Department of Computer Science

Computer Security

Exam 1

February 24, 2025

Exam 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. Which scenario is an example of a *privacy violation* but not a confidentiality breach?
   (a) A social media company collecting user location data without consent.
   (b) A hacker gaining access to classified government files.
   (c) An employee leaking customer financial data to a competitor.
   (d) A hospital's database being accessed by an unauthorized doctor.

   (A) **Correct**: Privacy is about controlling data collection and use, and this case involves unauthorized collection rather than unauthorized access.
   (B) Incorrect: Unauthorized access to classified files is a confidentiality breach.
   (C) Incorrect: Leaking financial data involves breaking confidentiality.
   (D) Incorrect: Unauthorized access to medical records is a breach of confidentiality.

2. The *CIA Triad* is a set of principles for information security. The letters stand for Confidentiality, Integrity, and:
   (a) Authorization.
   (b) Authentication.
   (c) Availability.
   (d) Accountability.

   The CIA Triad consists of Confidentiality (protecting data from unauthorized access), Integrity (ensuring data accuracy), and Availability (keeping systems accessible to authorized users).

3. Which of the following describes an *integrity attack*?
   (a) An employee sends sensitive company data to a competitor.
   (b) A hacker alters log files to erase evidence of a breach.
   (c) A power outage disrupts access to an online banking system.
   (d) A virus spreads across a corporate network, causing system slowdowns.

   (A) Incorrect: This is a confidentiality breach.
   (B) **Correct**: Integrity is compromised when log files or other critical data are altered.
   (C) Incorrect: This affects availability, not integrity.
   (D) Incorrect: This might affect availability, but does not necessarily change the accuracy of data.

4. Which of the following best describes an *attack vector*?
   (a) A security control designed to protect against attacks.
   (b) A strategy for recovering from a cyberattack.
   (c) A report listing all vulnerabilities in an organization's network.
   (d) The method or pathway an attacker uses to exploit a vulnerability.

   (A) Incorrect: Security controls defend against attacks but are not attack vectors themselves.
   (B) Incorrect: Recovery strategies deal with incident response, not attack vectors.
   (C) Incorrect: Vulnerability reports identify security flaws but do not describe how they are exploited.
   (D) **Correct**: An attack vector is the method or route attackers use to exploit vulnerabilities, such as phishing emails or open ports.

5. Which of the following is an example of an *attack surface*?
   (a) All publicly accessible web applications of an organization.
   (b) An email with a malicious attachment sent to an employee.
   (c) A exploit used against an unpatched system.
   (d) A security patch that fixes a software flaw.

(A) **Correct**: The attack surface includes all exposed points where an attacker can attempt to breach a system.

(B) Incorrect: A malicious email is an attack vector, not the attack surface.

(C) Incorrect: An exploit is a tool used in an attack, not an entire attack surface.

(D) Incorrect: A security patch reduces an attack surface but is not an attack surface itself.

6. Which of the following violates *Kerckhoffs's Principle*?
   (a) Using AES with a 128-bit key instead of a 256-bit key
   (b) Using a proprietary encryption algorithm.
   (c) Relying on keeping the keys secret rather than keeping the algorithm secret.
   (d) Publishing encryption protocols for peer review and security analysis.

(A) Incorrect: AES is a well-tested public algorithm that aligns with the principle, with either 128 or 256-bit keys.

(B) **Correct**: Keeping an encryption algorithm secret violates Kerckhoffs's Principle because security should not depend on obscurity.

(C) Incorrect: This aligns with the principle, as secrecy should reside in the key.

(D) Incorrect: Public cryptographic review strengthens security and follows Kerckhoffs's Principle.

7. A cryptanalyst observes that the frequency distribution of letters in the ciphertext closely matches the typical frequency distribution of those letters in French text. Which type of cipher was most likely used for encryption?
   (a) Monoalphabetic substitution cipher.
   (b) Polyalphabetic substitution cipher.
   (c) One-time pad cipher.
   (d) Transposition cipher.

(A) Incorrect: A monoalphabetic substitution cipher replaces letters but does not alter their frequency distribution, so frequency analysis can break it.

(B) Incorrect: A polyalphabetic cipher (e.g., Vigenère) spreads letter frequencies across multiple substitutions, making the distribution different from English plaintext.

(C) Incorrect: A one-time pad results in completely random ciphertext, with no resemblance to the original plaintext's frequency distribution.

(D) **Correct**: A transposition cipher rearranges the positions of letters without changing their individual frequencies, meaning the frequency distribution remains the same as the original plaintext.

