DUMPSARENA

Certified Ethical Hacker Exam (CEHv13)

ECCouncil 312-50v13

Version Demo

Total Demo Questions: 99

Total Premium Questions: 1477

Buy Premium PDF

https://dumpsarena.co

sales@dumpsarena.co

sales@dumpsarena.co
dumpsarena.co

Topic Break Down

Topic	No. of Questions
Topic 1, Information Security and Ethical Hacking Overview	185
Topic 2, Reconnaissance Techniques	269
Topic 3, System Hacking	225
Topic 4, Network and Perimeter Hacking	240
Topic 5, Web Application Hacking	222
Topic 6, Wireless Network Hacking	76
Topic 7, Mobile Platform, IoT, and OT Hacking	98
Topic 8, Cloud Computing	62
Topic 9, Cryptography	100
Total	1477

QUESTION NO: 1

A Windows environment has legacy name-resolution protocols enabled. A security assessment identifies that an attacker could respond to LLMNR and NetBIOS Name Service queries, capture challenge-response traffic, and attempt credential-relay attacks.

Which of the following actions reduce the attack surface associated with this condition? (Choose three.)

A. Disable LLMNR through policy.

B. Disable NetBIOS over TCP/IP where it is not needed.

C. Require SMB signing on clients and servers.

D. Enable anonymous access to IPC$ shares.

E. Use a shared local administrator password on all hosts.

ANSWER: A B C

Explanation:

Disabling LLMNR and disabling NetBIOS over TCP/IP where they are not required remove the broadcast or multicast name-resolution mechanisms that attackers commonly abuse when normal DNS resolution fails. Requiring SMB signing is also an important protection because it helps prevent an attacker from relaying captured NTLM authentication traffic to SMB services that require signed communications.

These controls should be implemented with appropriate testing because some legacy applications may depend on NetBIOS naming. Organizations should use properly configured DNS for host-name resolution and remove obsolete protocols rather than relying on unauthenticated fallback mechanisms. MITRE ATT&CK documents LLMNR and NetBIOS poisoning as part of Adversary-in-the-Middle: LLMNR/NBT-NS Poisoning and SMB Relay. Microsoft also documents SMB-signing protections in its SMB signing guidance.

QUESTION NO: 2

A telecommunications provider in Toronto operates a monitoring platform that analyzes inbound traffic streams during suspected denial-of-service conditions. The system converts traffic measurements into signal components and evaluates their energy across multiple frequency ranges to distinguish abnormal traffic bursts from background network noise.

Rather than focusing on traffic baselines or identifying the exact statistical breakpoint where behavior changes, the platform identifies anomalies by decomposing traffic signals into spectral components for analysis.

Which DDoS detection technique is being used in this scenario?

A. Traffic Pattern Analysis

B. Sequential Change-Point Detection

C. Activity Profiling

D. Wavelet-Based Signal Analysis

ANSWER: D

Explanation:

Wavelet-Based Signal Analysis is correct because the defining feature in the scenario is treating network-traffic measurements as a signal, decomposing that signal into components at different scales or frequency ranges, and examining the resulting energy to identify anomalous bursts. Wavelet methods provide simultaneous time- and frequency-oriented analysis: they can expose short-lived, high-volume traffic changes while retaining information about when those changes occurred. This is useful for denial-of-service and distributed denial-of-service monitoring, where attack traffic may create abrupt changes in traffic volume or structure that are difficult to isolate from normal background variation using a single aggregate measurement.

In CEH-aligned terminology, wavelet-based analysis is a signal-analysis approach for detecting DoS/DDoS-related anomalies in traffic data. The platform's focus on spectral components and multiscale energy evaluation directly matches this technique. Wavelet transforms are specifically designed for multiresolution signal analysis; the PyWavelets discrete wavelet transform documentation describes decomposition into approximation and detail coefficients across levels. For broader security context on detecting and handling denial-of-service incidents, see NIST's Computer Security Incident Handling Guide.

QUESTION NO: 3

Which of the following tools can be used to perform a zone transfer?

A. NSLookup

B. Finger

C. Dig

D. Sam Spade

E. Host

F. Netcat

G. Neotrace

ANSWER: A C D E

Explanation:

A DNS zone transfer is an authorized replication operation in which a secondary DNS server obtains DNS zone data from the primary server. During security testing, attempting an AXFR request can reveal whether a DNS server has been incorrectly configured to permit transfers to arbitrary clients. **NSLookup**, **Dig**, **Sam Spade**, and **Host** are appropriate tools for requesting or enumerating DNS zone-transfer data when the target name server permits it. In particular, `dig @server zone AXFR` explicitly requests a full zone transfer, while the host utility supports zone listing with its `-l` mode. NSLookup has historically provided the `ls` command for listing a domain from a selected DNS server where zone transfers are allowed. Sam Spade, a DNS and network-investigation utility, also includes DNS enumeration capabilities that can be used for zone-transfer checks. In a properly administered environment, authoritative DNS servers restrict AXFR/IXFR to designated secondary servers, typically using source-address ACLs and/or transaction signatures, so these requests should be performed only with authorization. See the BIND 9 command documentation for dig and host usage, and Microsoft's nslookup documentation for interactive DNS-query commands.

QUESTION NO: 4

Nicolas just found a vulnerability on a public-facing system that is considered a zero-day vulnerability. He sent an email to the owner of the public system describing the problem and how the owner can protect themselves from that vulnerability. He also sent an email to Microsoft informing them of the problem that their systems are exposed to. What type of hacker is Nicolas?

A. Red hat

B. white hat

C. Black hat

D. Gray hat

ANSWER: D

Explanation:

Gray hat is correct because Nicolas identified a zero-day flaw in a public-facing system without any stated authorization to test that system, then disclosed the finding to the affected system owner and the vendor whose products may be exposed. In common ethical-hacking terminology, gray-hat activity describes security research or vulnerability discovery performed without explicit permission but with a non-malicious intent to notify parties who can address the risk. The key detail is that Nicolas provides remediation information and alerts Microsoft, supporting a responsible, defensive outcome rather than exploitation for harm or personal gain.

Zero-day vulnerabilities are particularly important to report promptly because they may be unknown to the vendor and therefore have no available patch when discovered. Coordinated vulnerability disclosure gives vendors and affected organizations the opportunity to investigate, develop mitigations or updates, and communicate protective guidance. Microsoft's Security Response Center provides a vulnerability-reporting channel for researchers to submit potential security vulnerabilities in Microsoft products and services. CISA likewise promotes coordinated vulnerability disclosure as a process for receiving, validating, and remediating reported weaknesses. See Microsoft Security Response Center and CISA Vulnerability Disclosure Policy resources.

QUESTION NO: 5

Which information CANNOT be directly obtained from DNS interrogation?

A. Usernames and passwords

B. Server geolocation (via IPs)

C. Subdomains of the organization

D. IP addresses of mail servers

ANSWER: A

Explanation:

Usernames and passwords cannot be directly obtained through DNS interrogation. DNS is a distributed naming system that answers queries about domain-related resource records, such as host address records, mail-exchanger records, name-server records, aliases, and service-related records. Interrogating a target's DNS infrastructure can therefore reveal information useful during reconnaissance, including mail-server hostnames and associated IP addresses, as well as discovered subdomains when relevant records are publicly resolvable. IP addresses obtained from DNS can also be used as input to a separate IP-geolocation lookup, although the geographic result itself is not a DNS record.

Authentication credentials are outside DNS's intended data model. Usernames and passwords are managed by identity providers, directory services, operating systems, or individual applications, and are not returned as part of standard DNS resource-record queries. Even where DNS supports security mechanisms such as DNSSEC, those mechanisms protect the authenticity and integrity of DNS data; they do not expose account credentials. This distinction is important in ethical-hacking reconnaissance: DNS enumeration helps map externally visible naming and service infrastructure, while credential discovery requires authorization and different assessment methods. See the DNS overview in RFC 1034 and the resource-record definitions in RFC 1035.

QUESTION NO: 6

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

A. A zone transfer is accomplished with the DNS

B. A zone transfer is accomplished with the nslookup service

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

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

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

F. Zone transfers cannot occur on the Internet

ANSWER: A C E

Explanation:

A zone transfer is a DNS protocol operation used to replicate authoritative zone data between DNS servers, normally from a primary server to an authorized secondary server. “A zone transfer is accomplished with the DNS” is correct because the transfer is performed through DNS zone-transfer mechanisms, principally full transfers (AXFR) and incremental transfers (IXFR). A full AXFR transfer provides the complete contents of a zone, making “A zone transfer passes all zone information that a DNS server maintains” correct in the context of the transferred authoritative zone. This replication enables secondary authoritative servers to answer consistently with the primary server and improves DNS availability.

DNS zone transfers use TCP port 53; RFC 5936 specifies the DNS AXFR protocol over TCP. Consequently, blocking all inbound TCP port 53 connections prevents a server from accepting inbound zone-transfer sessions. In operational environments, administrators generally restrict TCP/53 zone transfers to explicitly authorized secondary servers rather than broadly exposing the service. This protects zone records from unauthorized disclosure while preserving necessary replication. See the IETF RFC 5936 DNS Zone Transfer Protocol (AXFR) and Microsoft’s DNS zone transfer guidance.

QUESTION NO: 7

At a private aerospace research facility in Mesa, Arizona, an executive raises concerns after sensitive discussion points from speakerphone meetings begin surfacing externally. The device shows no indicators of active audio recording, and application permission history does not reflect recent camera or microphone authorization changes.

A forensic mobile analysis identifies that an installed application has been continuously reading motion sensor output while the phone’s loudspeaker is active. The collected sensor data was later transmitted to a remote server, where acoustic characteristics were reconstructed from the recorded measurements.

Identify the attack technique responsible for this compromise.

A. Camfecting

B. StormBreaker Abuse

C. Android Camera Hijack Attack

D. Spearphone Attack

ANSWER: D

Explanation:

Spearphone Attack is correct because it is an acoustic side-channel attack that uses a smartphone’s motion sensors, particularly the accelerometer, to infer speech or other audio emitted by the device’s own loudspeaker. When the loudspeaker operates, its vibrations can propagate through the phone chassis and create measurable accelerometer signals. An application can collect those signals without requesting microphone access, then transmit the measurements for processing and reconstruction by an external server.

The key forensic indicators in the scenario align directly with this technique: continuous motion-sensor reads, loudspeaker activity, no apparent recording indicator, no recent microphone-permission grant, and later reconstruction of acoustic characteristics from sensor data. Motion sensors have historically been accessible to applications with fewer protections than microphone data, making this an important privacy risk when sensor output can be correlated with device-generated speech.

The original Spearphone research demonstrates that speech played through a smartphone loudspeaker may be inferred from accelerometer measurements and describes the relevant side channel and threat model. See the research paper at Spearphone: A Novel Remote Eavesdropping Attack via Motion Sensors. For Android’s current treatment of motion sensors and related privacy considerations, see Android Developers: Sensors overview.

QUESTION NO: 8

During an authorized penetration test of an organization ' s Operational Technology (OT) environment, the tester has already identified exposed industrial assets and now begins actively probing controllers, services, and interfaces to identify exploitable weaknesses. No exploitation attempts or persistence mechanisms have been performed yet.

According to the OT hacking methodology, which phase is currently being carried out?

A. Gain Remote Access

B. Information Gathering

C. Launch Attacks

D. Vulnerability Scanning

ANSWER: D

Explanation:

Vulnerability Scanning is the current phase because the tester is actively assessing already identified OT assets to discover weaknesses that could later be used for exploitation. In an OT assessment, this may include carefully interrogating exposed services and management interfaces, enumerating controller or gateway characteristics, identifying firmware and software versions, checking authentication exposure, and correlating observed configurations with known vulnerabilities. The defining goal is to identify and validate potential attack paths, not to use them to obtain access or maintain a foothold.

The scenario also establishes the sequence clearly: asset identification has already occurred, while exploitation and persistence have explicitly not begun. This places the work squarely in vulnerability assessment. Because industrial devices can be sensitive to unexpected traffic, this activity must be conducted within the authorized rules of engagement and with OT-safe techniques, rates, and maintenance constraints. NIST emphasizes that industrial control systems have unique safety, reliability, and availability requirements that must guide security testing. CISA likewise provides ICS-focused vulnerability information to support identification and remediation of weaknesses. See NIST SP 800-82 Rev. 3 and CISA ICS Advisories.

QUESTION NO: 9

A web application allows users to upload files and later include them in pages dynamically. Attackers exploit this to execute code. Which vulnerability exists?

A. LFI

B. RFI

C. CSRF

D. XSS

ANSWER: A

Explanation:

LFI is the correct classification because the application dynamically includes a file that resides on its own local filesystem. A file-upload feature can make this especially dangerous when an attacker can place a file in a server-accessible upload directory and then influence an include mechanism to load that local file. In platforms or configurations that parse the included content as server-side code, the attacker-controlled contents can execute under the web application's security context.

The essential issue is not merely that a file was uploaded; it is that application-controlled dynamic inclusion trusts an attacker-influenced local path or file without strict allowlisting and safe handling. Secure implementations should avoid constructing include paths from user input, restrict inclusion to a fixed allowlist of known templates, ensure upload directories are non-executable, and store uploaded content outside the web root where feasible. OWASP describes file-inclusion testing and the risk created when applications use user-controlled input to select files in Testing for File Inclusion. PHP also cautions that included files are parsed as PHP code, which explains the potential execution impact when a malicious local file reaches an include operation: PHP include documentation.

QUESTION NO: 10

Chandler works as a pen-tester in an IT firm in New York. As part of detecting viruses in the systems, he uses a detection method where the antivirus executes the malicious code on a virtual machine to simulate CPU and memory activities.

Which type of virus detection method did Chandler use in this context?

A. Heuristic Analysis

B. Code Emulation

C. Scanning

D. Integrity checking

ANSWER: B

Explanation:

Code Emulation is the correct virus-detection method because it executes suspicious instructions in a controlled, software-simulated environment rather than allowing them to run directly on the host operating system. The antivirus emulator supplies a virtual CPU, memory, and related execution context, enabling the detection engine to observe what the code attempts to do at runtime. This is particularly useful when malware is packed, encrypted, obfuscated, or polymorphic, since its harmful behavior may not be apparent from a static file inspection alone. During emulation, the security product can safely monitor for behaviors such as self-decryption, attempts to alter memory, process manipulation, persistence actions, or other malicious activity, while isolating the real endpoint from those actions. The question explicitly describes execution of malicious code on a virtual machine to simulate CPU and memory activity, which directly matches code emulation. MITRE describes malware analysis sandboxes as systems that execute suspicious files in an isolated environment and monitor their behavior; code emulation is a closely related analysis approach used by endpoint security products. See MITRE ATT&CK: Virtualization/Sandbox Evasion and Microsoft Defender advanced hunting overview for context on behavior-focused malware analysis and detection telemetry.

QUESTION NO: 11

You are attempting to run an Nmap port scan on a web server. Which of the following commands would result in a scan of common ports with the least amount of noise In order to evade IDS?

A. nmap -sP- -p-65535-T5

B. nmap-A-host-time 99-T1

C. nmap -A -Pn

D. nmap -sT-O- To

E. nmap -sS -T0

ANSWER: E

Explanation:

`nmap -sS -T0 <target>` is the best answer because it combines Nmap's default common-port selection with a comparatively low-noise scan method and its most conservative timing template. When no port-selection argument is supplied, Nmap scans its default set of 1,000 most commonly used TCP ports, satisfying the requirement to scan common ports rather than every possible port. The `-sS` SYN scan sends a SYN probe and normally does not complete the TCP three-way handshake, unlike a TCP connect scan. This reduces connection-level logging and is the standard scan type used when the scanner has the required privileges. The `-T0` paranoid timing template introduces very long delays between probes, substantially reducing the scan rate and making rate-based detection more difficult. It is therefore appropriate where minimizing probe volume and timing regularity is more important than scan speed. This command must be used only against systems for which authorization has been obtained. Nmap documents SYN scanning behavior in its TCP SYN scan reference and describes the timing templates, including paranoid timing, in its performance and timing documentation.

QUESTION NO: 12

In ethical hacking, what is black box testing?

A. Testing using only publicly available information

B. Testing without any prior knowledge of the system

C. Testing with full system knowledge

D. Testing knowing only inputs and outputs

ANSWER: B

Explanation:

Testing without any prior knowledge of the system is the correct characterization of black box testing in an ethical-hacking engagement. The tester begins from an outsider's perspective and is not supplied with internal details such as network diagrams, source code, application architecture, credentials, configurations, or a list of known assets. Consequently, the engagement starts with reconnaissance and discovery to identify the target's exposed attack surface, followed by enumeration and appropriately authorized vulnerability validation. This model most closely simulates the position of a real external attacker who must independently discover useful information before attempting access or exploitation. The scope, targets, rules of engagement, and written authorization still must be defined in advance; "no prior knowledge" describes the tester's knowledge of the target environment, not an absence of permission or engagement controls. This approach is useful for assessing what an organization exposes to the public internet and how effectively perimeter-facing defenses detect and resist an external threat. NIST describes penetration testing as a method for evaluating security through simulated attacks in SP 800-115. EC-Council's CEH program also covers reconnaissance, enumeration, and vulnerability assessment as core phases supporting this external-attacker simulation: CEH program overview.

