PassQueen
a.k.a dumpsqueen.com

Certified Ethical Hacker Exam (CEHv12)

ECCouncil 312-50v12

Version Demo

Total Demo Questions: 68

Total Premium Questions: 685

Buy Premium PDF

https://passqueen.com

support@passqueen.com

Topic Break Down

Topic	No. of Questions
Topic 1, Information Security and Ethical Hacking Overview	96
Topic 2, Reconnaissance Techniques	125
Topic 3, System Hacking	113
Topic 4, Network and Perimeter Hacking	125
Topic 5, Web Application Hacking	84
Topic 6, Wireless Network Hacking	41
Topic 7, Mobile Platform, IoT, and OT Hacking	24
Topic 8, Cloud Computing	21
Topic 9, Cryptography	56
Total	685

QUESTION NO: 1

John wants to send Marie an email that includes sensitive information, and he does not trust the network that he is connected to. Marie gives him the idea of using PGP. What should John do to communicate correctly using this type of encryption?

A. Use his own public key to encrypt the message.

B. Use Marie's public key to encrypt the message.

C. Use his own private key to encrypt the message.

D. Use Marie's private key to encrypt the message.

ANSWER: B

Explanation:

Use Marie's public key to encrypt the message. PGP uses public-key cryptography to establish confidentiality between communicating parties. Marie distributes her public key so that others, including John, can encrypt data intended only for her. After John encrypts the email with Marie's public key, the resulting encrypted session key and message can be decrypted only with the corresponding private key, which Marie must keep under her exclusive control. This protects the sensitive contents from interception or disclosure while the message traverses an untrusted network. In typical OpenPGP operation, the email content is encrypted with a newly generated symmetric session key for efficiency; that session key is then encrypted to Marie's public key and included with the ciphertext. The practical security requirement is that John obtain and validate the authentic public key belonging to Marie, such as by checking its fingerprint through a trusted channel or verifying an appropriate certification. This prevents an attacker from substituting a different public key and receiving the confidential message. PGP can also add a digital signature using John's private key to provide authenticity and integrity, but encryption for confidentiality is directed to the recipient's public key. See the OpenPGP standard (RFC 9580) and GnuPG documentation on encryption and decryption.

QUESTION NO: 2

Which of the following actions should be performed to preserve the integrity of digital evidence collected from a compromised workstation? (Choose three.)

A. Calculate a cryptographic hash of the collected evidence

B. Maintain a documented chain of custody

C. Create the forensic image using a write blocker

D. Restart the workstation before acquiring data

E. Perform analysis directly on the original drive

ANSWER: A B C

Explanation:

Calculating cryptographic hashes of collected evidence, maintaining a documented chain of custody, and creating a forensic image using a write blocker are appropriate evidence-preservation practices. A cryptographic hash, such as SHA-256, provides a value that can be recalculated later to demonstrate whether the acquired evidence has changed. Chain-of-custody documentation records who collected, handled, transferred, and stored the evidence, along with the relevant dates and times.

A hardware or software write blocker helps prevent accidental modification of the source storage media during acquisition. Investigators should perform analysis on verified copies rather than on the original evidence whenever possible. Restarting a live system can alter volatile memory, log files, temporary data, and timestamps, so it is not an evidence-preservation step. NIST provides forensic collection guidance in SP 800-86, Guide to Integrating Forensic Techniques into Incident Response.

QUESTION NO: 3

What is the common name for a vulnerability disclosure program opened by companies In platforms such as HackerOne?

A. Vulnerability hunting program

B. Bug bounty program

C. White-hat hacking program

D. Ethical hacking program

ANSWER: B

Explanation:

Bug bounty program is the common name for a company-run initiative that invites independent security researchers to identify and responsibly report security vulnerabilities, often through a platform such as HackerOne. These programs define the systems that researchers are authorized to test, set rules of engagement and safe-harbor terms, establish report-handling procedures, and typically offer rewards based on the validity and severity of reported findings. The primary purpose is to supplement an organization's internal security testing with the diverse skills and perspectives of an external researcher community. HackerOne describes bug bounty programs as a way for organizations to incentivize hackers to find and report vulnerabilities, with bounties paid for eligible discoveries. A bug bounty program is therefore a structured form of vulnerability disclosure that commonly includes financial recognition, although the precise scope, reward model, and disclosure requirements are determined by each organization. For more information, see HackerOne's overview of bug bounty programs and its guidance on creating a vulnerability disclosure program.

QUESTION NO: 4

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

A. NSLookup

B. Finger

C. Dig

D. Sam Spade

E. Host

F. Netcat

G. Neotrace

ANSWER: A C D E

Explanation:

A DNS zone transfer retrieves DNS records from an authoritative name server using the AXFR operation. **NSLookup** can request a full zone listing in implementations that support its `ls` command, such as the classic interactive nslookup workflow. **Dig** directly supports AXFR queries; for example, an authorized test can use `dig @nameserver domain AXFR`. **Host** can request a zone transfer through its zone-listing functionality, commonly expressed with the `-l` option. These command-line utilities are standard tools for DNS enumeration during an authorized security assessment and can reveal hosts, mail exchangers, name servers, and other DNS data when a server is incorrectly configured to allow transfers to unauthorized clients.

Sam Spade is also a DNS reconnaissance utility that historically provided zone-transfer capabilities alongside DNS lookup and other network-enumeration functions. Whether any of these tools succeeds depends on the target DNS server permitting AXFR for the requesting client; modern DNS administration should restrict transfers to explicitly authorized secondary servers. The DNS AXFR operation is defined in RFC 5936. Command syntax and AXFR support for dig are documented in the dig manual page.

QUESTION NO: 5

Ethical backer jane Doe is attempting to crack the password of the head of the it department of ABC company. She Is utilizing a rainbow table and notices upon entering a password that extra characters are added to the password after submitting. What countermeasure is the company using to protect against rainbow tables?

A. Password key hashing

B. Password salting

C. Password hashing

D. Account lockout

ANSWER: B

Explanation:

Password salting is the countermeasure being used. A salt is a unique, randomly generated value added to a password before the password is processed by a cryptographic password-hashing function. The resulting stored value is therefore based on both the user's password and the additional salt characters. This matches the observed behavior in which extra characters are appended or otherwise incorporated when the password is submitted.

Rainbow tables rely on precomputed mappings between likely plaintext passwords and their corresponding hash values. Salting defeats the practical value of these precomputed tables because the same plaintext password produces a different stored hash for each distinct salt. An attacker would need to generate or compute results for the specific salt associated with each target account, rather than reuse a single precomputed table across many accounts. Salts should be unique per password and generated using a cryptographically secure random source. Modern password-storage implementations also combine salting with deliberately slow, adaptive password-hashing algorithms such as Argon2, bcrypt, PBKDF2, or scrypt. See the NIST Digital Identity Guidelines and the OWASP Password Storage Cheat Sheet.

QUESTION NO: 6

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

A. Red hat

B. white hat

C. Black hat

D. Gray hat

ANSWER: D

Explanation:

Gray hat is correct because Nicolas discovered a vulnerability in a public-facing system without any stated authorization to test that system, then disclosed the issue to the affected owner and Microsoft rather than exploiting it for harm or personal gain. In CEH terminology, gray-hat activity commonly falls between authorized ethical testing and malicious compromise: the researcher may identify a weakness without prior permission but acts with an intention to alert affected parties and enable remediation. The key facts are the absence of stated authorization and the responsible notification of the system owner and relevant vendor. A zero-day is a previously unknown vulnerability for which the vendor has had no time to prepare or distribute a fix. Reporting such an issue privately with protective details supports coordinated vulnerability disclosure, allowing the parties responsible for the affected technology to investigate, develop mitigations, and protect users before broad public disclosure. CISA describes coordinated vulnerability disclosure as a process through which researchers report vulnerabilities to an organization so they can be addressed appropriately. Microsoft also provides channels and guidance for reporting security vulnerabilities affecting its products and services. CISA Vulnerability Disclosure Policy; Microsoft Security Response Center.

QUESTION NO: 7

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

A. Use port security on his switches.

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

C. Use a firewall between all LAN segments.

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

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

ANSWER: A B D

Explanation:

Using port security on his switches, using a tool like ARPwatch to monitor for strange ARP activity, and using static ARP entries on a small network are appropriate protective measures. Switch port security restricts which MAC addresses are permitted on an access port, helping administrators control endpoint access and detect or shut down ports when an unauthorized MAC address appears. This reduces opportunities for an unapproved device to connect and perform local-layer attacks. ARPwatch provides visibility by recording ARP mappings and alerting administrators when a host's MAC-to-IP relationship unexpectedly changes; prompt detection enables containment and remediation before a poisoning event causes prolonged interception or disruption. For a small, stable network, static ARP entries prevent hosts from accepting forged dynamic ARP replies for the configured peers because the intended IP-to-MAC mappings are fixed. In a production switched environment, these measures are commonly supplemented with DHCP snooping and Dynamic ARP Inspection, which validates ARP traffic against trusted address-binding information. Cisco describes this validation approach in its Dynamic ARP Inspection documentation; ARPwatch behavior and alerting are documented in the ARPwatch manual page.

QUESTION NO: 8

You are an ethical hacker tasked with conducting an enumeration of a company's network. Given a Windows system with NetBIOS enabled, port 139 open, and file and printer sharing active, you are about to run some nbtstat commands to enumerate NetBIOS names. The company uses |Pv6 for its network. Which of the

following actions should you take next?