8. Which of the following ciphers most closely attempts to simulate the functionality of a *one-time pad*?
   (a) ChaCha20, a stream cipher.
   (b) DES, a Feistel cipher.
   (c) AES, a substitution-permutation network cipher.
   (d) Caesar, a shift cipher.

(A) **Correct**: ChaCha20, a stream cipher, because stream ciphers generate a pseudorandom key stream, similar to a one-time pad (though weaker, since it is not truly random and can be reused).

(B) Incorrect: DES, a Feistel cipher, because DES is a block cipher, not a stream cipher, and does not attempt to simulate a one-time pad.

(C) Incorrect: AES, an SP-network cipher – Incorrect, because AES is also a block cipher and does not generate a continuous key stream like a one-time pad.

(D) Incorrect: Caesar, a shift cipher – Incorrect, because the Caesar cipher uses a fixed shift, making it completely deterministic and far from the randomness of a one-time pad.

9. Which of the following ciphers does *NOT* require *padding*?
   (a) Camellia, a Japanese cipher with a Feistel network structure.
   (b) ARIA, a South Korean standard cipher that uses a substitution-permutation network.
   (c) Playfair, a classical digraph substitution cipher.
   (d) RC4, a stream cipher that was previously used in Wi-Fi protocols.

(A) Incorrect: because Camellia is a block cipher and must use padding for messages that are not a multiple of the block size.

(B) Incorrect: because ARIA is a block cipher and requires padding when message lengths are not aligned with the block size.

(C) Incorrect: because Playfair encrypts letter pairs (digraphs) and pads single-letter endings with filler characters (e.g., 'X') if needed, making it a form of manual padding.

(D) **Correct**: because stream ciphers encrypt data bit-by-bit or byte-by-byte and do not require padding like block ciphers do.

10. Why is the *one-time pad* (OTP) not commonly used in practical cryptographic applications?
    (a) Its key size makes key distribution impractical.
    (b) It can be broken using frequency analysis.
    (c) It is vulnerable to quantum computing attacks.
    (d) Ciphers such as AES provide better security.

(A) **Correct**: Its key size makes key distribution impractical. – OTP requires a key as long as the message, and securely distributing such large keys is impractical for most real-world applications.

(B) It can be broken using frequency analysis. – Incorrect – OTP provides perfect secrecy, meaning it does not reveal any statistical patterns that could be exploited by frequency analysis.

(C) It is vulnerable to quantum computing attacks. – Incorrect – Unlike RSA and ECC, OTP remains secure even against quantum computers because it does not rely on mathematical complexity but rather on true randomness.

(D) Ciphers such as AES provide better security. – Incorrect – AES is widely used because it is practical, but OTP, when used correctly, provides perfect secrecy, which no modern cipher (including AES) can claim. However, OTP is impractical due to key management issues, not because of weaker security.

11. If an attacker can break a 64-bit key in 292 years using brute force, approximately how much longer would it take to brute-force a *128-bit key*, assuming the same computational power?
    (a) Approximately 600 years.
    (b) 2^{64} times longer, or roughly 5.4×10^{28} years.
    (c) 64 times longer, or about 18,688 years.
    (d) It would take the same amount of time with quantum computers.

(A) and (C) Incorrect: the time increase is not linear; it is exponential.

(D) Incorrect: Even quantum computers would not make brute-force feasible for 128-bit keys (Grover's algorithm only halves the exponent, making AES-128 roughly as hard as brute-forcing a 64-bit key today).

(B) **Correct**: 2^{64} times longer, or roughly 5.4×10^{28} years. Since 128-bit keys are twice as long as 64-bit keys, the key space increases by 2^{64} times. Given that breaking a 64-bit key takes 292 years, breaking a 128-bit key would take $292\times2^{64}\approx5.4\times10^{28}$ years, which is far longer than the age of the universe (~13.8 billion years).

12. What is the primary purpose of an Initialization Vector (IV) in encryption?
    (a) Maximize confusion and diffusion in the ciphertext.
    (b) Ensure that the same plaintext encrypted with the same key produces different ciphertexts.
    (c) Extend the plaintext so it aligns with the block size of the cipher.
    (d) Build a table of substitutions for the substitution-permutation network of a cipher.

(A) Incorrect: While IVs contribute to randomness, their primary purpose is not to maximize confusion and diffusion (this is more related to substitution-permutation networks and key scheduling).

(B) **Correct**: IVs prevent identical plaintexts from encrypting to the same ciphertext under the same key, reducing patterns and improving security.

(C) Incorrect: Padding, not IVs, ensures plaintext aligns with block size. IVs do not alter plaintext length.