QUESTION NO: 13

During a compliance audit at a logistics company in Columbus, Ohio, the mobile security team discovers that several field-issued Android devices are responding to remote commands from an unknown external system. The affected devices are not connected via USB, and no enterprise mobility policies were recently modified.

Network monitoring reveals that the devices have remote debugging enabled and are accepting connections over the wireless network on a specific high-numbered port commonly associated with remote device communication. Investigators determine that the external system was able to capture screenshots, list installed applications, forward ports, and install additional packages without requiring physical access to the devices.

Which attack technique most accurately explains this compromise?

A. Device Administration API Abuse

B. FRP Bypass

C. Android Rooting

D. ADB Exploitation via TCP 5555

ANSWER: D

Explanation:

ADB Exploitation via TCP 5555 most accurately accounts for the observed compromise. Android Debug Bridge (ADB) is Android's command-line debugging and device-management interface. Although it is frequently used over USB during development, it can be configured to listen for network connections using TCP port 5555. If wireless debugging is enabled and an attacker can reach and authenticate to the exposed ADB service, the attacker can remotely issue ADB commands without physical possession of the device.

The stated activities closely match ADB capabilities: querying installed packages through the package manager, capturing the device display using screen-capture commands, installing application packages, and configuring port forwarding. The

decisive indicators are remote debugging being enabled, network-based command acceptance, and the specific association with TCP 5555. Together, these establish an exposed ADB listener as the relevant remote-access mechanism rather than merely a general Android device-management condition.

Android's ADB documentation describes connecting to devices over Wi-Fi and identifies port 5555 as the traditional TCP/IP debugging port. Android's security guidance also emphasizes that debugging interfaces must be controlled because they expose powerful device-management functions. See the Android Debug Bridge documentation and Android Open Source Project ADB security documentation.

QUESTION NO: 14

Study the snort rule given below:

```
alert tcp $EXTERNAL_NET any -> $HOME_NET 135
(msg: "NETBIOS DCERPC ISystemActivator bind attempt";
flow:to_server, established; content: "|05|"; distance: 0; within: 1;
content: "|ob|"; distance: 1; within: 1; byte_test: 1, &, 1, 0, relative;
content: "|A0 01 00 00 00 00 00 C0 00 00 00 00 00 00 46|";
distance: 29; within: 16; reference: cve, CAN-2003-0352;
classtype: attempted-admin; sid: 2192; rev: 1;)

alert tcp $EXTERNAL_NET any -> $HOME_NET 445 (msg: "NETBIOS SMB
DCERPC ISystemActivator bind attempt"; flow: to_server, established;
content: "|FF|SMB|25|"; nocase; offset:4, depth:5; content: "|26 00|";
nocase; distance:5; within: 12; content: "|05|"; distance:0; within:1;
content: "|ob|"; distance: 1; within: 1; byte_test: 1, &, 1, 0, relative;
content: "|A0 01 00 00 00 00 00 C0 00 00 00 00 00 00 46|";
distance: 29; within: 16; reference: cve, CAN-2003-0352;
classtype: attempted-admin; sid: 2193; rev: 1;)
```

From the options below, choose the exploit against which this rule applies.

A. WebDav

B. SQL Slammer

C. MS Blaster

D. MyDoom

ANSWER: C

Explanation:

MS Blaster is correct because the rule identifies traffic associated with exploitation of the Microsoft Windows DCOM RPC vulnerability used by the Blaster worm. Blaster, first widely observed in 2003, scanned for systems exposing Microsoft RPC services and sent a specially crafted RPC request to trigger the buffer-overflow condition. Snort signatures for this activity commonly inspect RPC/DCERPC traffic and match the distinctive interface, operation, and payload characteristics used during the exploit attempt, often in traffic involving TCP port 135 and subsequent RPC-related communications. This detection is therefore focused on the network behavior of the MS Blaster attack rather than merely identifying a generic Windows service connection. The underlying vulnerability is tracked as CVE-2003-0352, which describes a buffer overrun in the Windows RPC service that could permit remote code execution. Microsoft's security bulletin for the issue explains the RPC flaw and its impact, while the CVE record provides the vulnerability identifier and references. See CVE-2003-0352 and Microsoft Security Bulletin MS03-026.

QUESTION NO: 15

During a scheduled red team engagement at a regional investment firm in Phoenix, Arizona, security consultants were permitted limited after-hours access to employee workstations. As part of the evaluation, a small intermediary device was placed inline between a keyboard and its connected desktop system.

Over time, the device began forwarding captured keystroke activity through the company's established wireless environment, allowing the assessment team to collect periodic log data without interacting further with the workstation.

What type of keylogger does this scenario describe?

A. Hardware Keylogger

B. Acoustic/CAM Keylogger

C. Wi-Fi Keylogger

D. Bluetooth Keylogger

ANSWER: C

Explanation:

Wi-Fi Keylogger is correct because the described device both captures keystrokes through a physical inline connection and transmits the recorded data over a Wi-Fi network. This is a hardware keylogger with wireless network capability, commonly characterized as a Wi-Fi keylogger. Its inline placement allows it to intercept keyboard input before the data reaches the desktop system, while its wireless connectivity enables remote retrieval of logged activity without requiring additional physical access to the workstation.

In penetration-testing and CEH terminology, keyloggers can be software- or hardware-based. Hardware devices may be placed between a keyboard and computer, embedded in keyboard hardware, or connected through USB/PS/2 interfaces. Some modern hardware keyloggers include Wi-Fi functionality for remote log access and management. The scenario's explicit emphasis on forwarding captured keystrokes through the organization's wireless environment identifies the wireless communications capability as the distinguishing feature, making "Wi-Fi Keylogger" the most specific classification.

For background on keylogging as a credential-access technique, see the MITRE ATT&CK Input Capture: Keylogging entry. General guidance on assessing and managing hardware-related cybersecurity risks is available from CISA Cyber Threats and Advisories.

QUESTION NO: 16

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. LM challenge-response data can be captured from network traffic for offline password-cracking attacks.

C. Makes use of only 32-bit encryption.

D. Effective length is 7 characters.

ANSWER: A B D

Explanation:

LM authentication is weak because its password-processing design drastically reduces the effort required to recover a password. "Converts passwords to uppercase" is correct: LM normalizes alphabetic characters to uppercase before hashing, eliminating case as a source of password complexity. "Effective length is 7 characters" is also correct because LM pads or truncates a password to 14 characters and processes it as two independent seven-character halves. An attacker can therefore attack each half separately rather than having to search the full password as one value.

"LM challenge-response data can be captured from network traffic for offline password-cracking attacks" is correct because legacy LM authentication uses a challenge-response exchange. Although the password itself is not transmitted as clear text, an observed challenge and response can be used to validate password guesses offline, making capture and cracking of weak LM-derived credentials practical. Microsoft recommends disabling LM-related compatibility where possible and using stronger modern authentication protections. Microsoft's NTLM overview describes the legacy NTLM authentication family

and its security considerations. Microsoft also documents the security risks of legacy credential protocols and the importance of moving to stronger authentication controls.

QUESTION NO: 17

You have been asked to perform a penetration test for a local company. You have had several meetings with the client and are now almost ready to begin the assessment. Which of the following is the document that would contain verbiage which describes what type of testing is allowed and when you will perform testing and limits your liabilities as a penetration tester?

A. Project scope

B. Nondisclosure agreement

C. Service-level agreement

D. Rules of engagement

ANSWER: D

Explanation:

Rules of engagement is correct because it establishes the mutually agreed operational and legal boundaries for the penetration-testing engagement before assessment activities start. It documents the authorization for the tester to act and specifies permitted testing methods, the approved targets and environments, testing windows, prohibited actions, escalation procedures, points of contact, and processes for handling unexpected events. Clear authorization and boundaries are essential because testing activities can otherwise be interpreted as unauthorized access or disruption. When the tester performs only the activities authorized in the rules of engagement, this documented agreement helps demonstrate that those actions were conducted with the client's consent and within defined limits, supporting liability management for both parties. Rules of engagement commonly work alongside the scope and a formal authorization or statement of work, but they are the document focused on how testing will be conducted safely and under what constraints. The National Institute of Standards and Technology's guidance on security testing emphasizes defining rules of engagement, including test constraints and coordination requirements, prior to testing. See NIST SP 800-115: Technical Guide to Information Security Testing and Assessment and the CISA penetration testing resources.

QUESTION NO: 18

To determine if a software program properly handles a wide range of invalid input, a form of automated testing can be used to randomly generate invalid input in an attempt to crash the program.

What term is commonly used when referring to this type of testing?

A. Randomizing

B. Bounding

C. Mutating

D. Fuzzing

ANSWER: D

Explanation:

Fuzzing is the correct term for automated robustness and security testing that supplies an application with unexpected, malformed, invalid, or randomly generated input and observes its behavior. A fuzzing tool, commonly called a fuzzer, repeatedly exercises an input-handling interface—such as a file parser, network protocol, API endpoint, command-line argument, or web form—with large numbers of test cases. The objective is to expose failures such as crashes, hangs, unhandled exceptions, memory-safety defects, and other abnormal behavior that could signal a security vulnerability. Fuzzing may use entirely generated inputs or mutate valid seed inputs, but its defining purpose is systematic testing of how software handles abnormal data. This directly matches the scenario of randomly generating invalid input to attempt to crash

a program. OWASP describes fuzzing as providing invalid, unexpected, or random data as input and monitoring for exceptions such as crashes; see the OWASP Fuzzing guidance. The U.S. National Institute of Standards and Technology also identifies fuzz testing as a technique for discovering vulnerabilities by testing software with malformed or unexpected inputs in its Fuzz Testing glossary entry.

QUESTION NO: 19

Consider the following Nmap output:



what command-line parameter could you use to determine the type and version number of the web server?

A. -sV

B. -Pn

C. -V

D. -ss

ANSWER: A

Explanation:

`-sV` is the Nmap service-version detection option. When used in a scan, Nmap sends a set of protocol-specific probes to discovered open ports and analyzes the returned responses. This allows it to identify the service running on a port and, where the response provides enough evidence, report product details such as the web-server implementation and its version number. For example, a web service may be identified with output resembling `Apache httpd 2.4.x`, `nginx 1.x`, or a particular Microsoft IIS version. This information is valuable during authorized security assessments because accurate service identification helps determine the relevant attack surface, advisories, and configuration checks. The uppercase `V` is significant: Nmap command-line switches are case-sensitive, so the version-detection syntax is specifically `-sV`. Nmap also supports tuning the intensity of version detection with `--version-intensity` when an assessor needs to balance probe coverage against scan speed. The official Nmap reference describes `-sV` as enabling service and version detection and details how the probe-based identification process works. See the Nmap Version Detection documentation and the Nmap Reference Guide.

QUESTION NO: 20

A penetration tester suspects that a web application ' s product search feature is vulnerable to SQL injection. The tester needs to confirm this by manipulating the SQL query. What is the best technique to test for SQL injection?

A. Inject a malicious script into the search field to test for Cross-Site Scripting (XSS)

B. Use directory traversal syntax in the search field to access server files

C. Input 1 OR 1=1 in the search field to retrieve all products from the database

D. Insert admin ' — in the search field to attempt bypassing authentication

ANSWER: C

Explanation:

Input 1 OR 1=1 in the search field to retrieve all products from the database is the appropriate technique because it uses a Boolean tautology, a condition that evaluates as true. In an authorized test, the tester can submit this payload in a product-search parameter and observe whether the application's response changes in a way consistent with the supplied input altering the database query. For example, if the application concatenates the input directly into a query's filtering condition, the tautology can cause that condition to match every accessible row, potentially resulting in an unexpectedly broad product listing. This behavior provides evidence that user-controlled input is being interpreted as query syntax rather than handled solely as data.

The test should be performed only within the approved scope and with care to avoid unnecessary data exposure or production impact. A controlled test account, a non-production environment where possible, and recorded baseline results make the finding more reliable. OWASP identifies Boolean-based testing as a useful approach for detecting SQL injection, while emphasizing that the underlying remediation is parameterized queries (prepared statements), rather than string concatenation of untrusted input. See the OWASP SQL Injection overview and the OWASP SQL Injection Prevention Cheat Sheet.

QUESTION NO: 21

A Linux administrator is hardening a server that processes confidential records. The administrator wants to improve the detection and investigation of unauthorized activity.

Which of the following logging controls should be implemented? (Choose three.)

A. Forward logs to a centralized logging system.

B. Restrict access to log files and logging configuration.

C. Synchronize system time with an approved time source.

D. Disable audit logging to reduce disk usage.

E. Allow all local users to modify security logs.

ANSWER: A B C

Explanation:

Forwarding logs to a centralized logging system, restricting access to log files, and synchronizing system time are appropriate controls. Centralized logging helps retain records if a local system is compromised and makes correlation across multiple systems easier. Access restrictions protect logs from unauthorized reading, deletion, and modification, preserving their value for monitoring and forensic review.

Accurate time synchronization is also essential because event correlation depends on reliable timestamps across servers, endpoints, network devices, and security tools. Disabling audit logging or allowing all users to alter logs weakens accountability and makes incident investigation more difficult. See NIST SP 800-92, Guide to Computer Security Log Management and CISA Logging Made Easy.

QUESTION NO: 22

What is the correct order of the five phases of ethical hacking?

A. Gaining Access → Maintaining Access → Covering Tracks → Reconnaissance → Scanning

B. Maintaining Access → Covering Tracks → Reconnaissance → Scanning → Gaining Access

C. Reconnaissance → Scanning → Gaining Access → Maintaining Access → Covering Tracks

D. Scanning → Reconnaissance → Gaining Access → Covering Tracks → Maintaining Access

ANSWER: C

Explanation:

Reconnaissance → Scanning → Gaining Access → Maintaining Access → Covering Tracks is the classic five-phase ethical-hacking sequence commonly taught in CEH-aligned materials. Reconnaissance comes first because an authorized tester initially gathers intelligence about the target, such as its public-facing infrastructure, domains, personnel information, and technology footprint. That intelligence guides scanning, where the tester actively identifies live systems, exposed ports, services, vulnerabilities, and configuration details. Gaining access follows only after sufficient information has been collected and validated, representing the controlled demonstration that a discovered weakness can be exploited. Maintaining access then evaluates the potential impact of persistence or continued access within the agreed scope, helping demonstrate the practical business risk associated with compromise. Finally, covering tracks represents understanding attacker techniques for concealing activity, such as altering or removing evidence; in an ethical engagement, this phase must remain explicitly authorized and carefully documented so the organization can improve detection and logging. This ordered model maps the progression from information gathering through validation of compromise and post-exploitation behavior. EC-Council describes CEH as training focused on identifying and addressing security weaknesses through ethical-hacking methodologies: EC-Council CEH. NIST also provides guidance on structured security testing in SP 800-115.

QUESTION NO: 23

During an external security review of a manufacturing firm in Detroit, Michigan, you ' re asked to prioritize patch baselines for internet-facing servers without logging in or establishing full sessions. To achieve this, you analyze network-level responses and capture application output in order to determine the underlying system and its software release. Which technique best fits this objective?

A. Service Version Discovery

B. Port Scanning

C. OS Discovery

D. Vulnerability Scanning

ANSWER: A

Explanation:

Service Version Discovery is the appropriate technique because its purpose is to identify the network service listening on a reachable port and determine the product and release information exposed by that service. This can be performed externally, without user credentials or an interactive login, by sending protocol-specific probes and evaluating returned banners, headers, greeting messages, handshake fields, error responses, and other application-level output. For example, an HTTP response may disclose a web-server product in headers, while an SSH or SMTP service may return a greeting containing implementation and release details.

That information is directly useful for establishing patch priorities. Knowing the precise service implementation and version enables the reviewer to compare the observed release against vendor-supported versions, security advisories, and known-vulnerability records. The technique may also use response behavior when a service does not explicitly reveal a banner, allowing the assessor to identify or narrow the likely software version through fingerprinting.

Nmap documents version detection as querying open ports to determine the application and version running there, using probes and response matching: Nmap Reference Guide: Service and Version Detection. Its documented service-detection workflow also illustrates the external, unauthenticated nature of this activity: Nmap Network Scanning: Service and Version Detection.

QUESTION NO: 24

A company is replacing an older wireless network that uses a shared WPA2-Personal passphrase. The security architect wants to strengthen authentication and reduce the impact of compromised wireless credentials.

Which of the following recommendations are appropriate? (Choose three.)

A. Deploy WPA3-Enterprise where supported.

B. Use 802.1X with unique user or device credentials.

C. Segment wireless clients from sensitive internal resources.

D. Disable wireless encryption to simplify support.

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

F. Publish one shared passphrase for all employees.

ANSWER: A B C

Explanation:

Deploying WPA3-Enterprise where supported, using 802.1X authentication with unique user or device credentials, and segmenting wireless clients from sensitive internal resources are appropriate recommendations. WPA3-Enterprise supports stronger enterprise authentication capabilities and is designed for managed wireless environments. 802.1X provides port-based network access control and enables organizations to authenticate individual users or devices instead of relying on one shared passphrase.

Network segmentation limits the systems that a wireless client can reach after connecting, reducing the impact of a compromised endpoint or credential. Disabling encryption or publishing a common password broadly increases risk rather than reducing it. MAC-address filtering alone is not a sufficient authentication control because MAC addresses can be observed and spoofed. See CISA guidance on WPA2 security and the NIST guidelines for securing wireless LANs.

QUESTION NO: 25

A compromised admin account is used to disable logging services. What is the attacker attempting?

