PassQueen
a.k.a dumpsqueen.com

Certified Ethical Hacker CEH v11

ECCouncil CEH-v11

Version Demo

Total Demo Questions: 54

Total Premium Questions: 546

Buy Premium PDF

https://passqueen.com

support@passqueen.com

Topic Break Down

Topic	No. of Questions
Topic 1, Introduction to Ethical Hacking	37
Topic 2, Footprinting and Reconnaissance	53
Topic 3, Scanning Networks	37
Topic 4, Enumeration	24
Topic 5, Vulnerability Analysis	23
Topic 6, System Hacking	54
Topic 7, Malware Threats	23
Topic 8, Sniffing	31
Topic 9, Social Engineering	25
Topic 10, Denial of Service	8
Topic 11, Session Hijacking	11
Topic 12, Evading IDS, Firewalls, and Honeypots	43
Topic 13, Hacking Web Servers	11
Topic 14, Hacking Web Applications	31
Topic 15, SQL Injection	10
Topic 16, Hacking Wireless Networks	29
Topic 17, Hacking Mobile Platforms	9
Topic 18, IoT Hacking	9
Topic 19, Cloud Computing	14
Topic 20, Cryptography	55
Topic 21, Mix Questions	9
Total	546

QUESTION NO: 1

John is investigating web-application firewall logs and observers that someone is attempting to inject the following:

char buff[10];

buff[>o] - 'a':

What type of attack is this?

A. CSRF

B. XSS

C. Buffer overflow

D. SQL injection

ANSWER: C

Explanation:

Buffer overflow is the intended answer because the code declares a fixed-size character buffer, `buff[10]`. A buffer overflow vulnerability occurs when a program writes data beyond the memory allocated for such a buffer, potentially overwriting adjacent memory. Attackers commonly exploit unsafe handling of fixed-length buffers by supplying crafted input that exceeds the expected bounds or manipulates indexing so that writes occur outside the valid range. Depending on the affected memory and platform protections, this may cause a crash, data corruption, altered control flow, or execution of attacker-controlled code.

The key security issue is not merely that a character is being assigned, but that a fixed buffer is being targeted in a way that suggests an attempt to cause an out-of-bounds write. In secure C/C++ development, applications must validate indexes and input lengths before copying or writing data into arrays, and should use bounds-aware APIs where appropriate. MITRE classifies this weakness as CWE-120, “Buffer Copy without Checking Size of Input,” and describes how such errors can enable memory corruption. See MITRE CWE-120 and OWASP Buffer Overflow.

QUESTION NO: 2

Abel, a cloud architect, uses container technology to deploy applications/software including all its dependencies, such as libraries and configuration files, binaries, and other resources that run independently from other processes in the cloud environment. For the containerization of applications, he follows the five-tier container technology architecture. Currently. Abel is verifying and validating image contents, signing images, and sending them to the registries. Which of the following tiers of the container technology architecture Is Abel currently working in?

A. Tier-1: Developer machines

B. Tier-4: Orchestrators

C. Tier-3: Registries

D. Tier-2: Testing and accreditation systems

ANSWER: D

Explanation:

Tier-2: Testing and accreditation systems is correct because this tier is the control point where container images are assessed and approved before they become available for deployment. Abel’s activities—verifying and validating image contents, applying image signatures, and then publishing approved images to a registry—are image assurance and accreditation functions. This workflow establishes trust in the artifact: the image is checked against the organization’s requirements, signed to support integrity and provenance verification, and only then promoted to the registry for later retrieval by orchestration and runtime environments.

Container-image security guidance emphasizes validating images and protecting their integrity throughout the build, distribution, and deployment lifecycle. NIST identifies image validation, provenance, and secure registry practices as key controls for container environments. Image signing provides cryptographic evidence that an image originates from a trusted publisher and has not been modified since signing. Docker's image-signing documentation similarly describes signing as a mechanism for publishers and consumers to establish trust in image content. Therefore, the described pre-registry verification, validation, and signing process belongs to the testing and accreditation systems tier. See NIST SP 800-190, Application Container Security Guide and Docker Content Trust documentation.

QUESTION NO: 3

Which is the first step followed by Vulnerability Scanners for scanning a network?

A. OS Detection

B. Firewall detection

C. TCP/UDP Port scanning

D. Checking if the remote host is alive

ANSWER: D

Explanation:

Checking if the remote host is alive is the first step because a vulnerability scanner must first identify which target systems are reachable before it can perform deeper assessment activities. This initial host-discovery phase commonly uses ICMP echo requests, ARP discovery on local networks, TCP probes, or other methods to determine whether an address is active and responsive. Establishing the live-host set prevents the scanner from wasting time attempting port, service, operating-system, and vulnerability checks against unused or unreachable IP addresses. After reachable hosts are identified, the scanner can enumerate exposed TCP or UDP ports, identify services and versions, and correlate discovered information with vulnerability data. Host discovery is therefore a foundational reconnaissance and scan-optimization activity in network vulnerability assessment workflows. NIST describes discovery as identifying assets and their characteristics, which provides the information needed for subsequent security assessment work; see NIST SP 800-115, Technical Guide to Information Security Testing and Assessment. For practical scanner behavior, Tenable also documents host discovery as the process of determining whether systems are live before conducting further scans: Tenable Nessus Host Discovery.

QUESTION NO: 4

How is the public key distributed in an orderly, controlled fashion so that the users can be sure of the sender's identity?

A. Hash value

B. Private key

C. Digital signature

D. Digital certificate

ANSWER: D

Explanation:

A digital certificate distributes a public key in a controlled, trustworthy manner by binding that key to an identified subject, such as an individual, server, or organization. The certificate is digitally signed by a certificate authority or another trusted issuer. Before relying on the public key, a user or application validates the issuer's signature, the certificate's validity period, its intended key usage, and the certification path to a trusted root. This validation provides assurance that the public key belongs to the entity named in the certificate, rather than to an attacker impersonating that entity. In public-key infrastructures, certificates are commonly formatted according to the X.509 standard and contain the subject identity, public-key information, issuer identity, serial number, and signature-related data. This trusted binding is what enables authenticated

public-key distribution for services such as TLS-protected websites and signed communications. See the NIST definition of a digital certificate and RFC 5280, Internet X.509 Public Key Infrastructure Certificate and CRL Profile.

QUESTION NO: 5

User A is writing a sensitive email message to user B outside the local network. User A has chosen to use PKI to secure his message and ensure only user B can read the sensitive email. At what layer of the OSI layer does the encryption and decryption of the message take place?

A. Application

B. Transport

C. Session

D. Presentation

ANSWER: D

Explanation:

Presentation is correct because the Presentation layer, Layer 6 of the OSI model, is responsible for translating data into a format the receiving application can interpret and for data transformation services such as encryption and decryption. When User A uses PKI to protect an email for User B, the message is encrypted using cryptographic material associated with User B, so that User B can decrypt and read the protected content with the corresponding private key. In OSI-model terms, this encryption/decryption function is conventionally assigned to the Presentation layer. This layer also encompasses related transformations such as character-code conversion, serialization, compression, and formatting, ensuring that data has an appropriate representation for secure transmission and processing at the endpoints. PKI supplies the certificate and public/private key infrastructure used to support confidentiality and authentication, while the OSI-layer question specifically asks where the data encryption and decryption service is categorized. The ISO/IEC 7498-1 OSI Basic Reference Model defines the OSI architecture, and the RFC 1009 OSI Routing Framework describes the Presentation layer as providing representation services, including encryption and decryption.