A. Use nbtstat -c to get the contents of the NetBIOS name cache

B. use nbtstat -a followed by the IPv6 address of the target machine

C. Utilize Nmap Scripting Engine (NSE) for NetBIOS enumeration

D. Switch to an enumeration tool that supports IPv6

ANSWER: D

Explanation:

Switch to an enumeration tool that supports IPv6 is the appropriate next action because an assessment method must be compatible with the target network's addressing protocol. NetBIOS over TCP/IP (NetBT), including the traditional services associated with port 139, is an IPv4-based mechanism. Microsoft's NetBIOS documentation identifies the NetBIOS-over-TCP transport in terms of IPv4 addressing, so IPv6 addressing cannot simply be substituted into a NetBIOS-oriented command and expected to provide equivalent results. In an IPv6 environment, the tester should first identify IPv6-reachable hosts and exposed services with tooling that explicitly supports IPv6, then select service-specific enumeration techniques appropriate to the protocols actually available on those hosts. This keeps enumeration accurate and avoids drawing conclusions from an IPv4-only discovery mechanism in an IPv6 deployment. Nmap supports IPv6 target specification and can be used for IPv6 host and service discovery, providing a practical starting point for adapting the enumeration workflow. See Microsoft's NetBIOS over TCP/IP documentation and the Nmap IPv6 guide.

QUESTION NO: 9

You are the lead cybersecurity analyst at a multinational corporation that uses a hybrid encryption system to secure inter-departmental communications. The system uses RSA encryption for key exchange and AES for data encryption, taking advantage of the strengths of both asymmetric and symmetric encryption. Each RSA key pair has a size of 'n' bits, with larger keys providing more security at the cost of slower performance. The

time complexity of generating an RSA key pair is O(n*2), and AES encryption has a time complexity of O(n). An attacker has developed a quantum algorithm with time complexity O((log n)*2) to crack RSA encryption. Given *n=4000' and variable 'AES key size', which scenario is likely to provide the best balance of security and

performance? which scenario would provide the best balance of security and performance?

A. Data encryption with 3DES using a 168-bit key: Offers high security but slower performance due to 3DES's inherent inefficiencies.

B. Data encryption with Blowfish using a 448-bit key: Offers high security but potential compatibility issues due to Blowfish's less widespread use.

C. Data encryption with AES-128: Provides moderate security and fast encryption, offering a balance between the two.

D. Data encryption with AES-256: Provides high security with better performance than 3DES, but not as fast as other AES key sizes.

ANSWER: C

Explanation:

Data encryption with AES-128: Provides moderate security and fast encryption, offering a balance between the two. is the best choice within the stated scenario. The RSA-4000 component is fixed and is used only to establish or protect the short-lived symmetric session key; the selectable parameter is the algorithm and key size used for bulk data encryption. AES is specifically designed for efficient processing of substantial volumes of data, and AES-128 uses fewer transformation rounds than AES-256, resulting in lower computational overhead while retaining strong conventional security for general enterprise data encryption. NIST specifies AES with 128-, 192-, and 256-bit keys, and all are standardized AES variants; selecting the 128-bit key size prioritizes throughput and latency where the scenario asks for a practical security-performance tradeoff rather than the highest possible symmetric security margin. See NIST FIPS 197, Advanced Encryption Standard. The specified quantum attack is directed at RSA key exchange, not at AES, so increasing the AES key size would not remedy the described RSA exposure. In this constrained comparison, AES-128 provides efficient linear-time bulk encryption and avoids the additional rounds and processing cost associated with larger AES keys. NIST's key-management guidance also recognizes 128-bit symmetric security strength as an approved security-strength category for appropriate use cases; see NIST SP 800-57 Part 1 Rev. 5.

QUESTION NO: 10

Which of the following are security benefits provided by a correctly implemented digital signature? (Choose three.)

A. Integrity

B. Authentication of the signer

C. Nonrepudiation support

D. Confidentiality of the signed message

E. Compression of the signed message

ANSWER: A B C

Explanation:

Integrity, **authentication of the signer**, and **nonrepudiation support** are the principal security benefits associated with a digital signature. The signer creates the signature using a private key, and a recipient verifies it using the corresponding

public key. If signed data is modified after signing, signature verification fails, providing evidence that integrity has not been preserved. A valid signature also demonstrates that the holder of the associated private key created the signature, subject to the trustworthiness of the certificate and key-management process.

Digital signatures do not provide confidentiality because the signed content remains readable unless it is separately encrypted. NIST describes the security services of digital signatures in FIPS 186-5, Digital Signature Standard.

QUESTION NO: 11

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 controls that reduce the attack surface and improve the resilience of DNS infrastructure. Hardening includes securely configuring the DNS software and host operating system, promptly applying patches, minimizing enabled services, using least-privilege administrative access, and enabling appropriate logging and monitoring. Split-horizon DNS separates internal and external name-resolution views, allowing internal records and addresses to remain unavailable to unauthenticated Internet clients while still serving necessary public records externally. Restricting zone transfers is essential because zone data can reveal hostnames, addressing schemes, and other useful reconnaissance information; transfers should be authorized only to designated secondary servers, commonly through explicit access controls and transaction security. Subnet diversity prevents authoritative DNS servers from sharing the same network dependency. Placing servers on distinct networks helps preserve name resolution if a routing failure, denial-of-service event, or compromise affects one subnet. Together, these practices support confidentiality of internal DNS information, integrity of DNS administration, and availability of name resolution. NIST provides detailed DNS security guidance in Secure Domain Name System Deployment Guide; DNS server placement and redundancy considerations are also documented in RFC 2182.

QUESTION NO: 12

Based on the below log, which of the following sentences are true?

Mar 1, 2016, 7:33:28 AM 10.240.250.23 - 54373 10.249.253.15 - 22 tcp_ip

A. Application is FTP and 10.240.250.23 is the client and 10.249.253.15 is the server.

B. Application is SSH and 10.240.250.23 is the server and 10.249.253.15 is the client.

C. SSH communications are encrypted; it's impossible to know who is the client or the server.

D. Application is SSH and 10.240.250.23 is the client and 10.249.253.15 is the server.

ANSWER: D

Explanation:

Application is SSH and 10.240.250.23 is the client and 10.249.253.15 is the server. The log records a TCP/IP connection from source address `10.240.250.23` using source port `54373` to destination address `10.249.253.15` using

destination port `22`. Port 54373 is a high-numbered ephemeral port, which is normally selected by the initiating host for an outbound client connection. Port 22 is the well-known assigned port for the Secure Shell (SSH) service, so the host receiving traffic at `10.249.253.15:22` is the SSH server and the host originating from `10.240.250.23:54373` is the client. Encryption protects SSH session contents, such as authentication exchanges and commands, but it does not conceal the IP addresses, TCP ports, and connection direction present in this network log. IANA maintains the authoritative service-name and port-number registry, including the assignment of TCP port 22 to SSH, at IANA Service Name and Port Number Registry. For SSH protocol context, see RFC 4251: The Secure Shell Protocol Architecture.

QUESTION NO: 13

which of the following protocols can be used to secure an LDAP service against anonymous queries?

A. SSO

B. RADIUS

C. WPA

D. NTLM

ANSWER: D

Explanation:

NTLM is correct because it is an authentication protocol that can be used by LDAP clients when binding to Microsoft Active Directory. Rather than allowing an unauthenticated, anonymous LDAP bind, a client can authenticate with domain credentials using an NTLM-based SASL bind. The directory service can then identify the requesting security principal and apply the permissions associated with that account, including restrictions on which directory objects and attributes may be queried. This changes LDAP access from anonymous access to authenticated, authorization-controlled access.

In an Active Directory environment, LDAP security is strengthened by requiring authenticated binds and, where appropriate, LDAP signing and channel binding. Microsoft documents that LDAP signing helps protect LDAP traffic from certain man-in-the-middle attacks, while channel binding strengthens authentication for LDAP connections protected with TLS. NTLM is one of the Windows authentication mechanisms relevant to LDAP authentication, although modern deployments should prefer Kerberos where it is available and enforce LDAP signing or LDAPS/TLS according to organizational requirements. See Microsoft's LDAP signing and channel binding guidance and Microsoft's NTLM documentation.

QUESTION NO: 14

An attacker has installed a RAT on a host. The attacker wants to ensure that when a user attempts to go to "www.MyPersonalBank.com", the user is directed to a phishing site.

Which file does the attacker need to modify?

A. Boot.ini

B. Sudoers

C. Networks

D. Hosts

ANSWER: D

Explanation:

Hosts is correct because the local hosts file can statically map a hostname, such as `www.MyPersonalBank.com`, to an IP address selected by the attacker. Operating systems commonly consult this local name-resolution file before making a DNS query, so a matching entry can cause the browser to connect to the specified phishing-server address even though the user typed the legitimate domain name. On Windows, this file is normally located at

`%SystemRoot%\System32\drivers\etc\hosts`; on Linux and macOS, it is generally `/etc/hosts`. This technique is commonly called hosts-file poisoning or hosts-file redirection and is a relevant persistence and traffic-redirection technique after a host has been compromised. The mapping affects name resolution on that individual endpoint, making it appropriate for redirecting a particular user on the infected host. Microsoft documents the hosts file as a mechanism for manual hostname-to-IP-address mappings, and the Linux manual page similarly describes it as a static table of hostnames and addresses. See Microsoft's hosts file documentation and the Linux hosts(5) manual page.

QUESTION NO: 15

Which of the following practices help reduce the risk of credential theft from a web application? (Choose three.)

