the path to the destination.

```
Juggyboy$ traceroute www.eccouncil.org
traceroute to www.eccouncil.org (64.147.99.90), 30 hops max, 52 byte packets
 1  * * *
 2  * * *
 3  ras.beamtele.net (183.82.15.69)  1.579 ms  1.513 ms  1.444 ms
 4  115.113.205.29.static-hyderabad.vsnl.net.in (115.113.205.29)  2.093 ms  1.963 ms  1.948 ms
 5  59.163.16.54.static.vsnl.net.in (59.163.16.54)  13.062 ms  13.094 ms  13.102 ms
 6  if-5-0-0-550.core2.cfo-chennai.as6453.net (116.0.84.69)  13.371 ms  13.103 ms  13.285 ms
 7  if-10-1-1-0.tcore2.cxr-chennai.as6453.net (180.87.37.18)  183.760 ms  165.805 ms  165.756 ms
 8  if-9-2.tcore2.mlv-mumbai.as6453.net (180.87.37.10)  172.479 ms  162.924 ms  162.835 ms
 9  if-6-2.tcore1.178-london.as6453.net (80.231.130.5)  151.203 ms  156.257 ms  150.901 ms
10  vlan704.icore1.ldn-london.as6453.net (80.231.130.10)  151.268 ms  152.167 ms  161.829 ms
11  * * *
12  ae-34-52.ebr2.london1.level3.net (4.69.139.97)  157.454 ms  151.607 ms  151.777 ms
13  ae-23-23.ebr2.frankfurt1.level3.net (4.69.148.194)  162.926 ms
    ae-22-22.ebr2.frankfurt1.level3.net (4.69.148.190)  170.020 ms
    ae-21-21.ebr2.frankfurt1.level3.net (4.69.148.186)  166.144 ms
14  ae-43-43.ebr2.washington1.level3.net (4.69.137.58)  236.524 ms
    ae-44-44.ebr2.washington1.level3.net (4.69.137.62)  246.080 ms  254.330 ms
15  ae-3-3.ebr1.newyork2.level3.net (4.69.132.90)  237.647 ms  252.050 ms
    ae-5-5.ebr2.washington12.level3.net (4.69.143.222)  258.821 ms
16  4.69.148.49 (4.69.148.49)  240.058 ms
    ae-4-4.ebr1.newyork1.level3.net (4.69.141.17)  242.545 ms
    4.69.148.49 (4.69.148.49)  240.874 ms
17  ae-61-61.csw1.newyork1.level3.net (4.69.134.66)  250.844 ms
    ae-71-71.csw2.newyork1.level3.net (4.69.134.70)  256.370 ms  242.690 ms
18  ae-34-89.car4.newyork1.level3.net (4.68.16.134)  250.200 ms
    ae-24-79.car4.newyork1.level3.net (4.68.16.70)  236.524 ms
    ae-14-69.car4.newyork1.level3.net (4.68.16.6)  255.573 ms
19  the-new-yor.car4.newyork1.level3.net (63.208.174.50)  249.250 ms  247.363 ms  243.364 ms
20  cs-nyi-gigalan-114.nyinternet.net (64.147.101.114)  240.236 ms  241.212 ms  240.654 ms
21  * * *         Request timed out
22  * * *         Request timed out
23  * * *         Request timed out
24  * * *         Request timed out
25  * * *         Request timed out
26  * * *         Request timed out
27  * * *         Request timed out
28  * * *         Request timed out
29  * * *         Request timed out
30  * * *         Request timed out

Destination Reached in 251 ms. Connection established to 64.147.99.90
Trace complete.
```

How would you overcome the Firewall restriction on ICMP ECHO packets?

A. Firewalls will permit inbound TCP packets to specific ports that hosts sitting behind the firewall are listening for connections. By sending out TCP SYN packets instead of ICMP ECHO packets, traceroute can bypass the most common firewall filters.

B. Firewalls will permit inbound UDP packets to specific ports that hosts sitting behind the firewall are listening for connections. By sending out TCP SYN packets instead of ICMP ECHO packets, traceroute can bypass the most common firewall filters.

C. Firewalls will permit inbound UDP packets to specific ports that hosts sitting behind the firewall are listening for connections. By sending out TCP SYN packets instead of ICMP ECHO packets, traceroute can bypass the most common firewall filters.