(D) Incorrect: IVs do not build substitution tables; they introduce randomness to block cipher modes like CBC or CTR.

13. What is the primary purpose of the *Diffie-Hellman algorithm*?
    (a) To encrypt and decrypt messages using asymmetric keys.
    (b) To allow two parties to securely establish a shared secret key over an insecure channel.
    (c) To hash messages to ensure integrity.
    (d) To provide digital signatures for authentication.

(A) Incorrect: Diffie-Hellman is not an encryption algorithm; it only establishes a shared secret.

(B) **Correct**: Diffie-Hellman enables two parties to create a shared secret that can be used for symmetric encryption, even if an attacker is eavesdropping.

(C) Incorrect: Hashing functions (e.g., SHA-256) are used for integrity, not key exchange.

(D) Incorrect: Digital signatures (e.g., RSA, ECDSA) are used for authentication, not key exchange.

14. How are public and private keys used where Alice wants to send a message securely to Bob that only he can read?
    (a) Alice encrypts a message using Bob's public key, and Bob decrypts it with his private key.
    (b) Alice encrypts the message with her private key and Bob decrypts it with his public key.
    (c) Alice and Bob agree on a key to use and then use that chosen key for encryption and decryption.
    (d) Alice encrypts the message with her private key and Bob decrypts it with Alice's public key.

(A) **Correct**: This ensures confidentiality, as only the recipient (who holds the private key) can decrypt the message.

(B) Incorrect: Decrypting with Bob's public key makes no sense.

(C) Incorrect: This describes symmetric encryption, not public-key cryptography.

(D) Incorrect: Anyone may have access to Alice's public key. Encrypting with a private key is used in digital signatures, not for confidentiality.

15. What is the primary characteristic of a *hybrid cryptosystem*?
    (a) It encrypts data with multiple keys to guard against the theft of a subset of those keys.
    (b) It enhances security by adding message authentication codes (MACs) to encrypted data.
    (c) It uses public key cryptography for key exchange and symmetric cryptography for data encryption.
    (d) It uses multiple layers of encryption for greater security in case one algorithm is found to be weak.

(A) Incorrect: Hybrid cryptosystems do not necessarily use multiple keys for the same data; instead, they use asymmetric encryption to protect the symmetric key.

(B) Incorrect: While MACs enhance integrity, they are not a defining feature of hybrid cryptosystems

(C) **Correct**: Hybrid cryptosystems use symmetric encryption for data encryption (because it is fast) and asymmetric encryption for secure key exchange. This approach combines efficiency with strong security.

(D) Incorrect: Hybrid cryptosystems do not simply layer encryption multiple times—they use asymmetric encryption for key exchange and symmetric encryption for data protection

16. How is *quantum computing* expected to impact modern cryptographic algorithms?
    (a) It will render symmetric encryption algorithms like AES and 3DES completely insecure.
    (b) It will make RSA and ECC public key algorithms ineffective.
    (c) It will make both symmetric and public-key encryption equally vulnerable.
    (d) It will have no significant impact on cryptographic security.

(A) Incorrect: Symmetric encryption (e.g., AES) remains secure, but Grover's algorithm reduces the effective security by half, meaning AES-128 security drops to AES-64 equivalent. However, doubling the key size (e.g., AES-256) mitigates this risk, unlike with RSA and ECC.

(B) **Correct**: Quantum computers running Shor's algorithm will be able to efficiently break RSA, ECC, and other asymmetric encryption methods, making them insecure.

(C) Incorrect: Public-key cryptography is far more vulnerable than symmetric cryptography to quantum attacks.

(D) Incorrect: Quantum computing will significantly impact cryptography, especially in public-key encryption and digital signatures.

17. What is the main benefit of *forward secrecy* in cryptographic communications?
    (a) It ensures that past encrypted sessions remain secure even if the long-term private key is compromised
    (b) It allows previously encrypted messages to be decrypted if the private key is later recovered.
    (c) It prevents attackers from performing a brute-force attack on a symmetric encryption key.
    (d) It guarantees protection against future quantum computing attacks.

(A) **Correct**: Forward secrecy ensures that session keys are generated uniquely for each session, preventing an attacker from decrypting past communications even if they later obtain the server's private key.

(B) Incorrect: Forward secrecy does not allow past messages to be decrypted—it prevents this from happening if a key is compromised.

(C) Incorrect: Brute-force resistance is a general property of strong encryption, not a direct result of forward secrecy.

(D) Incorrect: Forward secrecy does not eliminate key exchange; instead, it ensures that key exchange methods generate unique session keys.