QUESTION NO: 6

SQL injection (SQLi) attacks attempt to inject SQL syntax into web requests, which may Bypass authentication and allow attackers to access and/or modify data attached to a web application.

Which of the following SQLI types leverages a database server's ability to make DNS requests to pass data to

an attacker?

A. Union-based SQLI

B. Out-of-band SQLI

C. In-band SQLI

D. Time-based blind SQLI

ANSWER: B

Explanation:

Out-of-band SQLI is correct because it uses a separate communication channel from the normal web application request-and-response path to exfiltrate information. In this technique, an attacker causes the database server to initiate an outbound network request—commonly a DNS lookup, and sometimes an HTTP request—to infrastructure the attacker controls. Data can be encoded into a DNS subdomain, allowing the attacker to recover it from DNS queries received by the authoritative server or a monitoring service. This is particularly useful where application responses do not expose query results and where timing-based inference is impractical or unreliable. The technique depends on the database platform offering a function or procedure capable of generating outbound requests, and on network egress controls allowing the request to leave the

environment. For example, SQL Server environments have historically been susceptible to DNS-based data retrieval through features that resolve UNC paths, while other database products may expose network-capable packages. OWASP identifies out-of-band SQL injection as a blind SQL injection technique that relies on the database server making DNS or HTTP requests to attacker-controlled systems. See the OWASP Blind SQL Injection overview and PortSwigger's blind SQL injection documentation.

QUESTION NO: 7

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 DNS, restricting zone transfers, and providing subnet diversity between DNS servers are complementary DNS-security recommendations. DNS-server hardening reduces the exposed attack surface through secure configuration, prompt patching, minimal enabled services, restricted administrative access, and appropriately configured recursion and logging. NIST's DNS guidance emphasizes operating DNS infrastructure securely and protecting authoritative server functions: NIST SP 800-81 Revision 3.

Split-horizon operation separates internal DNS responses and namespace information from externally visible responses. This limits unnecessary disclosure of internal hostnames and addressing while allowing public authoritative DNS to serve only records intended for external users. Zone transfers can disclose an entire zone's records, so transfers should be permitted only to explicitly authorized secondary DNS servers, ideally using authenticated transfer mechanisms such as TSIG. Microsoft documents configuring zone transfers only to designated servers in its DNS zone-transfer guidance.

Finally, placing authoritative DNS servers on diverse subnets improves resilience and availability. A routing failure, subnet outage, or localized denial-of-service event is less likely to remove every authoritative server at once. Together, these practices provide layered protection for DNS confidentiality, integrity, and availability.

QUESTION NO: 8

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 is a human-focused attack technique in which an attacker manipulates, deceives, or persuades a person into disclosing sensitive information, granting access, or performing an action that benefits the attacker. Rather than initially exploiting a software vulnerability or defeating a technical control directly, the attacker takes advantage of human factors such as trust, urgency, fear, authority, curiosity, or helpfulness. For example, an attacker may impersonate an IT

support employee to persuade a user to reveal credentials, or send a convincing phishing message that directs the recipient to enter information into a fraudulent website. This makes social engineering highly relevant to ethical hacking because security assessments must consider people and organizational processes alongside networks, applications, and systems. The U.S. Cybersecurity and Infrastructure Security Agency describes social engineering as using deception to manipulate individuals into divulging confidential or personal information, and NIST identifies social engineering as a method used to gain unauthorized access through human interaction. See CISA: Social Engineering and NIST CSRC: Social Engineering.

QUESTION NO: 9

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

A. BA810DBA98995F1817306D272A9441BB

B. 44EFCE164AB921CAAAD3B435B51404EE

C. 0182BD0BD4444BF836077A718CCDF409

D. CEC52EB9C8E3455DC2265B23734E0DAC

E. B757BF5C0D87772FAAD3B435B51404EE

F. E52CAC67419A9A224A3B108F3FA6CB6D

ANSWER: B E

Explanation:

LM hashing processes a password as two independent seven-character halves. If the original password contains fewer than eight characters, the second half has no password characters and is therefore made entirely from the fixed padding used by the LM algorithm. Encrypting the relevant constant with that padding produces the recognizable second 16-hex-character value `AAD3B435B51404EE`. Consequently, an LM hash whose trailing half is `AAD3B435B51404EE` indicates that the password was seven characters or fewer. Both `44EFCE164AB921CAAAD3B435B51404EE` and `B757BF5C0D87772FAAD3B435B51404EE` have that exact trailing value, so each represents a password of less than eight characters. The first hash was normalized by replacing the non-hexadecimal `Q` with `C`; an LM hash must contain exactly 32 hexadecimal characters. This recognizable empty-second-half pattern is commonly used during password-audit analysis to identify short legacy LM passwords. Microsoft documents the legacy LM-hash storage behavior and associated security concerns at Microsoft Learn. A reference list of LM hash formats and examples is available from Hashcat example hashes.

QUESTION NO: 10

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

Explanation:

The version value in an SOA record is its *serial number*, which is **200302028**. The serial number is the first numeric field enclosed in the SOA record's parentheses, immediately after the primary authoritative name server and the responsible party's mailbox field. Secondary DNS servers use this value when checking whether the zone has changed: when the primary zone's serial is greater than the serial held by a secondary server, the secondary requests a zone transfer to obtain the updated data.

Although administrators commonly format a serial using a date-based convention, such as `YYYYMMDDnn`, DNS treats the field as an unsigned 32-bit serial value and compares it according to DNS serial-number arithmetic. In this record, the remaining numeric values are timing parameters—refresh, retry, expire, and negative-cache TTL—not the zone version. RFC 1035 defines the SOA RDATA field order and identifies SERIAL as the first numeric field; RFC 1912 provides operational guidance for incrementing SOA serial values whenever zone data changes. See RFC 1035, Section 3.3.13 and RFC 1912, Section 2.2.

QUESTION NO: 11

Tony is a penetration tester tasked with performing a penetration test. After gaining initial access to a target system, he finds a list of hashed passwords.

Which of the following tools would not be useful for cracking the hashed passwords?

A. John the Ripper

B. Hashcat

C. netcat

D. THC-Hydra

ANSWER: C

Explanation:

netcat is the correct answer because it is a general-purpose networking utility, not a password-hash recovery or password-auditing tool. Netcat is commonly used during authorized penetration testing to create TCP or UDP connections, transfer data, listen on ports, relay traffic, and perform basic network interaction. Those capabilities can support reconnaissance or post-exploitation workflows, but they do not provide hash-format recognition, candidate-password generation, dictionary processing, rule-based mutation, or cryptographic hash computation and comparison—the functions required to recover a plaintext password from a captured password hash.

For a list of hashes obtained after access, the relevant task is normally offline password auditing: candidate passwords are transformed with the relevant algorithm and compared to the captured hash value. Netcat has no built-in cracking engine or hash-verification capability, so it cannot meaningfully perform that task. The Netcat project describes the utility as a networking tool for reading from and writing to network connections, which aligns with its role in connectivity and data movement rather than cryptographic password recovery. See the Ncat Users' Guide and the Ncat project page.