D. Do not use traceroute command to determine the path packets take to reach the destination instead use the custom hacking tool JOHNTHETRACER and run with the command

E. \ > JOHNTHETRACER www.eccouncil.org -F -evade

ANSWER: A

Explanation:

Firewalls will permit inbound TCP packets to specific ports that hosts sitting behind the firewall are listening for connections. By sending out TCP SYN packets instead of ICMP ECHO packets, traceroute can bypass the most common firewall filters. This describes TCP-based traceroute, a technique that retains the essential TTL-expiry mechanism while making probes resemble permitted application connection attempts. The traceroute sends TCP SYN probes with progressively increasing TTL values toward a TCP port that is likely to be allowed, such as a web-service port. Each intermediate router where the TTL reaches zero can return an ICMP Time Exceeded message, allowing the probing host to identify that hop. When a probe reaches the destination, the destination's TCP response—commonly SYN/ACK for an open port or RST for a closed port—shows that the endpoint has been reached. Because perimeter firewalls frequently allow TCP connections to published services, TCP SYN probes can traverse paths where ICMP Echo probes are filtered. This is a legitimate diagnostic approach when performed only against systems and networks for which authorization has been obtained. RFC 1393 describes traceroute's use of TTL and ICMP messages: RFC 1393. The implementation behavior and TCP SYN probe method are also documented by tcptraceroute documentation.

QUESTION NO: 87

Which of the following statement correctly defines ICMP Flood Attack? (Select 2 answers)