A. Enforce multifactor authentication

B. Use HTTPS for all authentication pages

C. Store passwords with a salted adaptive hash

D. Use one shared administrator account

E. Log user passwords for troubleshooting

ANSWER: A B C

Explanation:

Enforcing multifactor authentication, **using HTTPS for all authentication pages**, and **storing passwords with a salted adaptive hash** are appropriate controls. Multifactor authentication reduces the value of a stolen password by requiring an additional authentication factor. HTTPS protects credentials from interception and modification while they travel between the user and the application. Passwords must be stored using a salted, computationally expensive password-hashing algorithm so that a database compromise does not expose recoverable plaintext credentials.

Reusing a shared administrator account removes accountability and creates a single highly valuable credential. Logging passwords for troubleshooting is also insecure because logs can be broadly accessible, retained for long periods, or copied into monitoring systems. OWASP provides guidance on authentication controls in its Authentication Cheat Sheet and password storage in its Password Storage Cheat Sheet.

QUESTION NO: 16

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 of an SOA record is its **serial number**, which is the first numeric field inside the parenthesized timing-and-control section. In the supplied record, that value is **200302028**. DNS secondary servers use the SOA serial to determine whether the zone data they hold is current. During a zone refresh check, a secondary server compares its stored serial

number with the serial published by the primary authoritative server. If the primary server's serial is greater, the secondary knows that the zone has changed and initiates a zone transfer to obtain the updated data. The serial must therefore be incremented whenever the zone is modified. Although administrators often encode a date and revision number in the serial, such as a YYYYMMDDnn-style convention, DNS interprets it as a 32-bit sequence value subject to serial-number arithmetic. RFC 1035 defines the SOA RDATA field order and identifies SERIAL as the first numeric field after the responsible-party mailbox. RFC 1912 also describes the operational need to increment the SOA serial after every zone-file change. See RFC 1035, section 3.3.13 and RFC 1912, section 2.2.

QUESTION NO: 17

Which of the following characteristics are associated with a properly configured Web Application Firewall (WAF)? (Choose three.)

A. Inspecting HTTP and HTTPS application traffic

B. Enforcing rules designed to identify malicious web requests

C. Providing virtual patching for certain known application flaws

D. Replacing secure coding and application patching

E. Providing antivirus protection for endpoint files

ANSWER: A B C

Explanation:

Inspecting HTTP and HTTPS application traffic, **enforcing rules designed to identify malicious web requests**, and **providing virtual patching for certain known application flaws** are characteristics associated with a WAF. A WAF operates at the web-application layer and can inspect requests for attack patterns, protocol anomalies, and policy violations. When a vulnerable application cannot be immediately updated, carefully deployed WAF rules can provide temporary compensating protection by blocking exploit attempts.

A WAF does not replace secure software development, code review, patching, or application testing. It also does not provide endpoint antivirus functionality or eliminate the need for network firewalls and access controls. NIST describes web application firewalls and their use as application-layer controls in SP 800-44 Version 2, Guidelines on Securing Public Web Servers. OWASP provides additional defensive guidance in the Web Application Firewall overview.

QUESTION NO: 18

An organization is preparing an incident-response plan for ransomware. Which of the following actions are appropriate during the containment and recovery phases? (Choose three.)

A. Isolate affected systems from the network

B. Preserve relevant evidence and logs

C. Restore systems from verified clean backups

D. Delete logs to free disk space

E. Reconnect infected systems before remediation

ANSWER: A B C

Explanation:

Isolating affected systems from the network, **preserving relevant evidence and logs**, and **restoring systems from verified clean backups** are appropriate actions. Isolation limits further propagation and can disrupt attacker command-and-control communications. Preserving volatile and nonvolatile evidence, including logs, helps investigators establish scope,

timeline, and root cause. Recovery should use backups that have been verified as clean and should occur only after the organization addresses the compromise path that enabled the incident.

Immediately deleting logs destroys valuable evidence and complicates incident analysis. Reconnecting infected systems before they have been remediated can permit reinfection or continued lateral movement. NIST describes containment, eradication, recovery, and post-incident activity in SP 800-61 Rev. 2, Computer Security Incident Handling Guide. CISA also provides ransomware response guidance at StopRansomware Guide.

QUESTION NO: 19

Bob received this text message on his mobile phone: “Hello, this is Scott Smelby from the Yahoo Bank. Kindly contact me for a vital transaction on: scottsmelby@yahoo.com”. Which statement below is true?

A. This is a scam as everybody can get a @yahoo address, not the Yahoo customer service employees.

B. This is a scam because Bob does not know Scott.

C. Bob should write to scottmelby@yahoo.com to verify the identity of Scott.

D. This is probably a legitimate message as it comes from a respectable organization.

ANSWER: A

Explanation:

This is a scam as everybody can get a @yahoo address, not the Yahoo customer service employees. The key warning sign is that the sender claims to represent an organization but directs Bob to a generic, consumer email address. An address ending in *@yahoo.com* can be registered by members of the public and does not establish that its user is an authorized representative of a purported bank or other business. Attackers commonly use this type of impersonation to create trust, initiate a conversation, and then seek sensitive information, payment, or further interaction outside legitimate support channels.

The appropriate security practice is to treat the unsolicited message as suspicious and independently locate contact information through a trusted source, such as the organization’s official website or an existing account statement. Bob should not use the contact address supplied in the message to validate the sender, because that would communicate directly with the unverified party. The FTC explains that phishing messages often impersonate trusted organizations and use unexpected communications to lure recipients into responding or disclosing information. CISA similarly advises users to verify suspicious requests through known, independently obtained contact methods. See the FTC phishing guidance and CISA phishing-recognition guidance.

QUESTION NO: 20

From the following table, identify the wrong answer in terms of Range (ft).

Standard Range (ft)

802.11a 150-150

802.11b 150-150

802.11g 150-150

802.16 (WiMax) 30 miles

A. 802.16 (WiMax)

B. 802.11g

C. 802.11b

D. 802.11a

ANSWER: A

Explanation:

802.16 (WiMax) is the incorrect table entry for a column explicitly labeled “Range (ft)” because its value is stated in miles rather than feet. Thirty miles is equivalent to 158,400 feet, using the standard conversion of 5,280 feet per mile. Therefore, even if “30 miles” is intended as an illustrative long-distance WiMAX deployment range, it is not presented in the unit requested by the table and is inconsistent with the other entries.

IEEE 802.16 is the standard family commonly associated with WiMAX broadband wireless access. Its deployments can be designed for substantially longer links than local-area Wi-Fi technologies, but actual coverage is not a fixed guarantee: it depends on spectrum, transmit power, antenna height, line of sight, terrain, modulation, and regulatory constraints. In exam-table terminology, the issue is specifically the unit mismatch, not a claim that WiMAX cannot support multi-mile coverage. A consistent table would express the value as 158,400 ft or change the range column to use an appropriate common unit. See the IEEE 802.16 standards information at IEEE 802.16 and the unit relationship defined by NIST.

QUESTION NO: 21

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:

Sending a customized SNMP set request with a spoofed source IP address in the permitted 192.168.1.0 range can bypass a source-address-based SNMP access control list during an authorized assessment. Cisco devices can use the CISCO-CONFIG-COPY-MIB to initiate configuration-copy operations through SNMP, including copying running or startup configuration to a TFTP server. The request must use a valid write-capable community string and set the copy parameters, such as source file type, destination file type, server address, and destination filename. Cisco documents this MIB and its configuration-copy workflow in the CISCO-CONFIG-COPY-MIB documentation.

Running a network sniffer and capturing the returned traffic with the configuration file from the router is also appropriate when the resulting copy uses TFTP. TFTP transfers data in cleartext over UDP, so an assessor positioned on a network segment capable of observing the transfer can capture the configuration contents as the router sends the file. This demonstrates why Cisco recommends securing management-plane access and replacing insecure SNMPv1/v2c and TFTP workflows with stronger controls. Cisco’s SNMP security guidance emphasizes restricting SNMP access and using SNMPv3 security features.

QUESTION NO: 22

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

A. Platform as a service

B. Software as a service

C. Functions as a

D. Infrastructure as a service

ANSWER: D

Explanation:

Infrastructure as a service is the appropriate cloud service model because it gives an organization control over the cloud resources it deploys and maintains, including its operating systems, applications, data, virtual networking configuration, and security settings. The hosting provider supplies and operates the underlying physical facilities, servers, storage, and virtualization layer, while the customer is responsible for administering the provisioned infrastructure and keeping workloads securely configured, patched, and monitored. This arrangement is common when an organization needs substantial administrative control but does not want to own or operate a physical data center. The National Institute of Standards and Technology defines Infrastructure as a Service as the capability to provision processing, storage, networks, and other fundamental computing resources on which the consumer can deploy and run software, including operating systems and applications. This directly aligns with an agreement that places maintenance responsibility for cloud-based resources on the organization. See the NIST definition of cloud computing and the AWS Shared Responsibility Model for further details on customer-managed responsibilities in infrastructure services.

QUESTION NO: 23

Annie, a cloud security engineer, uses the Docker architecture to employ a client/server model in the application she is working on. She utilizes a component that can process API requests and handle various Docker objects, such as containers, volumes. Images, and networks. What is the component of the Docker architecture used by Annie in the above scenario?

