PassQueen
a.k.a dumpsqueen.com

EC-Council Certified Security Analyst (ECSA) V10

ECCouncil 412-79v10

Version Demo

Total Demo Questions: 27

Total Premium Questions: 279

Buy Premium PDF

https://passqueen.com

support@passqueen.com

Topic Break Down

Topic	No. of Questions
Topic 1, Introduction to Penetration Testing and Methodologies	24
Topic 2, Penetration Testing Scoping and Engagement	29
Topic 3, Open Source Intelligence (OSINT)	16
Topic 4, Social Engineering Penetration Testing	9
Topic 5, Network Penetration Testing	100
Topic 6, Vulnerability Assessment	13
Topic 7, System Hacking	30
Topic 8, Web Application Penetration Testing	40
Topic 9, Wireless Penetration Testing	15
Topic 10, Mix Questions	3
Total	279

QUESTION NO: 1

What is the following command trying to accomplish?

A. Verify that NETBIOS is running for the 192.168.0.0 network

B. Verify that TCP port 445 is open for the 192.168.0.0 network

C. Verify that UDP port 445 is open for the 192.168.0.0 network

D. Verify that UDP port 445 is closed for the 192.168.0.0 networks

ANSWER: C

Explanation:

“Verify that UDP port 445 is open for the 192.168.0.0 network” is the intended result when the omitted command uses a UDP scan against port 445 across that target network. UDP scanning sends probes to the specified UDP port and evaluates responses, including UDP replies or ICMP unreachable messages, to determine whether the port appears open, closed, or open|filtered. Port 445 is registered for the Microsoft-DS service over both TCP and UDP, so explicitly selecting UDP changes the protocol being assessed; it does not test the TCP service. This type of targeted scan is useful during an authorized assessment to identify systems exposing UDP services on the selected port and to determine whether additional service enumeration or validation is appropriate. Nmap’s reference guide describes UDP scanning behavior and the response analysis used to infer UDP port states: Nmap UDP Scan documentation. The IANA Service Name and Transport Protocol Port Number Registry lists the Microsoft-DS assignment for port 445 and its supported transport protocols: IANA port 445 registry entry.

QUESTION NO: 2

A penetration tester is preparing a Rules of Engagement document before beginning an assessment of a production network.

Which of the following items should be included in the Rules of Engagement? (Select 3)

A. Approved target systems and network ranges

B. Testing hours and operational restrictions

C. Communication and escalation contacts

D. Final vulnerability severity ratings

E. Recovered password hashes

ANSWER: A B C

Explanation:

The Rules of Engagement should define the approved systems and networks to be tested, the testing schedule and maintenance-window restrictions, and the communication and escalation procedures to be followed if an issue is discovered. These details establish the tester's authorization and help prevent testing activities from causing unintended business disruption.

Password-cracking results and final vulnerability ratings are produced during or after the technical assessment, not established as pre-engagement operating rules. NIST recommends that security-testing plans document scope, constraints, points of contact, and incident-handling procedures before testing begins. See NIST SP 800-115.

QUESTION NO: 3

An assessor is reviewing the results of an Active Directory security assessment.

Which of the following conditions can increase the risk of Kerberoasting attacks? (Select 2)

A. Service accounts with registered Service Principal Names

B. Service accounts using weak or easily guessed passwords

C. Multi-factor authentication required for interactive users

D. Anonymous LDAP binds disabled

E. DNS zone transfers restricted to approved servers

ANSWER: A B

Explanation:

Service accounts with registered Service Principal Names (SPNs) and **service accounts using weak or easily guessed passwords** are correct. A domain user can generally request Kerberos service tickets for accounts associated with SPNs. The ticket is encrypted using material derived from the service account's secret, allowing an attacker to attempt offline password cracking against the captured ticket data.

Multi-factor authentication for interactive users and disabling anonymous LDAP binds do not directly eliminate the risk created by weak service-account credentials associated with SPNs. Strong, long, unique service-account passwords or managed service accounts help reduce this exposure. Microsoft documents SPNs and their relationship to Kerberos service authentication at Service Principal Names.

QUESTION NO: 4

Vulnerability assessment is an examination of the ability of a system or application, including the current security procedures and controls, to withstand assault.