QUESTION NO: 12

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|"; msG. "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, and its payload contains |00 01 86 a5|.

ANSWER: D

Explanation:

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, provided that its payload contains the specified byte sequence. This follows directly from the Snort rule header: `alert` is the action, `tcp` limits inspection to TCP traffic, and `any any` denotes any source IP address and source port. The directional operator `-->` indicates traffic flowing toward the destination, while `192.168.1.0/24 111` specifies any host in that subnet using destination port 111. Port 111 is commonly associated with the RPC port mapper service, and the message identifies the event as potential mount daemon access. The rule option `content:"|00 01 86 a5|"` is also required: Snort must find those four hexadecimal bytes in the inspected packet payload before producing the alert. Thus, the rule is not merely a port-based alert; it detects a particular TCP payload pattern in traffic directed to TCP port 111 on the stated network. Snort rule headers and payload options are documented in the Snort Rule Writing Guide: Headers and the Snort content option documentation.

QUESTION NO: 13

Windows LAN Manager (LM) hashes are known to be weak.

Which of the following are known weaknesses of LM? (Choose three.)

A. Converts passwords to uppercase.

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

C. Makes use of only 32-bit encryption.

D. Effective length is 7 characters.

ANSWER: A B D

Explanation:

LM password hashing has several design characteristics that make passwords substantially easier to recover than modern Windows credential formats. “Converts passwords to uppercase.” is correct because LM eliminates case sensitivity before deriving the hash, reducing the number of possible password combinations an attacker must test. “Effective length is 7 characters.” is also correct in practical cryptanalytic terms: LM divides a password into two independent seven-character portions and derives each portion separately. An attacker can therefore attack each half independently rather than having to search the full password as one value. “Hashes are sent in clear text over the network.” reflects the legacy LM security model’s inadequate protection of LM-derived credentials during older network authentication and communication scenarios; exposure of reusable credential material enables offline password-cracking attacks. Microsoft recommends disabling LM hash storage and avoiding legacy LM/NTLM authentication where possible, because these protocols do not provide the protection expected from modern authentication mechanisms. See Microsoft’s guidance on preventing LM hash storage and its overview of NTLM authentication.

QUESTION NO: 14

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 the appropriate choice because obstructed-URL attacks commonly use alternate encodings to disguise dangerous path characters and traversal sequences from simplistic string matching. For example, percent-encoded or nonstandard Unicode representations may be interpreted differently by an intermediary device and the destination web server. An intrusion detection system can be configured with signatures or normalization-aware inspection rules to recognize anomalous Unicode encodings, encoded directory-traversal patterns, and requests whose decoded form targets sensitive files or paths. Alerting on those events gives defenders visibility into attempted evasion and enables rapid investigation, blocking through an integrated IPS or WAF workflow, and tuning of web-server request filtering. Effective detection should normalize and decode input before applying inspection logic, while preserving the original request for forensic analysis; this avoids treating an obfuscated representation as harmless merely because its raw text does not contain an obvious traversal string. OWASP describes how encoded input can be used in path-traversal attacks and recommends validation and canonicalization controls: OWASP Path Traversal. NIST also provides guidance on using intrusion detection and prevention technologies to monitor and respond to suspicious network and application activity: NIST SP 800-94.

QUESTION NO: 15

The collection of potentially actionable, overt, and publicly available information is known as

A. Open-source intelligence

B. Real intelligence

C. Social intelligence

D. Human intelligence

ANSWER: A

Explanation:

Open-source intelligence is the correct term for intelligence derived from publicly available sources that can be collected openly and analyzed to produce actionable information. Common OSINT sources include public websites, news articles, social-media posts, search-engine results, public records, technical documentation, internet-facing service information, and published reports. In an ethical hacking engagement, OSINT is central to the reconnaissance phase because it helps the tester build an accurate picture of the target's people, systems, domains, technologies, and external exposure before conducting more active assessment activities. The defining characteristics in the question are that the information is overtly collected and publicly available; no covert human source, privileged access, or classified collection method is required. Although individual public facts may appear harmless, correlating and validating them can reveal meaningful security-relevant insights. The U.S. Intelligence Community describes open-source intelligence as intelligence derived from publicly available information, while CISA emphasizes using publicly accessible information to understand and manage cyber risk. See the Office of the Director of National Intelligence overview of intelligence disciplines and CISA's cyber threats and advisories resources.

QUESTION NO: 16

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

G. UDP 137 and UDP 138

ANSWER: C E G

Explanation:

Blocking **139**, **445**, and **UDP 137 and UDP 138** prevents the Windows NetBIOS/SMB service paths relevant to legacy Windows NT, Windows 2000, and Windows XP hosts from crossing a network firewall. NetBIOS over TCP/IP uses UDP port 137 for NetBIOS Name Service, which supports name registration and resolution, and UDP port 138 for NetBIOS Datagram Service. TCP port 139 carries NetBIOS Session Service, historically used for SMB file and printer sharing over NetBIOS. Microsoft documents these NetBIOS-over-TCP/IP port assignments in its service overview and network port requirements.

TCP port 445 should also be blocked at a perimeter firewall when the goal is to prevent externally reachable Windows file-sharing exposure. On Windows 2000 and later systems, SMB can use direct hosting over TCP 445 rather than using the NetBIOS session path on TCP 139. Therefore, filtering only the classic NetBIOS ports can still leave SMB services accessible through port 445. Microsoft's SMB security guidance recommends limiting SMB exposure and allowing it only where explicitly required. Applying these filters at untrusted boundaries reduces opportunities for reconnaissance, share enumeration, and unauthorized access to legacy Windows services.

QUESTION NO: 17

Which of the following viruses tries to hide from anti-virus programs by actively altering and corrupting the chosen service call interruptions when they are being run?

A. Macro virus

B. Stealth/Tunneling virus

C. Cavity virus

D. Polymorphic virus

ANSWER: B

Explanation:

Stealth/Tunneling virus is correct. A tunneling virus evades detection by intercepting, altering, or corrupting low-level operating-system service calls and interrupt handlers that an antivirus product relies on to inspect files, memory, and system activity. Rather than allowing the security software to receive an unmodified result from a requested operation, the malware tunnels beneath or around normal monitoring mechanisms and can provide a deceptive view that conceals its presence. This behavior is a form of stealth: the virus manipulates the interface between the operating system and defensive tools so malicious changes are not readily observable during scanning or execution. The defining clue is the active manipulation of service-call interruptions while they run, which specifically describes tunneling techniques. MITRE ATT&CK documents that adversaries may impair or evade defenses by modifying system-level mechanisms and security-tool visibility; see Impair Defenses. For broader authoritative background on malware behavior and defensive evasion, consult the U.S. Cybersecurity and Infrastructure Security Agency's malware guidance.

QUESTION NO: 18

A zone file consists of which of the following Resource Records (RRs)?