A. Docker client

B. Docker objects

C. Docker daemon

D. Docker registries

ANSWER: C

Explanation:

Docker daemon is the Docker architecture component that receives and processes Docker API requests, then manages Docker objects such as images, containers, networks, and volumes. The daemon runs as the `dockerd` service on the Docker host and performs the underlying work requested through the Docker client or another API consumer. In Docker's client-server architecture, a client sends commands or REST API calls to the daemon through a Unix socket, network interface, or another configured communication endpoint. The daemon is responsible for building images, creating and running containers, attaching containers to networks, and creating or mounting volumes. This directly matches the scenario's description of a server-side component that processes API requests and handles the lifecycle of Docker resources. Docker documentation describes the daemon as listening for Docker API requests and managing Docker objects, which is the defining responsibility identified in the question. See the Docker overview and Docker's dockerd reference for the daemon's role and API-based architecture.

QUESTION NO: 24

Upon establishing his new startup, Tom hired a cloud service provider (CSP) but was dissatisfied with their service and wanted to move to another CSP.

What part of the contract might prevent him from doing so?

A. Virtualization

B. Lock-in

C. Lock-down

D. Lock-up

ANSWER: B

Explanation:

Lock-in is correct because cloud-provider lock-in describes a situation in which contractual, technical, operational, or financial dependencies make it difficult for a customer to move workloads, applications, and data to a different cloud service provider. A CSP contract can create lock-in through restrictive data-export terms, proprietary service interfaces or formats, migration fees, minimum commitments, long notice periods, or terms that limit assistance with transition and portability. Although a customer may be unhappy with service quality, these provisions can make switching costly, slow, or impractical. For a startup, the impact can be especially significant: limited staff, budget constraints, and reliance on provider-specific managed services may turn migration into a major business and security project. Sound cloud procurement therefore evaluates exit clauses, data portability, migration support, deletion procedures, and the allocation of costs and responsibilities before signing. The U.S. National Institute of Standards and Technology identifies portability and interoperability as important cloud-computing concerns, including the ability to move data and applications across providers. The UK National Cyber Security Centre also recommends considering exit arrangements and avoiding dependencies that impede changing providers. See NIST Cloud Computing Standards Roadmap and NCSC guidance on using cloud services securely.

QUESTION NO: 25

Which of the following controls help protect an organization from VLAN hopping attacks? (Choose three.)

A. Disable Dynamic Trunking Protocol on access ports

B. Configure endpoint-facing ports as static access ports

C. Prune unnecessary VLANs from trunk links

D. Enable port mirroring on all access ports

E. Use the default VLAN as the native VLAN on every trunk

ANSWER: A B C

Explanation:

Disabling Dynamic Trunking Protocol, configuring access ports as nontrunking ports, and pruning unused VLANs from trunk links are effective VLAN-hopping defenses. DTP can negotiate a switch port into trunking mode when it is enabled or left in a dynamic mode; disabling DTP and explicitly configuring an endpoint-facing port as an access port prevents an unauthorized device from negotiating a trunk. Unused switch ports should also be placed into an unused VLAN and administratively shut down where practical.

Pruning prevents VLANs that do not need to traverse a trunk from being carried over that trunk. This reduces the opportunity for a system connected elsewhere on the switched infrastructure to reach a VLAN unnecessarily. Using an unused native VLAN is an additional common hardening practice for 802.1Q trunks. Cisco documents trunk configuration and DTP behavior in its Dynamic Trunking Protocol guidance.

QUESTION NO: 26

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

What Is the best Linux pipe to achieve your milestone?

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

B. curl -s https://site.com | grep '

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

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

ANSWER: B

Explanation:

curl -s https://site.com | grep 'is the best pipeline because it retrieves the main page's HTML directly to standard output, then filters the document for absolute HTTP(S) anchor references associated with the target site. The final cut operation uses the double-quote delimiter to extract the URL value from the HTML href attribute, producing a practical list of links for initial reconnaissance. The -s flag is also useful in a pipeline because it suppresses curl's progress meter and other non-content output, leaving the HTML body available for downstream text-processing commands. This is a lightweight first-pass approach for collecting obvious in-scope links before deeper crawling, parsing, normalization, or validation. In a real engagement, the tester should ensure authorization and should account for relative URLs, JavaScript-generated links, redirects, and malformed HTML, which may require a dedicated crawler or parser. Curl's output and silent-mode behavior are documented in the curl manual; the filtering semantics used here are documented in the GNU grep manual.

QUESTION NO: 27

Which of the following practices help secure an organization against removable-media threats? (Choose three.)

A. Disable AutoRun and AutoPlay where not required

B. Use endpoint protection to scan removable media

C. Restrict removable-media use through device-control policies

D. Allow unrestricted USB devices for all users

E. Disable endpoint logging for USB activity

ANSWER: A B C

Explanation:

Disabling AutoRun and AutoPlay where not required, using endpoint protection to scan removable media, and **restricting removable-media use through device-control policies** are appropriate controls. Disabling automatic execution prevents code from launching merely because a device is connected. Endpoint protection can identify known malicious files before users interact with them. Device-control policies can block unauthorized removable devices or permit only approved and encrypted media.

Allowing unrestricted USB devices increases the opportunity for malware introduction and data exfiltration. Disabling endpoint logging removes visibility that may be necessary to investigate suspicious device activity. NIST provides removable-media guidance in SP 800-111, Guide to Storage Encryption Technologies for End User Devices, and CISA provides practical recommendations for removable media at Using Caution with USB Drives and Other External Devices.

QUESTION NO: 28

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:

"Converts passwords to uppercase." is a fundamental LM weakness because the password is normalized to uppercase before hashing. This removes case sensitivity and substantially reduces the possible password search space available to an attacker. "Effective length is 7 characters." is also correct because LM processes a maximum 14-character password by splitting it into two independent seven-character blocks. Each block can therefore be cracked separately, making recovery much more practical than attacking a modern salted password verifier. "Hashes are sent in clear text over the network." is treated as a weakness in the context of legacy LAN Manager environments, where obsolete authentication and password-handling mechanisms can expose reusable credential material or permit interception-based attacks. LM's outdated design is why Microsoft recommends avoiding LM compatibility and disabling LM hash storage wherever possible. Microsoft documents that the LM hash uses the password converted to uppercase and split into two seven-character portions, characteristics that directly enable efficient password-cracking attacks. See Microsoft's guidance on preventing LM hash storage and the Microsoft NTLM protocol specification for legacy authentication context.

QUESTION NO: 29

When configuring wireless on his home router, Javik disables SSID broadcast. He leaves authentication "open" but sets the SSID to a 32-character string of random letters and numbers.

What is an accurate assessment of this scenario from a security perspective?

A. Since the SSID is required in order to connect, the 32-character string is sufficient to prevent brute-force attacks.

B. Disabling SSID broadcast prevents 802.11 beacons from being transmitted from the access point, resulting in a valid setup leveraging "security through obscurity".

C. It is still possible for a hacker to connect to the network after sniffing the SSID from a successful wireless association.

D. Javik's router is still vulnerable to wireless hacking attempts because the SSID broadcast setting can be enabled using a specially crafted packet sent to the hardware address of the access point.

ANSWER: C

Explanation:

It is still possible for a hacker to connect to the network after sniffing the SSID from a successful wireless association. Suppressing SSID broadcast does not make the wireless network undiscoverable or provide meaningful access control. An access point configured with a hidden SSID still transmits management traffic, and the SSID can be exposed when a legitimate client probes for, associates with, or reconnects to the network. A nearby attacker using wireless monitoring tools can capture that traffic and identify the network name, regardless of its length or randomness.

More importantly, the router uses open authentication, so the SSID is only an identifier, not a credential or cryptographic secret. Once the identifier is learned, a wireless client

can attempt to join without needing a protected authentication exchange. Effective Wi-Fi protection requires modern encryption and authentication, such as WPA3-Personal or WPA2-Personal with a strong, unique passphrase; hiding an SSID may reduce casual visibility but is not a security boundary. NIST specifically cautions that SSID cloaking does not prevent determination of the SSID and should not be relied upon for wireless security. See NIST SP 800-153: Guidelines for Securing Wireless Local Area Networks and CISA guidance on securing wireless networks.

QUESTION NO: 30

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-auditing and password-cracking tools. L0phtcrack was developed to assess the strength of Windows account passwords, particularly through recovery and auditing of LAN Manager and NTLM password hashes. Its capabilities include importing hashes, running dictionary and brute-force-style attacks, and reporting weak credentials so they can be remediated.

John the Ripper is a widely used, cross-platform password-security auditing tool. It supports many password-hash formats and can perform dictionary-based, incremental, and rule-based cracking attempts. Security professionals use it during authorized assessments to identify passwords that are susceptible to guessing or offline hash-cracking attacks. In a CEH context, both tools illustrate the password-cracking phase of an ethical hacking engagement: after authorized hash acquisition, the tester evaluates whether password choices and password-hash protections withstand realistic attack methods. Documentation for John the Ripper is available from the Openwall John the Ripper project, while L0phtcrack product information is available at L0phtCrack.

QUESTION NO: 31

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

A. A zone transfer is accomplished with the DNS

B. A zone transfer is accomplished with the nslookup service

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

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

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

F. Zone transfers cannot occur on the Internet

ANSWER: A C E

Explanation:

A zone transfer is a DNS function used to replicate DNS zone data between authoritative servers, such as a primary server and its authorized secondary servers. Therefore, "A zone transfer is accomplished with the DNS" is correct: the transfer uses DNS mechanisms and is performed by DNS servers. "A zone transfer passes all zone information that a DNS server maintains" correctly describes a full zone transfer (AXFR), which transfers the resource records comprising a zone so that a secondary can maintain an authoritative copy. Full zone transfers use TCP, conventionally on port 53; consequently, blocking inbound TCP port 53 connections prevents remote servers from establishing the TCP connection needed to request a zone transfer. In operational environments, the preferred security control is to allow TCP/53 only from explicitly authorized secondary DNS servers, rather than exposing transfers to arbitrary hosts. RFC 5936 specifies DNS zone-transfer procedures and the TCP connection used for AXFR. See RFC 5936: DNS Zone Transfer Protocol (AXFR) and ISC Knowledgebase: Restricting DNS zone transfers.