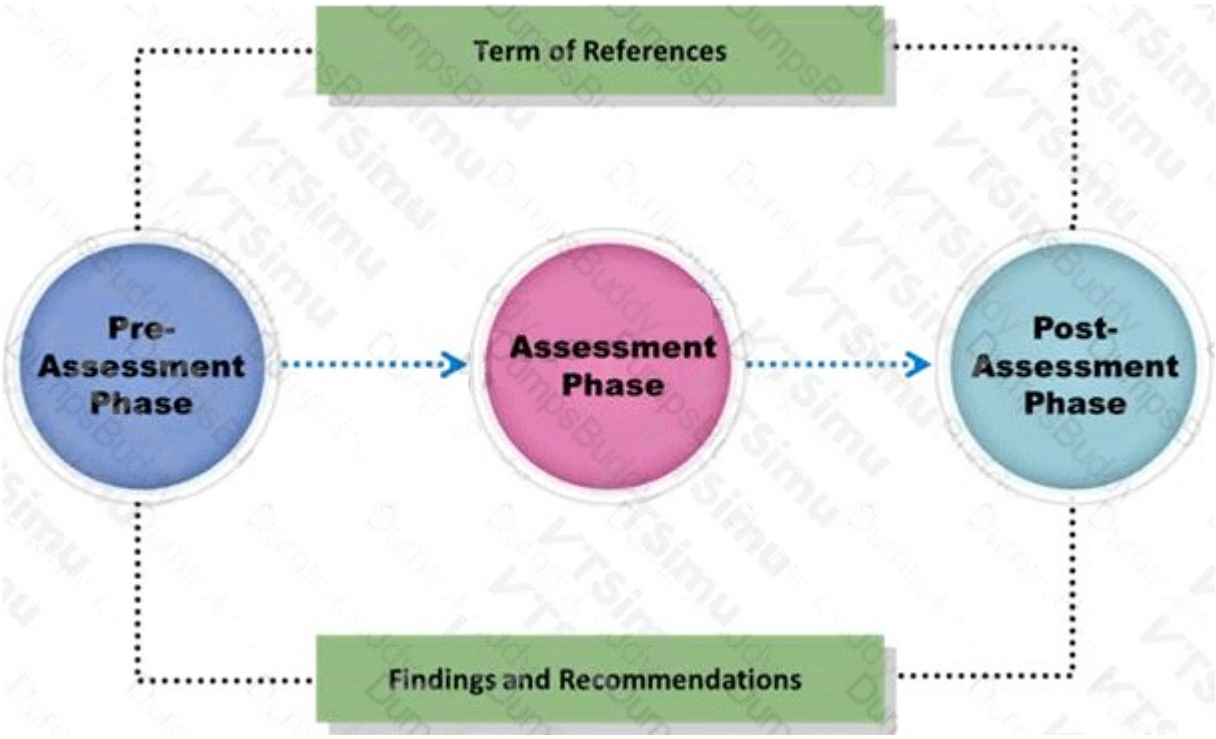

What does a vulnerability assessment identify?

A. Disgruntled employees

B. Weaknesses that could be exploited

C. Physical security breaches

D. Organizational structure

ANSWER: B

Explanation:

Vulnerability assessment identifies **weaknesses that could be exploited**. It is a structured process for discovering, cataloging, and prioritizing security weaknesses in systems, applications, networks, configurations, and implemented controls. The assessment typically uses automated scanners, configuration reviews, manual validation, and threat intelligence to identify conditions such as unpatched software, insecure services, weak authentication settings, excessive permissions, and known flaws associated with installed products.

The key outcome is an actionable view of exposure: each identified weakness can be evaluated according to its technical severity, likelihood of exploitation, affected assets, and potential business impact. This enables the organization to remediate, mitigate, or formally accept risk in a prioritized manner. Although an assessment may consider a broad range of technical and procedural controls, its defining purpose is to find exploitable vulnerabilities before an attacker can use them. NIST describes vulnerability scanning as identifying hosts and associated vulnerabilities, while CISA emphasizes remediation of known exploitable vulnerabilities as a core risk-reduction activity. See NIST SP 800-115: Technical Guide to Information Security Testing and Assessment and CISA Known Exploited Vulnerabilities Catalog.

QUESTION NO: 5

A tester is conducting passive reconnaissance of an organization's Internet presence before active testing begins.

Which of the following are passive reconnaissance techniques? (Select 3)

A. Reviewing public RDAP registration records

B. Querying publicly available DNS records

C. Reviewing public social-media profiles

D. Running a TCP SYN port scan

E. Performing directory brute forcing against a web server

ANSWER: A B C

Explanation:

Reviewing public WHOIS or RDAP records, searching public DNS information, and examining publicly accessible social-media profiles are passive reconnaissance techniques because they use information already published by registries, DNS infrastructure, or individuals without directly probing the target systems.

Port scanning and directory brute forcing actively send requests or probes to target infrastructure and are therefore active reconnaissance activities. Passive collection is useful during the early stages of an assessment because it can help identify domains, address allocations, personnel roles, and public technology information without generating traffic to the target environment. See ICANN RDAP information and NIST SP 800-115.

QUESTION NO: 6

A penetration tester has been authorized to assess a public web application that accepts file uploads from authenticated users.

Which of the following controls help reduce risks associated with file-upload functionality? (Select 3)

A. Allowlist permitted file types

B. Generate server-side file names

C. Store uploaded files outside the web root

D. Trust the Content-Type value supplied by the browser

E. Allow uploaded files to retain executable permissions

ANSWER: A B C

Explanation:

Allowlisting permitted file types, **generating server-side file names**, and **storing uploaded files outside the web root** are correct. These controls reduce the chance that an attacker can upload an executable file, control a dangerous file path, or directly request an uploaded file for execution through the web server.

Trusting the client-provided Content-Type header alone is insufficient because it can be modified. Allowing uploaded files to retain executable permissions creates unnecessary risk. OWASP provides detailed recommendations in the File Upload Cheat Sheet.

QUESTION NO: 7

A security analyst is assessing a Linux server that exposes SSH to the Internet.

Which of the following measures help reduce the risk associated with SSH remote access? (Select 3)

A. Use public-key authentication

B. Disable direct root login

C. Restrict access to approved source addresses

D. Enable Telnet as a fallback service

E. Use one shared administrator account

ANSWER: A B C

Explanation:

Using public-key authentication, disabling direct root login, and restricting SSH access with firewall rules or an allowlist are appropriate controls. Public-key authentication can reduce dependence on passwords, disabling direct root login requires administrative access to be performed through an accountable user account, and network restrictions reduce the number of systems that can attempt to connect to the service.

Enabling Telnet would replace encrypted SSH communication with an insecure cleartext protocol, and sharing a common administrator account removes accountability. OpenSSH documents configuration directives such as `PermitRootLogin` and public-key authentication in its server configuration reference. See OpenSSH sshd_config documentation.

QUESTION NO: 8

The IP protocol was designed for use on a wide variety of transmission links. Although the maximum length of an IP datagram is 64K, most transmission links enforce a smaller maximum packet length limit, called a MTU.