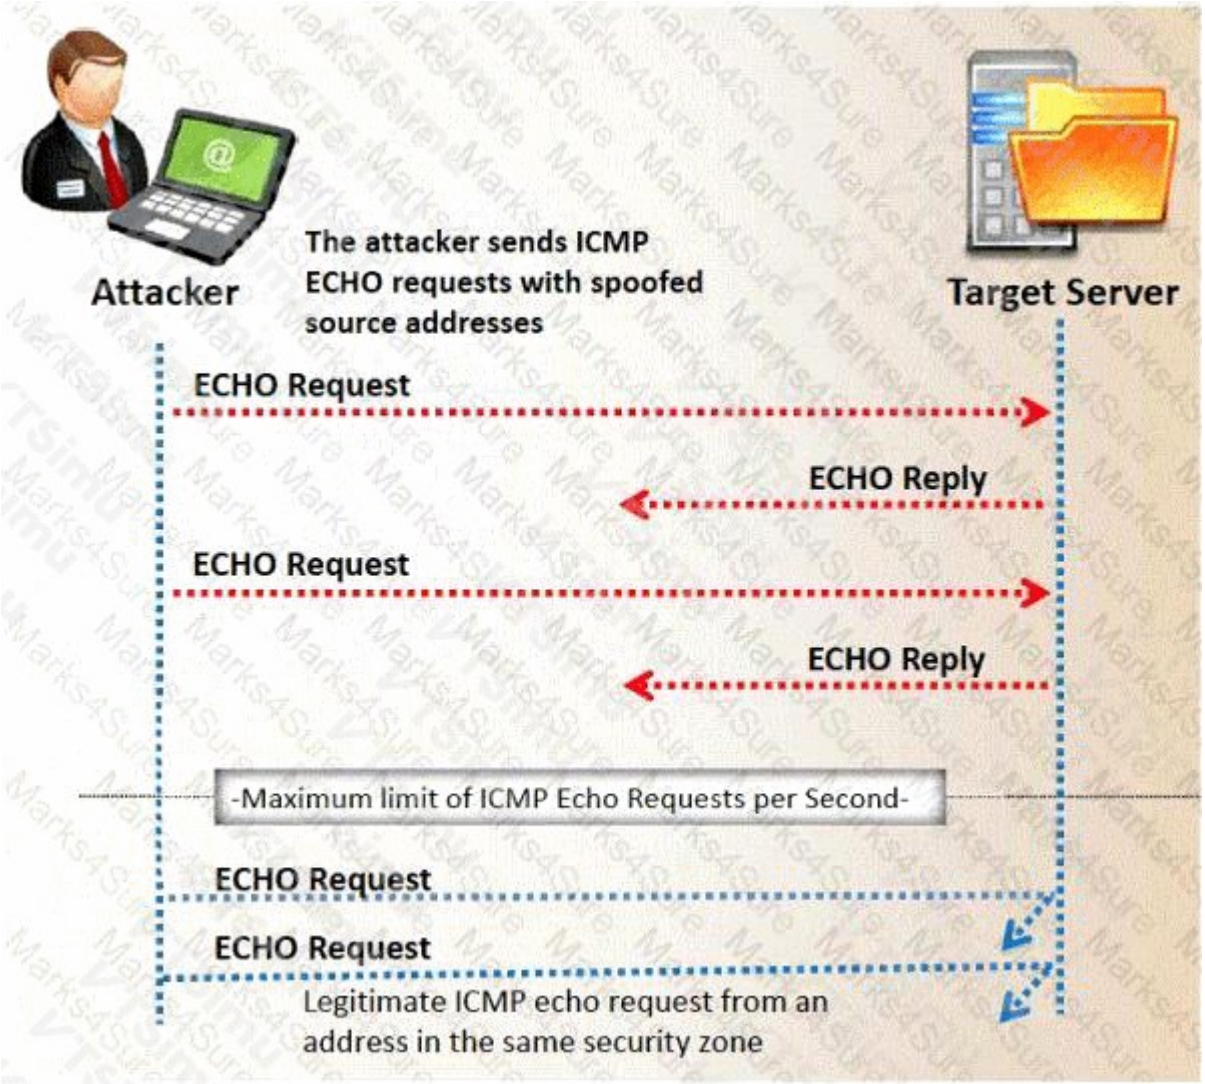

A. Bogus ECHO reply packets are flooded on the network spoofing the IP and MAC address

B. The ICMP packets signal the victim system to reply and the combination of traffic saturates the bandwidth of the victim ' s network

C. ECHO packets are flooded on the network saturating the bandwidth of the subnet causing denial of service

D. A DDoS ICMP flood attack occurs when the zombies send large volumes of ICMP_ECHO_REPLY packets to the victim system.

ANSWER: B D

Explanation:

An ICMP flood is a denial-of-service technique that abuses the Internet Control Message Protocol by sending ICMP traffic at a volume high enough to exhaust a target's available network capacity or processing resources. "The ICMP packets signal the victim system to reply and the combination of traffic saturates the bandwidth of the victim ' s network" correctly captures the common Echo-request form of the attack: the victim receives large quantities of ICMP Echo Requests and generates Echo Replies, so both incoming attack traffic and outgoing reply traffic consume resources and bandwidth. This amplification of work at the target is why sustained ICMP traffic can impair normal network services.

"A DDoS ICMP flood attack occurs when the zombies send large volumes of ICMP_ECHO_REPLY packets to the victim system." is also valid in the context of distributed ICMP flooding. A botnet can direct high volumes of ICMP Echo Reply traffic at a victim; regardless of whether flood traffic consists of Echo Requests or Echo Replies, the distributed sources collectively impose packet-processing and bandwidth pressure on the target. ICMP Echo message formats, including Echo and Echo Reply, are defined in RFC 792. For an overview of how ICMP/ping floods exhaust a victim's resources through high-rate ICMP traffic, see Cloudflare's ICMP flood explanation.

QUESTION NO: 88

A company has made the decision to host their own email and basic web services. The administrator needs to set up the external firewall to limit what protocols should be allowed to get to the public part of the company's network. Which ports should the administrator open? (Choose three.)

A. Port 22

B. Port 23

C. Port 25

D. Port 53

E. Port 80

F. Port 139

G. Port 445

ANSWER: C D E

Explanation:

Port 25, Port 53, and Port 80 are the appropriate inbound firewall openings for a public network segment providing basic email and web services. Port 25 is assigned to SMTP, which enables other Internet mail-transfer servers to deliver email to the organization's mail server. This is the standard server-to-server email delivery service and is listed in the IANA Service Name and Port Number Registry.

Port 53 supports DNS, which lets Internet users and external mail systems resolve the organization's public host names, including web host records and mail-exchanger records. DNS normally uses both UDP and TCP on this port: UDP handles most ordinary queries, while TCP is needed in circumstances such as large responses and zone transfers. Port 80 is the well-known port for HTTP and permits public access to the company's basic unencrypted web service. HTTP's use of TCP port 80 is specified by RFC 9110.

These openings should be restricted to the specific public-facing systems in a DMZ or equivalent segmented network. Firewall rules should permit only the needed transport protocols, log connections, and be paired with hardened, maintained DNS, mail, and web servers.

QUESTION NO: 89

What is the correct command to run Netcat on a server using port 56 that spawns command shell when connected?

