Department of Computer Science

Computer Security

Exam 3

April 28, 2025

Solutions & 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. Generative AI has made *spear phishing* attacks easier to carry out by:
   a. Crafting exploits based on newly disclosed vulnerabilities.
   b. Correlating personal data from leaked databases and crafting custom messages.
   c. Creating scripts that enable sending the same message out to tens of thousands of users.
   d. Discovering new vulnerabilities.

Correct answer: b
Generative AI can analyze personal data and generate highly personalized, convincing phishing messages, which is the hallmark of spear phishing.

a is incorrect because crafting exploits is not specific to spear phishing.

c is incorrect because that describes generic phishing, not spear phishing.

d is incorrect because discovering vulnerabilities is not the primary goal of spear phishing attacks.

2. Which of the following best describes a *backdoor* in the context of computer security?
   a. A program that records keystrokes and sends them to an attacker.
   b. A covert method for bypassing normal authentication or security mechanisms.
   c. A vulnerability in outdated software that allows code execution.
   d. A malicious web page that mimics a login form to steal passwords.

Correct answer: b
A backdoor is a hidden method, often intentionally placed, to bypass standard security controls and gain access to a system.

a: is incorrect because that describes a keylogger.

c: is incorrect because it describes an exploit, not a backdoor.

d: is incorrect because that is phishing, not a system-level access method.

3. A *signature* in anti-malware software is:
   a. A cryptographic hash used to verify file integrity.
   b. An encrypted identifier used to track malware infection.
   c. A digital signature to detect if the anti-malware database has been corrupted.
   d. A byte sequence believed to be unique to a piece of malware.

Correct answer: d
Malware signatures are typically unique byte patterns used to identify known malicious code.

a: is incorrect because a hash alone is too brittle for detection (e.g., small changes defeat it).

b: is incorrect because signatures are not encrypted.

c: is incorrect because that refers to verifying integrity of the antivirus software itself, not detecting malware.

4. What core lesson does Ken Thompson's *Reflections on Trusting Trust* illustrate?
   a. A program's source code must be inspected not only for malicious actions but also for vulnerabilities.
   b. Inspecting source code is sufficient only if the compiler source code has also been verified.
   c. A compiler can be malicious even if neither it nor the target program shows malicious code in source form.
   d. Virtualization layers can introduce undetectable changes in a program's execution environment.

Correct answer: c
Thompson's paper demonstrates that trusting source code is not enough — a compiler can be maliciously modified to insert backdoors into any program it compiles, including a clean version of itself. This creates a self-propagating vulnerability that cannot be detected by source inspection alone.t.

a: True, but misses the point of the paper – the danger is not in the program source but in the tools that compile it.

b: Still misses the point because even verifying the compiler source may not be enough if the existing compiler binary is malicious.

d: is unrelated to the topic and describes a different class of attack. The paper doesn't focus on runtime environments.

5. A *rootkit* is:
   a. Software designed to hide certain processes or files from detection.
   b. Software that installs itself in a bootloader to run before the operating system boots.
   c. A framework for building viruses and worms.
   d. Software that exploits privilege escalation vulnerabilities.

Correct answer: a

Rootkits are primarily designed for stealth, to conceal the presence of malware or unauthorized activity from system monitoring tools.

b: describes a bootkit, which is a type of rootkit, but more specific

c: refers to a malware development framework

d: describes a privilege escalation attack, not the concealment function of a rootkit

6. *Polymorphic* malware:
   a. Contacts an external server for instructions on what actions to take.
   b. Infects a system with a small piece of software that then downloads the full malware package from a server.
   c. Is a technique used to evade detection by modifying the code before propagating it.
   d. Uses multiple exploits to attempt to infiltrate a system.

Correct answer: c
Polymorphic malware rewrites or obfuscates its code each time it spreads, making signature-based detection more difficult.

a: describes command-and-control behavior, not polymorphism.

b: refers to a dropper or downloader, not polymorphic behavior.

d: describes a multi-vector or blended threat, not polymorphism.