The value of the MTU depends on the type of the transmission link. The design of IP accommodates MTU differences by allowing routers to fragment IP datagrams as necessary. The receiving station is responsible for reassembling the fragments back into the original full size IP datagram.

IP fragmentation involves breaking a datagram into a number of pieces that can be reassembled later. The IP source, destination, identification, total length, and fragment offset fields in the IP header, are used for IP fragmentation and reassembly.

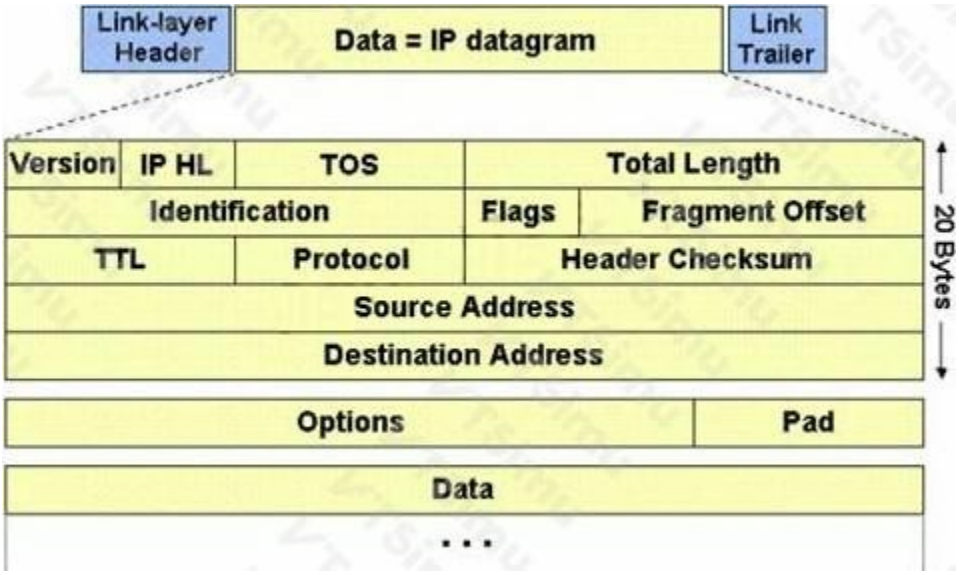

The fragment offset is 13 bits and indicates where a fragment belongs in the original IP datagram. This value is a:

A. Multiple of four bytes

B. Multiple of two bytes

C. Multiple of eight bytes

D. Multiple of six bytes

ANSWER: C

Explanation:

Multiple of eight bytes is correct because the IPv4 Fragment Offset field does not store a byte-for-byte position. Instead, its 13-bit value specifies the fragment's starting position within the original IPv4 datagram in units of eight octets (eight-byte blocks). For example, a Fragment Offset value of 185 means that the fragment's data begins at byte 1,480 of the original datagram (185 × 8). This unit size lets IPv4 represent offsets across a full datagram while using only 13 header bits.

Consequently, when a datagram is fragmented, every fragment other than the final fragment normally carries a payload length that is divisible by eight bytes. This alignment ensures the next fragment can begin at a position that the Fragment Offset field can represent. The destination system uses the Identification value together with source and destination addresses, protocol, the More Fragments flag, and these eight-byte-based offsets to place each fragment's payload in the proper location during reassembly.

RFC 791 defines the IPv4 Fragment Offset as measured in "units of 8 octets," and the IANA IPv4 header description likewise identifies it as a 13-bit field measured in eight-octet units. See RFC 791, Internet Protocol and IANA IPv4 Header Parameters.

QUESTION NO: 9

Windows stores user passwords in the Security Accounts Manager database (SAM), or in the Active Directory database in domains. Passwords are never stored in clear text; passwords are hashed and the results are stored in the SAM.

NTLM and LM authentication protocols are used to securely store a user's password in the SAM database using different hashing methods.

The SAM file in Windows Server 2008 is located in which of the following locations?

A. c:\windows\system32\config\SAM

B. c:\windows\system32\drivers\SAM

C. c:\windows\system32\Setup\SAM

D. c:\windows\system32\Boot\SAM

ANSWER: A

Explanation:

The correct location is **c:\windows\system32\config\SAM**. On Windows Server 2008, the Security Accounts Manager is implemented as a Registry hive named `SAM`. Registry hives that are backed by files are stored under the system root's `System32\Config` directory; when Windows is installed in its usual default location, the system root is `C:\Windows`. Therefore, the on-disk hive file is normally `C:\Windows\System32\Config\SAM`.

While Windows is running, the SAM hive is loaded into the Registry under `HKEY_LOCAL_MACHINE\SAM`. Its contents are protected by the operating system and cannot ordinarily be copied directly from the live system because the hive is in use and access-controlled. In an authorized security assessment or forensic workflow, analysts may instead work from an offline disk image, an appropriate backup, or a Volume Shadow Copy. The stored credential material is protected as password-derived hash data, and access to the SAM alone is not equivalent to ordinary readable password storage.

Microsoft documents that Registry hive files reside in the `%SystemRoot%\System32\Config` directory, which establishes this location. See Microsoft's Reg Save documentation and Microsoft's Reg Load documentation.

QUESTION NO: 10

An assessor has obtained a forensic image of a compromised server and must preserve it for possible legal proceedings.

Which of the following practices support proper evidence handling? (Select 3)

A. Calculate and record a cryptographic hash of the image

B. Maintain a chain-of-custody record

C. Analyze a verified working copy of the evidence

D. Delete unneeded files from the original evidence

E. Modify timestamps to simplify timeline analysis

ANSWER: A B C

Explanation:

Calculating and documenting cryptographic hashes, maintaining a chain-of-custody record, and storing the original evidence securely while examining a verified copy all support evidence integrity. Hashes provide a means to demonstrate that data has not changed, chain-of-custody records identify who possessed evidence and when, and working from a copy helps preserve the original artifact.