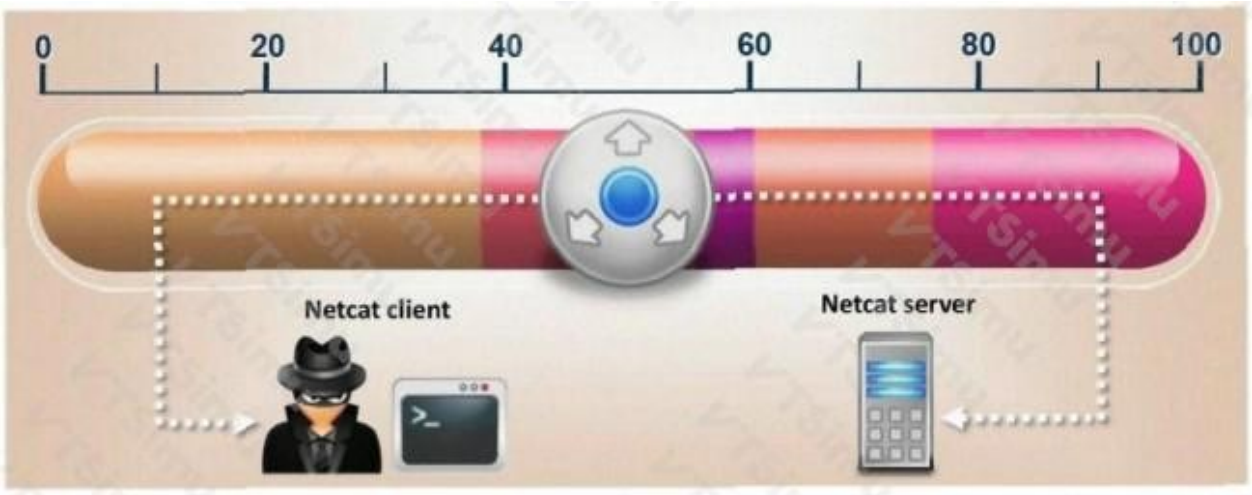

A. nc -port 56 -s cmd.exe

B. nc -p 56 -p -e shell.exe

C. nc -r 56 -c cmd.exe

D. nc -L -p 56 -t -e cmd.exe

ANSWER: D

Explanation:

`nc -L -p 56 -t -e cmd.exe` is the appropriate command for the legacy Windows Netcat syntax implied by this question. The `-p 56` parameter assigns local TCP port 56 for the listener. The uppercase `-L` creates a persistent listening service: after a client disconnects, Netcat continues listening for another connection. The `-t` setting enables Telnet negotiation support, which is commonly included in historical Windows Netcat listener examples. Finally, `-e cmd.exe` executes the Windows command interpreter and attaches its standard input, output, and error streams to the connected network session, thereby providing the requested command shell.

This syntax is associated with older Netcat implementations that expose the `-e` execution feature and the persistent `-L` listener option. Such functionality is highly dangerous outside a tightly controlled, authorized lab because it effectively creates a remote command-execution backdoor. Modern Ncat documentation similarly describes execution of a program in connection with listener operation, although its command-line syntax differs by implementation. Consult the relevant implementation's documentation before use: Ncat Users' Guide: Executing Commands and netcat-traditional manual page.

QUESTION NO: 90

Buffer X in an Accounting application module for Brownies Inc. can contain 200 characters. The programmer makes an assumption that 200 characters are more than enough. Because there were no proper boundary checks being conducted, Bob decided to insert 400 characters into the 200-character buffer. (Overflows the buffer). Below is the code snippet:

```
Void func (void)
{
int I; char buffer [200];
for (I=0; I<400; I++)
buffer [I]= 'A';
return;
}
```

How can you protect/fix the problem of your application as shown above?

A. Because the counter starts with 0, we would stop when the counter is less than 200

B. Because the counter starts with 0, we would stop when the counter is more than 200

C. Add a separate statement to signify that if we have written less than 200 characters to the buffer, the stack should stop because it cannot hold any more data

D. Add a separate statement to signify that if we have written 200 characters to the buffer, the stack should stop because it cannot hold any more data

ANSWER: A D

Explanation:

Because the counter starts with zero, the valid positions in a 200-character buffer run from zero through 199. The loop or write operation must therefore continue only while the counter is less than 200; this ensures that no write is attempted at position 200 or beyond. This is a fundamental bounds-checking control: the program validates the destination capacity before every access rather than relying on an assumption about the amount of supplied input.

Adding a condition that stops further writes once 200 characters have been written also correctly enforces the buffer's maximum capacity. In a production implementation, this limit should be applied before each copy or append operation, and the application should either reject, truncate according to an explicit policy, or otherwise safely handle excess input. Where available, developers should favor memory-safe string and collection APIs that accept the destination size and report truncation or failure. Such validation prevents input from overwriting adjacent memory and is a core mitigation for buffer-overflow weaknesses. See the MITRE CWE-120 buffer-copy guidance and the OWASP Buffer Overflow overview.

QUESTION NO: 91

Which of the following controls can help prevent Cross-Site Request Forgery (CSRF) attacks? (Choose three.)