A. Anti-forensics

B. Exfiltration

C. Recon

D. Privilege escalation

ANSWER: A

Explanation:

Anti-forensics is correct because disabling logging services is intended to reduce the evidence generated by malicious activity and hinder subsequent detection, incident response, and forensic investigation. After obtaining administrative access, an attacker can alter or stop audit mechanisms so that actions such as account changes, process execution, persistence installation, and data access are not recorded. This behavior aligns with the covering-tracks objective commonly taught in ethical-hacking methodology: concealing the attacker’s presence and limiting the artifacts defenders can use to reconstruct events. In the MITRE ATT&CK framework, impairing defenses includes inhibiting or disabling logging and monitoring capabilities, while indicator removal on hosts includes clearing or otherwise eliminating forensic artifacts. These techniques support evasion and concealment after access has been obtained. Maintaining protected, centralized, and tamper-resistant logs is therefore a core defensive practice, since an adversary with elevated privileges may try to manipulate local logging controls. See the MITRE ATT&CK descriptions for Impair Defenses and Indicator Removal.

QUESTION NO: 26

Which of the following tools are used for enumeration? (Choose three.)

A. SolarWinds

B. USER2SID

C. Cheops

D. SID2USER

E. DumpSec

ANSWER: B D E

Explanation:

Enumeration is the active collection and interpretation of target information that can identify users, groups, accounts, shares, services, hosts, and security configuration. **USER2SID** and **SID2USER** are classic Windows enumeration utilities focused on translating between Windows account names and security identifiers (SIDs). This is valuable during an authorized assessment because a SID is the identifier Windows uses for a security principal, while resolving it to a user or group name makes collected access-control and account information actionable. Microsoft describes SIDs as unique values used to identify security principals such as users and groups in access-control contexts.

DumpSec is also an enumeration tool for Windows environments. It is designed to retrieve and report security-related information, including account and group details, policy information, permissions, and shared resources. Such output supports the enumeration phase by organizing exposed Windows security data for review within the defined scope of an ethical-hacking engagement. Together, USER2SID, SID2USER, and DumpSec directly support Windows user, group, and security-information enumeration. See Microsoft's Security identifiers documentation and EC-Council's Certified Ethical Hacker overview.

QUESTION NO: 27

An organization is moving several applications to a public cloud provider. The security team wants to reduce the risk caused by overly broad identities and accidental exposure of cloud resources.

Which of the following recommendations follow cloud-security best practices? (Choose three.)

A. Apply least-privilege IAM permissions.

B. Enable multifactor authentication for privileged accounts.

C. Log cloud management activity.

D. Make storage containers public by default.

E. Share one administrator account among all cloud operators.

ANSWER: A B C

Explanation:

Applying least-privilege IAM permissions, enabling multifactor authentication for privileged accounts, and logging cloud management activity are appropriate controls. Least privilege limits identities, workloads, and services to only the permissions required for their defined function. This reduces the impact if an account, access key, or workload identity is compromised.

Multifactor authentication adds a second authentication factor and is especially important for administrative and high-impact accounts. Logging management-plane activity provides an audit trail for actions such as identity changes, security-group modifications, and creation of new resources, enabling detection and investigation of unauthorized activity. Making storage containers public by default and sharing one administrator account remove accountability and unnecessarily expand exposure. See the AWS Well-Architected Security Pillar and CISA Cloud Security Technical Reference Architecture.

QUESTION NO: 28

An ethical hacker needs to enumerate user accounts and shared resources within a company ' s internal network without raising any security alerts. The network consists of Windows servers running default configurations. Which method should the hacker use to gather this information covertly?

A. Deploy a packet sniffer to capture and analyze network traffic

B. Perform a DNS zone transfer to obtain internal domain details

C. Exploit null sessions to connect anonymously to the IPC$ share

D. Utilize SNMP queries to extract user information from network devices

ANSWER: C

Explanation:

Exploit null sessions to connect anonymously to the IPC$ share is the correct choice for the scenario's intended Windows-enumeration technique. A null session is an unauthenticated Server Message Block connection that historically allowed a tester to establish a session to the Interprocess Communication share and query certain remote information. Where anonymous enumeration is permitted, this can expose information useful for reconnaissance, such as account, group, and share details, without first obtaining a valid user credential. The technique is especially relevant when assessing legacy Windows hosts or environments where anonymous-access restrictions have not been hardened. It is relatively low interaction compared with authentication attempts because it uses standard Windows file-sharing and remote-procedure-call mechanisms. In a legitimate ethical-hacking engagement, the tester would validate whether an anonymous connection can be established and then limit queries to the authorized scope. Modern supported Windows versions commonly restrict anonymous access and anonymous SAM/share enumeration by default, so successful null-session enumeration is a finding that indicates legacy behavior or weakened security policy rather than an assumption that should be made for every current Windows deployment. Microsoft documents the relevant anonymous-access security settings in its security policy reference and provides background on null sessions.

QUESTION NO: 29

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:

Using port security on switches helps restrict which MAC addresses may connect through an access port. Limiting or locking permitted MAC addresses reduces the opportunity for an unauthorized system to attach to the LAN and participate in an ARP-poisoning attack. Switch port-security capabilities are documented by Cisco at Cisco Catalyst port-security documentation.

Using a tool such as ARPwatch provides an important compensating control. ARPwatch observes IP-to-MAC mappings and reports unexpected changes, duplicate addresses, or new MAC addresses. Although monitoring does not itself discard forged ARP replies, timely alerts enable the administrator to investigate, remove a rogue device, and correct poisoned mappings before interception persists.

For a small, stable network, static ARP entries directly prevent hosts from dynamically accepting a changed MAC address for a protected IP address, because the IP-to-MAC association is administratively fixed. Static mappings require careful maintenance when hardware changes, but they are effective where their operational overhead is practical. The ARP protocol's role in resolving IP addresses to link-layer addresses is described in RFC 826; this unauthenticated resolution behavior is why enforcing known mappings and detecting mapping changes are valuable protections.

QUESTION NO: 30

You are using a public Wi-Fi network inside a coffee shop. Before surfing the web, you use your VPN to prevent intruders from sniffing your traffic. If you did not have a VPN, how would you identify whether someone is performing an ARP spoofing attack on your laptop?

A. You should check your ARP table for an unexpected or changing MAC address associated with the default gateway IP address.

B. You should scan the network using Nmap to check the MAC addresses of all the hosts and look for duplicates.

C. You should use netstat to check for any suspicious connections with another IP address within the LAN.

D. You cannot identify such an attack and must use a VPN to protect your traffic, r

ANSWER: A

Explanation:

Check the ARP table for an unexpected or changing MAC address associated with the default gateway IP address. ARP resolves a local IPv4 address to a layer-two MAC address, and the operating system stores those learned mappings in its ARP cache. In an ARP-spoofing attack, an attacker sends forged ARP replies intended to poison that cache, commonly causing the victim to associate the gateway's IP address with the attacker's MAC address. This can place the attacker in the traffic path and enable interception or manipulation of unprotected local-network traffic.

On a public Wi-Fi network, inspect the current ARP cache and pay particular attention to the default gateway entry. A gateway MAC address that changes unexpectedly during the session, differs from a value obtained through a trusted network source, or appears suspicious when monitored over time is a strong indication that ARP poisoning may be occurring. This approach directly examines the protocol state that ARP spoofing attempts to alter. ARP's address-to-hardware-address mapping behavior is defined in RFC 826, and Cisco provides practical discussion of recognizing and mitigating ARP poisoning at Cisco's ARP poisoning guidance.

QUESTION NO: 31

A tester is performing passive and low-impact reconnaissance against an authorized web application. The tester wants to identify server technologies, frameworks, and client-side libraries exposed by the application.

Which of the following tools can be used for web technology fingerprinting? (Choose three.)

A. WhatWeb

B. Wappalyzer

C. BuiltWith

D. Wireshark

E. John the Ripper

ANSWER: A B C

Explanation:

WhatWeb, **Wappalyzer**, and **BuiltWith** can be used to identify web technologies. These tools examine information such as HTTP response headers, cookies, HTML structure, script references, metadata, and other observable characteristics to infer web servers, content-management systems, JavaScript frameworks, analytics platforms, and related components.

Technology fingerprinting is useful during authorized reconnaissance because an identified component can guide later validation for configuration weaknesses, unsupported versions, or known vulnerabilities. It should not be treated as conclusive proof of a product version because banners and observable indicators can be changed or obscured. Wireshark is a packet analyzer, while John the Ripper is a password-auditing tool. See the WhatWeb project and Wappalyzer.

QUESTION NO: 32

An attacker has partial root access to a mobile application. What control best prevents further exploitation?

A. Secure coding and automated reviews

B. Certificate pinning

C. Regular penetration testing

D. Mobile Application Management (MAM)

ANSWER: D

Explanation:

Mobile Application Management (MAM) is the most appropriate control because it is designed to enforce security policy at the application and application-data layer, including when a device presents an elevated-risk condition. A MAM platform can apply managed-app policies such as restricting data sharing between managed and unmanaged applications, requiring application authentication, selectively wiping corporate application data, blocking access when device-risk conditions are detected, and revoking a user or application session. These capabilities provide containment after a compromise is suspected: they reduce the compromised application's useful access to protected organizational data and let administrators rapidly remove access without necessarily wiping the entire personally owned device.

For an attacker who has gained partial elevated access, the immediate goal is to limit what the affected application can access, move, or retain, while allowing the organization to enforce policy and terminate access centrally. MAM supports that response through granular application-level controls and isolation of managed work data. Microsoft describes app protection policies as controls that help protect organizational data within managed applications, including restricting data transfer and selectively wiping corporate data; see Microsoft Intune app protection policies. The NIST mobile-device security guidance also recognizes enterprise mobility management capabilities as important for applying and enforcing mobile security policy: NIST SP 800-124 Rev. 2.

QUESTION NO: 33

Robin, a professional hacker, targeted an organization's network to sniff all the traffic. During this process.

Robin plugged in a rogue switch to an unused port in the LAN with a priority lower than any other switch in the network so that he could make it a root bridge that will later allow him to sniff all the traffic in the network.

What is the attack performed by Robin in the above scenario?

A. ARP spoofing attack

B. VLAN hopping attack

C. DNS poisoning attack

D. STP attack

ANSWER: D

Explanation:

STP attack is correct because the rogue switch is deliberately configured with a lower bridge priority to influence Spanning Tree Protocol root-bridge election. In STP, the switch with the lowest bridge ID becomes the root bridge; bridge priority is a principal component of that ID. By presenting a superior BPDU and becoming root, the attacker can influence the Layer 2 topology and cause traffic to traverse a switch under the attacker's control, creating an opportunity to capture or manipulate frames. This is commonly called an STP root-bridge takeover attack.

The scenario's key indicators are the connection of an unauthorized switch to an unused LAN port and the intentional use of a lower priority to win the root election. Network administrators mitigate this risk using features such as Root Guard on ports where an unexpected root bridge must not be accepted, BPDU Guard on edge/access ports, and appropriate port-security controls. Cisco's STP overview explains that root selection is based on the lowest bridge ID, while its Root Guard documentation describes preventing a port from accepting superior BPDUs that could change the root bridge: Cisco: Spanning Tree Protocol Root Bridge Election and Cisco: Configuring STP Root Guard.

QUESTION NO: 34

When you are testing a web application, it is very useful to employ a proxy tool to save every request and response. You can manually test every request and analyze the response to find vulnerabilities. You can test parameter and headers manually to get more precise results than if using web vulnerability scanners.

What proxy tool will help you find web vulnerabilities?

A. Maskgen

B. Dimitry

C. Burpsuite

D. Proxychains

ANSWER: C

Explanation:

Burpsuite is correct because Burp Suite is an integrated web-application security testing platform whose Proxy component intercepts traffic between a browser and the target application. When interception is enabled, a tester can view, retain, modify, and forward individual HTTP and HTTPS requests and responses. This supports the manual testing workflow described in the question: examining requests, changing parameter values, altering headers, replaying requests, and observing how the application responds. Capturing this traffic also provides a detailed record that can be reviewed while testing authentication, session handling, input validation, access controls, and application logic.

Burp Suite's proxy works with a browser configured to send its web traffic through the local Burp listener. For HTTPS applications, the browser can trust Burp's generated certificate so encrypted requests and responses can be inspected during an authorized assessment. The Proxy tool integrates with other Burp features, including HTTP history, Repeater, and Scanner editions where available, allowing intercepted requests to be selected for deeper analysis. PortSwigger documents the Proxy as the feature used to intercept, inspect, and modify web traffic: Burp Suite Proxy documentation. The product overview is available at PortSwigger Burp Suite.

QUESTION NO: 35

A systems administrator must provide secure remote administration for Linux servers. The organization wants to reduce exposure of privileged access while retaining the ability to administer systems remotely.

Which of the following are appropriate SSH hardening measures? (Choose three.)

A. Disable direct SSH login for root.

B. Restrict SSH access to authorized management networks.

C. Use public-key authentication and multifactor authentication.

D. Enable Telnet for administrative access.

E. Use a shared root password for all administrators.

ANSWER: A B C

Explanation:

Disabling direct root login forces administrators to authenticate with individual accounts before using controlled privilege escalation, improving accountability and reducing direct attacks against the root account. Restricting SSH access to authorized management networks or approved source addresses limits the systems that can attempt remote administration. Using public-key authentication together with multifactor authentication, where supported, provides stronger protection than passwords alone.

Telnet should not be used as a replacement because it transmits credentials and session traffic in clear text. Allowing root logon with a shared password removes accountability and increases the impact of password compromise. OpenSSH documents relevant settings, including `PermitRootLogin`, `AllowUsers`, and public-key authentication in the sshd_config manual page. Additional remote-access guidance is available in NIST SP 800-46 Rev. 2.

QUESTION NO: 36

A web app fails to restrict API request frequency. What risk exists?

A. Data scraping

B. CSRF

C. XSS

D. SQLi

ANSWER: A

Explanation:

Data scraping is correct because an API that accepts unlimited or insufficiently controlled requests can be queried automatically at high volume. An attacker can enumerate identifiers, paginate through results, or repeatedly invoke search and lookup endpoints to harvest data that the application exposes through normal API functionality. This can include customer profiles, product catalogs, pricing, public records, metadata, or other information that becomes sensitive when collected comprehensively and at scale. Rate limits, quotas, pagination controls, and monitoring help make automated extraction more difficult and provide a means to detect abnormal usage patterns. OWASP identifies missing limits on API interactions as a resource-consumption risk: excessive requests can consume bandwidth, compute capacity, and other operational resources. In practice, the same absence of request-frequency controls also enables systematic collection of API responses, making data scraping the most direct risk described by the question. A sound defensive implementation applies rate limits appropriate to the endpoint and caller identity, uses authentication where appropriate, restricts returned data to what is necessary, and alerts on anomalous request volume or enumeration patterns. See the OWASP API Security Top 10 guidance on unrestricted resource consumption and the OWASP Denial of Service Cheat Sheet for practical request-limiting guidance.

QUESTION NO: 37

During a penetration test at a financial services firm in Boston, ethical hacker Daniel simulates a DDoS against the customer portal. To handle the surge, the IT team sets a rule that caps the number of requests a single user can make per second; aggressive connections are delayed or dropped while most legitimate customers continue to use the service.

Which countermeasure strategy is the IT team primarily using?

A. Rate Limiting

B. Shutting Down Services

C. Absorb the Attack

D. Degrading Services

ANSWER: A

Explanation:

Rate Limiting is the correct countermeasure because the team is enforcing a defined maximum number of requests that one user or client can submit during a specified interval, in this case per second. Once a client exceeds that threshold, the system delays or rejects additional requests. This constrains the amount of application and network capacity any individual source can consume, helping preserve availability for ordinary customers during a request-flooding or distributed-denial-of-service event.

Rate limiting is commonly applied at web servers, API gateways, load balancers, web application firewalls, and content delivery networks. It can use identifiers such as source IP address, authenticated account, session, API key, or request path, depending on the service and threat model. The described behavior—throttling aggressive connections while allowing most normal traffic to proceed—is the practical purpose of rate limiting: prioritizing fair resource use and reducing exhaustion of finite server resources. Cloudflare describes rate limiting as restricting the number of requests a client may make within a specified time period, and OWASP identifies rate limiting as a control for protecting services from excessive automated requests. See Cloudflare's rate-limiting overview and the OWASP Denial of Service Cheat Sheet.

QUESTION NO: 38

Which of the following is a low-tech way of gaining unauthorized access to systems?

A. Social Engineering

B. Eavesdropping

C. Scanning

D. Sniffing

ANSWER: A

Explanation:

Social Engineering is a low-tech means of obtaining unauthorized access because it targets the human element of security rather than relying primarily on exploiting a software vulnerability or conducting a technical attack against a device. An attacker uses deception, impersonation, persuasion, urgency, or pretexting to induce a person to disclose credentials, approve a fraudulent request, open a malicious attachment, or bypass an established security process. For example, an attacker may claim to be from the help desk and persuade an employee to reveal a password or approve a multifactor-authentication prompt. Once the attacker obtains valid credentials or convinces an authorized user to take an action, they can access systems while appearing to operate as a legitimate user. This makes social engineering especially significant in ethical-hacking assessments: testing must evaluate whether organizational personnel consistently verify identity, follow authorization procedures, and report suspicious requests. CISA identifies social engineering as the use of psychological manipulation to get people to perform actions or reveal confidential information, and recommends user awareness and verification practices as important defenses. See CISA's social engineering guidance and EC-Council's CEH program overview.