Modifying timestamps or deleting files from the original system changes potential evidence and undermines the ability to establish integrity. NIST guidance for incident response emphasizes preserving evidence and documenting handling activities. See NIST SP 800-86: Guide to Integrating Forensic Techniques into Incident Response.

QUESTION NO: 11

What will the following URL produce in an unpatched IIS Web Server?

A. Execute a buffer flow in the C: drive of the web server

B. Insert a Trojan horse into the C: drive of the web server

C. Directory listing of the C:\windows\system32 folder on the web server

D. Directory listing of C: drive on the web server

ANSWER: D

Explanation:

The intended result is **Directory listing of C: drive on the web server**. This question refers to the historical IIS Unicode/canonicalization directory-traversal vulnerability. In affected, unpatched IIS configurations, specially encoded path separators could bypass normal path validation and traverse outside the published web root. When the crafted request reached the command interpreter and included a command such as `dir c:\`, IIS returned the command output in the HTTP response. That output is a listing of the root of the server's C: drive, subject to the permissions of the IIS process account and the exact command included in the omitted URL. This was a serious vulnerability because it could expose filesystem structure and, in exploitable configurations, allow commands to be executed through the web service. Microsoft documented the issue as an IIS canonicalization error and released a patch to correct the improper handling of encoded directory traversal sequences. See Microsoft's MS00-078 security bulletin and the associated IIS canonicalization patch information.

QUESTION NO: 12

An organization is hardening its web servers against clickjacking attacks.

Which of the following HTTP response headers can be used to restrict whether a page may be displayed in a frame? (Select 2)

A. X-Frame-Options

B. Content-Security-Policy: frame-ancestors

C. Strict-Transport-Security

D. X-Content-Type-Options

E. Referrer-Policy

ANSWER: A B

Explanation:

X-Frame-Options and a Content Security Policy **frame-ancestors** directive are correct. X-Frame-Options can instruct browsers to deny framing or allow framing only by the same origin. CSP frame-ancestors provides a modern and more flexible mechanism for defining which origins are permitted to frame a page.

Strict-Transport-Security enforces HTTPS use, while X-Content-Type-Options controls MIME-type sniffing. Neither header prevents a page from being embedded in a frame. OWASP documents clickjacking defenses in the Clickjacking Defense Cheat Sheet.

QUESTION NO: 13

TCP/IP provides a broad range of communication protocols for the various applications on the network. The TCP/IP model has four layers with major protocols included within each layer. Which one of the following protocols is used to collect information from all the network devices?

A. Simple Network Management Protocol (SNMP)

B. Network File system (NFS)

C. Internet Control Message Protocol (ICMP)

D. Transmission Control Protocol (TCP)

ANSWER: A

Explanation:

Simple Network Management Protocol (SNMP) is the protocol designed to monitor and manage network-connected devices. SNMP-enabled devices, such as routers, switches, firewalls, servers, printers, and wireless equipment, expose management information through structured Management Information Bases (MIBs). A network-management system can query a device's SNMP agent to retrieve operational data including interface status and counters, bandwidth utilization, processor and memory statistics, device uptime, routing-related values, and error conditions. Devices can also send unsolicited SNMP notifications, called traps or informs, when defined events occur. This makes SNMP a foundational protocol for centralized network discovery, inventory, health monitoring, fault detection, and performance collection across an organization's infrastructure. In the TCP/IP suite, SNMP operates at the application layer and commonly uses UDP for transport, conventionally ports 161 for requests and 162 for notifications. The IETF architecture description explains that SNMP provides a framework for monitoring and managing resources in an internetwork, while the NIST glossary similarly identifies SNMP as a protocol for managing devices on IP networks. See RFC 3411: An Architecture for Describing SNMP Management Frameworks and NIST: Simple Network Management Protocol.

QUESTION NO: 14

A security consultant is preparing the final technical report for a penetration-testing engagement.

Which of the following information should be included for each confirmed vulnerability? (Select 3)

A. Affected assets and locations

B. Evidence supporting the finding

C. Remediation recommendations

D. All recovered passwords in clear text

E. Unverified scanner findings presented as confirmed exploits

ANSWER: A B C

Explanation:

Affected assets, **evidence supporting the finding**, and **remediation recommendations** are correct. A technical finding should identify the affected system or application, provide sufficient evidence to support the conclusion, explain the security impact, and give practical remediation guidance so the client can reproduce, prioritize, and correct the issue.

Credentials and raw sensitive data should not be unnecessarily reproduced in a report. The report should instead protect sensitive evidence and include only the minimum information necessary to communicate risk and support remediation. NIST describes reporting and analysis activities for technical assessments in NIST SP 800-115.

QUESTION NO: 15

An assessor is reviewing an Nmap TCP SYN scan report. Which of the following port states can indicate that packet filtering is affecting the scan results? (Select 2)

A. Filtered

B. Open|filtered

C. Open

D. Closed

ANSWER: A B

Explanation:

Filtered indicates that Nmap cannot determine whether the port is open because packet filtering prevents probes from reaching the port or prevents responses from returning. **Open|filtered** means that Nmap cannot distinguish between an open port and a filtered port for the scan technique used because no response was received.

An open state means that an application is accepting connections, while closed means that the host is reachable but no application is listening on that port. Nmap documents filtered and open|filtered states as conditions in which filtering or lack of a response prevents a definitive determination. See Nmap Port Scanning Basics.