A. DNS, NS, AXFR, and MX records

B. DNS, NS, PTR, and MX records

C. SOA, NS, AXFR, and MX records

D. SOA, NS, A, and MX records

ANSWER: D

Explanation:

SOA, NS, A, and MX records is the correct answer because these are standard DNS resource-record types commonly contained in a forward DNS zone file. The Start of Authority (SOA) record defines authoritative administrative information for the zone, including its primary authoritative server, the responsible party, the zone serial number, and timing values used by secondary servers. Authoritative name-server (NS) records identify the DNS servers that are authoritative for the zone. Address (A) records associate host names with IPv4 addresses, allowing clients to locate hosts and services by name. Mail exchanger (MX) records specify the mail servers responsible for accepting email for a domain, along with their delivery preference values.

DNS zones are text representations of DNS data and can contain additional record types depending on operational requirements, such as AAAA, CNAME, TXT, SRV, or PTR records. However, SOA, NS, A, and MX are the intended conventional set in this question and are all valid DNS RRs used in zone data. RFC 1035 defines the DNS database and core resource-record formats, including SOA, NS, A, and MX: RFC 1035. The current registry of DNS RR type assignments is maintained by IANA: IANA DNS Parameters.

QUESTION NO: 19

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 accomplished with DNS because it is a DNS protocol operation used to replicate authoritative zone data from a primary server to one or more secondary servers. Full zone transfers are commonly referred to as AXFR transfers and are handled through DNS messaging, normally over TCP port 53. A zone transfer passes all zone information that a DNS server maintains because AXFR is intended to provide the receiving secondary server with a complete copy of the zone, including its resource records and associated zone content needed to answer authoritatively. This is why exposed or unrestricted transfers can reveal useful DNS information during reconnaissance. A zone transfer can be prevented by blocking all inbound TCP port 53 connections because full zone transfers require a TCP connection to the authoritative DNS server. In a production environment, the more targeted best practice is to allow TCP port 53 only from explicitly authorized secondary DNS servers and to configure zone-transfer access controls; however, blocking all inbound TCP port 53 traffic does prevent external AXFR connections to that server. The DNS zone-transfer specification is documented in RFC 5936. Operational controls for restricting transfers are also described in the BIND 9 Administrator Reference Manual.

QUESTION NO: 20

Which of the following practices help protect a Linux system from unauthorized SSH access? (Choose three.)

A. Disable direct root login

B. Use public-key authentication

C. Restrict access with firewall rules or AllowUsers settings

D. Enable Telnet as a fallback remote-access service

E. Rely only on changing the default SSH port

ANSWER: A B C

Explanation:

Disable direct root login, **use public-key authentication**, and **restrict SSH access with firewall rules or AllowUsers settings** are effective controls. Disabling direct root access requires administrators to authenticate through an individual account and then use approved privilege escalation, improving accountability and reducing direct attacks against the root account. Public-key authentication provides stronger protection than password-only authentication when private keys are protected appropriately, especially when combined with passphrases and multifactor authentication.

Network filtering and SSH daemon access restrictions reduce the number of systems and accounts that can initiate SSH sessions. Changing the default port may reduce opportunistic scanning noise, but it is not a substitute for authentication or access controls. Telnet should not be enabled as an alternative because it sends credentials and session traffic without encryption. OpenSSH documents access controls and authentication settings in its sshd_config manual page.

QUESTION NO: 21

Henry is a penetration tester who works for XYZ organization. While performing enumeration on a client

organization, he queries the DNS server for a specific cached DNS record. Further, by using this cached record, he determines the sites recently visited by the organization's user. What is the enumeration technique used by Henry on the organization?

A. DNS zone walking

B. DNS cache snooping

C. DNS SEC zone walking

D. DNS cache poisoning

ANSWER: B

Explanation:

DNS cache snooping is the correct enumeration technique because it tests a recursive DNS resolver’s cache to infer whether a domain name was resolved recently. A tester can send a query for a particular name and inspect characteristics of the response, especially the remaining time-to-live value or whether the server provides a cached recursive answer. If the record is already present in the resolver cache, that indicates a user or system using that resolver likely accessed the associated domain within the record’s cache lifetime. This can reveal browsing patterns, external services in use, update infrastructure, or other useful reconnaissance information without directly inspecting endpoint browser histories. The technique depends on a DNS server permitting recursive queries and exposing information sufficient to distinguish cached responses. DNS cache behavior and TTL handling are defined in the DNS standards; RFC 1034 describes resolvers using cached data to answer queries, while RFC 2308 further documents caching behavior for DNS responses. See RFC 1034 and RFC 2308.

QUESTION NO: 22

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:

USER2SID, **SID2USER**, and **DumpSec** are enumeration tools because they assist in identifying Windows accounts, security identifiers, and security-relevant configuration information. USER2SID converts a known Windows account or user name into its associated Security Identifier (SID), while SID2USER resolves a SID to the corresponding account name. These mappings are valuable during Windows and Active Directory enumeration because SIDs are central to Windows authorization and can reveal account identities, domain associations, and well-known principals. Microsoft describes SIDs as unique values used to identify security principals and control access decisions in Windows environments: Microsoft: Security Identifiers.

DumpSec is designed to collect and report security-related information from Windows systems, including users, groups, permissions, policies, shares, and related configuration data. Gathering this information is a core enumeration activity: it builds a detailed picture of accessible identities, privileges, and exposed resources before further security testing. The tool's documentation and product information describe its role in auditing and reporting Windows security settings: SystemTools/Somarsoft. Together, these tools specifically support account and security-information discovery, making them the appropriate three selections.

QUESTION NO: 23

A network administrator wants to reduce the attack surface of IPv6 on an enterprise LAN. Which of the following are appropriate controls? (Choose three.)

A. Enable RA Guard on untrusted access ports

B. Enable DHCPv6 Guard on untrusted access ports

C. Apply IPv6-aware firewall and ACL policies

D. Rely only on MAC address filtering

E. Disable all switch-port security features

ANSWER: A B C

Explanation:

RA Guard, **DHCPv6 Guard**, and **IPv6 access-control policies** are appropriate controls. RA Guard helps block unauthorized IPv6 Router Advertisement messages on untrusted switch ports, reducing the risk of rogue routers advertising malicious default gateways or prefixes. DHCPv6 Guard limits unauthorized DHCPv6 server messages. IPv6-aware firewall and ACL policies restrict traffic according to approved business requirements.

Disabling IPv6 without assessing business dependencies is not a general security control, and MAC filtering does not validate IPv6 router advertisements or DHCPv6 messages. Cisco documents first-hop security capabilities including RA Guard and DHCPv6 Guard in its IPv6 First-Hop Security overview.

QUESTION NO: 24

A security team is collecting evidence after a suspected compromise of a Windows server. Which of the following sources can provide useful authentication and account-activity evidence? (Choose three.)

A. Windows Security event log

B. Domain controller security logs

C. PowerShell operational logs

D. Recycle Bin contents

E. Browser bookmarks

ANSWER: A B C

Explanation:

The **Windows Security event log**, **domain controller security logs**, and **PowerShell operational logs** can provide valuable evidence. Windows Security logs may record successful and failed logons, account changes, and privilege-related events. Domain controllers provide centralized evidence of Active Directory authentication and account activity. PowerShell logs can reveal command execution and script-related events when the relevant logging policies are enabled.