QUESTION NO: 39

Study the snort rule given below and interpret the rule. alert tcp any any --> 192.168.1.0/24 111

(content:"|00 01 86 a5|"; ms

G. "mountd access";)

A. An alert is generated when a TCP packet is generated from any IP on the 192.168.1.0 subnet and destined to any IP on port 111

B. An alert is generated when any packet other than a TCP packet is seen on the network and destined for the 192.168.1.0 subnet

C. An alert is generated when a TCP packet is originated from port 111 of any IP address to the 192.168.1.0 subnet

D. An alert is generated when a TCP packet originating from any IP address is seen on the network and destined for any IP address on the 192.168.1.0 subnet on port 111

E. "mountd access";)

ANSWER: D

Explanation:

The correct interpretation is: “An alert is generated when a TCP packet originating from any IP address is seen on the network and destined for any IP address on the 192.168.1.0 subnet on port 111.” The rule header begins with `alert`, which instructs Snort to generate an alert when the entire rule matches. The protocol field is `tcp`. The source address and source port are both `any`, so the traffic may originate from any host and any source port. The directional operator `-->` specifies traffic flowing toward the destination specified on the right side. That destination is any host within `192.168.1.0/24`, and the destination service port is `111`, the well-known port commonly associated with RPC portmapper/rpcbind services. In addition to the header conditions, the packet payload must contain the hexadecimal byte sequence `00 01 86 a5`. The message text identifies the matched activity as “mountd access,” providing a human-readable alert description. Snort evaluates rule headers and rule options together, so an alert requires both the specified TCP traffic direction, addresses, and port, plus the stated content match. See Snort’s rule-header documentation and its content-option documentation.

QUESTION NO: 40

During a forensic log review at a satellite communications provider in Denver, Colorado, cybersecurity analyst Kevin Morales identified subtle timestamp irregularities in archived telemetry records. Although the discrepancies were minor, regulatory reporting standards required confirmation that the system clock was synchronizing correctly with its configured time sources.

Kevin needed to interact directly with the hostâ€™s running time service to review its current associations and operational state. He was not attempting to reset the clock or trace the hierarchy of upstream time authorities, but rather to query the active service for detailed status information from the target machine.

Identify the command Kevin should execute to obtain this information.

A. ntptrace [-n] [-m maxhosts] [servername/IP address]

B. ntpq [-inp] [-c command] [host] [...]

C. ntpdc [-ilnps] [-c command] [host] [...]

D. ntpq -p [host]

ANSWER: D

Explanation:

`ntpq -p [host]` is the appropriate command because it queries the Network Time Protocol daemon on the specified host and requests its peer-association display. This output provides the configured and discovered time peers along with operational indicators that help an analyst assess synchronization health. In particular, it shows peer addresses or names, the selected synchronization source, reachability, delay, offset, jitter, and related association status values. Those details directly support investigation of minor timestamp anomalies: Kevin can determine whether the target is currently reaching its time sources and whether its clock discipline appears stable without changing the system clock or configuration.

The `-p` request is commonly used as the concise peer-summary query for an NTP service. It communicates with the running daemon using NTP’s control-query interface, making it suitable for checking the target host’s live NTP operational state during a forensic review. The NTP Project documentation describes `ntpq` as a utility for querying and monitoring `ntpd`, and documents the peers display as a list of peers known to the server with a summary of their state. See the NTP Project ntpq documentation and the ntpq manual page.

QUESTION NO: 41

What type of analysis is performed when an attacker has partial knowledge of inner-workings of the application?

A. Black-box

B. Announced

C. White-box

D. Grey-box

ANSWER: D

Explanation:

Grey-box is correct because grey-box testing models an assessment in which the tester or attacker has limited, partial knowledge of the target application's internals. This knowledge may include selected design documents, authentication details, API information, network diagrams, user roles, or limited source-code and architectural information. The assessor combines this internal context with an external attacker's perspective to identify vulnerabilities and evaluate how effectively they can be discovered, exploited, and traversed in a realistic environment. In penetration-testing terminology, grey-box engagements are useful for efficiently testing important attack paths while retaining uncertainty and discovery work that would occur during a real attack. The scope and degree of knowledge should be explicitly agreed upon before testing, since “partial knowledge” can range from basic credentials to substantial environment documentation. The National Institute of Standards and Technology describes security testing as a means to identify vulnerabilities and verify security controls, including testing approaches appropriate to the assessment context. OWASP likewise identifies grey-box testing as an approach where the tester has partial knowledge of the application or environment. See OWASP Web Security Testing Guide and NIST SP 800-115: Technical Guide to Information Security Testing and Assessment.

QUESTION NO: 42

You start performing a penetration test against a specific website and have decided to start from grabbing all the links from the main page.

What Is the best Linux pipe to achieve your milestone?

A. dirb https://site.com | grep "site"

B. curl -s https://sile.com | grep ‘’< a href-\’http" | grep "Site-com- | cut -d "\V" -f 2

C. wget https://stte.com | grep "< a href=*http" | grep "site.com"

D. wgethttps://site.com | cut-d"http-

E. wget -qO- https://site.com | grep -oP '(?<=href=")[^"]*(?=")'

ANSWER: E

Explanation:

`wget -qO- https://site.com | grep -oP '(?<=href=")[^"]*(?=")'` is the best answer because it retrieves the target page’s HTML and sends it directly through the pipe for link extraction without saving a local file. The `-q` switch suppresses routine status output, while `-O-` directs the downloaded response body to standard output. This makes the command suitable for quick, passive link collection during authorized reconnaissance. The `grep -oP` portion enables Perl-compatible regular expressions and prints only the text matched by the expression. Its lookbehind and lookahead isolate the value within an `href="..."` attribute, producing a clean list of linked URLs or paths rather than entire HTML lines. This is useful for identifying the application’s reachable resources from its main page before deeper mapping and validation. In practice, the tester should account for relative links, single-quoted attributes, JavaScript-generated navigation, redirects, and modern HTML variations; a dedicated parser or crawler may be more comprehensive. However, for the stated goal of piping the main page into a command that extracts standard double-quoted `href` values, this command correctly combines retrieval and filtering. GNU documents `wget` output handling at the GNU Wget manual, and GNU grep documents Perl-compatible regular expressions at the GNU Grep manual.

QUESTION NO: 43

During a red team assessment at a retail bank in New York, ethical hacker Aisha launches a flood of TCP connection initiation packets against the bank ' s online portal. The target accepts each initial handshake packet but never receives the final ACK to complete the three-way handshake, exhausting the server ' s backlog of half-open connections and preventing legitimate users from establishing new sessions.

Which type of DoS attack is Aisha most likely simulating?

A. ACK Flood

B. TCP SACK Panic

C. APT Attack

D. SYN Flood Attack

ANSWER: D

Explanation:

SYN Flood Attack is correct because it abuses the TCP three-way handshake at the connection-initiation stage. An attacker transmits a large number of TCP SYN packets to a listening service. For each SYN received, the server ordinarily reserves state for the pending connection and returns a SYN-ACK response, awaiting the client's final ACK. When that ACK is deliberately withheld, each request remains in the half-open state until the server's timeout expires. A sufficiently high volume of such incomplete requests fills the server's pending-connection backlog and consumes connection-handling resources, leaving little or no capacity for legitimate clients to initiate sessions.

The scenario explicitly identifies the defining technical indicators: TCP connection-initiation packets, absence of the final handshake acknowledgment, and exhaustion of half-open connections in the backlog. Source addresses are frequently spoofed in real SYN floods, which makes completion of the handshake unlikely and can make mitigation more difficult, but spoofing is not required for the described behavior. This attack is a transport-layer denial-of-service technique directed at TCP services, such as an online banking portal. For further technical background, see the TCP specification (RFC 9293) and CISA's guidance on denial-of-service attacks.

QUESTION NO: 44

Which defense MOST disrupts ransomware spread?

A. Backup

B. IDS

C. Network segmentation

D. AV

ANSWER: C

Explanation:

Network segmentation is correct because it limits an attacker's ability to move laterally after gaining an initial foothold. Ransomware operators commonly seek to reach file shares, domain controllers, backup infrastructure, and additional endpoints before triggering encryption. Dividing the environment into controlled network segments, with tightly restricted and monitored traffic between them, reduces the number of systems and services reachable from a compromised host. This containment can prevent a local compromise from becoming an organization-wide ransomware event and can preserve critical systems for response and recovery. Effective segmentation applies least-privilege access controls between user networks, server networks, administrative systems, and backup environments; it should also restrict common lateral-movement paths such as remote administration protocols to only explicitly authorized management sources. The U.S. Cybersecurity and Infrastructure Security Agency identifies network segmentation as a ransomware resilience measure because it can reduce the impact of an incident by preventing or limiting propagation. Microsoft similarly recommends

segmentation and isolation of high-value assets as part of a Zero Trust strategy to contain breaches. See CISA's Ransomware Guide and Microsoft's Zero Trust network guidance.

QUESTION NO: 45

Abel, a security professional, conducts penetration testing in his client organization to check for any security loopholes. He launched an attack on the DHCP servers by broadcasting forged DHCP requests and leased all the DHCP addresses available in the DHCP scope until the server could not issue any more IP addresses. This led to a Dos attack, and as a result, legitimate employees were unable to access the clients network. Which of the following attacks did Abel perform in the above scenario?

A. VLAN hopping

B. DHCP starvation

C. Rogue DHCP server attack

D. STP attack

ANSWER: B

Explanation:

DHCP starvation is the correct answer because the described activity exhausts the address pool maintained by a DHCP server. The attacker sends a large number of forged DHCP client requests, typically varying the source MAC address so that the server treats each request as a separate client. The DHCP server then creates leases for those apparent clients until every address in the configured DHCP scope has been allocated.

Once the available pool is depleted, genuine endpoints cannot obtain an IP configuration when they join the network or attempt to renew connectivity. Since DHCP commonly supplies not only an IP address but also the subnet mask, default gateway, DNS servers, and lease information, preventing a legitimate host from receiving a lease effectively denies that host normal network access. This is therefore a denial-of-service condition caused specifically by DHCP lease exhaustion.

In a penetration test, demonstrating DHCP starvation identifies a need for controls such as DHCP snooping, switch-port rate limits, port security, appropriately sized scopes, and monitoring of abnormal DHCP request volumes. Cisco describes DHCP snooping and rate limiting as protections against DHCP-related attacks in its DHCP Snooping Configuration Guide. The DHCP protocol's address-allocation and lease behavior is specified in RFC 2131.

QUESTION NO: 46

Your organization has signed an agreement with a web hosting provider that requires you to take full responsibility of the maintenance of the cloud-based resources. Which of the following models covers this?

A. Platform as a service

B. Software as a service

C. Functions as a service

D. Infrastructure as a service

ANSWER: D

Explanation:

Infrastructure as a service is the appropriate model because it gives the customer responsibility for administering and maintaining the cloud resources deployed above the provider's underlying infrastructure. In an IaaS arrangement, the hosting provider delivers the foundational computing environment, such as physical datacenter facilities, hardware, virtualization, and the basic infrastructure service. The organization then has substantial control over its provisioned virtual machines, operating systems, installed applications, configurations, access controls, and the ongoing maintenance required for those resources. This model is commonly selected when an organization needs flexibility to build, configure, patch,

monitor, and operate its own server workloads while avoiding ownership of the physical datacenter. The precise division of duties is defined by the provider's shared-responsibility model and the service agreement, but IaaS places the greatest operational responsibility on the customer among the listed cloud service models. NIST describes IaaS as providing processing, storage, networks, and other fundamental computing resources on which consumers deploy and run software, while Microsoft provides a practical overview of the customer-management responsibilities in IaaS. See NIST SP 800-145 and Microsoft Azure: What is IaaS?.

QUESTION NO: 47

Which tool dumps Windows hashes?

A. Mimikatz

B. John

C. Hydra

D. Aircrack-ng

ANSWER: A

Explanation:

Mimikatz is the correct answer because it is a post-exploitation credential-access tool widely used to extract authentication material from Windows systems. With appropriate privileges, its credential modules can read sensitive security-process memory and retrieve credential artifacts, including NTLM password hashes, Kerberos keys and tickets, and other logon-session secrets. This capability is commonly described as credential dumping or OS credential dumping. In an authorized security assessment, obtaining such artifacts demonstrates the impact of local-administrator or SYSTEM-level compromise and helps defenders validate protections such as Credential Guard, LSASS protection, privileged-access controls, and monitoring for suspicious access to the Local Security Authority Subsystem Service.

MITRE ATT&CK documents OS Credential Dumping as a technique for obtaining account credentials from operating-system sources, including Windows security-process memory. Its software entry for Mimikatz specifically identifies credential dumping as a core capability. The Mimikatz project also documents its credential modules and their use in examining Windows authentication data. See MITRE ATT&CK: OS Credential Dumping and MITRE ATT&CK: Mimikatz. Use such tooling only within explicit authorization and a defined testing scope.

QUESTION NO: 48

A company uses S/MIME to protect sensitive email communications between employees and business partners.

Which of the following statements about S/MIME are correct? (Choose three.)

A. S/MIME can encrypt email content for confidentiality.

B. S/MIME digital signatures can provide message integrity.

C. S/MIME uses public-key certificates.

D. A digital signature automatically encrypts the message body.

E. S/MIME requires all recipients to share the same private key.

ANSWER: A B C

Explanation:

S/MIME can encrypt email content to provide confidentiality, and it can apply digital signatures to provide message integrity and authenticate the signer. S/MIME relies on public-key certificates and asymmetric cryptography to bind a public key to an identity and to support encryption and signature verification.

A digital signature does not by itself encrypt a message, and symmetric encryption alone does not establish the sender's identity. In typical S/MIME use, the sender generates a random symmetric content-encryption key, encrypts the email content with that key, and protects the key using the recipient's public key. S/MIME message specifications are defined in RFC 8551. NIST describes digital-signature security services in FIPS 186-5, Digital Signature Standard.

QUESTION NO: 49

A company has observed repeated password-spraying attempts against its cloud email service. The attackers try a small number of commonly used passwords against many accounts in an effort to avoid traditional account-lockout thresholds.

Which of the following controls are most appropriate to reduce the risk of successful password spraying? (Choose three.)

A. Require multifactor authentication.

B. Disable legacy authentication where possible.

C. Use smart lockout and monitor distributed failed logons.

D. Use the same administrator password on every system.

E. Allow unlimited authentication attempts from any location.

ANSWER: A B C

Explanation:

Multifactor authentication is effective because a stolen or guessed password alone is insufficient for access. Disabling legacy authentication protocols reduces exposure because older protocols may not support modern authentication controls such as MFA and conditional access. Smart lockout, risk-based access policies, and monitoring of distributed authentication failures help identify and slow password-spraying behavior without unnecessarily locking large numbers of legitimate users.

Increasing password reuse or using one shared password among administrators would make password spraying more damaging, not less. NIST recommends resistance to online guessing attacks through rate limiting and additional authentication protections in NIST SP 800-63B. Microsoft also describes password-spray defenses in its Microsoft Entra smart lockout guidance.

QUESTION NO: 50

A penetration tester suspects that a web application ' s login form is vulnerable to SQL injection due to improper sanitization of user input. What is the most appropriate approach to test for SQL injection in the login form?

A. Inject JavaScript into the input fields to test for Cross-Site Scripting (XSS)

B. Enter ' OR ' 1 ' = ' 1 in the username and password fields to bypass authentication

C. Perform a directory traversal attack to access sensitive files

D. Use a brute-force attack on the login page to guess valid credentials

ANSWER: B

Explanation:

Entering `' OR '1'='1` in the username and password fields is an appropriate initial SQL-injection test payload for an authorized assessment. The payload attempts to alter the logic of a backend query by introducing a Boolean expression that evaluates as true. For example, if an application unsafely concatenates user input into a query used for authentication, the injected condition may cause the query's authentication predicate to be satisfied independently of a legitimate supplied credential. A tester should perform this test only within the agreed scope and should record the exact request, response status, application behavior, and any database-error indicators. The result helps establish whether user-controlled input is being interpreted as SQL syntax rather than being handled solely as data. Secure applications prevent this by using parameterized queries (prepared statements), which keep query structure separate from user-supplied values, along with

server-side input validation and appropriate error handling. OWASP identifies SQL injection as a major injection risk and recommends parameterized interfaces as the primary defense. See the OWASP SQL Injection Prevention Cheat Sheet and OWASP SQL Injection overview.

QUESTION NO: 51

Which of the following LM hashes represent a password of less than 8 characters? (Choose two.)

A. BA810DBA98995F1817306D272A9441BB

B. 44EFCE164AB921CQAAD3B435B51404EE

C. 0182BD0BD4444BF836077A718CCDF409

D. CEC52EB9C8E3455DC2265B23734E0DAC

E. B757BF5C0D87772FAAD3B435B51404EE

F. E52CAC67419A9A224A3B108F3FA6CB6D

ANSWER: B E

Explanation:

LM hashing converts a password to uppercase, pads it with null bytes to fourteen bytes, and divides the result into two seven-byte sections. Each section is used independently to produce an eight-byte DES-based hash value, yielding the familiar 32-hex-character LM hash. When the original password is fewer than eight characters, all password characters fit in the first seven-byte section. The second section therefore consists entirely of null bytes and always produces the well-known constant value `AAD3B435B51404EE`. Consequently, `44EFCE164AB921CQAAD3B435B51404EE` and `B757BF5C0D87772FAAD3B435B51404EE` exhibit the required second-half marker and represent passwords shorter than eight characters. The first listed value contains a non-hex character in its first half, which appears to be a transcription defect, but its displayed second half is still the conventional LM empty-second-block marker used by the question. LM hashes are obsolete and insecure because the split into seven-character blocks substantially reduces cracking effort; Windows security guidance recommends preventing their storage. See Microsoft's policy documentation for disabling LAN Manager hash storage and Openwall's technical description of LAN Manager hashes.