QUESTION NO: 32

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

A. BA810DBA98995F1817306D272A9441BB

B. 44EFCE164AB921CQAAD3B435B51404EE

C. 0182BD0BD4444BF836077A718CCDF409

D. CEC52EB9C8E3455DC2265B23734E0DAC

E. B757BF5C0D87772FAAD3B435B51404EE

F. E52CAC67419A9A224A3B108F3FA6CB6D

ANSWER: B E

Explanation:

LM hashes consist of two 16-hex-character halves because the legacy LM algorithm uppercases the password, pads it as needed, and processes it in two seven-character blocks. When a password contains fewer than eight characters, the second seven-character block contains only padding. Its resulting, well-known LM hash half is AAD3B435B51404EE. Therefore, 44EFCE164AB921CQAAD3B435B51404EE and B757BF5C0D87772FAAD3B435B51404EE identify passwords shorter than eight characters: each ends with that standard empty-second-block value. The first value contains a non-hexadecimal Q, which appears to be a transcription error, but its intended identifying suffix remains unambiguous. This recognition pattern is useful during credential auditing because it reveals the legacy LM split-block behavior without needing to crack the password first. LM is obsolete and should not be retained or generated in modern environments; Microsoft specifically recommends preventing storage of LAN Manager hash values, and modern password-storage guidance favors adaptive, salted password-hashing mechanisms. See Hashcat example hashes for the LM hash format and Microsoft's LM-hash storage policy documentation for the security guidance.

QUESTION NO: 33

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 Windows-focused enumeration tools. Enumeration follows initial discovery and obtains actionable details about identified systems, including accounts, group memberships, security identifiers, shares, permissions, and policy-related information. USER2SID is used to derive the Windows security identifier associated with a known user account. This allows an assessor to identify account SID values that may be useful when mapping Windows identities and domains. SID2USER performs the complementary task of resolving a SID to its corresponding user or group name, supporting the identification of accounts from enumerated SID data. Windows uses SIDs as unique identity values for security principals; Microsoft describes their role in access control and account identification in its Security Identifiers documentation.

DumpSec is designed to collect and report Windows security information, including users, groups, shares, permissions, and related security settings. These outputs are classic enumeration results because they reveal the target environment's identity and authorization structure. This use aligns with the information-gathering and enumeration activities described in NIST's Technical Guide to Information Security Testing and Assessment.

QUESTION NO: 34

An attacker with access to the inside network of a small company launches a successful STP manipulation attack. What will he do next?

A. He will create a SPAN entry on the spoofed root bridge and redirect traffic to his computer.

B. He will activate OSPF on the spoofed root bridge.

C. He will repeat this action so that it escalates to a DoS attack.

D. He will repeat the same attack against all L2 switches of the network.

ANSWER: A

Explanation:

He will create a SPAN entry on the spoofed root bridge and redirect traffic to his computer. A successful Spanning Tree Protocol manipulation attack typically means the attacker has caused a rogue device or switch under their control to be elected as the STP root bridge, usually by sending superior Bridge Protocol Data Units with a more favorable bridge ID. Because STP calculates active Layer 2 forwarding paths toward the elected root bridge, the network reconverges and traffic from other switches may traverse the attacker-controlled bridge. This placement creates an opportunity to observe traffic that would otherwise remain on switched segments. Configuring Switched Port Analyzer (SPAN), also called port mirroring, copies selected ingress and/or egress traffic from source ports or VLANs to a destination port connected to the attacker's capture system. The attacker can then collect packets for analysis while retaining the forwarding role needed to avoid immediately disrupting connectivity. Cisco documents that SPAN copies traffic from configured sources to a monitor destination, and its STP guidance explains the importance of root-bridge selection in determining Layer 2 topology. See Cisco SPAN and RSPAN configuration guidance and Cisco STP design guidance.

QUESTION NO: 35

_____ is a tool that can hide processes from the process list, can hide files, registry entries, and intercept keystrokes.

A. Trojan

B. RootKit

C. DoS tool

D. Scanner

E. Backdoor

ANSWER: B

Explanation:

RootKit is correct because a rootkit is designed to maintain stealthy, privileged persistence on a compromised system by concealing evidence of its presence and activity. A defining capability is subverting operating-system mechanisms that report system state, allowing malicious processes to be omitted from process listings and selected files, directories, services, or Registry keys to be hidden from normal administrative tools. Rootkits can operate in user mode, kernel mode, or at lower layers such as boot or firmware, with deeper placement generally providing more effective concealment. They may also include surveillance functions such as keystroke interception, enabling collection of credentials and other sensitive input while avoiding detection. These capabilities make rootkits particularly dangerous: system output can no longer necessarily be trusted because the malware may alter what security tools and administrators are shown. MITRE ATT&CK documents rootkit use as an adversary technique for hiding artifacts and processes, and CISA describes rootkits as malware that enables unauthorized access while hiding its existence or the existence of other software. See MITRE ATT&CK: Rootkit and CISA: What is a Rootkit?.

QUESTION NO: 36

in an attempt to increase the security of your network, you Implement a solution that will help keep your wireless network undiscoverable and accessible only to those that know It. How do you accomplish this?

A. Delete the wireless network

B. Remove all passwords

C. Lock all users

D. Disable SSID broadcasting

ANSWER: D

Explanation:

Disabling SSID broadcasting is the configuration that best matches the stated goal of making a wireless network less visible to casual users. The service set identifier is the network name ordinarily included in Wi-Fi beacon frames, allowing nearby devices to display the network in their available-network lists. When SSID broadcast is disabled, the access point omits the network name from those beacons. Users who already know the exact SSID can manually configure a wireless profile and attempt to connect, which is why this is commonly described as a “hidden” wireless network.

This measure should be understood as a limited visibility control rather than a standalone security boundary. Wireless clients configured for a hidden network may still send probe requests that reveal the SSID, and wireless monitoring tools can identify the network through management or data traffic. Effective protection therefore also requires modern authentication and encryption, such as WPA3-Personal or WPA2-AES with a strong, unique passphrase, along with appropriate access-point hardening. Nevertheless, for the specific objective of preventing ordinary beacon-based network discovery while permitting

connections by users who know the network name, disabling SSID broadcasting is the appropriate answer. See Cisco's wireless access-point guidance and CISA's wireless network security guidance.

QUESTION NO: 37

Peter, a Network Administrator, has come to you looking for advice on a tool that would help him perform SNMP enquires over the network.

Which of these tools would do the SNMP enumeration he is looking for? Select the best answers.

A. SNMPUtil

B. SNScan

C. SNMPScan

D. Solarwinds IP Network Browser

E. NMap

ANSWER: A B D

Explanation:

SNMPUtil, **SNScan**, and **Solarwinds IP Network Browser** are appropriate tools for SNMP-focused discovery and enumeration. SNMP enumeration involves querying an SNMP-enabled device—normally using a known or discovered community string in SNMPv1/v2c, or valid SNMPv3 credentials—to retrieve management information exposed through Management Information Base (MIB) object identifiers. Useful results can include interface details, IP addressing, routing information, device descriptions, installed software, running services, and system-contact data. SNMPUtil is a Windows command-line utility designed to issue SNMP GET, GETNEXT, and WALK-style queries against specified MIB object identifiers, making it suited to direct information retrieval. SNScan is intended to scan hosts for accessible SNMP services and enumerate SNMP information at scale. SolarWinds IP Network Browser provides network-discovery and IP-management functionality that can use SNMP polling to identify managed devices and collect their device and interface data. These tools support the SNMP-enumeration phase described in CEH methodology, where the assessor identifies exposed management services and documents only the information authorized by the engagement scope. Microsoft describes SNMP management concepts and usage in its Windows Server documentation, while SolarWinds documents IP network discovery and management capabilities in its IP Network Browser documentation.

QUESTION NO: 38

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

A. 110

B. 135

C. 139

D. 161

E. 445

F. 1024

ANSWER: B C E

Explanation:

For a Windows NT, Windows 2000, and Windows XP environment, blocking **135**, **139**, and **445** at network boundaries is an appropriate defensive control for preventing the legacy Windows networking exposure intended by the question. Port 135 is used by the Microsoft RPC Endpoint Mapper, which helps clients locate RPC services and is commonly associated with Windows remote-management and file-sharing attack paths. Port 139 is the NetBIOS Session Service port used for NetBIOS-based SMB sessions. Port 445 is used by SMB directly over TCP, without the NetBIOS session layer, and must also be filtered because Windows systems can use it for file and printer sharing even when NetBIOS-based connectivity is not used. Blocking these ports inbound and outbound where such services are not explicitly required reduces exposure to unauthenticated discovery, share enumeration, lateral movement, and SMB/RPC vulnerabilities. Microsoft's service and port requirements identify RPC Endpoint Mapper and SMB as relevant Windows network services, while the IANA registry documents the registered NetBIOS Session Service and Microsoft-DS port assignments. See Microsoft's service overview and network port requirements and the IANA Service Name and Transport Protocol Port Number Registry.

QUESTION NO: 39

Which of the following activities are examples of passive reconnaissance? (Choose three.)