A. Use unpredictable anti-CSRF tokens

B. Set cookies with an appropriate SameSite attribute

C. Require re-authentication or a separate user confirmation for sensitive transactions

D. Disable JavaScript on the web server

E. Use only GET requests for sensitive operations

ANSWER: A B C

Explanation:

Use unpredictable anti-CSRF tokens is correct because the application can require a secret, per-session or per-request value in each state-changing request. A malicious external site cannot normally read that token because of the browser same-origin policy, so it cannot construct a valid forged request.

Set cookies with an appropriate SameSite attribute is also correct. SameSite cookie settings can limit whether browser cookies are automatically attached to cross-site requests, reducing the ability of a forged request to execute in the victim's authenticated session. **Require re-authentication or a separate user confirmation for sensitive transactions** provides an additional safeguard for high-impact actions, such as changing payment details or transferring funds.

Input validation is important for many web vulnerabilities, but it does not by itself establish that a request originated from the legitimate application. CSRF protections should be applied to all state-changing functions. See the OWASP CSRF overview and the Mozilla SameSite cookie documentation.

QUESTION NO: 92

Which of the following represent weak password? (Select 2 answers)

A. Passwords that contain letters, special characters, and numbers ExamplE. ap1$%##f@52

B. Passwords that contain only numbers ExamplE. 23698217

C. Passwords that contain only special characters ExamplE. &*#@!(%)

D. Passwords that contain letters and numbers ExamplE. meerdfget123

E. Passwords that contain only letters ExamplE. QWERTYKLRTY

F. Passwords that contain only special characters and numbers ExamplE. 123@$45

G. Passwords that contain only letters and special characters ExamplE. bob@&ba

ANSWER: B E

Explanation:

“Passwords that contain only numbers ExamplE. 23698217” is weak because an eight-digit, numeric-only password has a very limited search space: only 100 million possible combinations. Modern password-cracking tools can test such a constrained format efficiently, particularly when an attacker knows or can infer that the password policy permits digits alone. Numeric-only passwords also commonly resemble memorable sequences, dates, or identifiers, which makes guessing attacks more effective.

“Passwords that contain only letters ExamplE. QWERTYKLRTY” is weak because the example is based on a highly recognizable keyboard pattern. Attackers do not rely solely on exhaustive brute force; they use dictionaries, password lists, and rule-based guessing techniques that prioritize common strings and keyboard walks. A password that is easily recognizable or predictable can therefore be compromised far faster than its raw character count might suggest. Current guidance emphasizes password length and resistance to common, expected, or compromised values rather than relying on a simple character-type requirement. See NIST SP 800-63B and CISA guidance on strong passwords.

QUESTION NO: 93

Snort has been used to capture packets on the network. On studying the packets, the penetration tester finds it to be abnormal. If you were the penetration tester, why would you find this abnormal?

05/20-17:0645.061034 192.160.13.4:31337 --> 172.16.1.101:1

TCP TTL:44 TOS:0x10 ID. 242

***FRP** Seq:0xA1D95 Ack:0x53 Win: 0x400

What is odd about this attack? (Choose the most appropriate statement)

A. This is not a spoofed packet as the IP stack has increasing numbers for the three flags.

B. This is back orifice activity as the scan comes from port 31337.

C. The attacker wants to avoid creating a sub-carrier connection that is not normally valid.

D. There packets were created by a tool; they were not created by a standard IP stack.

ANSWER: D

Explanation:

There packets were created by a tool; they were not created by a standard IP stack. Snort’s TCP flag display shows the packet has the FIN, RST, and PSH flags set together (`FRP`). This is an abnormal combination for ordinary TCP communications. A reset is used to immediately terminate or refuse a connection, whereas FIN is used for the orderly closing sequence of an established connection. PSH requests prompt delivery of data to the receiving application. Combining these control intents in one packet is characteristic of deliberately crafted traffic rather than traffic emitted through the normal state transitions of a conventional TCP implementation.

Penetration testers and attackers may use raw-packet tools to construct unusual flag combinations for reconnaissance, firewall/IDS testing, evasion experiments, or to observe how a target’s TCP stack responds to malformed or unexpected

segments. The source port value may be suggestive, but it is not sufficient evidence by itself to identify the activity as a specific Back Orifice communication; ports are easily selected or forged in crafted packets. The meaningful anomaly is the contradictory TCP control-flag pattern, which indicates packet construction outside typical IP/TCP stack behavior. TCP reset and connection-closing behavior is defined in RFC 9293; Snort documentation is available from Snort Documents.