QUESTION NO: 52

A red team operator wants to obtain credentials from a Windows machine without touching LSASS memory due to security controls and Credential Guard. They use SSPI to generate NetNTLM responses in the logged-in user context and collect those responses for offline cracking. Which attack technique is being used?

A. Internal Monologue attack technique executed through OS authentication protocol manipulations

B. Replay attack attempt by reusing captured authentication traffic sequences

C. Hash injection approach using credential hashes for authentication purposes

D. Pass-the-ticket attack method involving forged tickets for network access

ANSWER: A

Explanation:

Internal Monologue attack technique executed through OS authentication protocol manipulations is correct because it abuses Windows' integrated authentication mechanisms rather than reading credential material directly from LSASS process memory. The technique invokes SSPI functions in the security context of the currently logged-on user and causes the NTLM authentication package to calculate a challenge-response value. The resulting NetNTLM response can be captured by the operator and subjected to offline password cracking. In practical implementations, the operator controls or supplies the

challenge used in the authentication exchange, allowing a response to be generated locally without needing to coerce the user to authenticate to a remote host.

This is especially relevant where direct credential dumping is constrained by endpoint defenses or protections designed to isolate credentials from LSASS memory. The key characteristic is obtaining a crackable NTLM challenge-response through legitimate Windows authentication interfaces, not extracting a reusable password hash from memory. Microsoft describes SSPI as the Windows interface through which applications use security packages for authentication, including NTLM: Microsoft SSPI documentation. Credential Guard's role in protecting credential secrets and authentication material is described at Microsoft Credential Guard documentation.

QUESTION NO: 53

An organization uses SHA-256 for data integrity checks but still experiences unauthorized data modification. Which cryptographic tool can help resolve this issue?

A. Asymmetric encryption

B. SSL/TLS certificates

C. Symmetric encryption

D. Digital signatures

ANSWER: D

Explanation:

Digital signatures are the appropriate cryptographic tool because a standalone SHA-256 hash does not authenticate the source of a hash value. If an attacker can alter both the data and its accompanying unhashed integrity value, they can calculate a new SHA-256 digest for the modified content, causing a simple comparison check to succeed. A digital signature binds the data's hash to the signer's private key. The recipient verifies the signature with the corresponding public key and independently calculates the hash of the received data. A valid verification establishes that the signed data has not changed since it was signed and that the signature was produced by the holder of the associated private key. This provides integrity together with origin authentication; in appropriate key-management circumstances, it also supports non-repudiation. In practice, the private signing key must be protected, and the verifier must trust the public-key binding through a certificate or another validated distribution mechanism. NIST describes digital signatures as mechanisms that detect unauthorized modification and authenticate the signatory, while RFC 8017 specifies widely deployed RSA signature operations. See NIST FIPS 186-5, Digital Signature Standard and RFC 8017: PKCS #1.

QUESTION NO: 54

Peter, a Network Administrator, has come to you looking for advice on a tool that would help him perform SNMP inquiries 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 C D

Explanation:

SNMPUtil, SNScan, SNMPScan, and SolarWinds IP Network Browser are appropriate tools for performing SNMP-focused discovery and enumeration. SNMP enumeration consists of sending authorized SNMP queries to managed devices and interpreting the returned management information. Depending on the configured SNMP version and available credentials or community strings, these queries can retrieve system identification, interface details, IP addressing data, routing information, installed software, running services, and other Management Information Base (MIB) objects. SNMPUtil is a Microsoft command-line utility that can issue SNMP requests for specific object identifiers, making it useful for direct inquiry and validation. SNScan and SNMPScan are purpose-built utilities for locating SNMP-capable hosts and gathering information exposed through SNMP. SolarWinds IP Network Browser provides a graphical way to browse SNMP-enabled devices and inspect available MIB data, which is useful when reviewing managed network infrastructure. These tools support the CEH enumeration objective because they concentrate on identifying and querying SNMP-exposed information rather than merely detecting that a host or port is present. The protocol behavior underlying these inquiries is defined in the SNMP RFC 1157 specification; Microsoft also documents the availability and use of SNMP utility tooling for SNMP requests.

QUESTION NO: 55

During a red team engagement at a technology startup in Austin, ethical hacker Priya simulates an internal attacker by connecting a laptop to the corporate LAN. Within minutes, nearby workstations begin receiving incorrect network settings such as altered gateways and DNS servers. Employees trying to access the intranet are redirected to fake login portals hosted on Priyaâ€™s machine. Security tools record temporary IP conflicts, but no alerts are triggered against the altered traffic paths.

Which attack technique did Priya most likely use?

A. DHCP Starvation Attack

B. DNS Cache Poisoning

C. Rogue DHCP Server Attack

D. Packet Sniffing

ANSWER: C

Explanation:

Rogue DHCP Server Attack is correct because the defining evidence is that workstations receive unauthorized network-configuration values shortly after the attacker connects to the LAN. DHCP provides clients with more than an IP address: its lease configuration can include the default gateway and DNS server addresses. An attacker operating an unauthorized DHCP service can answer client DHCP requests and supply values that direct victims' network traffic or name resolution through attacker-controlled systems. This directly accounts for the altered gateway and DNS settings described in the scenario.

With a malicious DNS-server setting, the attacker can resolve an intranet hostname to the IP address of a credential-harvesting site hosted on the attacker's laptop. A malicious default-gateway setting can similarly place the attacker in the traffic path, supporting interception or manipulation. Temporary address conflicts are also consistent with a rogue service issuing leases that overlap with the legitimate DHCP server's address scope. DHCP client-server behavior and the configuration parameters delivered in DHCP messages are defined in RFC 2131. For practical enterprise guidance on detecting unauthorized DHCP servers and protecting DHCP environments, see Cisco's rogue DHCP server guidance.

QUESTION NO: 56

One of your team members has asked you to analyze the following SOA record. What is the version?

Rutgers.edu. SOA NS1.Rutgers.edu ipad.college.edu (200302028 3600 3600 604800 2400.) (Choose four.)

A. 200302028

B. 3600

C. 604800

D. 2400

E. 60

F. 4800

ANSWER: A B C D

Explanation:

The values selected are the four distinct SOA timing and version values represented in the record: **200302028**, **3600**, **604800**, and **2400**. In an SOA record, the first numeric field is the zone serial number. It functions as the zone-data version: secondary DNS servers compare this number with their local copy to determine whether a zone transfer is needed. Here, that version is **200302028**; the original option contained a transcription error and has been corrected to match the record.

The remaining selected values are also explicit operational fields of this SOA record. **3600** is used for both the refresh interval and retry interval in the supplied data. **604800** is the expire interval, which limits how long a secondary server may continue serving data when it cannot refresh from the primary. **2400** is the minimum/negative-caching TTL value under the modern SOA interpretation. These fields collectively identify the SOA versioning and synchronization settings that are useful during DNS enumeration. The SOA resource-record layout is specified in RFC 1035, section 3.3.13; negative caching use of the SOA minimum field is defined by RFC 2308.

QUESTION NO: 57

In the neon-lit sprawl of Las Vegas, Nevada, a luxury hotelâ€™s smart room control system suffered a breach, allowing an intruder to manipulate guest room settings. The incident investigation revealed that the IoT devices lacked any mechanism to verify the integrity or authenticity of software prior to execution, allowing tampered instructions to run unchecked. As Emna Ruza, a cybersecurity consultant brought in to assess the breach, you recommend a solution that ensures only authorized, validated code is executed on the devices.

Which secure development practice are you advising the hotel to implement?

A. Implement code signing

B. Ensure secure boot

C. Secure firmware or software updates

D. Utilize secure communication protocols

ANSWER: A

Explanation:

Implement code signing is the appropriate practice because it establishes cryptographic proof that software, firmware, or update packages originated from an authorized publisher and have not been modified since signing. Before accepting or executing a software image, the IoT device verifies its digital signature using a trusted public key or certificate. A valid verification confirms both authenticity of the publisher and integrity of the code; any alteration to the signed content causes verification to fail. The device can then reject the unauthorized or tampered image rather than execute it.

This control is especially important for smart-room IoT equipment, where a malicious firmware image or altered application logic could directly affect physical room controls and guest safety or privacy. Code signing should be enforced throughout the device lifecycle: during manufacturing, software deployment, firmware updates, and recovery procedures. The device must protect its trusted verification keys and must be designed to refuse unsigned or invalidly signed code. NIST describes digital signatures as mechanisms that provide assurance of data integrity and origin authentication, which are precisely the security properties required here. Guidance on connected-device security likewise identifies authenticated software updates and verification as key protections for IoT products. See NIST FIPS 186-5: Digital Signature Standard and NISTIR 8259A: IoT Device Cybersecurity Capability Core Baseline.

QUESTION NO: 58

Which system consists of a publicly available set of databases that contain domain name registration contact information?

A. WHOIS

B. CAPTCHA

C. IANA

D. IETF

ANSWER: A

Explanation:

WHOIS is correct because it is the Internet query-and-response service historically used to retrieve registration information for Internet resources, most notably domain names. A WHOIS query is sent to the appropriate registry or registrar service and can return registration data associated with the queried domain. Depending on the applicable registry policy and privacy or proxy services, the returned record may include registrant, administrative, technical, and abuse-contact details, as well as the sponsoring registrar, registration dates, nameservers, and domain status. This makes WHOIS a valuable reconnaissance resource for ethical hackers when identifying domain ownership relationships, registrar details, DNS infrastructure, and potential points of contact during an authorized assessment. The precise data exposed can vary because modern privacy requirements and registration-data policies may redact personal information. ICANN explains that its Registration Data Lookup Tool provides access to publicly available registration data for generic top-level domains, while the WHOIS protocol itself is specified in RFC 3912. See ICANN Registration Data Lookup and RFC 3912: WHOIS Protocol Specification.

QUESTION NO: 59

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 because they send SNMP requests to managed network devices and present the returned management information in a usable form. Enumeration commonly includes querying object identifiers from a device's Management Information Base, identifying system details, discovering network interfaces, and retrieving operational or configuration-related values when the available SNMP credentials and access controls permit it. SNMPUtil is a command-line utility suited to making targeted SNMP queries. SNScan is designed to probe hosts for SNMP and collect exposed SNMP information across a network range. Solarwinds IP Network Browser provides a graphical approach to discovering devices and browsing available network-management data. These capabilities directly support an administrator who needs to inventory and assess SNMP-enabled infrastructure. SNMP request and response operations are defined by the SNMP protocol specifications, including retrieval operations used to obtain managed-object values; see RFC 3416: Version 2 of the Protocol Operations for SNMP. For practical monitoring and discovery context, see SolarWinds SNMP monitoring overview.

QUESTION NO: 60

Why would you consider sending an email to an address that you know does not exist within the company you are performing a Penetration Test for?

A. To elicit a response back that will reveal information about email servers and how they treat undeliverable mail

B. To create needless SPAM

C. To determine who is the holder of the root account

D. To test for virus protection

E. To perform a DoS

ANSWER: A

Explanation:

“To elicit a response back that will reveal information about email servers and how they treat undeliverable mail” is correct because an intentionally addressed message to a nonexistent recipient can generate a non-delivery report or other SMTP-level response. In an authorized penetration test, this is a limited reconnaissance technique: the returned information may disclose externally exposed mail-exchange infrastructure, relay or gateway naming, mail-routing behavior, recipient-validation behavior, and the organization’s handling of undeliverable messages. These details can help a tester accurately map the email attack surface and identify systems that should be included in the agreed scope of later assessment activities.

SMTP defines how mail transfer agents report delivery failures and may produce diagnostic information for the sending system or sender. The precise content varies by configuration, forwarding rules, and security controls, so testing must be explicitly authorized and performed with a controlled test mailbox to avoid impacting real users or creating unnecessary mail traffic. The result should be documented as reconnaissance evidence, with any exposed infrastructure details handled as sensitive assessment data. See the SMTP specification in RFC 5321 and OWASP guidance on information gathering in the Web Security Testing Guide.

QUESTION NO: 61

A cybersecurity research team identifies suspicious behavior on a user’s Android device. Upon investigation, they discover that a seemingly harmless app, downloaded from a third-party app store, has silently overwritten several legitimate applications such as WhatsApp and SHAREit. These fake replicas maintain the original icon and user interface but serve intrusive advertisements and covertly harvest credentials and personal data in the background. The attackers achieved this by embedding malicious code in utility apps like video editors and photo filters, which users were tricked into installing. The replacement occurred without user consent, and the malicious code communicates with a command-and-control (C & C) server to execute further instructions. What type of attack is being carried out in this scenario?

A. Simjacker attack

B. Man-in-the-Disk attack

C. Agent Smith attack

D. Camfecting attack

ANSWER: C

Explanation:

Agent Smith attack is correct because it describes an Android malware campaign that disguises its initial payload in apparently legitimate utility applications, commonly distributed through third-party app stores. After installation, the malware exploits weaknesses in Android’s application-installation and update mechanisms to replace installed apps with malicious versions. Those replacements can retain the target application’s expected name, icon, and user-facing functionality, making the compromise difficult for a user to notice. The altered application can then inject advertisements, collect sensitive information such as credentials, and receive additional instructions from attacker-controlled command-and-control infrastructure. This combination of a Trojanized utility app, silent replacement of legitimate apps, preserved branding and interface, intrusive advertising, and background data harvesting is the defining pattern associated with Agent Smith. Check Point Research documented the Agent Smith campaign and its method of replacing installed applications with malicious

variants: Agent Smith: A New Species of Mobile Malware. Android's application-signing model, which is central to understanding why legitimate update paths normally require the same signing certificate, is described in the Android Open Source Project APK signing documentation.

QUESTION NO: 62

A penetration tester evaluates an industrial control system (ICS) that manages critical infrastructure. The tester discovers that the system uses weak default passwords for remote access. What is the most effective method to exploit this vulnerability?

A. Perform a brute-force attack to guess the system ' s default passwords

B. Execute a Cross-Site Request Forgery (CSRF) attack to manipulate system settings

C. Conduct a denial-of-service (DoS) attack to disrupt the system temporarily

D. Use the default passwords to gain unauthorized access to the ICS and control system operations

ANSWER: D

Explanation:

Use the default passwords to gain unauthorized access to the ICS and control system operations is the correct choice because the identified weakness is the continued use of known factory-set or vendor-supplied credentials. Where those credentials remain active on a remotely accessible ICS component, authenticating with them is the most direct way to validate the exposure during an authorized penetration test. Successful authentication demonstrates a real compromise path and allows the tester to assess the permissions associated with that account, the reachable management functions, and the potential operational impact while following the engagement's rules of engagement. ICS environments require exceptional care: testing should use approved accounts and windows where possible, minimize interaction with live processes, avoid unsafe commands, and document the evidence needed for remediation. The appropriate corrective action is to replace all default credentials with unique, strong credentials, restrict remote administrative access, apply least privilege, and monitor authentication activity. CISA specifically recommends changing default passwords and securing remote access for operational technology assets. NIST's industrial-control-system guidance likewise emphasizes strong authentication and access control as foundational protections for ICS environments. See CISA Industrial Control Systems and NIST SP 800-82 Rev. 3.

QUESTION NO: 63 - (SIMULATION)

In Atlanta, Georgia, ethical hacker James Patel is hired by Southern Retail, a major e-commerce chain, to test the security of their online shopping platform. During his penetration test, James aims to simulate a session hijacking attack by setting up a proxy to intercept HTTP traffic between customers and the platform, log the requests, and perform advanced searches on the captured data to identify session tokens. He needs a lightweight tool specifically designed for security research that can handle these tasks in a controlled environment to demonstrate vulnerabilities to the company ' s security team.

Which tool should James use to perform this session hijacking simulation?

ANSWER: See the explanation for the answer

Explanation:

Use Caido.

Caido is a lightweight web-security interception proxy suited to a controlled session-hijacking demonstration. Configure the test browser to use Caido as its HTTP/HTTPS proxy, capture the platform requests and responses, then use its request history and search/filter features to locate session-related values such as `Cookie`, `Set-Cookie`, authorization headers, or token parameters.

This directly supports the required workflow: intercept HTTP traffic, log requests, and perform advanced searches on captured web traffic to identify session tokens.

QUESTION NO: 64

An authorized security assessment is performed on a public-sector services portal in Madison, Wisconsin. After authenticating with a controlled test account, the assessor captures the authentication identifier issued by the application.

Under controlled lab conditions, she attempts to reuse the captured identifier from a separate machine connected through a different encrypted channel. Although the identifier remains valid and within its lifetime, the application rejects the request when presented from the alternate environment.

Analysis indicates that the server evaluates characteristics associated with the original secure exchange before allowing continued use of the issued identifier.

Which defensive mechanism most likely explains this behavior?