QUESTION NO: 16

SQL injection attack consists of insertion or "injection" of either a partial or complete SQL query via the data input or transmitted from the client (browser) to the web application. A successful SQL injection attack can:

i) Read sensitive data from the database

iii) Modify database data (insert/update/delete)

iii) Execute administration operations on the database (such as shutdown the DBMS)

iV) Recover the content of a given file existing on the DBMS file system or write files into the file system

v) Issue commands to the operating system

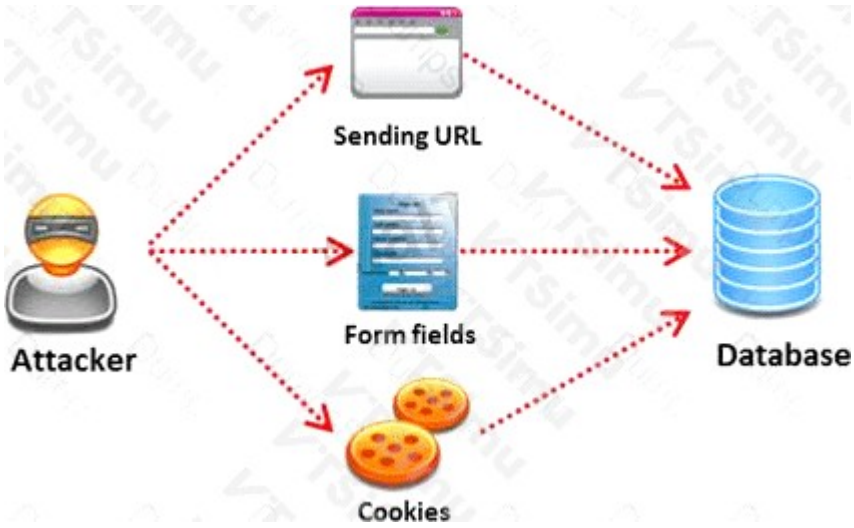

Pen tester needs to perform various tests to detect SQL injection vulnerability. He has to make a list of all input fields whose values could be used in crafting a SQL query, including the hidden fields of POST requests and then test them separately, trying to interfere with the query and to generate an error.

In which of the following tests is the source code of the application tested in a non-runtime environment to detect the SQL injection vulnerabilities?

A. Automated Testing

B. Function Testing

C. Dynamic Testing

D. Static Testing

ANSWER: D

Explanation:

Static Testing is correct because it examines an application's source code, bytecode, or other non-executing representations without running the application. For SQL injection detection, a static analysis process follows data originating at potentially untrusted input points—such as HTTP parameters, form fields, cookies, headers, and hidden POST values—and determines whether that data can reach SQL-query construction or database-execution functions without safe parameter binding or appropriate validation. This permits a tester or security tool to identify vulnerable data flows early, including paths that may be difficult to trigger during a live test. It is therefore explicitly a non-runtime approach, unlike testing that sends crafted values to an executing web application and observes its responses. Static testing is especially useful during secure code review and continuous integration because developers can remediate unsafe query construction before deployment. OWASP describes static application security testing as analysis performed without executing the application, and identifies static analysis as a useful way to find code-level security weaknesses. The OWASP SQL Injection Prevention Cheat Sheet also identifies parameterized queries as the primary defense, a control that source-code analysis can verify. See OWASP: Static Application Security Testing and OWASP: SQL Injection Prevention Cheat Sheet.

QUESTION NO: 17

A wireless intrusion detection system (WIDS) monitors the radio spectrum for the presence of unauthorized, rogue access points and the use of wireless attack tools.

The system monitors the radio spectrum used by wireless LANs, and immediately alerts a systems administrator whenever a rogue access point is detected. Conventionally it is achieved by comparing the MAC address of the participating wireless devices.

Which of the following attacks can be detected with the help of wireless intrusion detection system (WIDS)?

