RUTGERS UNIVERSITY

Department of Computer Science

Computer Security

Exam 3

December 1, 2025

Review – Answers Discussion

Prepared with tender loving care by Paul Krzyzanowski

100 POINTS – 25 QUESTIONS – 4 POINTS EACH – For each statement, select the *most* appropriate answer.

1. How does a *virus* differ from other types of malware?
   a. It provides attackers with remote access to a system.
   b. It attaches itself to other programs or files to replicate.
   c. Users install it willingly because it disguises its true purpose.
   d. It spreads automatically to other systems without user action.

Correct answer: b

A virus attaches itself to other files or programs and relies on them for replication.

Incorrect choices:

(b) This describes remote access trojans, not viruses.

(c) This behavior describes trojans that rely on user installation.

(d) Autonomous propagation describes worms, not viruses.

2. How does a *Trojan* differ from other types of malware?
   a. It attaches itself to other programs or files to replicate.
   b. It provides attackers with remote access to a system.
   c. Users install it willingly because it disguises its true purpose.
   d. It spreads automatically to other systems without user action.

Correct answer: c

A trojan is installed because it masquerades as something legitimate while hiding its malicious purpose.

Incorrect choices:

(a) This describes a virus attaching to files.

(b) Some trojans offer remote access, but that is not what defines them.

(d) Worms propagate automatically, but trojans do not.

3. What is the primary reason attackers use *packers*?
   a. To evade static detection by obscuring code structure.
   b. To reduce the size of the malware in memory.
   c. To increase the network throughput of command traffic.
   d. To automatically gain administrative privileges.

Correct answer: a

Packers obfuscate malware code to evade signature-based and static analysis.

Incorrect choices:

(b) Packers do not primarily aim to reduce memory usage.

(c) Packers do not affect network performance.

(d) Packers do not perform privilege escalation.

4. What distinguishes an *N-day exploit* from a zero-day exploit?
   a. It targets a vulnerability for which a patch or workaround already exists.
   b. It relies on social engineering instead of technical flaws.
   c. It targets a vulnerability that is expected to become public within a predicted number of days.
   d. It targets a theoretical vulnerability that has not yet been weaponized.

Correct answer: a

An N-day exploit abuses a known vulnerability that already has a fix but has not been applied everywhere.

Incorrect choices:
(b) Social engineering is not part of the N-day definition.

(c) Disclosure timing prediction is unrelated to N-day terminology.

(d) A theoretical flaw is unrelated to N-days and zero-days.

5. Why are *rootkits* particularly difficult for standard antivirus tools to detect?
   a. They exist only in GPU (Graphical Processing Unit) memory and avoid system RAM entirely.
   b. They disguise themselves as useful applications to trick users.
   c. They intercept or modify system calls to hide files and processes.
   d. They frequently move between different systems to evade detection.

Correct answer: c

Rootkits hook or filter system calls so tools never see their processes or files.

Incorrect choices:

(a) Rootkits generally reside in normal system memory.

(b) Disguising malicious programs is a trojan tactic.

(d) Rootkits do not evade detection by migrating across systems.

6. What was the central lesson of Ken Thompson's *Reflections on Trusting Trust*?
   a. You cannot trust code you did not write yourself.
   b. Open-source software is safer because you can audit it and search for malicious code.
   c. Password-based authentication is inherently unsafe.
   d. A program may contain a backdoor even if both the program and the compiler source are audited.

Correct answer:d

Thompson showed that a malicious compiler can inject a backdoor even when source appears clean.

Incorrect choices:

(a) His point was specifically about compilers, not all outside code.

(b) Open-source availability does not stop compiler-inserted backdoors.

(c) Password handling was not the focus of the paper.

7. A *zero-day vulnerability* gets its name because:
   a. Attackers built an exploit immediately after the vulnerability was disclosed.
   b. Software vendors have zero days to fix or patch the vulnerability.
   c. It spreads to multiple systems in zero days.
   d. The vulnerability was found and exploited on the same day.

Correct answer: b

A zero-day vulnerability is unknown to the vendor, meaning they have had zero days to prepare or release a fix.

Incorrect choices:

(a) Zero-day vulnerabilities exist before disclosure, not after.

(c) The term has nothing to do with propagation speed.

(d) Development timing does not define a zero-day.

8. Which attack does *NOT* provide the attacker with opportunities for *eavesdropping* on traffic?
   a. CAM table overflow.
   b. SYN flooding.
   c. Rogue DHCP server.
   d. ARP spoofing.

Correct answer: b

A SYN flood exhausts server resources without intercepting traffic.

Incorrect choices:

(a) CAM overflow causes switches to flood frames, enabling sniffing.

(c) Rogue DHCP servers redirect traffic through an attacker-controlled gateway.

(d) ARP spoofing enables interception by forging MAC mappings.

9. Why are UDP protocols appealing to attackers for *DDoS reflection attacks*?
   a. UDP services often generate responses larger than the incoming request.
   b. UDP supports small payloads, allowing high packet-per-second volume.
   c. UDP is a connectionless protocol that can be easily spoofed.
   d. UDP supports broadcast packets, increasing the attacker's reach.