The Recycle Bin is not a reliable audit source for authentication events, and browser bookmarks do not provide authoritative account-activity records. Microsoft documents Windows security auditing at Advanced Security Auditing and PowerShell logging at about_Logging.

QUESTION NO: 25

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 the system historically used to query publicly available registration records for domain names and IP address resources. For domain registrations, WHOIS data can include the registrar, registration and expiration dates, authoritative name servers, domain status, and—where publication is permitted under applicable privacy rules—registrant, administrative, and technical contact details. During reconnaissance, ethical hackers may use WHOIS to identify domain ownership relationships, DNS infrastructure, registrar information, and dates that could help establish the scope and history of an organization's Internet-facing assets.

Modern registration-data access is often subject to privacy controls, redaction requirements, and registry-specific policies. As a result, a WHOIS response may not disclose all contact fields that were traditionally available. Nevertheless, WHOIS remains the recognized term for the publicly queryable domain-registration information service and its associated databases. ICANN describes Registration Data Directory Services, including WHOIS, as services that provide access to domain-name registration data. The IANA's Root Zone Database also provides registry-related information and links for top-level domains, supporting discovery of the authoritative registry for a particular domain suffix.

References: ICANN Registration Data Directory Services; IANA Root Zone Database.

QUESTION NO: 26

Becky has been hired by a client from Dubai to perform a penetration test against one of their remote offices.

Working from her location in Columbus, Ohio, Becky runs her usual reconnaissance scans to obtain basic information about their network. When analyzing the results of her Whois search, Becky notices that the IP was allocated to a location in Le Havre, France. Which regional Internet registry should Becky go to for detailed information?

A. ARIN

B. APNIC

C. RIPE

D. LACNIC

ANSWER: C

Explanation:

RIPE is the appropriate regional Internet registry because Le Havre is in France, which falls within the RIPE NCC service region. RIPE NCC is responsible for coordinating and registering Internet number resources—including IPv4 and IPv6 address allocations and Autonomous System Numbers—across Europe, the Middle East, and parts of Central Asia. A WHOIS lookup through RIPE NCC can provide registration information associated with the relevant IP-address allocation, such as the organization or network maintaining the resource, abuse-contact details, allocation status, and related route or network objects where published. During authorized reconnaissance, this information helps a penetration tester accurately identify the organization responsible for the target address space and locate appropriate operational or abuse contacts if needed. The client's location in Dubai and Becky's location in Ohio do not determine the applicable registry; the registry is selected based on the geographic region in which the IP resource was allocated or registered. RIPE NCC provides its WHOIS query service and documentation at RIPE Database. Its regional service area is described by RIPE NCC.

QUESTION NO: 27

Which of the following are security benefits provided by Kerberos authentication? (Choose three.)

A. Mutual authentication between clients and services

B. Single sign-on to supported services

C. Reduced transmission of user passwords over the network

D. Permanent authentication tickets that never expire

E. Automatic encryption of all application data

ANSWER: A B C

Explanation:

Kerberos provides **mutual authentication**, **single sign-on**, and **reduced transmission of passwords across the network**. A client and service can validate each other through tickets issued by the Key Distribution Center (KDC). After obtaining a ticket-granting ticket, a user can request service tickets without repeatedly entering credentials, supporting single sign-on. Kerberos also avoids sending the user password to each requested network service.

Kerberos does not rely on plaintext passwords being sent to services, and it does not inherently encrypt every item of application data. Application-layer confidentiality depends on the selected protocol and whether message protection is negotiated. Microsoft provides an overview at Kerberos authentication overview.

QUESTION NO: 28

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

A. Use port security on his switches.

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

C. Use a firewall between all LAN segments.

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

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

ANSWER: A B D

Explanation:

Use port security on his switches, **use a tool like ARPwatch to monitor for strange ARP activity**, and **if you have a small network, use static ARP entries** are appropriate complementary controls for ARP-spoofing risk. Switch port security restricts which and how many MAC addresses can be learned on an access port. This reduces the opportunity for an attacker to introduce unauthorized devices or rapidly change source MAC addresses as part of local-layer attacks. Cisco documents port-security MAC-address restrictions and violation actions at Cisco: Configuring Port Security.

Static ARP mappings bind a known IP address to its expected MAC address, preventing an endpoint from accepting a malicious ARP reply that attempts to substitute the attacker's MAC address for a critical host such as the default gateway. They are most manageable for small environments or a limited set of high-value systems. ARPwatch adds important detection and response coverage by observing ARP traffic and reporting unexpected IP-to-MAC mapping changes, allowing administrators to investigate and contain poisoning attempts quickly. Its purpose and change-monitoring behavior are described by the ARPwatch manual page. Together, access-layer restrictions, trusted static bindings where practical, and continuous ARP monitoring provide defense in depth.

QUESTION NO: 29

What hacking attack is challenge/response authentication used to prevent?

A. Replay attacks

B. Scanning attacks

C. Session hijacking attacks

D. Password cracking attacks

ANSWER: A

Explanation:

Challenge/response authentication is used to prevent **Replay attacks**. Rather than transmitting a reusable credential directly, the authenticating system issues a unique, unpredictable challenge, often called a nonce. The client calculates a response using that challenge and a secret, such as a password-derived key or cryptographic key. Because the challenge is fresh for each authentication attempt, a response captured from an earlier exchange is not valid for a later challenge. This prevents an attacker who monitors network traffic from simply retransmitting previously captured authentication data to impersonate the legitimate user. Effective implementations must ensure challenges are unique, unpredictable, and accepted only within an appropriate validity period; otherwise, replay protection can be weakened. This pattern is a core security property of modern authentication protocols and is explicitly described in NIST guidance on nonce-based challenge-response mechanisms. See NIST SP 800-63B: Digital Identity Guidelines and RFC 4949: Internet Security Glossary.

QUESTION NO: 30

Todd has been asked by the security officer to purchase a counter-based authentication system. Which of the following best describes this type of system?

A. A biometric system that bases authentication decisions on behavioral attributes.

B. A biometric system that bases authentication decisions on physical attributes.

C. An authentication system that generates one-time passwords from a shared secret key and a moving counter.

D. An authentication system that uses passphrases that are converted into virtual passwords.

ANSWER: C

Explanation:

An authentication system that generates one-time passwords from a shared secret key and a moving counter is a counter-based authentication system. This approach is standardized as the HMAC-Based One-Time Password algorithm (HOTP). The authenticating server and the user's token or authenticator share a secret key and maintain a counter value. To create a one-time password, the authenticator applies an HMAC function to the secret and current counter, then derives a short numeric value for the user to enter. After a successful authentication, the counter advances, ensuring that the password is not intended for repeated use.

The counter is the defining feature: unlike time-based one-time-password methods, which derive codes from the current time, this method advances according to authentication events. Implementations need a controlled resynchronization process in case the client and server counters become misaligned. RFC 4226 defines HOTP and specifies the use of a shared secret with a moving factor, usually a counter: RFC 4226: HOTP. NIST also recognizes one-time-password authenticators and describes their use in digital identity authentication guidance: NIST SP 800-63B.

QUESTION NO: 31

A company's security policy states that all Web browsers must automatically delete their HTTP browser cookies upon terminating. What sort of security breach is this policy attempting to mitigate?