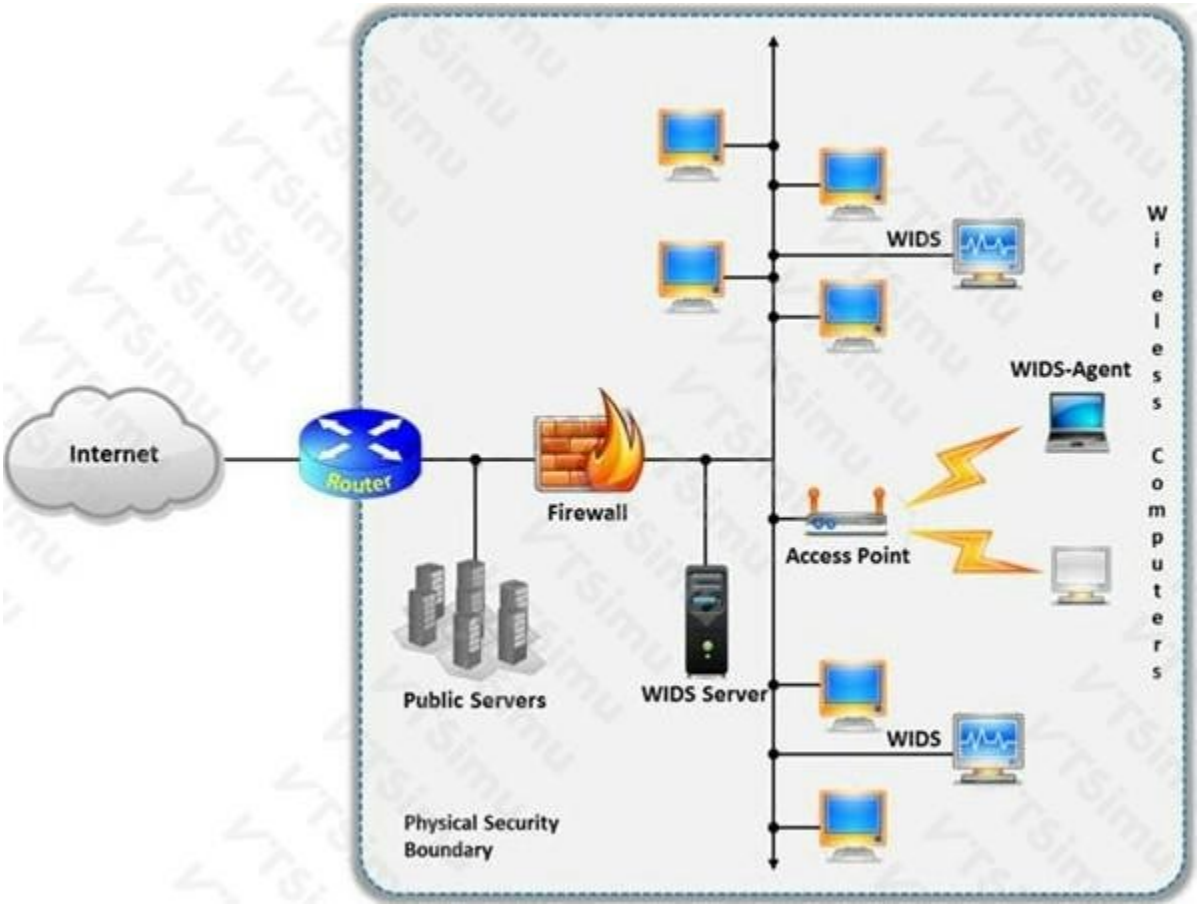

A. Social engineering

B. SQL injection

C. Parameter tampering

D. Man-in-the-middle attack

ANSWER: D

Explanation:

Man-in-the-middle attack is correct because a wireless intrusion detection system can observe management and data-related wireless activity for patterns associated with interception of client traffic. A common wireless man-in-the-middle scenario uses a rogue or evil-twin access point that impersonates a trusted network, persuading stations to associate through the attacker. The monitoring sensors can identify unauthorized access points, suspicious duplicate or spoofed device identities, unexpected associations, anomalous beacon characteristics, and other radio-frequency events that indicate an attacker is inserting itself between a client and its intended network.

Wireless monitoring is especially valuable because these attacks occur over the shared radio medium and may be visible even when the affected endpoint has not yet reported a problem. Alerts enable administrators to investigate the observed transmitter, compare its identity and configuration with authorized infrastructure, locate it physically, and contain the threat. NIST's guidance on securing wireless LANs describes the importance of wireless intrusion detection and prevention capabilities for detecting rogue devices and wireless attacks: NIST SP 800-153, Guidelines for Securing Wireless Local Area Networks. NIST also discusses wireless intrusion detection systems and their monitoring role in NIST SP 800-94, Guide to Intrusion Detection and Prevention Systems.

QUESTION NO: 18

Jessica works as systems administrator for a large electronics firm. She wants to scan her network quickly toÂ detect live hosts by using ICMP ECHO Requests. What type of scan is Jessica going to perform?

A. Smurf scan

B. Tracert

C. Ping trace

D. ICMP ping sweep

ANSWER: D

Explanation:

ICMP ping sweep is correct because it discovers active hosts by sending ICMP Echo Request packets to a range of target IP addresses and identifying systems that return ICMP Echo Reply packets. This is a host-discovery technique, not a port scan: its immediate purpose is to establish which addresses are currently reachable and likely have a live network stack. An administrator can use the resulting list of responsive hosts to define the scope for later enumeration, vulnerability assessment, or more detailed service scanning. The technique is commonly called a "ping sweep" because the Echo Request is swept across a subnet or other address range rather than sent to only one host. ICMP Echo is defined by the Internet Control Message Protocol specification, which describes the request and reply messages used to test reachability. Nmap also documents ICMP echo requests as a host-discovery probe and notes that host discovery helps determine which systems are online before further scanning. See RFC 792: Internet Control Message Protocol and Nmap Network Scanning: Host Discovery.

QUESTION NO: 19

In the context of penetration testing, what does blue teaming mean?

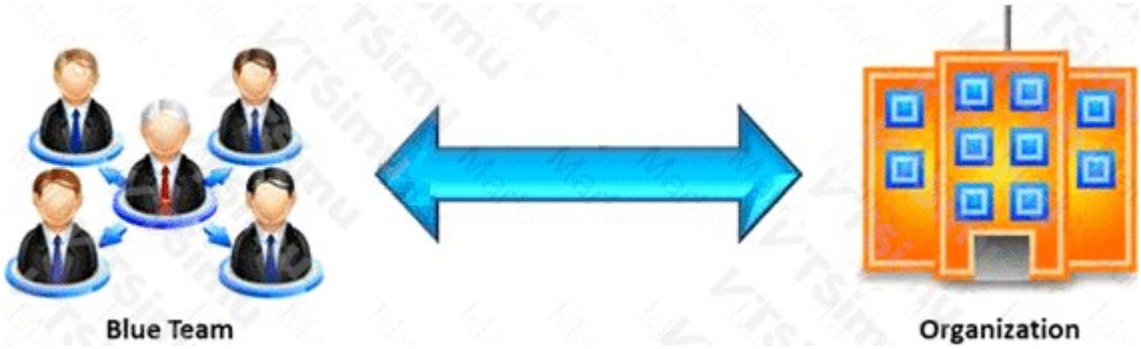

A. A penetration test performed with the knowledge and consent of the organization ' s IT staff

B. It is the most expensive and most widely used

C. It may be conducted with or without warning

D. A penetration test performed without the knowledge of the organization ' s IT staff but with permission from upper management