QUESTION NO: 94

Bill has successfully executed a buffer overflow against a Windows IIS web server. He has been able to spawn an interactive shell and plans to deface the main web page. He first attempts to use the "echo" command to simply overwrite index.html and remains unsuccessful. He then attempts to delete the page and achieves no progress. Finally, he tries to overwrite it with another page in which also he remains unsuccessful. What is the probable cause of Bill's problem?

A. You cannot use a buffer overflow to deface a web page

B. There is a problem with the shell and he needs to run the attack again

C. The HTML file has permissions of read only

D. The system is a honeypot

ANSWER: C

Explanation:

The HTML file has permissions of read only is correct. Successfully obtaining an interactive shell through a vulnerable IIS process does not automatically provide administrative or write-level access to website content. The shell runs under the security context assigned to the compromised process, commonly an IIS application-pool or service identity. That identity may be deliberately restricted to reading and serving files from the web root while being denied permission to modify, replace, or delete them. Consequently, each attempted change to `index.html` fails even though command execution itself has been achieved. This is an example of operating-system access control limiting the impact of a server-side compromise: effective permissions on the target file and its containing directory determine whether the process can write or delete content. On Windows, a read-only file attribute can contribute to this behavior, and—more importantly in IIS deployments—NTFS access-control entries can prevent the worker-process identity from changing the page. Microsoft documents that IIS application pools run under identities whose access must be explicitly granted to resources, and that Windows file permissions govern access to files and folders. See IIS Application Pool Identities and Microsoft documentation for Windows file permissions.

QUESTION NO: 95

What are the two basic types of attacks? (Choose two.

A. DoS

B. Passive

C. Sniffing

D. Active

E. Cracking

ANSWER: B D

Explanation:

The correct classifications are **Passive** and **Active**. These are the two foundational categories used in information-security literature to describe how an adversary interacts with a target system or its communications. A passive attack obtains information without altering system resources or interfering with normal operation. It is primarily an attack on confidentiality, such as observing communications or collecting transmitted data. An active attack involves an attempt to alter resources, modify data, disrupt services, inject content, or otherwise affect the operation or state of a system. This classification is

useful because it guides defensive priorities: passive attacks require strong confidentiality controls and detection of unauthorized observation, while active attacks require integrity, authentication, availability, and recovery protections. NIST defines a passive attack as one that attempts to learn or use information from a system without affecting system resources, and defines an active attack as one that attempts to alter system resources or affect their operation. These definitions support treating Passive and Active as the basic attack types. See the NIST definition of passive attack and the NIST definition of active attack.

QUESTION NO: 96

Exhibit:

The following is an entry captured by a network IDS.You are assigned the task of analyzing this entry. You notice the value 0x90, which is the most common NOOP instruction for the Intel processor. You figure that the attacker is attempting a buffer overflow attack. You also notice "/bin/sh" in the ASCII part of the output. As an analyst what would you conclude about the attack?

A. The buffer overflow attack has been neutralized by the IDS

B. The attacker is creating a directory on the compromised machine

C. The attacker is attempting a buffer overflow attack and has succeeded

D. The attacker is attempting an exploit that launches a command-line shell

ANSWER: D

Explanation:

The attacker is attempting an exploit that launches a command-line shell. Repeated `0x90` bytes are commonly called a NOP sled: in Intel x86 machine code, byte `0x90` represents the `NOP` instruction, which performs no operation. Attack payloads have historically used a sequence of NOPs to increase the likelihood that redirected execution reaches injected shellcode. The visible string `/bin/sh` is a strong indicator that the shellcode's intended payload is to execute the Unix Bourne-compatible command shell. Together, these artifacts support the conclusion that the IDS observed an attempted memory-corruption exploit carrying a shell-launching payload. Importantly, a network IDS entry can identify the attack pattern and payload intent, but the entry by itself does not establish that exploitation succeeded on the destination host. The Intel instruction reference documents `NOP` as opcode `90`, and MITRE ATT&CK describes Unix shell execution as command and scripting interpreter activity. See the Intel x86 NOP instruction reference and MITRE ATT&CK: Unix Shell.