A. Attempts by attackers to access the user and password information stored in the company's SQL database.

B. Attempts by attackers to access Web sites that trust the Web browser user by stealing the user's authentication credentials.

C. Attempts by attackers to access password stored on the user's computer without the user's knowledge.

D. Attempts by attackers to determine the user's Web browser usage patterns, including when sites were visited and for how long.

ANSWER: B

Explanation:

Attempts by attackers to access Web sites that trust the Web browser user by stealing the user's authentication credentials is correct because cookies commonly hold a session identifier or authentication token after a user signs in. A server uses that value to associate subsequent browser requests with the authenticated session, so possession of a valid session cookie can allow an attacker to impersonate the user without needing the user's password. This is generally known as session hijacking or session theft.

Automatically deleting cookies at browser termination reduces the time that persistent authentication artifacts remain on an endpoint. This is especially useful on shared, public, or insufficiently protected computers, where another person or malicious software could retrieve a stored cookie after the legitimate user has closed the browser. Removing the cookie forces a new authentication flow when the browser is reopened and prevents an old, retained browser session from continuing to be trusted.

This control is part of sound session-management practice, alongside short session lifetimes, secure transport, and appropriate cookie protections. OWASP describes the risks associated with stolen session identifiers in its Session Management Cheat Sheet. Cookie lifetime and session-cookie behavior are also documented by MDN Web Docs.

QUESTION NO: 32

A security administrator wants to reduce the risk of sensitive information being exposed through a publicly accessible Git repository. Which of the following are appropriate actions? (Choose three.)

A. Remove secrets from source code and configuration files

B. Rotate credentials that were committed to the repository

C. Use pre-commit or CI scanning for secrets

D. Store credentials in source-code comments instead

E. Make the complete repository history publicly readable

ANSWER: A B C

Explanation:

Remove secrets from source code and configuration files, **rotate credentials that were committed**, and **use pre-commit or CI scanning for secrets** are appropriate actions. API keys, passwords, private keys, access tokens, and connection strings should not be stored in repositories, particularly public repositories. If a secret was committed, deleting the visible file is insufficient because repository history, clones, caches, and forks may retain the value. The exposed credential must be revoked or rotated.

Automated secret-scanning tools can identify likely credentials before code is committed or merged, reducing recurrence. Making repository history publicly readable or storing secrets in comments only increases disclosure risk. GitHub provides guidance on responding to exposed credentials and using secret-scanning capabilities in its Secret scanning documentation.

QUESTION NO: 33

Which of the following are characteristics of a properly implemented digital signature? (Choose three.)

A. Verifies that signed data has not been altered

B. Authenticates the signer through the associated public key

C. Supports non-repudiation evidence

D. Encrypts the message content for confidentiality

E. Eliminates the need to protect private keys

ANSWER: A B C

Explanation:

A properly implemented digital signature provides **integrity verification**, **authentication of the signer**, and evidence supporting **non-repudiation**. The signer creates the signature using a private key, while recipients validate it using the corresponding public key. If the signed content is changed after signing, signature verification fails, providing integrity evidence. Successful verification also establishes that the signer possessed the associated private key at the time of signing, subject to the trustworthiness of the certificate and key management process.

Digital signatures do not encrypt message content by themselves. Anyone with access to the message and the public key can verify a signature; confidentiality requires encryption or another approved protection mechanism. NIST describes the security services provided by digital signatures in FIPS 186-5, Digital Signature Standard.

QUESTION NO: 34

Which of the following actions help an organization defend against ransomware? (Choose three.)

A. Maintain tested offline or immutable backups

B. Apply security patches promptly

C. Use least privilege and restrict administrative access

D. Disable endpoint protection to improve performance

E. Depend on ransom payment as the recovery strategy

ANSWER: A B C

Explanation:

Maintain tested offline or immutable backups, **apply security patches promptly**, and **use least privilege with restricted administrative access** are important ransomware defenses. Tested backups provide a recovery path when data is encrypted or destroyed, and offline or immutable designs help prevent attackers from encrypting or deleting backup copies. Prompt patching reduces exposure to vulnerabilities that ransomware operators may exploit for initial access or lateral movement.

Least privilege limits the systems and data that a compromised account can affect. Restricting administrative privileges, segmenting access, and using separate privileged accounts can reduce the scope of an intrusion. Disabling endpoint security tools weakens detection and prevention, while paying a ransom does not ensure decryption or data recovery and does not prevent further extortion. CISA recommends resilient backups, patching, and access-control measures in its Ransomware Guide.

QUESTION NO: 35

Louis, a professional hacker, had used specialized tools or search engines to encrypt all his browsing activity and navigate anonymously to obtain sensitive/hidden information about official government or federal databases. After gathering the Information, he successfully performed an attack on the target government organization without being traced. Which of the following techniques is described in the above scenario?

A. Dark web footprinting

B. VoIP footpnnting

C. VPN footprinting

D. website footprinting

ANSWER: A

Explanation:

Dark web footprinting is the correct technique because the scenario describes gathering intelligence from hidden or non-indexed sources while preserving anonymity during reconnaissance. In CEH terminology, footprinting is the information-gathering phase used to develop a target profile before attempting an attack. Dark web footprinting specifically involves using anonymity-oriented networks and specialized search resources to locate exposed, leaked, restricted, or otherwise hidden information that may not be available through conventional web searches. The key indicators are the use of specialized tools or search engines, anonymous navigation, encrypted traffic, and the collection of sensitive information about government or federal systems before the attack. Tor is a common technology associated with this activity: it is designed to route traffic through multiple relays and provide privacy protections for users accessing services, including onion services that are not part of the ordinary indexed web. This intelligence can help an attacker identify useful target details while reducing direct attribution to the reconnaissance source. The Tor Project provides an overview of its anonymity network at https://www.torproject.org/about/overview/, and CISA discusses the risks and operational considerations surrounding dark-web exposure at https://www.cisa.gov/news-events/news/dark-web-and-cybercrime.

QUESTION NO: 36

What would be the purpose of running "wget 192.168.0.15 -q -S" against a web server?

A. Performing content enumeration on the web server to discover hidden folders

B. Using wget to perform banner grabbing on the webserver

C. Flooding the web server with requests to perform a DoS attack

D. Downloading all the contents of the web page locally for further examination

ANSWER: B

Explanation:

Using wget to perform banner grabbing on the webserver is correct. When wget requests the specified web server, the `-S` option instructs it to print the HTTP response headers received from the server. HTTP headers can disclose useful service-identification details, such as the HTTP protocol version, server software and version through a `Server` header when exposed, redirect locations, cookies, supported behavior, and other application or infrastructure clues. This is a form of banner grabbing because the tester is collecting identifying information supplied by the service in its response rather than attempting to enumerate directories, mirror site content, or generate high request volume. The `-q` option enables quiet operation, suppressing wget's normal progress and informational output; it does not prevent the explicitly requested server-response headers from being displayed with `-S`. In an authorized CEH-style reconnaissance workflow, reviewing these headers helps identify the web service and informs subsequent, appropriately scoped assessment activities. GNU documents `-S` as printing server headers and `-q` as turning off wget's output: GNU Wget Manual. HTTP response header semantics are standardized in RFC 9110.