ANSWER: A

Explanation:

Blue teaming refers to the defensive side of an authorized security exercise. In this context, the organization's IT and security personnel are aware of the test and participate with their knowledge and consent, allowing them to monitor activity, apply defensive processes, investigate alerts, and assess how effectively the environment detects and responds to the simulated attack. The exercise is still governed by an agreed scope, rules of engagement, and authorization, but defender awareness is a defining operational characteristic of this blue-team scenario. It is useful when the goal extends beyond identifying technical vulnerabilities: the organization can evaluate logging, alert triage, incident-handling procedures, communications, and defensive capabilities while the test is underway. This aligns with the general distinction between adversarial testing functions, where defensive personnel protect systems and respond to authorized simulated attacks. NIST's technical guide emphasizes that testing must be planned, authorized, and coordinated to obtain meaningful security results safely; see NIST SP 800-115. EC-Council also describes blue-team work as the defensive function that detects, responds to, and strengthens organizational security during adversarial assessments; see EC-Council: Red Team vs. Blue Team.

QUESTION NO: 20

An external intrusion test and analysis identify security weaknesses and strengths of the client's systems and networks as they appear from outside the client's security perimeter, usually from the Internet.

The goal of an external intrusion test and analysis is to demonstrate the existence of known vulnerabilities that could be exploited by an external attacker.

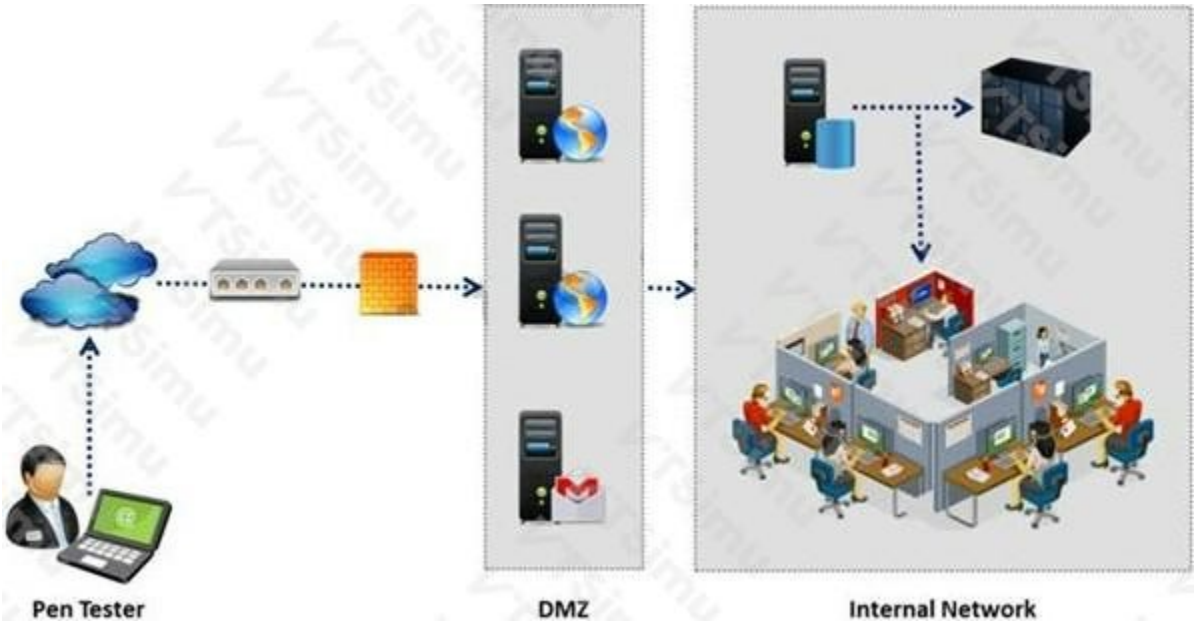

During external penetration testing, which of the following scanning techniques allow you to determine a port's state without making a full connection to the host?

A. XMAS Scan

B. SYN scan

C. FIN Scan

D. NULL Scan

ANSWER: B

Explanation:

SYN scan is the correct technique because it determines a target port's likely state by initiating, but deliberately not completing, the TCP three-way handshake. The scanner sends a SYN packet to the target port. A SYN/ACK response indicates that the port is accepting connection attempts and is therefore open; a reset response indicates that the port is closed. Once a SYN/ACK is received, the scanner sends a reset rather than the final ACK that would establish a normal application connection. This is why the method is commonly called a half-open scan.

For an external penetration test, this approach provides efficient service discovery while avoiding the full TCP session establishment used by a conventional connect scan. It is widely used by port-scanning tools when the tester has the required packet-capture and packet-generation privileges, because direct inspection of returned TCP flags gives useful evidence of a reachable and listening service. Nmap documents this behavior as "TCP SYN scan," noting that it is relatively unobtrusive because it never completes TCP connections. See the Nmap TCP SYN scan documentation and the TCP specification in RFC 793.

QUESTION NO: 21

When you are running a vulnerability scan on a network and the IDS cuts off your connection, what type of IDS is being used?

A. Passive IDS

B. Active IDS

C. Progressive IDS

D. NIPS

ANSWER: B

Explanation:

Active IDS is correct because the defining characteristic in this scenario is an automated response that interrupts the suspected hostile activity. A vulnerability scan can generate patterns associated with reconnaissance or attack attempts, such as repeated probes across ports and services. When the detection system identifies that activity and cuts off the scanner's connection, it is doing more than observing and alerting: it is actively taking a protective action. Depending on the product and deployment, that response may involve terminating a TCP session, injecting reset packets, dynamically blocking the source address, or updating a firewall rule to prevent continued communication.