7. Why can a *hypervisor rootkit* be more difficult to detect than traditional kernel-mode rootkits?
   a. It runs at the same privilege level as the kernel but hides in encrypted memory regions.
   b. It runs beneath the operating system, making the OS unaware of its manipulation.
   c. It installs as a user-mode program and hides its actions through process injection.
   d. It subverts the boot process to hijack system calls from the BIOS.

Correct answer: b

A hypervisor rootkit virtualizes the running OS, executing beneath it at a lower privilege level (ring -1), allowing it to intercept and modify system behavior invisibly.

a is incorrect because it misstates the privilege level and technique.

c is incorrect as it describes a user-mode stealth technique, not hypervisor-level.

d confuses BIOS attacks with hypervisor-based virtualization.

8. In a *CAM overflow attack*, what is the attacker's goal?
   a. Crash a switch by flooding it with malformed packets.
   b. Exhaust the switch's ability to track device locations, causing it to broadcast traffic.
   c. Reconfigure switch ports to forward traffic to a specific port.
   d. Redirect traffic by overwhelming the router's ARP cache.

Correct answer: b

In a CAM overflow attack, the attacker floods the switch with frames from many different (spoofed) source addresses. This exhausts the switch's memory for tracking which device is on which port. As a result, the switch acts like a hub, broadcasting frames, making it possible for the attacker to intercept traffic.

a: Describes a crash or DoS attack, but that's a different attack (e.g., malformed frame flooding or control plane attacks).

c: Implies reconfiguration or VLAN misbehavior (not CAM overflow). Attackers can cause behavior like port mirroring, but CAM overflow doesn't reconfigure ports.

d: ARP cache poisoning is a different attack on end-hosts, not the switch itself.

9. Which of the following best explains a common vulnerability in both ARP and DHCP protocols?
   a. Both are legacy protocols that encrypt messages using outdated algorithms.
   b. Both use a challenge-response mechanism that can be bypassed by attackers.
   c. Both produce responses without using cryptographic checks.
   d. Both accept responses without authenticating or validating the sender's identity.

Correct answer: d

Both ARP and DHCP lack authentication and assume that responses received on the network are trustworthy. There's no built-in mechanism to verify whether the sender is trustworthy. This makes them vulnerable to spoofing (ARP poisoning and rogue DHCP servers).

a: Neither protocol uses encryption at all, so the issue isn't weak encryption — it's the lack of any authentication.

b: Neither uses a challenge-response protocol. This applies to other protocols (like some authentication systems), but not ARP or DHCP.

c: While it's technically true that they don't use cryptographic integrity, this is more of a symptom than the core problem.

10. What is the key mechanism behind a DNS *rebinding attack*?
    a. Tricking the user into visiting a malicious site that returns an incorrect IP address for a trusted domain.
    b. Modifying the victim's hosts file to change domain resolution permanently.
    c. Changing a domain's IP address after the initial resolution to bypass same-origin restrictions.
    d. Returning invalid DNS records to crash the resolver or cause denial of service.

Correct answer: c

DNS rebinding works by changing the IP address associated with a domain name between successive DNS queries. This tricks the browser into allowing cross-origin requests to internal IPs.

a: Describes basic DNS spoofing, not rebinding.

b: Refers to hosts file tampering, not DNS. Hosts file modification is local and typically requires elevated privileges.

d: Describes a denial-of-service technique, not the purpose of rebinding. Unlikely in practice and unrelated to the goal of rebinding.

11. Why was BGP vulnerable to *prefix hijacking*?
    a. It was built on trust with no authentication for route advertisements.
    b. It relied on DNS for authentication.
    c. It lacked path length validation.
    d. It encrypted advertisements using outdated algorithms.

Correct answer: a
BGP trusts all announcements, allowing malicious ASes to hijack prefixes.
Wrong answers:
b – BGP does not use DNS for this purpose.
c – It does consider path length but prioritizes more specific prefixes.
d – BGP does not rely on encryption by default

12. How can DNS spoofing via *cache poisoning* be avoided?
    a. Use longer DNS TTLs.
    b. Redirect all DNS requests to local routers.
    c. Validate the UDP checksums.
    d. Use randomized query IDs and source ports.