QUESTION NO: 37

This type of injection attack does not show any error message. It is difficult to exploit as it returns information when the application is given SQL payloads that elicit a true or false response from the server. By observing the response, an attacker can extract sensitive information. What type of attack is this?

A. Time-based SQL injection

B. Union SQL injection

C. Error-based SQL injection

D. Blind SQL injection

ANSWER: D

Explanation:

Blind SQL injection is the correct answer because the application does not directly return database errors or query results to the attacker. Instead, the attacker submits carefully constructed SQL conditions that evaluate to true or false and observes a detectable difference in the application's behavior. For example, a true condition might cause a page to display normal content, return a different response length, redirect, or produce another observable change, whereas a false condition does not. By repeatedly testing such conditions, an attacker can infer database details one bit, character, or value at a time, including database names, table names, and sensitive stored data.

This technique is often called boolean-based blind SQL injection because information is inferred from true/false response behavior. It can be slower and more difficult to exploit than injection that directly exposes query output, but it remains a serious risk wherever application input is incorporated into SQL statements without parameterized queries and proper input handling. OWASP describes blind injection as an inference technique in which an attacker learns information from application behavior rather than visible error output. See the OWASP Blind SQL Injection page and the OWASP SQL Injection Prevention Cheat Sheet.

QUESTION NO: 38

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 C D E

Explanation:

SNMPUtil, SNScan, SNMPScan, Solarwinds IP Network Browser, and NMap can each be used to obtain SNMP-related information from reachable devices when the appropriate SNMP version, credentials, and network access are available. SNMPUtil is a command-line utility for sending SNMP requests and retrieving management information. SNScan and SNMPScan are purpose-built SNMP discovery/enumeration tools that query SNMP-enabled hosts and collect exposed management data. Solarwinds IP Network Browser provides a graphical SNMP browsing capability for viewing managed-device information and object identifiers. NMap also supports practical SNMP enumeration through its Nmap Scripting Engine; for example, the `snmp-info` script retrieves system details, while additional SNMP scripts can identify accessible community strings and query device data. In an authorized assessment, the administrator should use valid SNMP credentials or approved community strings and query only in-scope systems. SNMP enumeration is valuable because management interfaces can reveal device identity, operating-system details, interfaces, routing information, and other inventory data needed for administration and security validation. Nmap documents its SNMP NSE scripts in the Nmap SNMP script category, and Microsoft documents command-line SNMP tooling in its SNMPUTIL reference.

QUESTION NO: 39

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:

DNS zone transfer functionality is implemented through the AXFR query type, which requests a complete copy of a DNS zone from an authoritative name server. NSLookup can request zone information using its interactive `ls` command when the queried server permits transfers. Dig directly supports AXFR queries, such as `dig AXFR example.com @server`, and is a standard utility for testing authorized DNS transfer exposure. Host also supports complete zone transfer requests through its zone-listing capability, for example `host -l example.com server`. Sam Spade, a legacy network reconnaissance suite commonly referenced in CEH-era materials, includes DNS-query and zone-transfer capabilities. These tools can retrieve the records only when the DNS server is configured to allow the requesting client to transfer the zone; properly configured authoritative servers restrict transfers to approved secondary DNS servers. AXFR is formally specified as the DNS full zone-transfer mechanism in RFC 5936. Current BIND documentation also describes using `dig` to make DNS queries and is useful for authorized DNS administration and assessment: BIND 9 manual pages.

QUESTION NO: 40

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 password-cracking programs used in authorized password-security auditing and recovery scenarios. L0phtcrack is designed primarily to audit Windows account-password hashes. It can obtain or import Windows password hashes and perform dictionary, brute-force, hybrid, and related cracking attacks to identify weak credentials. Its use in a sanctioned assessment helps administrators find passwords that can be guessed or recovered too easily and then improve password controls.

John the Ripper is a widely used open-source password-auditing and recovery tool. It supports many password-hash formats and uses configurable wordlists, rules, incremental modes, and other techniques to test password strength. The John the Ripper project describes the tool as being intended to detect weak Unix passwords, while its broader implementations support numerous platforms and hash types. Both tools therefore directly fit the category of password crackers and are commonly referenced in ethical-hacking and credential-auditing contexts. They should be used only with explicit authorization and within the approved scope of a security assessment. See the L0phtCrack website and the Openwall John the Ripper project page.

QUESTION NO: 41

Which of the following incident handling process phases is responsible for defining rules, collaborating human workforce, creating a back-up plan, and testing the plans for an organization?

A. Preparation phase

B. Containment phase

C. Identification phase

D. Recovery phase

ANSWER: A

Explanation:

The **Preparation phase** is responsible for establishing an organization's ability to respond effectively before an incident occurs. This includes defining incident-response policies, procedures, roles, escalation paths, and rules for handling security events. It also involves organizing and training the incident-response team and ensuring coordination among relevant personnel, such as IT, security, legal, management, and communications staff. Preparation includes planning for business continuity through backups and recovery capabilities, then exercising and testing those plans so gaps can be identified and corrected before a real security incident creates operational pressure. Well-prepared organizations maintain appropriate tools, asset information, communications mechanisms, and documented playbooks, enabling a faster and more consistent response when suspicious activity is detected. NIST's incident-response lifecycle identifies preparation as the foundational activity that develops policies, procedures, resources, and readiness required for incident handling. The related contingency-planning guidance also emphasizes testing plans and maintaining recovery capabilities, including backups. See NIST SP 800-61 Rev. 2, Computer Security Incident Handling Guide and NIST SP 800-34 Rev. 1, Contingency Planning Guide.

QUESTION NO: 42

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

access-list 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” is appropriate because an SNMP access control list commonly permits requests only from designated management-network addresses. If the discovered community has write access, a request appearing to originate from an allowed management address can invoke Cisco configuration-copy functionality. Cisco devices expose configuration-copy operations through the CISCO-CONFIG-COPY-MIB, enabling a running or startup configuration to be copied to a specified TFTP destination. The relevant operation is performed by the router after it accepts the authorized-looking SNMP request, so the configuration transfer can be directed to a test-controlled TFTP service.

“Run a network sniffer and capture the returned traffic with the configuration file from the router” is also appropriate because traditional TFTP transfers carry configuration data without encryption. Capturing the resulting router-to-TFTP traffic allows the tester to reconstruct the transferred configuration, including the contents sent during the copy operation. These two actions work together: the spoofed, permitted-source SNMP write triggers the copy, and passive packet capture obtains the plaintext configuration transfer. Cisco documents SNMP support and its management functions in its SNMP configuration guide; the underlying SNMP protocol behavior is defined in RFC 1157.

QUESTION NO: 43

Insecure direct object reference is a type of vulnerability where the application does not verify if the user is authorized to access the internal object via its name or key. Suppose a malicious user Rob tries to get access to the account of a benign user Ned.

Which of the following requests best illustrates an attempt to exploit an insecure direct object reference vulnerability?