A. Encrypting DNS resolution traffic using DNS over HTTPS

B. Cryptographically binding authentication tokens to the TLS connection context

C. Applying IPsec protection at the network layer

D. Enforcing HTTP Strict Transport Security

ANSWER: B

Explanation:

Cryptographically binding authentication tokens to the TLS connection context is the mechanism that best matches the observed behavior. A normal bearer session identifier can generally be replayed by anyone who possesses it while it is valid. Token binding changes that model by associating the authenticated token with cryptographic key material or channel-specific properties established during the original TLS-protected exchange. When the application later receives the identifier, it can verify that the request demonstrates possession of the binding key or is associated with the expected secure-channel context.

Consequently, copying only the captured identifier to a separate machine or a newly established encrypted connection is insufficient. The alternate client does not possess the private key or equivalent connection-bound proof needed to satisfy the server's validation, so the server rejects the otherwise unexpired identifier. This provides a meaningful defense against token theft and replay because a stolen value alone is not usable outside its intended TLS context. The IETF Token Binding Protocol describes how application-layer tokens can be bound to a TLS connection, while TLS channel binding concepts similarly support tying authentication to the underlying secure channel. See RFC 8471: The Token Binding Protocol Version 1.0 and RFC 5929: Channel Bindings for TLS.

QUESTION NO: 65

During a cloud security assessment, it was discovered that a former employee still had access to critical resources months after leaving the organization. Which practice would have most effectively prevented this issue?

A. Using multi-cloud deployment models

B. Implementing real-time traffic analysis

C. Conducting regular penetration tests

D. Enforcing timely user de-provisioning

ANSWER: D

Explanation:

Enforcing timely user de-provisioning is the appropriate control because it addresses the employee identity lifecycle at the point of termination. When a person leaves an organization, a defined offboarding workflow should promptly disable or delete their cloud identity and revoke all associated authorization paths. This includes console and federated sign-in access, IAM role assignments, group memberships, access keys, API tokens, application credentials, SSH keys, and active sessions

where supported. Automating this workflow through integration between HR systems, identity providers, and cloud IAM platforms reduces the risk that an account remains active because of manual oversight or incomplete handoffs.

Cloud environments make prompt de-provisioning especially important because remotely accessible resources can often be reached from anywhere using valid credentials. Applying least privilege and periodically reviewing access complement offboarding, but immediate revocation upon departure is the control that would have prevented the former employee's continued access. NIST's digital identity guidance describes account management processes that include disabling accounts when access is no longer required, while CISA recommends removing accounts and credentials as part of secure employee offboarding. See NIST SP 800-63B and CISA insider-threat mitigation resources.

QUESTION NO: 66

A senior attacker uses OAuth tokens stolen from browser storage to access APIs. What attack does this represent?

A. SQL Injection

B. CSRF

C. XSS

D. Token replay

ANSWER: D

Explanation:

Token replay is correct because the attacker has obtained a valid OAuth token and presents that token to an API in order to impersonate its legitimate holder. OAuth access tokens are commonly bearer tokens: the resource server authorizes a request based on possession of a valid token, rather than on a separate demonstration that the requester is the original user or client that received it. Consequently, an attacker who steals an unexpired token from browser-accessible storage can submit it in the API request's `Authorization: Bearer` header and obtain the permissions associated with that token. This is replay because the attacker reuses previously issued authentication material to establish an unauthorized authenticated session or API call. The core security event described is the use of the stolen credential, irrespective of how it was initially taken from browser storage. OAuth security guidance specifically identifies token theft and replay as risks of bearer-token use and recommends controls such as sender-constrained tokens, short lifetimes, secure storage, and token rotation where applicable. See RFC 6750: OAuth 2.0 Bearer Token Usage and the OAuth 2.0 Security Best Current Practice (RFC 9700).

QUESTION NO: 67

Who are â€œscript kiddiesâ€□ in the context of ethical hacking?

A. Highly skilled hackers who write custom scripts

B. Novices who use scripts developed by others

C. Ethical hackers using scripts for penetration testing

D. Hackers specializing in scripting languages

ANSWER: B

Explanation:

"Novices who use scripts developed by others" is correct. In information-security terminology, script kiddies are attackers with relatively limited technical skill who use ready-made tools, exploit code, automated scanners, or scripts written and published by more capable individuals. Their defining characteristic is reliance on existing attack resources rather than a detailed ability to discover vulnerabilities, develop original exploits, or understand every technical mechanism used by the tool.

These users may obtain tools from public repositories, exploit frameworks, or online communities and run them against targets with little or no modification. The label describes skill and method rather than authorization: a person using another party's script without permission is still conducting an attack, regardless of whether the tool itself is readily available. Because automation makes known vulnerabilities easier to probe and exploit, organizations should prioritize prompt patching, secure configuration, monitoring, and controls that reduce exposure to publicly available exploits.

This usage is consistent with common security guidance describing "script kiddies" as unskilled or low-skilled actors who rely on tools produced by others. See the UK National Cyber Security Centre guidance for practical defensive context and the MITRE ATT&CK knowledge base for structured information on adversary techniques and defensive detection.

QUESTION NO: 68

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

A. Use the same machines for DNS and other applications

B. Harden DNS servers

C. Use split-horizon operation for DNS servers

D. Restrict Zone transfers

E. Have subnet diversity between DNS servers

ANSWER: B C D E

Explanation:

Hardening DNS servers, using split-horizon operation for DNS servers, restricting zone transfers, and maintaining subnet diversity between DNS servers are sound DNS-security recommendations because they address both the attack surface and service resilience. DNS server hardening includes minimizing exposed services, promptly applying security updates, using strong administrative controls, enabling logging, and securely configuring recursive resolution and authoritative service. These practices reduce the opportunity for unauthorized modification, abuse, and exploitation of the DNS infrastructure.

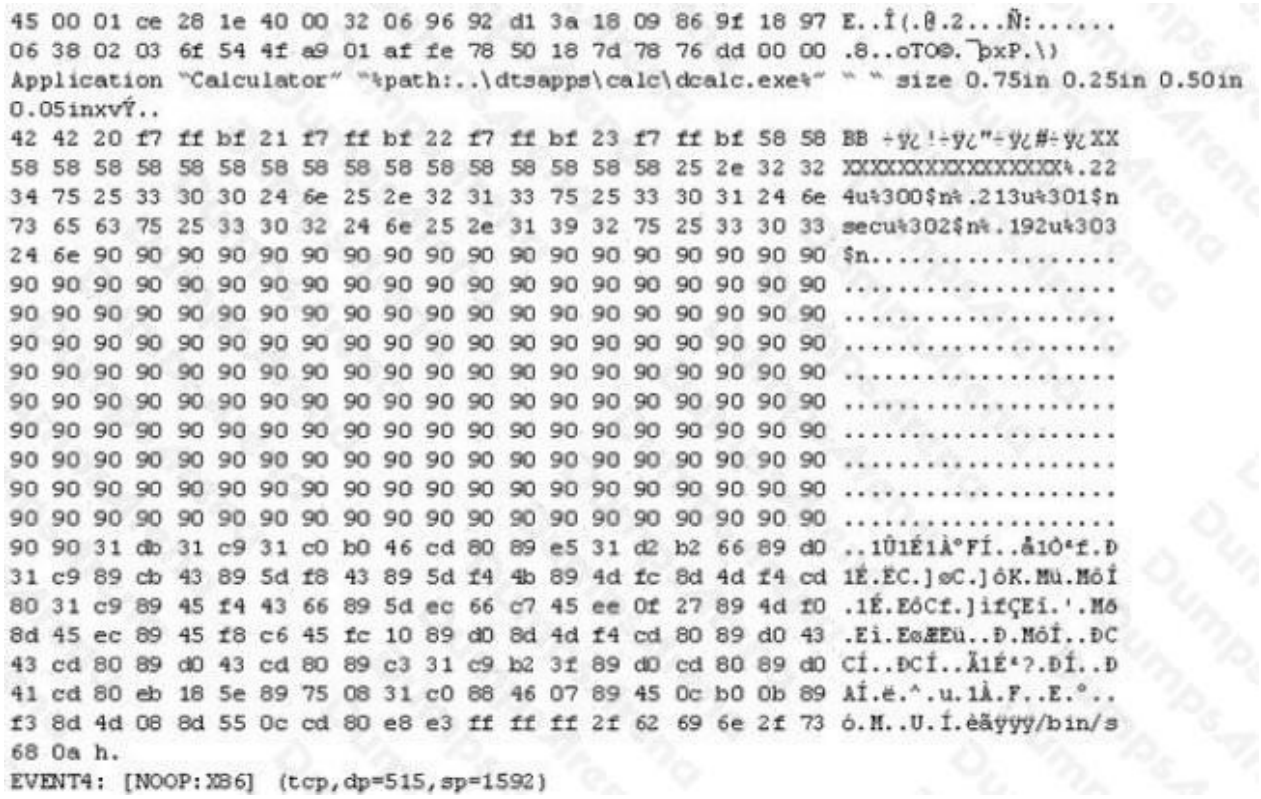
Split-horizon DNS provides different DNS responses according to where a query originates, allowing internal names and addresses to remain unavailable to external clients while preserving appropriate public DNS records. Restricting zone transfers limits replication of authoritative zone data to explicitly authorized secondary servers, reducing unnecessary disclosure of a company's namespace and host information. Finally, placing authoritative DNS servers on diverse subnets improves availability and fault tolerance: a routing, subnet, or localized network failure is less likely to take every authoritative server offline simultaneously. Together, these measures support confidentiality of internal naming data, integrity of DNS records, and continued name-resolution availability. See CISA DNS Best Practices and RFC 1918: Address Allocation for Private Internets.

QUESTION NO: 69

The following is an entry captured by a network IDS. You are assigned the task of analyzing this entry. You notice the value 0x90, which is the most common NOOP instruction for the Intel processor. You figure that the attacker is attempting a buffer overflow attack.

You also notice "/bin/sh" in the ASCII part of the output.

As an analyst what would you conclude about the attack?



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

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

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

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

ANSWER: D

Explanation:

The attacker is attempting an exploit that launches a command-line shell. A sequence containing many `0x90` bytes is characteristic of a NOP sled: it provides a broad landing area so that an overwritten instruction pointer can transfer execution into injected shellcode even when the precise jump address is uncertain. The `/bin/sh` string is a strong indicator that the payload is intended to invoke the standard Unix shell after code execution is obtained. Together, these artifacts identify a classic stack-based buffer-overflow exploit payload designed to yield an interactive command interpreter on the affected host. An IDS entry establishes that network traffic matching this payload was observed; it does not, by itself, establish whether the exploit succeeded, whether the target was vulnerable, or whether a shell session was actually created. This distinction is important during incident analysis: the evidence supports an attempted shell-launching exploit and should prompt validation through host, service, and network-session telemetry. MITRE ATT&CK describes Unix shell use under Command and Scripting Interpreter: Unix Shell, while CWE documents the underlying out-of-bounds write weakness commonly associated with buffer-overflow exploitation. See MITRE ATT&CK: Unix Shell and CWE-787: Out-of-bounds Write.

QUESTION NO: 70

Identify the UDP port that Network Time Protocol (NTP) uses as its primary means of communication?

A. 113

B. 69

C. 123

D. 161

ANSWER: C

Explanation:

Network Time Protocol uses UDP port 123 for its primary communication. NTP clients send requests to an NTP server on destination UDP port 123, and servers listen on that port to provide time-synchronization responses. Accurate, consistent time is operationally and security relevant: it supports reliable event correlation in logs, certificate validation, authentication systems, scheduled tasks, and incident investigation. During network enumeration or firewall review, identifying UDP/123 commonly indicates an NTP service; however, service presence should be validated through appropriate, authorized testing rather than inferred solely from a port number. The port assignment is standardized by the Internet Assigned Numbers Authority, whose service-name registry lists *ntp* as using UDP port 123. NTP's protocol specification also defines UDP port 123 as the NTP port used by clients and servers. See the IANA Service Name and Transport Protocol Port Number Registry and RFC 5905: Network Time Protocol Version 4.

QUESTION NO: 71

Which of the following are well known password-cracking programs?

A. L0phtcrack

B. NetCat

C. Jack the Ripper

D. Netbus

E. John the Ripper

ANSWER: A E

Explanation:

L0phtcrack and John the Ripper are well-known tools used in authorized password auditing and recovery activities. L0phtcrack is a commercial password-recovery and password-auditing product with a long association with assessing the strength of Microsoft Windows credentials. It obtains and audits password hashes, helping security professionals identify weak passwords and demonstrate the risk posed by easily guessed or cracked credentials. Its current product information describes password auditing, recovery, and remediation capabilities for Windows environments.

John the Ripper is an established password-security auditing and recovery tool maintained by Openwall. It can test password hashes from many operating systems, applications, and encrypted file formats, and supports multiple cracking approaches, such as dictionary-based and brute-force-style candidate generation. In an ethical hacking engagement, these tools must only be used against systems, accounts, and password hashes for which the tester has explicit authorization. Their value is in measuring password policy effectiveness, identifying weak credentials, and guiding remediation such as stronger password requirements and multifactor authentication. See the Openwall John the Ripper page and L0phtCrack product site.

QUESTION NO: 72

Scenario:

Victim opens the attacker ' s website.

Attacker sets up a website containing interesting and attractive content such as â€œDo you want to make $1000 in a day?â€□.

Victim clicks the attractive content URL.

The attacker creates a transparent iframe in front of the URL that the victim attempts to click. The victim believes he/she is clicking the â€œDo you want to make $1000 in a day?â€□ link, but is actually clicking content or a URL hidden inside the transparent iframe controlled by the attacker.

What is the name of the attack mentioned in the scenario?

A. HTTP Parameter Pollution

B. Clickjacking Attack

C. HTML Injection

D. Session Fixation

ANSWER: B

Explanation:

Clickjacking Attack is correct because the scenario describes a user-interface redress technique in which an attacker overlays a transparent or invisible iframe on top of enticing, visible web content. The victim sees and intends to select the displayed “Do you want to make $1000 in a day?” link, but the browser sends the click to an element embedded in the attacker-controlled iframe. Thus, the victim performs an action that is different from the action they believe they are taking.

Clickjacking commonly relies on framing a target page or sensitive control and disguising it beneath a harmless-looking interface. Depending on the framed content and the victim’s authenticated state, an attacker may attempt to induce unintended actions such as changing settings, granting permissions, following accounts, or approving transactions. The defining detail is not merely that a malicious link exists, but that the victim’s click is intercepted or redirected through a concealed interface layer. OWASP identifies clickjacking as an attack that tricks users into clicking something other than what they perceive, and recommends framing protections such as the `frame-ancestors` Content Security Policy directive. See the OWASP Clickjacking overview and the MDN frame-ancestors documentation.

QUESTION NO: 73

A penetration tester is tasked with uncovering historical content from a companyâ€™s website, including previously exposed login portals or sensitive internal pages. Direct interaction with the live site is prohibited due to strict monitoring policies. To stay undetected, the tester decides to explore previously indexed snapshots of the organizationâ€™s web content saved by external sources. Which approach would most effectively support this passive information-gathering objective?

A. Search with intext: " login " site:target.com to retrieve login data

B. Use the link: operator to find backlinks to login portals

C. Use the Internet Archive Wayback Machine to view archived versions of target pages

D. Use the intitle:login operator to list current login pages

ANSWER: C

Explanation:

Using the Internet Archive Wayback Machine to view archived versions of target pages is the most effective passive approach because it retrieves historical material from an independent archival service rather than sending requests to the organization’s live web server. Archived captures can reveal prior site structures, retired login paths, old administrative interfaces, documents, exposed directories, and content that has since been removed or protected. This makes web archives especially valuable during reconnaissance when direct interaction is prohibited or likely to trigger monitoring.

The Wayback Machine lets an authorized tester search a domain or specific URL and select captures from particular dates. The resulting evidence can be used to map historical attack surface and identify legacy resources for later, explicitly authorized validation. Importantly, Google discontinued its public cached-pages feature, so referring to a current Google `cache:` capability would be inaccurate. The Internet Archive’s web collection remains an established source for historical snapshots and fits the stated requirement to rely on externally saved content. See the Internet Archive guide to using the Wayback Machine and its Wayback Machine search interface.

QUESTION NO: 74

A malware analyst is tasked with evaluating a suspicious PDF file suspected of launching attacks through embedded JavaScript. Initial scans using pdfid show the presence of /JavaScript and /OpenAction keywords. What should the analyst do next to understand the potential impact?

A. Upload the file to VirusTotal and rely on engine consensus

B. Disassemble the PDF using PE Explorer

C. Extract and analyze stream objects using PDFStreamDumper

D. Compute file hashes using HashMyFiles for signature matching

ANSWER: C

Explanation:

Extract and analyze stream objects using PDFStreamDumper is the appropriate next action because pdfid is a triage tool: it identifies structural indicators such as `/JavaScript` and `/OpenAction`, but it does not reveal the script's actual contents or execution logic. An `/OpenAction` entry can cause an action to occur when a PDF is opened, while JavaScript may be stored in compressed, encoded, or obfuscated object streams. The analyst therefore needs to locate the relevant objects, decode their streams, and inspect the recovered JavaScript and associated data for exploit logic, URL retrieval, command execution attempts, or secondary-payload behavior. PDFStreamDumper supports this static-analysis workflow by exposing PDF objects and decoded stream content, allowing the analyst to understand what the suspicious document is designed to do before any controlled dynamic execution is considered. This approach also preserves forensic evidence and enables targeted identification of the object references associated with automatic actions. Didier Stevens' PDF analysis tools and guidance describe using structural inspection and stream analysis to investigate suspicious PDFs: PDF Tools by Didier Stevens. Adobe's PDF reference documents the PDF object model and action mechanisms that underpin entries such as `/OpenAction`: PDF 1.7 Reference.