QUESTION NO: 97

This is an example of whois record.

```
Registrant:
Jason Springfield, Inc
11807 N.E. 99th Street, Suite 1100
New York, NY 98682
USA

Registrar: Jason Springfield (http://www. jspringfield.com)
Domain Name: jspringfield.com
Created on: 29-DEC-10
Expires on: 29-DEC-14
Last Updated on: 23-FEB-11

Administrative Contact:
Contact, Admin Jack_Smith@jspringfield.com
Jason Springfield, Inc
11807 N.E. 99th Street, Suite 1100
New York, NY 98682
USA
360.253.6744
360.253.3556

Technical Contact:
Contact, Technical Sheela_Ravin@jspringfield.com
Jason Springfield, Inc
11807 N.E. 99th Street, Suite 1100
New York, NY 98682
USA
360.253.3456
360.253.2675

Billing Contact:
Contact, Technical David_Bruce@jspringfield.com
Jason Springfield, Inc
11807 N.E. 99th Street, Suite 1100
New York, NY 98682
USA
360.253.6654
360.253.1256

Domain servers (DNS) in listed order:
NS1.jspringfield.com
NS2.jspringfield.com
```

Sometimes a company shares a little too much information on their organization through public domain records. Based on the above whois record, what can an attacker do? (Select 2 answers)

A. Search engines like Google, Bing will expose information listed on the WHOIS record

B. An attacker can attempt phishing and social engineering on targeted individuals using the information from WHOIS record

C. Spammers can send unsolicited e-mails to addresses listed in the WHOIS record

D. IRS Agents will use this information to track individuals using the WHOIS record information

ANSWER: B C

Explanation:

An attacker can attempt phishing and social engineering on targeted individuals using the information from WHOIS record because registration data may identify contacts responsible for a domain, along with organizational details, email addresses, telephone numbers, and physical addresses. Such details help an attacker create a convincing pretext, impersonate a registrar or internal support team, and tailor a message to a specific role or person. This is a common reconnaissance objective: publicly available information is correlated with other sources to make social-engineering attempts more credible.

Spammers can send unsolicited e-mails to addresses listed in the WHOIS record because publicly disclosed registrant or administrative contact addresses can be collected and used as delivery targets for bulk email, scams, or other unwanted communications. WHOIS privacy services and modern registration-data redaction reduce this exposure in many cases, but when contact data is publicly visible, it remains useful to malicious actors. ICANN describes registration data as information associated with domain-name registrations and explains its availability through Registration Data Lookup Service processes. The United States Cybersecurity and Infrastructure Security Agency also identifies phishing as a technique that uses deceptive communications to obtain information or induce harmful actions. See ICANN Registration Data Lookup FAQs and CISA phishing guidance.

QUESTION NO: 98

A developer for a company is tasked with creating a program that will allow customers to update their billing and shipping information. The billing address field used is limited to 50 characters. What pseudo code would the developer use to avoid a buffer overflow attack on the billing address field?

A. if (billingAddress = 50) {update field} else exit

B. if (billingAddress != 50) {update field} else exit

C. if (billingAddress >= 50) {update field} else exit

D. if (billingAddress <= 50) {update field} else exit

ANSWER: D

Explanation:

The appropriate validation is *if (billingAddress <= 50) {update field} else exit*, because input must be accepted only when its length does not exceed the maximum capacity allocated for the billing-address field. In practical code, the condition should explicitly test the string length, such as `if (billingAddress.length <= 50)`, before copying, storing, or otherwise processing the supplied value. This permits valid addresses from zero through 50 characters while rejecting values that would exceed the defined field boundary.

Enforcing a strict maximum length at the application boundary is an important input-validation control. The validation must be performed server-side, even if the user interface also limits entry length, because an attacker can bypass client-side controls and submit an oversized request directly. The receiving implementation should also use bounded memory-safe APIs or memory-safe language features so that data is never written beyond the destination buffer. OWASP recommends validating all untrusted input using allowlist rules and enforcing appropriate length constraints; see the OWASP Input Validation Cheat Sheet. This control addresses the risk described by CWE-120: Buffer Copy without Checking Size of Input.

QUESTION NO: 99

What does a type 3 code 13 represent?(Choose two.

A. Echo request

B. Destination unreachable

C. Network unreachable

D. Administratively prohibited

E. Port unreachable

F. Time exceeded

ANSWER: B D

Explanation:

“Destination unreachable” and “Administratively prohibited” are correct because ICMP uses a type-and-code structure: Type 3 identifies the overall ICMP error-message category as Destination Unreachable, while Code 13 gives the specific reason, Communication Administratively Prohibited. This response indicates that a router, firewall, access-control list, or other policy-enforcing device deliberately blocked the attempted communication based on administrative policy. During ethical hacking reconnaissance and troubleshooting, seeing this ICMP response is useful evidence that the target path is reachable enough for an intermediate device to evaluate the traffic, but a configured security rule prevents delivery. It can therefore reveal packet-filtering behavior and help distinguish policy-based blocking from failures such as an absent route or an unavailable service. The Internet Engineering Task Force’s router requirements document defines ICMP Type 3, Code 13 as “Communication Administratively Prohibited.” ICMP’s destination-unreachable message category is also defined in the core ICMP specification. See RFC 1812, section 5.2.7.1 and RFC 792.