Correct answer: c

Reflection attacks rely on spoofing the victim's address, which is trivial with a connectionless protocol like UDP.

Incorrect choices:

(a) Amplification alone does not enable reflection without spoofing.

(b) Packet-per-second rates do not determine reflection feasibility.

(d) Reflection relies on spoofed unicast replies, not broadcast behavior.

10. Why is *BGP hijacking* possible on today's Internet?
    a. BGP routers trust route advertisements without verifying ownership.
    b. BGP uses a Spanning Tree Protocol, enabling routing loops.
    c. It is easy to passively monitor route advertisements from other ASes.
    d. BGP uses simple cryptographic signatures that are easy to break.

Correct answer: a

BGP lacks built-in validation of which AS owns a prefix.

Incorrect choices:

(b) BGP does not use the spanning tree protocol (STP).

(c) Monitoring routes does not cause hijacking.

(d) Standard BGP does not use cryptographic signatures.

11. How do *SYN cookies* protect a server from SYN flooding attacks?
    a. They block UDP traffic to ensure only TCP connections are allowed.
    b. They increase the timeout for half-open connections.
    c. They encode connection state in the sequence number so memory allocation is deferred.
    d. They require client authentication before sending SYN packets.

Correct answer: c

Encoding state in the sequence number prevents the need to allocate memory during the handshake.

Incorrect choices:

(a) UDP is unrelated to SYN floods.

(b) Longer timeouts worsen half-open issues.

(d) TCP SYN packets are never authenticated.

12. What makes NTP (Network Time Protocol) an attractive target for *amplification*?
    a. It can generate large responses from small queries.
    b. It uses simple password-based authentication.
    c. Its responses are unencrypted.
    d. It can modify system clocks, which can affect time-dependent protocols.

Correct answer: a

NTP commands like monlist return large replies to small requests.

Incorrect choices:

(b) NTP amplification does not involve passwords.

(c) Encryption is irrelevant to amplification.

(d) Time manipulation is harmful but not part of amplification attacks.

13. A *CAM table overflow* attack results in:
    a. An Ethernet switch flooding packets out all ports.
    b. An Ethernet switch refusing to receive additional frames.
    c. An Ethernet switch being unable to forward any packets.
    d. A switch permanently disabling unused ports.

Correct answer: a

Overflowing the CAM table forces the switch to flood traffic to all ports.

Incorrect choices:

(b) Switches still receive frames.

(c) Switching continues; it is just no longer selective.

(d) CAM attacks do not disable ports.

14. How does the attacker use *DNS rebinding* to bypass the browser's same-origin policy?
    a. By embedding the target site inside an iframe to load it in the victim's browser.
    b. By stealing session cookies using cross-site scripting (XSS).
    c. By presenting a forged TLS certificate to impersonate the target site.
    d. By using a very short TTL so the same hostname resolves to different IP addresses.

Correct answer: d

Short TTLs force the browser to resolve the same hostname to a new IP, tricking same-origin checks.

Incorrect choices:

(a) Iframes do not bypass same-origin.

(b) XSS is unrelated to DNS behavior.

(c) TLS spoofing requires certificate compromise.

15. Why is *IPsec ESP* more commonly used than AH?
    a. ESP provides confidentiality in addition to integrity.
    b. ESP requires no key management.
    c. ESP is faster because it avoids authentication overhead.
    d. AH works only with IPv6 and not with IPv4.

Correct answer: a

ESP protects both confidentiality and integrity, which makes it broadly useful.

Incorrect choices:

(b) Both ESP and AH require key management.

(c) ESP includes authentication options; it is not "faster because unauthenticated."

(d) Both AH and ESP work with IPv4 and IPv6.

16. How is *WireGuard* designed to be simpler than other VPNs?
    a. It avoids its own encryption layer and relies on applications to provide security, such as using TLS.
    b. It removes negotiation complexity and uses a fixed set of cryptographic primitives.
    c. It relies only on pre-shared keys instead of using a key exchange protocol.
    d. It disables roaming support to simplify state management.

Correct answer: b

WireGuard simplifies configuration by using a small, fixed modern crypto suite without negotiation.

Incorrect choices:

(a) WireGuard always encrypts tunnel traffic and does not rely on application-layer security like TLS.

(c) WireGuard does use preshared keys optionally but also performs key exchange using the Noise protocol.

(d) WireGuard supports roaming and does not disable it.

17. What is a key difference between Transport Layer Security (TLS) and a Virtual Private Network (VPN)?
    a. TLS uses symmetric encryption, while VPNs use asymmetric encryption for data.
    b. TLS protects communication for specific applications, while a VPN protects all traffic.
    c. TLS requires certificates, while VPNs only require a password.
    d. TLS provides integrity but not confidentiality, while VPNs provide both.

Correct answer: b

TLS protects individual application streams; VPNs protect all network-layer traffic.