Correct answer: d
Random query IDs and ports make it harder for attackers to guess the right values.
Wrong answers:
a – This does the opposite of preventing spoofing. Long TTLs make poisoned records persist longer in the cache, increasing the impact of a successful attack.
b – This does nothing to prevent spoofing. If the local router is also vulnerable or compromised, spoofed responses could still be injected.
c – UDP checksums only detect transmission errors, not spoofing. An attacker can easily compute a valid checksum.

13. Why are *reflection amplification* attacks typically carried out over UDP rather than TCP?
    a. UDP allows services to respond only to encrypted traffic.
    b. UDP does not require a connection, so responses can be redirected to spoofed IP addresses.
    c. UDP services ignore malformed headers, making spoofing easier.
    d. TCP is incompatible with amplification due to smaller headers.

Correct answer: b
Because UDP is connectionless, attackers can spoof the source IP address, causing services to send large replies to a victim who never made the request.

a: UDP services like DNS or NTP do not generally require encryption, and encryption is irrelevant to the attack.

c: While some UDP services may not validate carefully, the key issue is that UDP is connectionless, not header validation.

d: TCP amplification is *harder* because TCP requires a handshake (SYN, SYN-ACK, ACK). The issue is not header size, but the fact that TCP expects valid stateful connections before sending data.

14. What is the primary purpose of the *TLS handshake*?
    a. To compress the application data before encryption.
    b. To verify firewall traversal capabilities.
    c. To establish a shared secret and authenticate the server.
    d. To encapsulate IP packets for tunneling.

Correct answer: c
The TLS handshake negotiates cryptographic parameters, performs authentication (typically of the server), and establishes a shared symmetric key.

15. How does TLS differ from a VPN in terms of protection scope?
    a. TLS protects only DNS traffic, while VPNs protect all network traffic.
    b. TLS encrypts at the network layer; VPNs encrypt at the application layer.
    c. TLS requires certificates, while VPNs do not.
    d. TLS protects individual application sessions; VPNs protect all traffic from the device.

Correct answer: d
TLS is used to secure individual application-level connections (e.g., HTTPS), whereas VPNs secure all IP traffic between endpoints.

16. What is a typical use case for *VPN tunneling* that TLS alone cannot handle?
    a. Providing access to a corporate network from a remote location.
    b. Sending encrypted email between clients.
    c. Securing file transfers between two specific applications.
    d. Encrypting credentials sent in a login form.

Correct answer: a
VPNs are often used to connect remote users to a private network, which TLS alone (being application-specific) cannot do.

17. What is the primary security purpose of a *DMZ (Demilitarized Zone)*?
    a. To isolate internet-facing services from the internal network.
    b. To connect untrusted networks directly to internal systems.
    c. To block outbound traffic from internal users.
    d. To host backup and archival data.

Correct answer: a

A DMZ places public-facing services (like web servers) in a separate zone, preventing direct access to internal networks.

a is the opposite of best practice.

c might be part of egress filtering but isn't DMZ-specific.

d is irrelevant

18. Which of the following is a limitation of *Deep Packet Inspection* (DPI)?
    a. It can only filter based on IP addresses and ports.
    b. It cannot detect known attack signatures in unencrypted traffic.
    c. It can only monitor traffic entering the network, not leaving it.
    d. It cannot inspect the content of traffic protected by end-to-end encryption.

Correct answer: d

DPI cannot inspect inside HTTPS (TLS) or other encrypted payloads without breaking the encryption, limiting its visibility and effectiveness.

a This describes basic filtering.

b DPI includes signature-based detection.

c DPI can apply to any traffic direction.

19. What is the primary security challenge posed by *deperimeterization*?
    a. Systems become more reliant on encrypted traffic, which firewalls cannot inspect.
    b. DNS and DHCP services are no longer functional in segmented networks.
    c. User authentication must be offloaded to cloud providers.
    d. Internal and external networks blend, making it harder to define trust boundaries.