QUESTION NO: 75

A security analyst is tasked with gathering detailed information about an organization ' s network infrastructure without making any direct contact that could be logged or trigger alarms. Which method should the analyst use to obtain this information covertly?

A. Examine leaked documents or data dumps related to the organization

B. Use network mapping tools to scan the organization ' s IP range

C. Initiate social engineering attacks to elicit information from employees

D. Perform a DNS brute-force attack to discover subdomains

ANSWER: A

Explanation:

Examine leaked documents or data dumps related to the organization is the appropriate method because it is a passive reconnaissance activity. Passive reconnaissance obtains intelligence from sources already available outside the target environment, rather than sending packets, queries, authentication attempts, or other interactions to the organization's systems. Publicly exposed breach data, leaked configuration files, technical documents, source-code repositories, cached materials, and document metadata can reveal useful infrastructure details such as internal hostnames, IP-addressing conventions, network diagrams, software and service references, email formats, and naming patterns. Because the analyst accesses information that has already been disclosed or published, this approach does not directly contact the target network and therefore avoids creating target-side logs or alerts. In an authorized ethical-hacking engagement, this intelligence helps establish the external footprint and prioritize later testing while preserving the required covert posture. This is consistent with the reconnaissance concept of using publicly available information and open-source intelligence before active enumeration. For background on reconnaissance and the distinction between passive and active information gathering, see CISA's OSINT resources and MITRE ATT&CK: Search Open Websites/Domains.

QUESTION NO: 76

Let's imagine three companies (A, B, and C), all competing in a challenging global environment.

Company A and B are working together in developing a product that will generate a major competitive advantage for them.

Company A has a secure DNS server while company B has a DNS server vulnerable to spoofing.

With a spoofing attack on the DNS server of company B, company C gains access to outgoing e-mails from company B.

How do you prevent DNS spoofing?

A. Install DNS logger and track vulnerable packets

B. Disable DNS timeouts

C. Install DNS Anti-spoofing

D. Disable DNS Zone Transfer

ANSWER: C

Explanation:

Install DNS Anti-spoofing is the best answer because DNS spoofing, often called DNS cache poisoning, succeeds when a resolver accepts a forged DNS response and caches or returns attacker-controlled address data. DNS anti-spoofing controls are designed to make forged replies difficult to inject and to reject responses that cannot be reliably associated with an outstanding query. In practice, this hardening includes DNS Security Extensions (DNSSEC) validation where supported, source-port and transaction-ID randomization, strict response/query matching, and securely configured recursive resolvers. DNSSEC is especially important because it adds cryptographic origin authentication and integrity protection to DNS data; a validating resolver can detect that a purported response was not signed by the authoritative zone owner or was altered in transit. These measures directly address the trust failure that lets an attacker redirect users or mail-related lookups to infrastructure under the attacker's control. NIST's secure DNS deployment guidance describes DNSSEC validation and resilient resolver configuration as central DNS security practices: NIST SP 800-81 Revision 2. ICANN also explains that DNSSEC protects DNS responses through digital signatures and enables resolvers to validate DNS data authenticity: ICANN DNSSEC overview.

QUESTION NO: 77

Within the context of Computer Security, which of the following statements describes Social Engineering best?

A. Social Engineering is the act of publicly disclosing information

B. Social Engineering is the means put in place by human resource to perform time accounting

C. Social Engineering is the act of getting needed information from a person rather than breaking into a system

D. Social Engineering is a training program within sociology studies

ANSWER: C

Explanation:

Social Engineering is the act of getting needed information from a person rather than breaking into a system. In computer security, social engineering refers to attacks that exploit human behavior, trust, authority, urgency, or curiosity to persuade a target to reveal sensitive information or perform an action that benefits the attacker. Rather than beginning with a technical vulnerability in software or network infrastructure, the attacker targets the human element—for example, convincing someone to disclose credentials, verify internal details, open a malicious attachment, or grant access to a restricted area or account.

This description correctly captures the core distinction: the information is obtained through interaction and manipulation of a person instead of directly compromising a system through a technical exploit. The interaction may occur through phishing email, phone calls, text messages, impersonation, pretexting, or in-person contact. A successful social-engineering attempt

can then provide the information or access needed for later technical compromise. CISA identifies social engineering as tactics used to deceive individuals into divulging sensitive information or taking unsafe actions; see CISA's phishing guidance. For further practical guidance on recognizing these manipulation-based attacks, see NCSC's phishing and scam guidance.

QUESTION NO: 78

You need to deploy a new web-based software package for your organization. The package requires three separate servers and needs to be available on the Internet. What is the recommended architecture in terms of server placement?

A. All three servers need to be placed internally

B. A web server facing the Internet, an application server on the internal network, a database server on the internal network

C. A web server and the database server facing the Internet, an application server on the internal network

D. All three servers need to face the Internet so that they can communicate between themselves

ANSWER: B

Explanation:

A web server facing the Internet, an application server on the internal network, a database server on the internal network is the recommended placement because it applies network segmentation and minimizes exposure of sensitive systems. The Internet-facing web tier should be deployed in a controlled perimeter network, commonly called a DMZ, where firewall policy permits only the required inbound web traffic, such as HTTPS. The web server can then communicate to the application tier through narrowly scoped firewall rules, and the application tier can access the database through only the required database port and account permissions.

This layered design keeps the application logic and, especially, sensitive database records out of direct reach from untrusted Internet clients. If the public web service is compromised, segmentation limits an attacker's ability to move directly into the internal application and data environment. It also supports monitoring and logging at each trust boundary, least-privilege network access, and separate patching and hardening responsibilities for each tier. NIST describes DMZs as perimeter network segments that provide controlled external access while helping protect internal networks. Cisco likewise describes the DMZ as an isolated network used for publicly accessible services. See NIST: DMZ definition and Cisco: Demilitarized Zone architecture.

QUESTION NO: 79

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?

A. 110

B. 135

C. 139

D. 161

E. 445

F. 1024

ANSWER: B C E

Explanation:

The appropriate firewall ports to block are 135, 139, and 445. On legacy Windows NT-family networks, these ports form an important part of the exposure associated with Windows remote services, file sharing, named pipes, and NetBIOS/SMB-related communications. TCP port 139 is the NetBIOS Session Service port and is used by SMB when it operates over

NetBIOS. TCP port 445 supports Microsoft-DS, which carries SMB directly over TCP without requiring NetBIOS; Windows 2000 and XP commonly use this direct-hosted SMB path. Blocking both ports at a network boundary prevents inbound SMB file-sharing and session traffic from reaching internal Windows hosts through either transport method. TCP port 135 is used by the Microsoft RPC Endpoint Mapper. Although it is not itself a NetBIOS service port, it is commonly filtered alongside 139 and 445 to reduce exposure to Windows RPC services and the dynamic RPC connections they coordinate on legacy Windows systems. This filtering approach is especially important when these services do not need to be accessible from untrusted networks. Microsoft documents the RPC Endpoint Mapper on port 135 and SMB-related ports 139 and 445 in its service overview and network port requirements. Microsoft also describes SMB direct hosting over TCP port 445 in its SMB documentation.

QUESTION NO: 80

An ethical hacker conducting an authorized assessment of a multinational advisory firm begins collecting intelligence exclusively from publicly accessible online platforms where employees share professional background details and engage in industry-related discussions.

By correlating individual role descriptions, publicly endorsed technical competencies, collaborative conversations referencing internal initiatives, and recurring terminology used to describe projects and departments, the tester develops a structured view of reporting relationships, identifies commonly deployed technologies, and infers internal naming conventions.

From a reconnaissance methodology perspective, which technique is being applied?

A. Footprinting through Social Networking Sites

B. Footprinting through Internet Research Services

C. Footprinting through Social Engineering

D. Footprinting through Search Engines

ANSWER: A

Explanation:

Footprinting through Social Networking Sites is the correct technique because the intelligence is gathered from publicly accessible professional profiles and industry-discussion platforms used by employees. These sources can reveal job titles, role histories, reporting and team relationships, technical skills, endorsements, project references, business vocabulary, and details about the organization's operational environment. Correlating this material is a passive reconnaissance activity: the assessor is not contacting personnel or attempting to induce disclosure, but is instead analyzing information that users and the organization have already made public.

In an authorized assessment, this form of open-source intelligence collection helps develop an organizational profile that can inform later, properly scoped testing. For example, recurring references to products, platforms, department names, and project terminology may reveal likely technology stacks, internal business functions, and naming patterns. Professional networking and social platforms are particularly valuable because employees often publish information that is individually harmless but becomes highly informative when aggregated and correlated.

This use of public employee content aligns with the general OSINT principle that publicly available social-media and professional-network information can support organizational reconnaissance. See the CISA overview of open-source intelligence and the MITRE ATT&CK Search Open Websites/Domains technique.

QUESTION NO: 81

An ethical hacker conducts testing with full knowledge and permission. What type of hacking is this?

A. Blue Hat

B. Grey Hat

C. White Hat

D. Black Hat

ANSWER: C

Explanation:

White Hat is correct because it describes authorized security testing performed with the organization's explicit knowledge and permission. A white-hat hacker works under a defined scope, such as a rules-of-engagement document or penetration-testing agreement, to assess systems, identify vulnerabilities, validate security controls, and report findings so that the organization can remediate risk. Authorization is fundamental: it establishes the tester's legal authority to interact with the target environment and defines boundaries such as permitted systems, testing windows, and allowed techniques. In CEH terminology, this activity is ethical hacking because the work is performed for a legitimate defensive purpose and is intended to improve the confidentiality, integrity, and availability of organizational assets. Ethical hackers should document authorization and communicate findings responsibly, including sufficient technical evidence and remediation guidance. This aligns with the NIST guidance that penetration testing is a security assessment methodology used to identify and validate exploitable weaknesses, and with EC-Council's description of CEH training as teaching authorized testing techniques used to assess and strengthen security posture. See NIST SP 800-115: Technical Guide to Information Security Testing and Assessment and EC-Council Certified Ethical Hacker (CEH).

QUESTION NO: 82

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. LM challenge-response authentication traffic can be captured and cracked offline.

C. Makes use of only 32-bit encryption.

D. Effective length is 7 characters.

ANSWER: A B D

Explanation:

LM hashing has structural weaknesses that substantially reduce the effort required to recover a password. "Converts passwords to uppercase" is correct because LM removes case sensitivity before calculating the hash. Consequently, users cannot increase password complexity through mixed case, and the original capitalization cannot be recovered from an LM hash. "Effective length is 7 characters" is also correct in the practical cracking sense: LM pads a password and divides it into two independent seven-character halves. Each half is processed separately, allowing an attacker to crack two much smaller search spaces rather than a single password of up to fourteen characters. "LM challenge-response authentication traffic can be captured and cracked offline" is correct because LM-based challenge-response authentication enables an attacker who captures the exchange to attempt offline password recovery against the response; modern Windows security guidance therefore recommends preventing LM hash storage and disabling legacy LM/NTLM usage where possible. These properties make LM materially weaker than modern password-verifier and authentication designs. Microsoft specifically documents the security reasons to prevent LM hash storage and provides policy guidance for controlling LAN Manager authentication levels. See Microsoft: Prevent Windows from storing an LM hash and Microsoft: LAN Manager authentication level.

QUESTION NO: 83

During a quarterly security audit at a financial services company in Charlotte, North Carolina, you are tasked with reviewing exposed services on legacy servers inherited from a third-party vendor. While scanning, you discover that TCP port 1434 is open on a database node that is not listed in the company ' s active inventory. The IT team has no records explaining why this service is running, and you are asked to determine whether the exposure of this port could indicate an unnecessary database-related risk. Based on standardized port assignments, which service is most likely running on this port and requires further review?

A. ms-sql-m

B. sqlsrv

C. sql*net

D. ms-sql-s

ANSWER: A

Explanation:

ms-sql-m is correct because TCP port 1434 is registered by the Internet Assigned Numbers Authority (IANA) under the service name *ms-sql-m*, historically referring to Microsoft SQL Server Monitor. The IANA Service Name and Transport Protocol Port Number Registry lists this assignment for both TCP and UDP, so the standardized port-to-service mapping requested by the question is ms-sql-m. See the IANA port-number registry entry for 1434.

From a security-review perspective, an unexpected listener on this port warrants investigation because it may indicate a legacy Microsoft SQL Server-related component, an undocumented application dependency, or an unintended service exposure. Microsoft SQL Server Browser commonly uses UDP port 1434 to provide instance-discovery information, while the IANA registration also reserves TCP 1434 for ms-sql-m. Therefore, a TCP scan result should be validated with service fingerprinting, host-process inspection, and configuration review rather than assuming it is necessarily SQL Server Browser. Microsoft describes SQL Server Browser behavior and its instance-discovery role in its SQL Server Browser documentation. In an unmanaged legacy environment, confirming the owning process, business purpose, exposure path, and need for the service is an appropriate next step before restricting or retiring it.

QUESTION NO: 84

Kevin and his friends are going through a local IT firm ' s garbage. Which of the following best describes this activity?

A. Intelligence gathering

B. Reconnaissance

C. Dumpster diving

D. Social engineering

ANSWER: C

Explanation:

Dumpster diving is the correct term for searching an organization’s discarded materials to obtain useful information or assets. In an ethical-hacking and security context, discarded paper records, notes, device labels, storage media, packaging, and improperly disposed electronics can expose details that help an attacker understand an organization or gain unauthorized access. For example, waste might contain network diagrams, employee contact information, passwords written on paper, customer data, or equipment configuration details. This is a physical information-gathering technique and illustrates why secure disposal procedures are an important part of an organization’s security program. Sensitive paper documents should be shredded or otherwise destroyed, while electronic media and devices require sanitization or destruction that prevents recovery of stored data. NIST’s media-sanitization guidance describes how organizations should render information on media inaccessible before disposal or reuse. The Cybersecurity and Infrastructure Security Agency also identifies proper handling and disposal of sensitive information as a practical protective measure. See NIST SP 800-88 Rev. 1, Guidelines for Media Sanitization and CISA Secure Our World.

QUESTION NO: 85

In Miami, Florida, Sarah Thompson, a security analyst at Apex Cyber Defense, is tasked with monitoring the wireless infrastructure at Coastal Healthcare, a busy urban hospital. One morning, nurse Emily Carter reports that her tablet used for accessing patient records is unexpectedly connecting to an access point broadcasting a name and signal similar to the hospital’s secure Wi-Fi. Upon investigation, Sarah’s log analysis reveals an unauthorized device on the network capturing sensitive traffic from connected systems. Suspecting a breach, she identifies that the attacker has deployed an access point to mimic the hospital’s legitimate network.

Based on this behavior, which wireless threat is the attacker executing?

A. Misconfigured AP

B. Rogue AP

C. Evil Twin AP

D. Honeypot AP

ANSWER: C

Explanation:

Evil Twin AP is correct because the attacker is deliberately impersonating the hospital's legitimate wireless network to cause a user device to associate with attacker-controlled infrastructure. An evil twin typically advertises a copied or confusingly similar SSID and may use a strong signal to make the fraudulent access point more attractive to nearby clients. After the tablet connects, the attacker can position their device between the victim and network services, enabling interception or capture of sensitive traffic. The facts that the access point imitates the hospital Wi-Fi and that an unauthorized device is capturing traffic are defining characteristics of this client-deception attack.

In a healthcare setting, this is especially serious because intercepted traffic can include credentials, session tokens, and electronic protected health information. Strong enterprise wireless authentication, particularly certificate-validated WPA2-Enterprise or WPA3-Enterprise deployments, helps clients distinguish authorized infrastructure from an impersonator. Security teams should also use wireless monitoring to detect duplicate SSIDs, unexpected BSSIDs, and anomalous signal patterns, while ensuring users report unanticipated connection prompts. The U.S. Cybersecurity and Infrastructure Security Agency discusses risks and protections for wireless networking at CISA: Securing Wireless Networks. Additional enterprise Wi-Fi security guidance is available from NSA Cybersecurity Advisories and Guidance.

QUESTION NO: 86

An attacker plans to compromise IoT devices to pivot into OT systems. What should be the immediate action?

A. Perform penetration testing

B. Secure IoT–OT communications with encryption and authentication

C. Deploy ML-based threat prediction

D. Deploy an IPS

ANSWER: B

Explanation:

Secure IoT–OT communications with encryption and authentication is the immediate priority because it directly protects the trust boundary that an attacker would use to move from an IoT foothold into operational technology. Strong, mutually authenticated communications ensure that OT assets accept commands, telemetry, and connections only from verified devices or services. Encryption protects data in transit from interception or alteration, while authentication prevents an unauthorized or impersonated IoT endpoint from establishing trusted access to the OT environment. This control is especially important in industrial environments, where compromised edge devices can otherwise become a path to safety- or availability-critical systems. In practice, this should be implemented with managed device identities and certificates, secure protocol configurations, key lifecycle management, and strict authorization for the minimum communications each IoT device requires. These measures reduce the opportunity for lateral movement at the point where IoT and OT systems interconnect. NIST guidance for IoT cybersecurity emphasizes device identification, secure communications, and protection of data and interfaces as foundational capabilities; see NISTIR 8259A. CISA also recommends securing industrial control system communications and controlling access paths in its Industrial Control Systems guidance.