A. Reviewing public job postings

B. Examining certificate-transparency logs

C. Searching public code repositories

D. Performing a TCP SYN scan against the target

E. Connecting to a service to retrieve its banner

ANSWER: A B C

Explanation:

Reviewing public job postings, examining certificate-transparency logs, and searching public code repositories are passive reconnaissance activities because they collect information without directly interacting with the target's systems. Job postings can disclose technologies, cloud platforms, security products, and organizational structure. Certificate-transparency logs can reveal publicly issued certificates and associated subdomains. Public code repositories may expose usernames, project names, deployment references, or accidentally committed secrets.

Port scanning and banner grabbing require direct interaction with target hosts and are active reconnaissance techniques. Passive reconnaissance is useful early in an authorized engagement because it can develop a target profile while generating little or no traffic visible to the organization. MITRE describes public-facing information gathering under its Search Open Websites/Domains reconnaissance technique.

QUESTION NO: 40

Which command can be used to show the current TCP/IP connections?

A. Netsh

B. Netstat

C. Net use connection

D. Net use

ANSWER: B

Explanation:

Netstat is the command used to display current network connections and protocol statistics, including active TCP connections. On Windows, running netstat shows active TCP connections by default, presenting local and remote endpoint addresses along with each connection's state, such as ESTABLISHED, LISTENING, or TIME_WAIT. This makes it a core command for enumerating network activity during troubleshooting and security assessments. Useful variations include netstat -a to include listening ports and all connections, netstat -n to display numeric addresses and port numbers without DNS name resolution, and netstat -o to associate connections with process IDs. In an ethical-hacking workflow, these views help an analyst identify exposed services, inspect unexpected remote sessions, and correlate a suspicious connection with the responsible local process. Microsoft's command reference documents that Netstat displays active TCP connections, ports on which the computer is listening, Ethernet statistics, and IP routing information: Microsoft Netstat command documentation. The equivalent utility is also standardized and documented for Unix-like environments in the netstat manual page.

QUESTION NO: 41

Sam, a professional hacker. targeted an organization with intention of compromising AWS IAM credentials. He attempted to lure one of the employees of the organization by initiating fake calls while posing as a legitimate employee. Moreover, he sent phishing emails to steal the AWS 1AM credentials and further compromise the employee's account. What is the technique used by Sam to compromise the AWS IAM credentials?

A. Social engineering

B. insider threat

C. Password reuse

D. Reverse engineering

ANSWER: A

Explanation:

Social engineering is correct because the attack relies on manipulating an employee into disclosing credentials or interacting with a deceptive message. The fake telephone calls impersonating a legitimate employee are a form of pretexting or vishing, while the fraudulent email messages are phishing. Both methods exploit trust, urgency, and perceived legitimacy rather than directly defeating a technical control. The intended result is theft of cloud identity credentials, which can then let an attacker authenticate as the victim and access resources according to the permissions assigned to that identity.

For cloud environments, credential phishing is particularly significant because access keys, passwords, session tokens, and multifactor-authentication approvals can enable access to highly valuable services. Effective defenses include verifying unexpected identity-related requests through an independent channel, using phishing-resistant multifactor authentication, applying least privilege, monitoring unusual sign-in activity, and promptly revoking or rotating credentials believed to be exposed. AWS specifically advises customers to protect credentials, avoid sharing them, and use multifactor authentication for additional account security. See the AWS IAM security best practices and CISA's guidance on recognizing and reporting phishing.

QUESTION NO: 42

An administrator is reviewing wireless security for a corporate WLAN. Which of the following practices are appropriate for securing the WLAN? (Choose three.)

A. Use WPA3-Enterprise

B. Use 802.1X authentication with a RADIUS server

C. Segment guest wireless users from internal resources

D. Rely on a hidden SSID as the primary access control

E. Use WEP for compatibility with legacy clients

ANSWER: A B C

Explanation:

Using WPA3-Enterprise, deploying 802.1X authentication with a RADIUS server, and separating guest wireless users from internal resources are appropriate security practices. WPA3-Enterprise provides strong Wi-Fi security for enterprise deployments and can use 802.1X authentication to validate individual users or devices through a RADIUS infrastructure. Per-user authentication improves accountability and avoids the shared-secret limitations associated with a single pre-shared key.

Guest networks should be segmented from internal corporate systems so that Internet access can be provided without granting guest devices access to sensitive network resources. Hiding an SSID is not a meaningful access control because the SSID can be observed in wireless management traffic. WEP is obsolete and vulnerable to practical attacks. NIST provides wireless security recommendations in SP 800-153, Guidelines for Securing Wireless Local Area Networks.

QUESTION NO: 43

Gerard, a disgruntled ex-employee of Sunglass IT Solutions, targets this organization to perform sophisticated attacks and bring down its reputation in the market. To launch the attacks process, he performed DNS footprinting to gather information about ONS servers and to identify the hosts connected in the target network. He used an automated tool that can retrieve information about DNS zone data including DNS domain names, computer names. IP addresses. DNS records, and network Who is records. He further exploited this information to launch other sophisticated attacks. What is the tool employed by Gerard in the above scenario?

A. Knative

B. zANTI

C. Towelroot

D. Bluto

ANSWER: D

Explanation:

Bluto is the correct tool because it is an automated reconnaissance utility designed to collect DNS-related intelligence about a target. Its purpose aligns directly with the scenario’s footprinting activity: identifying domains and host-related information, resolving associated IP addresses, obtaining DNS record data, and performing WHOIS queries. This information helps an assessor build a profile of an organization’s Internet-facing presence and understand relationships among domains, name servers, records, and potentially exposed systems before conducting authorized follow-on testing.

DNS footprinting is valuable because DNS records can reveal infrastructure details such as authoritative name servers, mail-routing hosts, aliases, and address mappings. WHOIS results can additionally provide registration and administrative context when that data is publicly available. Bluto consolidates these reconnaissance functions into an automated workflow, which is precisely the behavior described in the question. The project's documentation identifies it as a DNS reconnaissance tool and provides its usage context at Bluto on GitHub. For background on the DNS record system that such reconnaissance queries examine, see the IETF DNS implementation and specification document.

QUESTION NO: 44

You are performing a penetration test for a client and have gained shell access to a Windows machine on the internal network. You intend to retrieve all DNS records for the internal domain, if the DNS server is at 192.168.10.2 and the domain name is abccorp.local, what command would you type at the nslookup prompt to attempt a zone transfer?

A. list server=192.168.10.2 type=all

B. ls -d abccorp.local

C. Iserver 192.168.10.2-t all

D. List domain=Abccorp.local type=zone

ANSWER: B

Explanation:

ls -d abccorp.local is the interactive nslookup command used to request a full listing of the specified domain's DNS data. The ls command causes nslookup to attempt a DNS zone transfer, while -d requests all records for the domain rather than restricting the output to host, alias, or a particular subdomain record type. Before issuing this command, the tester would direct nslookup at the target name server with server 192.168.10.2, then run ls -d abccorp.local. In a successful assessment scenario, the DNS server responds to the request using the DNS zone-transfer mechanism and returns the zone contents, which can reveal internal hosts, service records, and naming structure. Whether the transfer succeeds depends on the DNS server's zone-transfer authorization configuration; properly managed DNS servers restrict transfers to approved secondary servers. Microsoft documents the ls command and its -d form in its nslookup command reference. The protocol behavior for full DNS zone transfers is standardized as AXFR in RFC 5936.

QUESTION NO: 45

Joseph was the Web site administrator for the Mason Insurance in New York, who's main Web site was located at www.masonins.com. Joseph uses his laptop computer regularly to administer the Web site. One night, Joseph received an urgent phone call from his friend, Smith. According to Smith, the main Mason Insurance web site had been vandalized! All of its normal content was removed and replaced with an attacker 's message ''Hacker Message: You are dead! Freaks!â€ From his office, which was directly connected to Mason Insurance's internal network, Joseph surfed to the Web site using his laptop. In his browser, the Web site looked completely intact.

No changes were apparent. Joseph called a friend of his at his home to help troubleshoot the problem. The Web site appeared defaced when his friend visited using his DSL connection. So, while Smith and his friend could see the defaced page, Joseph saw the intact Mason Insurance web site. To help make sense of this problem, Joseph decided to access the Web site using hisdial-up ISP. He disconnected his laptop from the corporate internal network and used his modem to dial up the same ISP used by Smith. After his modem connected, he quickly typed www.masonins.com in his browser to reveal the following web page:

After seeing the defaced Web site, he disconnected his dial-up line, reconnected to the internal network, and used Secure Shell (SSH) to log in directly to the Web server. He ran Tripwire against the entire Web site, and determined that every system file and all the Web content on the server were intact. How did the attacker accomplish this hack?

A. ARP spoofing

B. SQL injection

C. DNS poisoning

D. Routing table injection

ANSWER: C

Explanation:

DNS poisoning is correct because the evidence shows that the legitimate Mason Insurance web server and its web content were not modified, yet users connecting through external ISP DNS resolution reached a page containing the attacker's defacement. A poisoned DNS record can cause www.masonins.com to resolve to an attacker-controlled IP address rather than the legitimate web server's address. The victim's browser still displays the expected domain name, but the HTTP request is sent to the fraudulent host, which can serve a convincing or overtly defaced copy of the site.

Joseph's differing results are especially significant: while attached to the corporate network, his laptop likely used an internal DNS server or split-DNS zone that retained the legitimate address. When he connected through the same external ISP used by the other observer, DNS resolution returned the poisoned result and exposed the malicious page. Logging in to the genuine server through SSH and finding Tripwire integrity checks clean confirms that the attack redirected users before they reached that server, rather than altering server files. DNS cache-poisoning risks and the importance of unpredictable query identifiers and source ports are detailed in RFC 5452. DNS security concepts and DNSSEC validation are also described by ICANN.