Correct answer: d

Deperimeterization refers to the breakdown of traditional network boundaries due to mobile devices, cloud services, and remote access, making perimeter-based defenses less effective.

a is a secondary challenge, not the core issue

b is unrelated

c is not a required outcome of deperimeterization

20. Which is a core principle in a *Zero Trust architecture*?
    a. All internet-to-internal network traffic must pass through a firewall.
    b. Users are trusted once authenticated inside the internal network.
    c. No device or user is trusted by default, regardless of network location.
    d. Security products must be thoroughly tested and audited before they are deployed.

> Correct answer: c
>
> Zero Trust assumes that threats exist both outside and inside the network, so no user or device is trusted by default. Trust must be continually evaluated.
>
> a: A perimeter-based security approach, not Zero Trust.
>
> b: This reflects outdated implicit trust models.
>
> d: Good practice, but not a defining principle of Zero Trust.

21. The *same-origin policy* enforces which of the following restrictions?
    a. Scripts can access data only if both pages share the same protocol, host, and port.
    b. Scripts can only manipulate elements on the same page.
    c. Scripts are not allowed to store data in the browser.
    d. Scripts must be served from the same server as the page.

> Correct answer: a
>
> The Same-Origin Policy limits access to content unless the protocol, domain (host), and port match, protecting data across different web origins.
>
> b: JavaScript can interact with other elements, including across frames if same-origin.
>
> c: Scripts can use cookies, localStorage, etc.
>
> d: Scripts can be loaded cross-origin; restriction applies to access, not loading.

22. What is a primary goal of *Cross-Origin Resource Sharing* (CORS)?
    a. Prevent mixed content from being loaded on secure pages.
    b. Enable browsers to reject tracking cookies.
    c. Allow JavaScript from one origin to access data from another with server consent.
    d. Isolate iFrame content from its parent.

> Correct: c
>
> CORS allows a server to specify which other origins may access its resources.
>
> a: Mixed content is handled via HTTPS rules, not CORS.
>
> b: Cookies are managed via cookie flags.
>
> d: iFrame isolation is managed by the Same-Origin Policy.

23. Why is *Cross-Site Request Forgery* (*CSRF*) a security problem?
    a. It abuses the user's authenticated session to perform unwanted actions.
    b. It injects malicious scripts into a victim's browser.
    c. It tricks users into downloading malware via email links.
    d. It forces a site to execute system commands.

Correct: a

CSRF uses the victim's browser and existing authentication (e.g., cookies) to execute commands unknowingly.

b: That's XSS.

c: That describes phishing or drive-by downloads.

d: That's command injection.

24. Why are *tracking pixels* considered a privacy concern?
    a. They encrypt user sessions.
    b. They inject malicious JavaScript into pages.
    c. They prevent users from opting out of data collection.
    d. They allow servers to log visits and send cookies invisibly.

Correct: d

Tracking pixels are tiny, often invisible images embedded in emails or web pages. When the image loads, it sends a request to a remote server, which can log the visit, record user-agent info, IP address, and drop or read cookies — all without the user's awareness. This enables covert tracking of user behavior across sites and emails.

a: They don't encrypt anything. Encryption of a session is usually handled by https (HTTP over TLS).

b: They are passive content and do not contain JavaScript. While they can accompany scripts, the pixel itself is not responsible for JavaScript execution.

c: This is a result, not a mechanism. They don't inherently block opt-outs, though they may operate without consent or transparency. Opt-out mechanisms (like Do Not Track or cookie banners) are broader concerns.

25. Which condition would most likely enable a reflected xss (Cross-Site Scripting) attack?
    a. A login form stores submitted usernames in a database and displays them on a profile page.
    b. A web application includes a user-provided value in a search results page without escaping it.
    c. A content delivery network serves JavaScript files from multiple domains.
    d. A browser extension injects a script that disables third-party cookies.

Correct answer: b

Reflected XSS occurs when untrusted user input is immediately used in the HTML response without proper escaping, such as dynamically including it in the page without sanitization. If unescaped input is reflected in the response, JavaScript can be injected via the URL.

a: This describes stored XSS.

c: While third-party JavaScript can be dangerous, it's not a direct cause of reflected XSS.

d: This relates to browser behavior and privacy, not XSS.

The end.