Incorrect choices:

(a) Both TLS and VPNs use symmetric ciphers for bulk encryption.

(c) VPNs may use certificates, passwords, or both.

(d) TLS provides confidentiality and integrity.

18. What is the purpose of network *segmentation* with firewalls?
    a. To eliminate the need for VLANs.
    b. To enable Wi-Fi connectivity.
    c. To block specific applications from establishing outbound connections.
    d. To limit lateral movement within the organization.

Correct answer: d

Segmentation limits how far attackers can move within a compromised environment.

Incorrect choices:

(a) Segmentation often uses VLANs rather than eliminating them.

(b) Wi-Fi connectivity does not require segmentation.

(c) Outbound blocking is a firewall feature but not the goal of segmentation.

19. What is the defining characteristic of a *signature-based* Intrusion Detection System (IDS)?
    a. It establishes a baseline of normal traffic and alerts on deviations.
    b. It validates strict adherence to protocol standards.
    c. It sits inline and drops suspicious packets.
    d. It compares traffic against a database of known attack patterns.

Correct answer: d

Signature-based IDS detects attacks by matching known malicious patterns.

Incorrect choices:

(a) Baseline deviation describes anomaly-based IDS.

(b) Standards compliance checking is unrelated to signature matching.

(c) Inline traffic blocking describes IPS, not IDS.

20. What is a key advantage of *host-based firewalls* over network firewalls?
    a. They handle higher throughput than core routers.
    b. They cannot be disabled by malware with administrative access.
    c. They provide visibility into all network traffic across the enterprise.
    d. They can block traffic based on the specific application or executable generating it.

Correct answer: d

Host-based firewalls can enforce rules tied to specific applications or executables.

Incorrect choices:

(a) Host firewalls are not designed for high-throughput routing.

(b) Malware with admin rights can disable them.

(c) Enterprise-wide visibility requires network-level monitoring.

21. What is the core principle of Zero Trust regarding trust verification?
    a. Trust is granted once at the network edge and persists.
    b. Internal users are trusted by default; external users require MFA.
    c. Devices with valid company certificates are trusted for all services.
    d. Identity and context must be explicitly verified for every access request, regardless of location.

Correct answer: d

Zero Trust requires continuous verification of identity, device state, and context for each request.

Incorrect choices:

(a) Zero Trust does not provide implicit session-long trust.

(b) Zero Trust does not distinguish “internal” vs. “external” networks.

(c) Device certificates alone do not grant unconditional trust.

22. A site sets the `HttpOnly` flag on its session cookie. What risk is it trying to reduce?
    a. Exposure of the cookie over unencrypted HTTP connections.
    b. Cross-site cookie sharing by third-party pages.
    c. Theft of the cookie through malicious JavaScript.
    d. Unauthorized reuse of the cookie by related subdomains (for example, `api.example.com`).

Correct answer: c

HttpOnly prevents client-side scripts from reading the cookie, reducing theft via XSS.

Incorrect choices:

(a) Protection during transmission is controlled by the Secure flag, not HttpOnly.

(b) Third-party cookie restrictions are handled by SameSite and Domain attributes.

(d) Subdomain access is controlled by the Domain attribute, not HttpOnly.

23. What makes CSRF (Cross-Site Request Forgery) possible?
    a. Attackers can create lookalike domain names.
    b. Browsers send a user's cookies whenever a request is made on their behalf.
    c. Cookies store passwords.
    d. Attackers are restricted by the same-origin policy when sending requests.

Correct answer: b

CSRF occurs because browsers automatically attach the user's cookies even when another site triggers the request.

Incorrect choices:

(a) Lookalike domains facilitate phishing, not CSRF.

(c) Cookies contain session identifiers, not passwords.

(d) SOP limits reading responses, not sending requests, so it does not stop CSRF.

24. What does CORS (Cross-Origin Resource Sharing) accomplish?
    a. Allows a website to explicitly permit certain cross-origin reads.
    b. Prevents SQL injection attacks.
    c. Allows attackers to load malicious JavaScript.
    d. Enables malicious JavaScript to upload cookies to an attacker's service.

Correct answer: a

CORS lets servers specify which external origins may read their resources.

Incorrect choices:

(b) SQL injection is unrelated to browser origin rules.

(c) CORS is a defensive mechanism, not an attack vector.

(d) CORS does not allow JavaScript to steal cookies.

25. What is the basic idea behind a cross-site scripting (XSS) attack?
    a. An attacker forges requests to trick the browser into automatically sending cookies.
    b. An attacker poisons DNS records, causing browsers to load the wrong IP address.
    c. An attacker injects a malicious script into content that a victim's browser executes.
    d. An attacker intercepts TLS traffic by spoofing certificates.

Correct answer: c

XSS occurs when user-supplied data becomes executable script in a victim's browser.

Incorrect choices:

(a) That describes CSRF.

(b) That describes DNS poisoning.

(d) This describes TLS interception, not XSS.

The end.