In traditional security terminology, an active IDS combines detection with response capabilities, whereas a passive IDS primarily records events and sends alerts for a person or another system to act on. The key evidence provided by the question is the connection termination itself, which demonstrates active intervention. This aligns with NIST's description of intrusion detection and prevention technologies as systems that can monitor activity and, when prevention capabilities are enabled, react to detected incidents. See NIST SP 800-94: Guide to Intrusion Detection and Prevention Systems and the CISA overview of intrusion detection systems.

QUESTION NO: 22

Harold wants to set up a firewall on his network but is not sure which one would be the most appropriate. He knows he needs to allow FTP traffic to one of the servers on his network, but he wants to only allow FTP-PUT.

Which firewall would be most appropriate for Harold?

A. Application-level proxy firewall

B. Data link layer firewall

C. Packet filtering firewall

D. Circuit-level proxy firewall

ANSWER: A

Explanation:

An application-level proxy firewall is the appropriate choice because the requirement is to distinguish a particular FTP operation—FTP `PUT`, which uploads a file—from other FTP commands. This requires inspection and control at the application layer, where the firewall understands the FTP protocol and can examine commands carried on the FTP control connection. An application proxy terminates the client connection and establishes a separate connection to the server, enabling it to enforce policy based on protocol-aware content rather than merely IP addresses, ports, or connection state. FTP defines `STOR` as the command used to store (upload) a file; a proxy that recognizes FTP commands can permit that upload behavior according to policy. This level of inspection is especially useful for FTP because its control commands and data connections have protocol-specific behavior that must be interpreted correctly. The FTP command definitions are documented in RFC 959, and the role of application-layer gateways/proxies is described in NIST's Guidelines on Firewalls and Firewall Policy.

QUESTION NO: 23

The IP protocol was designed for use on a wide variety of transmission links. Although the maximum length of an IP datagram is 64K, most transmission links enforce a smaller maximum packet length limit, called a MTU.

The value of the MTU depends on the type of the transmission link. The design of IP accommodates MTU differences by allowing routers to fragment IP datagrams as necessary. The receiving station is responsible for reassembling the fragments back into the original full size IP datagram.

IP fragmentation involves breaking a datagram into a number of pieces that can be reassembled later. The IP source, destination, identification, total length, and fragment offset fields in the IP header, are used for IP fragmentation and reassembly.

The fragment offset is 13 bits and indicates where a fragment belongs in the original IP datagram. This value is a:

A. Multiple of four bytes

B. Multiple of two bytes

C. Multiple of eight bytes

D. Multiple of six bytes

ANSWER: C

Explanation:

Multiple of eight bytes is correct because the IPv4 Fragment Offset field does not store a byte-for-byte position. Instead, its 13-bit value specifies the fragment's starting position within the original datagram in units of eight octets (bytes). A receiver obtains the actual byte position by multiplying the field value by eight, then uses that position, together with the Identification value and source/destination/protocol information, to place the payload correctly during reassembly.

This eight-byte unit is an important format constraint in IPv4 fragmentation. Except for the final fragment, a fragmented datagram's payload must have a length that is divisible by eight bytes so the following fragment can begin at an offset representable by the Fragment Offset field. The last fragment may contain a remaining payload length that is not divisible by eight because no subsequent fragment needs to start after it. The field's 13-bit width therefore represents offsets from zero through 65,528 bytes, supporting placement within the maximum IPv4 datagram size. RFC 791 defines the Fragment Offset as measured in units of eight octets: RFC 791, Internet Protocol. See also the IPv4 header definition in RFC 8200 for related packet-format context.

QUESTION NO: 24

An organization is deploying SNMPv3 to monitor routers and switches.

Which of the following security features are provided by SNMPv3? (Select 2)

A. Authentication of SNMP messages

B. Encryption of management traffic

C. Automatic network segmentation

D. Stateful packet filtering

E. Anonymous community-string access

ANSWER: A B

Explanation:

Authentication and **encryption of management traffic** are correct. SNMPv3 supports the User-based Security Model (USM), which can authenticate the origin of SNMP messages and provide privacy by encrypting message contents. These capabilities address the clear-text community-string weaknesses associated with SNMPv1 and SNMPv2c.

SNMPv3 does not use community strings as its primary security mechanism, and it does not inherently provide firewall filtering or network segmentation. Those are separate security controls. The SNMPv3 architecture and security models are described in RFC 3411 and RFC 3414.

QUESTION NO: 25

You setup SNMP in multiple offices of your company. Your SNMP software manager is not receiving data from other offices like it is for your main office. You suspect that firewall changes are to blame.

What ports should you open for SNMP to work through Firewalls. (Select 2)

A. 162

B. 160

C. 161

D. 163

ANSWER: A C

Explanation:

SNMP normally uses UDP port 161 for communication between an SNMP manager and SNMP agents. The manager sends polling requests—such as requests for interface status, CPU utilization, or system information—to an agent on UDP 161, and the agent returns its response as part of that SNMP exchange. Therefore, firewall policy must permit the required UDP 161 traffic between the central manager and the monitored devices in remote offices.

UDP port 162 is used by the SNMP manager or trap receiver to accept unsolicited notifications from agents. These notifications include SNMP traps and, depending on configuration, informs. Allowing UDP 162 to the manager enables remote devices to report notable events without waiting for the manager's next scheduled poll. In a multi-office design, rules should be narrowly scoped to the known manager and agent IP addresses, with the necessary directionality based on whether polling, traps, or both are used. The IANA service-name registry lists *snmp* on port 161 and *snmptrap* on port 162: IANA Service Name and Port Number Registry. Net-SNMP also documents the standard agent port and trap port: Net-SNMP documentation.