QUESTION NO: 87

Which of the following hacking frameworks describes adversary tactics, techniques, and procedures (TTPs) used in cyberattacks?

A. NIST CSF 2.0

B. ISSF

C. MITRE ATT & CK

D. ISC 28901

ANSWER: C

Explanation:

MITRE ATT&CK is the appropriate framework because it provides a structured, publicly accessible knowledge base for describing real-world adversary behavior. ATT&CK organizes offensive activity by the adversary's tactical goals, such as gaining initial access, executing code, establishing persistence, escalating privileges, discovering systems, moving laterally, collecting data, and exfiltrating it. For each goal, it documents techniques and, where applicable, more granular sub-techniques that describe how attackers accomplish that objective. This makes it highly useful for threat intelligence, detection engineering, incident response, purple-team exercises, and security-control assessment. A defender can map observed evidence, such as suspicious credential dumping or remote-service activity, to ATT&CK techniques and then use that mapping to identify detection gaps and prioritize mitigations. MITRE maintains distinct ATT&CK matrices for enterprise, mobile, and industrial-control-system environments, enabling teams to apply the framework to the technologies they operate. The official ATT&CK website describes the knowledge base and provides the Enterprise matrix, while MITRE's introduction explains its use in modeling adversary tactics and techniques. See MITRE ATT&CK and the MITRE ATT&CK Enterprise matrix.

QUESTION NO: 88

A technology consulting firm in Denver, Colorado, recently experienced a wave of suspicious account compromise incidents. Several employees reported receiving an email that appeared identical to a legitimate cloud storage notification they had received earlier that week. The message reused the original branding, formatting, sender display name, and subject line. However, it informed recipients that the previously shared document had been â€œupdated due to synchronization errorsâ€□ and instructed them to reauthenticate using the embedded link. The link directed users to a convincing replica of the organizationâ€™s authentication portal. Investigation revealed that the attacker had reused content from a genuine prior communication and modified only the embedded hyperlink. Which type of social engineering attack does this scenario most accurately represent?

A. Clone Phishing

B. Consent Phishing

C. Search Engine Phishing

D. Tabnabbing

ANSWER: A

Explanation:

Clone Phishing is correct because the attacker copied a legitimate message that recipients had already received and trusted, preserving its recognizable appearance, branding, sender display name, and subject line. The malicious activity was introduced by changing an actionable component of that otherwise familiar communication: the embedded hyperlink. The altered link led to a replica authentication portal designed to capture credentials when employees attempted to reauthenticate.

This technique exploits the victim's memory of a genuine earlier email. Rather than creating a wholly new lure, the attacker clones the original communication and makes a small, credible modification, such as replacing a link or attachment and adding a plausible reason to act. The synchronization-error notice and reauthentication request provide the urgency and business context needed to make the changed link appear routine. This is the defining pattern of clone phishing and is particularly effective when the duplicated message relates to a service employees commonly use, such as cloud storage. For phishing concepts and defensive guidance, see CISA's phishing guidance and the UK NCSC phishing guidance.

QUESTION NO: 89

At a government research lab, cybersecurity officer Nikhil is compiling a vulnerability assessment report after scanning the internal subnet. As part of his documentation, he lists the IP addresses of all scanned hosts and specifies which machines are affected. He includes tables categorizing discovered vulnerabilities by type such as outdated software, default credentials, and open ports.

Which section of the vulnerability assessment report is Nikhil working on?

A. Findings

B. Risk Assessment

C. Supporting Information

D. Assessment Overview

ANSWER: A

Explanation:

Findings is the appropriate report section because it records the specific security conditions identified during the vulnerability assessment. Nikhil is documenting the assessed hosts by IP address, identifying the systems affected by each issue, and organizing the discovered weaknesses into categories such as outdated software, default credentials, and open ports. These details are the substantive assessment results: they establish what was found, where it was found, and which assets require attention.

A useful findings presentation links each vulnerability to the affected asset or assets and supplies enough technical detail for remediation teams to validate and address the condition. Tables are commonly used for this purpose because they make host-level results, vulnerability categories, and remediation tracking clear and actionable. The findings content can subsequently be used to prioritize remediation and communicate risk to stakeholders, but its central function is to document the discovered vulnerabilities and their affected systems.

This approach aligns with NIST guidance that vulnerability-management activities identify vulnerabilities in systems and use the resulting information to support remediation decisions. See NIST SP 800-40 Rev. 4 and the CISA Vulnerability Management guidance.

QUESTION NO: 90

During an internal assessment, a penetration tester gains access to a hash dump containing NTLM password hashes from a compromised Windows system. To crack the passwords efficiently, the tester uses a high-performance CPU setup with Hashcat, attempting millions of password combinations per second. Which technique is being optimized in this scenario?

A. Spoof NetBIOS to impersonate a file server

B. Leverage hardware acceleration for cracking speed

C. Dump SAM contents for offline password retrieval

D. Exploit dictionary rules with appended symbols

ANSWER: B

Explanation:

Leverage hardware acceleration for cracking speed is correct because the scenario focuses on maximizing the rate at which candidate passwords can be transformed into NTLM hashes and compared against the captured hash values. Password-hash cracking is an offline process once the hashes have been obtained: there is no need to send each guess to the target system. Consequently, performance is primarily determined by the cracking hardware, the cracking engine's optimized hash implementation, and the number of hash computations that can be performed per second. A high-performance CPU running Hashcat is being used to increase that throughput, allowing millions of candidate values to be tested rapidly. Hashcat is specifically designed to use available compute devices and optimized kernels for high-speed

recovery of passwords from supported hash types, including NTLM. In an authorized assessment, this capability helps demonstrate the practical risk posed by weak, reused, or predictable passwords stored as recoverable hash material. Hashcat's documentation describes its role as an advanced password-recovery tool and its use of CPU and accelerator back ends; see the Hashcat Wiki. For context on NTLM authentication, Microsoft documents the protocol in its NTLM protocol specification.

QUESTION NO: 91

A forensic examiner acquires a laptop that may contain evidence of unauthorized activity. The examiner needs to preserve the evidentiary value of a copied disk image.

Which actions should the examiner take? (Choose three.)

A. Calculate and record hashes of the acquired image.

B. Document every evidence transfer in the chain of custody.

C. Store the original evidence in controlled storage.

D. Perform analysis directly on the original disk.

E. Remove acquisition timestamps from the evidence record.

ANSWER: A B C

Explanation:

Calculating and recording cryptographic hashes of the acquired image, documenting each transfer of custody, and preserving the original evidence in controlled storage are essential evidence-handling practices. A hash value provides a means to verify that the copied image has not changed after acquisition. The hash should be recorded at acquisition and verified whenever the evidence is accessed or copied.

A chain-of-custody record identifies who handled evidence, when it was transferred, why it was transferred, and where it was stored. Controlled storage protects the original media from alteration, loss, or unauthorized access. An examiner should analyze verified working copies rather than modify the original evidence. See NIST SP 800-86, Guide to Integrating Forensic Techniques into Incident Response.

QUESTION NO: 92

A penetration tester gains access to a target system through a vulnerability in a third-party software application. What is the most effective next step to take to gain full control over the system?

A. Conduct a denial-of-service (DoS) attack to disrupt the system's services

B. Execute a Cross-Site Request Forgery (CSRF) attack to steal session data

C. Perform a brute-force attack on the system ' s root password

D. Use a privilege escalation exploit to gain administrative privileges on the system

ANSWER: D

Explanation:

Use a privilege escalation exploit to gain administrative privileges on the system. Initial access obtained through a vulnerable third-party application commonly runs in the security context of a limited user account or service account. That foothold may permit reconnaissance and limited actions, but it does not necessarily provide access to protected operating-system resources, security settings, credential stores, or other users' data. Privilege escalation is the post-exploitation activity used to move from that restricted context to a more powerful one, such as local administrator, root, or SYSTEM, subject to the engagement's scope and authorization.

A tester should first enumerate the host carefully to identify an appropriate and safe escalation path. Relevant evidence can include the operating-system version and patch level, service configurations, file and directory permissions, scheduled tasks, installed drivers, sudo or sudoers rules, token privileges, and credentials or secrets exposed in configuration files. Any exploit or configuration weakness used must be validated in a controlled manner, avoiding unnecessary disruption and documenting the resulting access level and remediation evidence. This approach aligns with established penetration-testing practice: exploit the initial foothold, enumerate the compromised host, and use authorized privilege-escalation techniques to demonstrate the real impact of the finding. See NIST SP 800-115, Technical Guide to Information Security Testing and Assessment and MITRE ATT&CK: Privilege Escalation.

QUESTION NO: 93

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

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” form the intended technique. Cisco devices can restrict SNMP management access with an ACL, which evaluates the source address of the incoming SNMP packet. If an assessment demonstrates that source-address spoofing is feasible, using an address permitted by that ACL can cause the router to accept the authenticated SNMP request. SNMP write access is significant because Cisco’s SNMP MIB support includes objects that can initiate configuration-file transfers, enabling a running or startup configuration to be copied to a TFTP server.

The capture component is needed because IP spoofing normally prevents ordinary request-response communication: replies will be routed to the spoofed address rather than the tester. A sniffer positioned where the relevant traffic can be observed can collect the configuration transfer or associated returned traffic. This scenario therefore combines bypass of a source-IP-based SNMP ACL with passive collection of the resulting configuration data, rather than relying on an interactive SNMP session. Cisco documents SNMP access restrictions and configuration-copy behavior in its SNMP community configuration guide and configuration file copy guide.

QUESTION NO: 94

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

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

How would you protect from these attacks?

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

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

C. Use SSL authentication on Web Servers

D. Enable Active Scripts Detection at the firewall and routers

ANSWER: B

Explanation:

Create rules in IDS to alert on strange Unicode requests is correct because obstructed-URL attacks commonly use unusual encodings to conceal directory-traversal sequences or sensitive resource paths from simplistic inspection. An intrusion detection system can inspect HTTP request URIs and query strings for anomalous Unicode, overlong UTF-8 encodings, repeated encoding, or traversal indicators after normalization. Alerting on these patterns gives defenders visibility into attempted exploitation, supports rapid incident response, and enables correlated investigation with web-server and application logs.

This detection should be implemented with signatures tailored to known malformed or suspicious request encodings and regularly reviewed to account for legitimate application traffic. It is especially useful where an attacker attempts to exploit differences between how an intermediary security control and the web server decode a URL. OWASP explains that traversal attacks can use encoded forms of traversal characters and recommends handling input only after appropriate decoding and validation: OWASP Path Traversal. For operational IDS/IPS monitoring guidance and the importance of inspecting traffic for malicious patterns, see CISA Intrusion Detection and Prevention Systems.

QUESTION NO: 95

While simulating a reconnaissance phase against a cloud-hosted retail application, your team attempts to gather DNS records to map the infrastructure. You avoid brute-forcing subdomains and instead aim to collect specific details such as the domainâ€™s mail server, authoritative name servers, and potential administrative information such as serial number and refresh interval.

Given these goals, which DNS record type should you query to extract both administrative and technical metadata about the target zone?

A. TXT

B. MX

C. NS

D. SOA

ANSWER: D

Explanation:

SOA is the DNS record type that supplies the requested administrative and technical metadata for a DNS zone. The Start of Authority record establishes core zone-management information, including the primary authoritative name server and the responsible party’s mailbox in DNS format. Crucially, it also contains the zone serial number and timing parameters used by secondary authoritative servers to synchronize their copies of zone data: refresh, retry, expire, and minimum/negative-cache TTL values. These fields allow a reconnaissance team to understand how the zone is administered and propagated without attempting subdomain brute force.

In practical DNS enumeration, separate lookups may be used to identify mail exchangers and the full set of authoritative name servers, while an SOA lookup provides the zone’s centralized operational metadata. The serial number can indicate when zone content has changed, and the refresh interval indicates how often secondary servers should check the primary server for updates. RFC 1035 defines the SOA resource record and its fields, including SERIAL, REFRESH, RETRY, EXPIRE, and MINIMUM. Cloud deployments still use these standard DNS semantics even when the authoritative DNS service is managed by a cloud provider. See RFC 1035, Section 3.3.13 and IANA DNS Resource Record Types.

QUESTION NO: 96

An organization is deploying a public web application that processes customer information. The application must be accessible from the Internet while the application and database tiers must remain protected from direct external access.

Which of the following controls should be included in the deployment? (Choose three.)

A. Place the public web tier in a DMZ.

B. Restrict application-to-database traffic to required ports and hosts.

C. Require TLS for Internet-facing client connections.

D. Expose the database server directly to the Internet.

E. Allow unrestricted connections from the web tier to internal networks.

ANSWER: A B C

Explanation:

Placing the Internet-facing web tier in a DMZ creates a controlled boundary between untrusted networks and internal resources. Restricting application-to-database connectivity to only the required ports and sources applies network least privilege and reduces lateral-movement opportunities. Deploying TLS for all client connections protects the confidentiality and integrity of customer traffic between browsers and the public service.

A database server should not be exposed directly to the Internet merely because the web application needs database access. Likewise, allowing unrestricted traffic from the web tier to all internal systems defeats the purpose of segmentation. NIST discusses firewall architecture and DMZ deployment in SP 800-41, Guidelines on Firewalls and Firewall Policy. OWASP provides additional guidance for protecting web applications in its Application Security Verification Standard.

QUESTION NO: 97 - (SIMULATION)

Lily, a network security analyst at a regional healthcare provider, is preparing defenses ahead of a scheduled external vulnerability assessment. During internal simulation drills, she observes that scanners are successfully identifying open ports and service banners across critical systems. Tasked with reducing exposure to such reconnaissance efforts, Lily is instructed to apply measures that specifically hinder port scanning activity without disrupting legitimate traffic.

Which of the following actions should Lily implement?

ANSWER: See the explanation for the answer

Explanation:

Implement stateful firewall/IPS port-scan protection: configure the perimeter firewall (or IPS) to detect abnormal sequential or high-rate connection attempts and rate-limit or temporarily block the source IP, while allowing approved connections to required services.

In practice, Lily should also ensure that inbound ACL/firewall rules expose only required ports and deny unsolicited access to all other ports. This minimizes what a scanner can identify without interrupting legitimate traffic to authorized services.

Allow only necessary service ports from authorized source networks.
Enable TCP/UDP scan detection (for example, SYN scan thresholds).
Rate-limit or automatically block sources performing scan-like behavior.
Suppress unnecessary service banner/version disclosure where supported.

QUESTION NO: 98

A new wireless client is configured to join a 802.11 network. This client uses the same hardware and software as many of the other clients on the network. The client can see the network, but cannot connect. A wireless packet sniffer shows that the

Wireless Access Point (WAP) is not responding to the association requests being sent by the wireless client. What is a possible source of this problem?

A. The WAP does not recognize the client's MAC address

B. The client cannot see the SSID of the wireless network

C. Client is configured for the wrong channel

D. The wireless client is not configured to use DHCP

ANSWER: A

Explanation:

The WAP does not recognize the client's MAC address is the correct answer. A wireless client that can discover the network has already received the access point's beacon or probe-response management frames, so radio visibility and SSID discovery are functioning. The failure occurs later in the 802.11 connection process, when the client sends an association request and the access point does not allow the association to proceed. One possible cause is a MAC address access-control list configured on the WAP. In this arrangement, the WAP compares the source MAC address in the client's management frames with a permitted or denied list. If the new device's MAC address is absent from an allow list, or present on a deny list, the WAP can refuse the device's attempt to associate. This is especially plausible where otherwise identical client hardware and software successfully connect, because the individual client's layer-2 MAC address remains unique. MAC filtering is an access-control measure that may be used as an additional administrative restriction, although it is not a strong standalone wireless security control because MAC addresses can be observed and spoofed. The association phase takes place before IP addressing, so this symptom is consistent with a layer-2 admission policy such as MAC filtering. See the Cisco access-point configuration guidance and the CAPWAP protocol specification for wireless access-point and station-management context.

QUESTION NO: 99

John, a professional hacker, performs a network attack on a renowned organization and gains unauthorized access to the target network. He remains in the network without being detected for a long time and obtains sensitive information without sabotaging the organization. Which of the following attack techniques is used by John?

A. Advanced persistent threat

B. threat Diversion theft

C. Spear-phishing sites

D. insider threat

ANSWER: A

Explanation:

Advanced persistent threat is correct because the defining characteristics in the scenario are unauthorized access to a specifically valuable target, a long-term covert presence, and the collection of sensitive information. An advanced persistent threat (APT) is not a single exploit or malware type; it is a sustained, targeted intrusion campaign. The attacker first obtains a foothold, then maintains persistence while moving through the environment, escalating privileges as needed, locating valuable data, and exfiltrating it while attempting to avoid detection. The word "persistent" directly reflects the attacker's ability to remain within the target network for an extended period, rather than conducting a brief opportunistic compromise. "Advanced" reflects the deliberate, capable, and often multi-stage nature of the operation, while "threat" describes the actor or campaign posing the risk. Sabotage can occur in some APT operations, but it is not required: cyberespionage and theft of confidential, proprietary, or strategic data are common APT objectives. NIST defines an APT as an adversary that uses sophisticated techniques to pursue objectives over an extended duration, adapting to defenders' efforts to resist it. See the NIST definition of Advanced Persistent Threat and MITRE ATT&CK's enterprise adversary tactics and techniques knowledge base.