QUESTION NO: 46

Geena, a cloud architect, uses a master component in the Kubernetes cluster architecture that scans newly generated pods and allocates a node to them. This component can also assign nodes based on factors such as the overall resource requirement, data locality, software/hardware/policy restrictions, and internal workload interventions.

Which of the following master components is explained in the above scenario?

A. Kube-controller-manager

B. Kube-scheduler

C. Kube-apiserver

D. Etcd cluster

ANSWER: B

Explanation:

Kube-scheduler is the Kubernetes control-plane component that watches for newly created Pods that do not yet have a node assigned, then selects an appropriate node for each Pod. Its scheduling process evaluates the Pod's declared resource requests and the available capacity of candidate nodes, ensuring that workloads are placed where they can run successfully. It also applies placement constraints and preferences, including node affinity and anti-affinity, taints and tolerations, topology considerations, and policies expressed through scheduling configuration. These mechanisms support the scenario's

references to resource requirements, data locality, software or hardware restrictions, policy restrictions, and workload-related placement decisions. After choosing the best feasible node, the scheduler creates the binding that assigns the Pod to that node; the node's kubelet can then begin the process of running the Pod's containers. Kubernetes documents the scheduler as the component responsible for assigning unscheduled Pods to nodes, using filtering and scoring to identify the most suitable placement. See the Kubernetes Scheduler documentation and the Assigning Pods to Nodes guide.

QUESTION NO: 47

Alice needs to send a confidential document to her coworker. Bryan. Their company has public key infrastructure set up. Therefore. Alice both encrypts the message and digitally signs it. Alice uses_____to encrypt the message, and Bryan uses_____to confirm the digital signature.

A. Bryan's public key; Bryan's public key

B. Alice's public key; Alice's public key

C. Bryan's private key; Alice's public key

D. Bryan's public key; Alice's public key

ANSWER: D

Explanation:

Bryan's public key; Alice's public key is correct because public key cryptography provides confidentiality by encrypting data for its intended recipient with that recipient's public key. In this case, Alice encrypts the confidential document using Bryan's public key. Only the holder of the corresponding private key—Bryan—can decrypt the protected content, which preserves confidentiality while the document is transmitted or stored.

For the digital signature, Alice creates the signature using her private key. Bryan confirms the signature using Alice's public key, typically obtained from and validated through her digital certificate in the organization's PKI. Successful verification demonstrates that the signature corresponds to Alice's key pair and that the signed data has not been altered since it was signed. In practical secure-message protocols, the document itself is commonly encrypted using a one-time symmetric session key, and that session key is then encrypted with Bryan's public key; however, the public-key roles tested by the question remain the same. See RFC 5280: Internet X.509 Public Key Infrastructure Certificate and CRL Profile and NIST's FIPS 186-5, Digital Signature Standard.

QUESTION NO: 48

An ethical hacker is testing the security of a website's database system against SQL Injection attacks. They discover that the IDS has a strong signature detection mechanism to detect typical SQL injection patterns.

Which evasion technique can be most effectively used to bypass the IDS signature detection while performing a SQL Injection attack?

A. Implement case variation by altering the case of SQL statements

B. Employ IP fragmentation to obscure the attack payload

C. Use Hex encoding to represent SQL string literals

D. Leverage string concatenation to break identifiable keywords

ANSWER: C

Explanation:

Use hexadecimal encoding to represent SQL string literals is correct because a database can interpret a hexadecimal literal as the same underlying value while the textual form no longer contains the ordinary quoted strings and character sequences that many SQL-injection signatures search for. For example, a value such as 'admin' can be represented as a hexadecimal literal in database dialects that support it. This preserves the query's intended semantics while changing the payload representation enough to evade a signature engine that does not fully canonicalize or decode alternate encodings before inspection. MySQL documents hexadecimal literals and their conversion to strings, demonstrating why this is a database-recognized representation rather than merely an encoded transport artifact: MySQL Reference Manual: Hexadecimal Literals. Encoding-based payload variation is a relevant concern in SQL-injection testing because defenses must validate input and avoid dynamically constructing SQL, rather than relying solely on pattern matching. OWASP emphasizes parameterized queries as the primary control that prevents attacker-supplied input from changing SQL syntax: OWASP SQL Injection Prevention Cheat Sheet. In an authorized assessment, this technique helps verify whether an IDS properly normalizes alternate representations before applying detection logic.

QUESTION NO: 49

The network in ABC company is using the network address 192.168.1.64 with mask 255.255.255.192. In the network the servers are in the addresses 192.168.1.122, 192.168.1.123 and 192.168.1.124. An attacker is trying to find those servers but he cannot see them in his scanning. The command he is using is: nmap 192.168.1.64/28.

Why he cannot see the servers?

A. He needs to add the command ""ip address"" just before the IP address

B. He needs to change the address to 192.168.1.0 with the same mask

C. He is scanning from 192.168.1.64 to 192.168.1.78 because of the mask /28 and the servers are not in that range

D. The network must be dawn and the nmap command and IP address are ok

ANSWER: C

Explanation:

He is scanning from 192.168.1.64 to 192.168.1.78 because of the mask /28 and the servers are not in that range. The prefix supplied to Nmap controls the target range, rather than merely describing the organization's actual subnet. A /28 prefix leaves four host bits, producing a block of 16 addresses. Starting at 192.168.1.64, that block spans 192.168.1.64 through 192.168.1.79. Within that block, .64 is the network address and .79 is the broadcast address, so the normal usable host range is .65 through .78. Consequently, targets .122, .123, and .124 are outside the scan target range and will not be probed by that command.

The stated mask of 255.255.255.192 corresponds to /26, whose network beginning at 192.168.1.64 spans .64 through .127 and includes all three server addresses. Nmap accepts CIDR notation for target specifications, and the selected prefix determines which addresses are expanded and scanned. The relevant subnetting behavior is defined by IPv4 CIDR addressing rules; see Nmap Target Specification and RFC 4632: Classless Inter-domain Routing.

QUESTION NO: 50

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

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

A. Converts passwords to uppercase.

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

C. Makes use of only 32-bit encryption.

D. Effective length is 7 characters.

E. Splits passwords into two 7-character halves and hashes each half independently.

ANSWER: A D E

Explanation:

LM password hashing has structural weaknesses that substantially reduce the work needed to recover passwords. "Converts passwords to uppercase." is correct because the LM process normalizes lowercase characters to uppercase before deriving the hash. This removes case sensitivity and reduces the effective password search space. "Effective length is 7 characters." is correct in the practical cracking sense: although LM accepts a password of up to 14 characters, it pads the password and divides it into two independent seven-character blocks. An attacker can therefore attack each block separately rather than searching all 14 characters as one unit. "Splits passwords into two 7-character halves and hashes each half independently." describes this same critical design flaw more explicitly. Each half is used to create separate DES-based values, so a recovered first half does not depend on recovering the second half. These characteristics make LM hashes especially susceptible to dictionary, brute-force, and precomputed-table attacks. Microsoft's NTLM protocol specification documents the LM one-way-function calculation, including uppercasing, padding, and the two seven-byte portions: Microsoft MS-NLMP: LMOWF. Microsoft also recommends eliminating legacy LM authentication where possible: Microsoft NTLM overview.

QUESTION NO: 51

What would be the fastest way to perform content enumeration on a given web server by using the Gobuster tool?

A. Performing content enumeration using the bruteforce mode and 10 threads

B. Shipping SSL certificate verification

C. Performing content enumeration using a wordlist

D. Performing content enumeration using the bruteforce mode and random file extensions

ANSWER: A

Explanation:

Performing content enumeration using the bruteforce mode and 10 threads is the correct choice because Gobuster accelerates web-content discovery by issuing requests concurrently rather than waiting for each candidate path to finish before starting the next one. In directory/content enumeration, Gobuster tests candidate directory and file names from a wordlist against the target web server. Its thread setting controls how many workers make these requests in parallel; using 10 threads allows multiple enumeration requests to be processed simultaneously, substantially reducing total scan time compared with a single-threaded approach while remaining a practical, commonly used level of concurrency. This approach is appropriate when the tester has authorization and should still be tuned to the target's capacity and the applicable rules of engagement. Gobuster's documentation describes its directory enumeration mode and the -t option for setting the number of concurrent threads. See the Gobuster project documentation and the Kali Linux Gobuster tool page for usage details.

QUESTION NO: 52

A company's policy requires employees to perform file transfers using protocols which encrypt traffic. You suspect some employees are still performing file transfers using unencrypted protocols because the employees do not like changes. You have positioned a network sniffer to capture traffic from the laptops used by employees in the data ingest department. Using Wireshark to examine the captured traffic, which command can be used as display filter to find unencrypted file transfers?

A. tcp.port == 21

B. tcp.port = 23

C. tcp.port = = 21 | | tcp.port = =22

D. tcp.port ! = 21

ANSWER: A

Explanation:

tcp.port == 21 is the appropriate Wireshark display filter for identifying FTP control traffic, which is a strong indicator that unencrypted FTP is being used. Traditional FTP uses TCP port 21 for its control channel: commands such as authentication, directory navigation, and transfer requests are sent there in cleartext unless an explicit TLS-protected FTP configuration has been negotiated. In the scenario, observing this traffic lets the analyst quickly locate hosts initiating FTP sessions and inspect the associated commands and credentials where authorized. Wireshark display-filter syntax uses the equality operator ==; therefore, the corrected expression filters packets where either the TCP source or destination port is 21. This is useful in packet captures because it finds both client-to-server and server-to-client control-channel packets. After identifying an FTP control session, the analyst can follow the TCP stream and correlate FTP transfer commands with any negotiated data connections to investigate the transfer activity. Wireshark documents tcp.port as a field that can be used in display filters, and its FTP dissector documentation identifies FTP traffic for protocol analysis. See the Wireshark display-filter reference and the Wireshark FTP protocol field reference.