QUESTION NO: 100

What ports should be blocked on the firewall to prevent NetBIOS traffic from not coming through the firewall if your network is comprised of Windows NT, 2000, and XP?(Choose all that apply.

A. 110

B. 135

C. 139

D. 161

E. 445

F. 1024

ANSWER: B C E

Explanation:

For legacy Windows NT, Windows 2000, and Windows XP environments, blocking **135**, **139**, and **445** at the network perimeter is an appropriate defensive control for reducing exposure to Windows file-sharing, NetBIOS-session, and related remote procedure call services. Port 139 is the TCP NetBIOS Session Service port used by SMB when it is carried over NetBIOS. Port 445 is used by Microsoft-DS/SMB directly over TCP, without the NetBIOS session layer, and is a major exposure point for Windows file and printer sharing. Port 135 is the Microsoft RPC Endpoint Mapper port; although it is not itself a NetBIOS port, it is commonly blocked with 139 and 445 in legacy Windows firewall filtering because RPC services can support remote enumeration and management paths associated with Windows-host attacks. Microsoft documents that SMB uses TCP 445 and, when NetBIOS over TCP/IP is used, TCP 139. Microsoft also identifies TCP 135 as the RPC Endpoint Mapper service. These blocks should be applied at untrusted boundaries while allowing only specifically required internal management or file-sharing traffic through tightly scoped rules. See Microsoft’s service and port requirements and Microsoft’s RPC documentation.

QUESTION NO: 101

What is the main disadvantage of the scripting languages as opposed to compiled programming languages?

A. Scripting languages are hard to learn.

B. Scripting languages are not object-oriented.

C. Scripting languages cannot be used to create graphical user interfaces.

D. Scripting languages are slower because they require an interpreter to run the code.

ANSWER: D

Explanation:

Scripting languages are slower because they require an interpreter to run the code. Traditionally, a compiled language is translated ahead of execution into native machine instructions for a particular target platform. A scripting language commonly executes through an interpreter or a runtime environment, which must process the program at run time. This additional execution layer can introduce overhead, especially for CPU-intensive tasks and repeated operations, so scripts may perform less quickly than equivalent programs compiled directly to native code.

Modern language implementations can reduce this gap substantially. For example, some runtimes use just-in-time compilation, caching, or optimized virtual machines; therefore, the statement describes a general trade-off rather than an absolute rule applicable to every workload. In ethical-hacking work, scripting remains highly useful for quickly automating reconnaissance, parsing output, calling web APIs, and building proof-of-concept tooling. When low-level speed, tight resource use, or direct native-code execution is the primary requirement, compiled implementations are often preferred. Python's documentation describes the Python interpreter and its execution model, while the Microsoft documentation explains how just-in-time compilation converts managed code into native instructions during execution: Python documentation: Using the Python Interpreter; Microsoft: Managed execution process.

QUESTION NO: 102

What do you conclude from the nmap results below?

Staring nmap V. 3.10ALPHA0 (www.insecure.org/map/)

(The 1592 ports scanned but not shown below are in statE. closed)

Port State Service

21/tcp open ftp

25/tcp open smtp

80/tcp open http

443/tcp open https

Remote operating system guess: Too many signatures match the reliability guess the OS. Nmap run completed – 1 IP address (1 host up) scanned in 91.66 seconds

A. The system is a Windows Domain Controller.

B. The system is not firewalled.

C. The system is not running Linux or Solaris.

D. The system is not properly patched.

ANSWER: B

Explanation:

The system is not firewalled is the intended conclusion from this scan output. Nmap reports four TCP ports as open and explicitly classifies the remaining 1,592 scanned ports as closed. A closed TCP port means that the target responded to the probe but that no service is listening on that port. This is materially different from Nmap's filtered state, which is used when a packet-filtering device prevents Nmap from determining whether a port is open or closed. Because the output accounts for the non-open ports as closed rather than filtered, the scan did not encounter filtering that concealed those ports. Thus, in the terminology and assumptions used by this question, the host is reachable without a firewall filtering the scanned TCP probes.

This conclusion is limited to the scan path, port range, protocol, and scan technique used. In a real assessment, a host could still be protected by a firewall that permits or actively rejects the tested traffic, or by controls affecting other ports and protocols. Nmap's port-state definitions distinguish closed ports, which are accessible but have no listening application, from filtered ports, where filtering obstructs Nmap's determination. See the Nmap port-scanning state documentation and the Nmap OS-detection documentation.