A. “GET /restricted/goldtransfer?to=Rob&from=1 or 1=1’ HTTP/1.1Host: westbank.com”

B. “GET /restricted/\r\n\%00account%00Ned%00access HTTP/1.1 Host: westbank.com”

C. “GET /restricted/accounts/?name=Ned HTTP/1.1 Host westbank.com”

D. “GET /restricted/ HTTP/1.1 Host: westbank.com

ANSWER: C

Explanation:

“GET /restricted/accounts/?name=Ned HTTP/1.1 Host: westbank.com” best illustrates an insecure direct object reference attempt because it directly supplies Ned’s account identifier in a request for a restricted account resource. In this scenario, Rob changes or chooses the object reference represented by the `name` parameter, attempting to make the application return an account that belongs to another user. The security flaw exists if the server uses that identifier to retrieve the account but fails to enforce an object-level authorization check confirming that the authenticated requester is permitted to access Ned’s

record. A secure implementation must not rely on an untrusted URL parameter as proof of ownership; it must validate authorization for the specific requested account on every request. This is commonly categorized as broken access control, including insecure direct object references. OWASP describes this risk as occurring when an application exposes object references and does not verify that the user is authorized to access the requested object. See the OWASP Broken Access Control guidance and the OWASP Authorization Cheat Sheet.

QUESTION NO: 44

Which of the following practices help protect an organization's cloud storage buckets from unintended public exposure? (Choose three.)

A. Enable public-access blocking where public access is not required

B. Apply least-privilege bucket policies

C. Review storage permissions and access logs regularly

D. Embed long-term access keys in public client-side code

E. Allow anonymous write access for easier collaboration

ANSWER: A B C

Explanation:

Blocking public access, **applying least-privilege bucket policies**, and **regularly reviewing access logs and permissions** help prevent unintended cloud-storage exposure. Public-access blocking prevents broad anonymous access when it is not explicitly required. Least-privilege policies limit identities to the specific storage actions and resources needed for their duties. Permission reviews and logging help identify accidental policy changes and anomalous access.

Embedding long-term cloud credentials in publicly accessible application code and granting anonymous write permissions increase risk rather than reduce it. AWS provides storage-security guidance in its Amazon S3 security best practices.

QUESTION NO: 45

Mary, a penetration tester, has found password hashes in a client system she managed to breach. She needs to use these passwords to continue with the test, but she does not have time to find the passwords that correspond to these hashes. Which type of attack can she implement in order to continue?

A. LLMNR/NBT-NS poisoning

B. Internal monologue attack

C. Pass the ticket

D. Pass the hash

ANSWER: D

Explanation:

Pass the hash is the appropriate technique because it allows a tester to authenticate to applicable Windows services by using a captured NTLM password hash directly, without recovering the original plaintext password. In Windows NTLM authentication, the secret is used in a challenge-response process; an attacker who possesses a usable NTLM hash may be able to present it as credential material and establish an authenticated session where NTLM is accepted. This is particularly useful during an authorized penetration test when time constraints make offline password cracking impractical and the objective is to validate the potential for lateral movement or access expansion. The technique depends on the target environment, protocol, account permissions, and whether protections such as Credential Guard or restrictions on NTLM use are in place. MITRE ATT&CK documents pass-the-hash under its “Use Alternate Authentication Material” technique and notes that adversaries can use stolen password hashes to authenticate to remote systems. Microsoft also describes NTLM

as a challenge-response authentication protocol, which explains why authentication can occur without transmitting or knowing the original password. See MITRE ATT&CK: Pass the Hash and Microsoft: NTLM overview.

QUESTION NO: 46

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

A. Isolate affected systems from the network

B. Block confirmed malicious indicators

C. Preserve relevant logs and forensic evidence

D. Delete all logs immediately

E. Restore every affected system before investigating

ANSWER: A B C

Explanation:

Isolating affected systems, **blocking known malicious indicators**, and **preserving relevant evidence** are appropriate containment actions. Isolation can prevent an attacker or malware from moving laterally or communicating with external infrastructure. Blocking confirmed malicious domains, IP addresses, hashes, or accounts can limit further activity. Evidence should be preserved before making disruptive changes so that investigators can establish scope and root cause.

Immediately deleting all logs destroys evidence, and restoring systems from backups is generally part of recovery after eradication and validation rather than initial containment. NIST describes containment, eradication, and recovery activities in SP 800-61 Rev. 2, Computer Security Incident Handling Guide.

QUESTION NO: 47

Which of the following controls help mitigate clickjacking attacks against a web application? (Choose three.)

A. Set X-Frame-Options to DENY or SAMEORIGIN

B. Use Content-Security-Policy frame-ancestors

C. Require server-side confirmation for sensitive actions

D. Disable JavaScript for all users

E. Use HTTPS without additional response headers

ANSWER: A B C

Explanation:

X-Frame-Options, a restrictive **Content-Security-Policy frame-ancestors directive**, and server-side confirmation for sensitive actions are appropriate controls. Clickjacking deceives a victim into clicking an interface element that is hidden or overlaid within an attacker-controlled page, commonly by loading the legitimate application in an iframe. The `X-Frame-Options` response header can prevent framing or limit framing to the same origin. The modern and more flexible `frame-ancestors` Content Security Policy directive specifies which origins may embed the application.

For high-impact operations, such as changing an email address or authorizing a payment, requiring explicit server-side confirmation provides defense in depth if a framing protection is absent or bypassed. Removing all JavaScript does not address the browser framing condition, and HTTPS protects transport confidentiality and integrity but does not by itself prevent a page from being framed. OWASP describes clickjacking defenses in its Clickjacking Defense Cheat Sheet.

QUESTION NO: 48

Yancey is a network security administrator for a large electric company. This company provides power for over 100, 000 people in Las Vegas. Yancey has worked for his company for over 15 years and has become very successful. One day, Yancey comes in to work and finds out that the company will be downsizing and he will be out of a job in two weeks. Yancey is very angry and decides to place logic bombs, viruses, Trojans, and backdoors all over the network to take down the company once he has left. Yancey does not care if his actions land him in jail for 30 or more years, he just wants the company to pay for what they are doing to him.

What would Yancey be considered?

A. Yancey would be considered a Suicide Hacker

B. Since he does not care about going to jail, he would be considered a Black Hat

C. Because Yancey works for the company currently; he would be a White Hat

D. Yancey is a Hacktivist Hacker since he is standing up to a company that is downsizing

ANSWER: A

Explanation:

“Yancey would be considered a Suicide Hacker” is correct. In CEH terminology, a suicide hacker is a malicious actor who is willing to accept severe personal consequences—including identification, prosecution, imprisonment, or loss of employment—in order to damage a target. The defining feature is not merely unauthorized access; it is the attacker’s indifference to the consequences of being caught while pursuing destructive or retaliatory objectives.

Yancey is an especially clear example because he is an insider with privileged knowledge and access to the organization’s environment. His planned deployment of logic bombs, viruses, Trojans, and backdoors is intended to cause harm after his departure, and he expressly states that he does not care about a potentially lengthy jail sentence. This combination of retaliatory motive, planned sabotage, and disregard for legal consequences matches the suicide-hacker classification. Security teams should treat this scenario as an insider-threat risk and promptly apply access revocation, monitoring, separation procedures, and incident-response controls. CISA describes insider threats as risks posed by individuals with authorized access who may misuse that access: CISA Insider Threat Mitigation. Intentional deployment of malicious code and unauthorized damage may also implicate the U.S. Computer Fraud and Abuse Act: U.S. DOJ CFAA overview.