QUESTION NO: 53

You want to do an ICMP scan on a remote computer using hping2. What is the proper syntax?

A. hping2 host.domain.com

B. hping2 --set-ICMP host.domain.com

C. hping2 -i host.domain.com

D. hping2 -1 host.domain.com

ANSWER: D

Explanation:

The proper syntax is hping2 -1 host.domain.com. In hping2, the -1 command-line switch selects ICMP mode, causing hping2 to construct and send ICMP packets rather than its default TCP packets. This is useful when testing a remote host's response to ICMP traffic, including host-discovery behavior and packet-filtering rules. Once ICMP mode is selected, hping2 can use its additional ICMP-specific options to choose an ICMP type and code, but those are not required for the basic ICMP scan syntax shown here. The destination hostname follows the mode switch, so host.domain.com is the remote target to which the ICMP probes are sent. Hping's documented usage identifies -1 as the ICMP mode option, and the tool's manual describes this mode as sending ICMP packets to the supplied target. See the hping3

manual page (the maintained successor with compatible mode syntax) and the Kali Linux hping3 tool documentation for command usage and ICMP-mode details.

QUESTION NO: 54

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:

For legacy Windows NT, Windows 2000, and Windows XP environments, firewall rules intended to prevent NetBIOS and Windows file-sharing exposure should block UDP ports 137 and 138, TCP port 139, and TCP port 445 at the network perimeter. UDP 137 is the NetBIOS Name Service used for name registration and resolution, while UDP 138 is the NetBIOS Datagram Service used for connectionless datagram traffic. TCP 139 is the NetBIOS Session Service port used to establish NetBIOS sessions and historically carries SMB file- and printer-sharing communications over NetBIOS over TCP/IP. Microsoft Windows also supports SMB directly over TCP port 445, without requiring the NetBIOS session layer. Blocking TCP 445 is therefore necessary alongside the traditional NetBIOS ports when the goal is to prevent inbound legacy Windows sharing and related service traffic from traversing the firewall. Microsoft's SMB documentation identifies TCP 445 as the direct-hosted SMB port, and its NetBIOS documentation identifies the Name Service, Datagram Service, and Session Service ports. See Microsoft's service overview and network port requirements and Microsoft NetBIOS documentation.

QUESTION NO: 55

You want to analyze packets on your wireless network. Which program would you use?

A. Wireshark with Airpcap

B. Airsnort with Airpcap

C. Wireshark with Winpcap

D. Ethereal with Winpcap

ANSWER: A

Explanation:

Wireshark with Airpcap is the appropriate choice for wireless packet analysis because AirPcap was a dedicated wireless capture adapter and driver family designed to provide Wi-Fi frames to Wireshark on Windows. Unlike a standard network interface operating in its usual mode, a wireless capture adapter must support monitor-mode capture so that

management, control, and data frames transmitted over the air can be collected and inspected. AirPcap integrates with Wireshark's capture mechanisms, allowing an analyst to select wireless channels, capture raw IEEE 802.11 traffic, and examine packet contents and protocol fields in Wireshark's detailed analysis interface. This makes the pairing suitable for tasks such as identifying access points and clients, reviewing beacon and probe traffic, observing authentication and association exchanges, and investigating wireless security events. The question reflects the historically common CEH toolset: AirPcap hardware was specifically marketed for use with Wireshark and was supported through the AirPcap capture library. Wireshark's documentation describes its packet-capture architecture and supported capture interfaces at Wireshark User's Guide: Capturing Live Network Data. Historical AirPcap technical documentation is available from Riverbed AirPcap Documentation.

QUESTION NO: 56

ViruXine.W32 virus hides their presence by changing the underlying executable code.

This Virus code mutates while keeping the original algorithm intact, the code changes itself each time it runs, but the function of the code (its semantics) will not change at all.

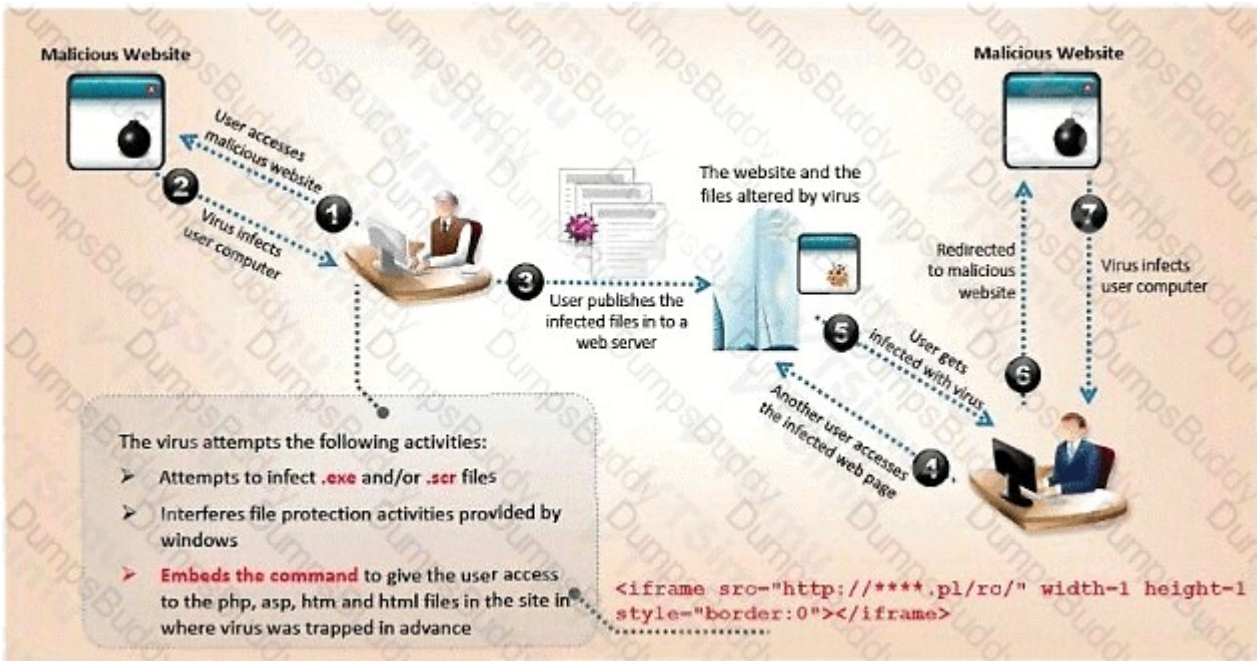

Here is a section of the Virus code:

```
1. lots of encrypted code
2. …
3. Decryption_Code:
4. C=C+1
5. A=Encrypted
6. Loop:
7. B=*A
8. C=3214*A
9. B=B XOR CryptoKey
10. *A=B
11. C=1
12. C=A+B
13. A=A+1
14. GOTO Loop IF NOT A=Decryption_Code
15. C=C^2
16. GOTO Encrypted
17. CryptoKey:
18. some_random_number
```

What is this technique called?

A. Polymorphic Virus

B. Metamorphic Virus

C. Dravidic Virus

D. Stealth Virus

ANSWER: B

Explanation:

Metamorphic Virus is correct because the defining characteristic of metamorphism is that the malware rewrites or transforms its own executable implementation while preserving its original behavior. Each newly generated version can have a substantially different instruction sequence, register allocation, control-flow layout, ordering of independent instructions, or equivalent code substitutions, yet it still executes the same malicious algorithm. The question specifically states that the underlying executable code changes each time it runs while the algorithm and semantics remain intact. That is a direct description of metamorphic code transformation.

This self-rewriting behavior is intended to produce successive variants that do not share a stable byte-level signature, complicating static signature-based identification. A metamorphic engine can regenerate functionally equivalent code repeatedly, making the executable appearance vary without changing the malware's operational purpose. This distinction is important in malware analysis because analysts must focus on behavioral and semantic properties, rather than expecting each generation to contain identical instruction bytes. Fortinet describes metamorphic malware as malware that changes its code while retaining its functionality; see Fortinet's metamorphic virus overview. MITRE also documents code-obfuscation and transformation techniques used to impede static analysis at MITRE ATT&CK: Obfuscated Files or Information.

QUESTION NO: 57

Which of the following controls are effective for reducing the risk of unauthorized access to an Active Directory environment? (Choose three.)

A. Use separate privileged administrative accounts

B. Enforce multifactor authentication for privileged access

C. Apply the principle of least privilege

D. Grant Domain Admin privileges to all support staff

E. Leave former employee accounts enabled

ANSWER: A B C

Explanation:

Using separate privileged administrative accounts, enforcing multifactor authentication for privileged access, and **applying the principle of least privilege** are effective controls. Dedicated administrative accounts reduce the exposure created when a daily-use account is compromised. Multifactor authentication adds a separate factor that makes password-only compromise less likely to result in account takeover. Least privilege limits administrative rights to the minimum permissions required for assigned duties.

Granting Domain Admin privileges to all support staff substantially increases the number of accounts capable of making forest-wide changes. Leaving former employee accounts enabled creates unnecessary opportunities for unauthorized use. Microsoft provides guidance on securing privileged access in its privileged access strategy documentation.