18. Why is *collision resistance* an important property of cryptographic hash functions?
    (a) It makes it impossible to find the original input given a hash value.
    (b) It prevents an attacker from modifying an encrypted message without knowing the encryption key.
    (c) It makes finding two inputs with the same hash infeasible, ensuring data integrity.
    (d) It ensures that hash functions are always faster than using symmetric encryption algorithms.

(A) Incorrect: This describes preimage resistance, not collision resistance.

(B) Incorrect: Hash functions do not provide encryption; they ensure data integrity, not confidentiality.

(C) **Correct**: It makes finding two inputs with the same hash infeasible, ensuring data integrity. Collision resistance is critical for data integrity, digital signatures, and authentication, as hash collisions can be exploited to forge messages.

(D) Incorrect: Hash functions are generally faster than encryption, but this is not related to collision resistance.

19. Which of the following is a key difference between *Message Authentication Codes* (MACs) and *digital signatures*?
    (a) MACs use asymmetric cryptography, while digital signatures use symmetric cryptography.
    (b) MACs generate a fixed-length hash of a message, while digital signatures encrypt the message.
    (c) MACs use a shared secret key, while digital signatures provide non-repudiation by using public-key cryptography.
    (d) MACs require a trusted third party (Certificate Authority), while digital signatures do not.

(A) Incorrect: This is reversed—MACs use symmetric cryptography, and digital signatures use asymmetric cryptography.

(B) Incorrect: Digital signatures do not provide encryption; they ensure authenticity and integrity.

(C) **Correct**: MACs rely on symmetric keys for authentication and message integrity, meaning only parties with the shared key can verify the message. Digital signatures use asymmetric cryptography (e.g., RSA, ECDSA) and provide non-repudiation, meaning the sender cannot deny signing the message.

(D) Incorrect: Digital signatures require a Certificate Authority (CA) in many cases, while MACs do not, but this is not their fundamental difference.

20. In the Needham-Schroeder protocol, what role does the *trusted third party* play?
    (a) It acts as a trusted central storage place for the public keys of all users.
    (b) It encrypts all messages between Alice and Bob.
    (c) It generates a session key and securely distributes it to both parties.
    (d) It provides non-repudiation for the established session.

(A) Incorrect: The KDC does not store public keys in this protocol; it manages symmetric session keys.

(B) Incorrect: The KDC does not encrypt all communications—Alice and Bob encrypt messages after receiving the session key.

(C) **Correct**: It generates a session key and securely distributes it to both parties. The trusted third party (key distribution center, KDC) creates a session key and sends encrypted versions of it to both Alice and Bob, ensuring only the intended recipient can decrypt it.

(D) Incorrect: The protocol does not provide non-repudiation, as it relies on symmetric encryption where both parties share the same key.

21. Why does Kerberos use *timestamps* instead of nonces in its authentication process?
    (a) To increase the randomness of session keys, making brute-force attacks harder.
    (b) To prevent replay attacks by ensuring authentication messages are valid only for a short time.
    (c) To prevent users from having to remember their session keys after authentication.
    (d) To ensure that messages are delivered in the correct sequence during the authentication protocol.

(A) Incorrect: Timestamps do not directly increase key randomness; they prevent replay attacks.

(B) **Correct**: Unlike Needham-Schroeder, Kerberos uses timestamps instead of nonces to prevent replay attacks, reducing the number of messages exchanged between parties.

(C) Incorrect: Kerberos uses session tickets to avoid frequent authentication, but timestamps do not store keys. The architecture of Kerberos that splits it into a Ticket Granting Service & Authentication Service allows users to authenticate once for an extended time (usually a login session) and cache a session key to the TGS. Timestamps have no bearing on this.

(D) Incorrect: Timestamps help detect replay attacks, not enable them.

22. Which of the following is an example of multi-factor authentication (MFA)?
    (a) A password and a security question.
    (b) A fingerprint scan and facial recognition
    (c) A PIN and a username.
    (d) A password and a one-time passcode sent via SMS.

(A) Incorrect: A password and security question are both "something you know," meaning this is not MFA.

(B) Incorrect: Fingerprints and facial recognition are both "something you are," making it not MFA.

(C) Incorrect: A PIN and username are both "something you know," so this is single-factor authentication.

(D) **Correct**: A password and a one-time passcode (OTP) sent via SMS. This method combines "something you know" (password) and "something you have" (a phone to receive the OTP), meeting MFA requirements.

23. What is the primary reason for *salting* passwords before hashing them?
    (a) To increase the randomness of the password itself.
    (b) To allow users to reset their passwords securely.
    (c) To prevent attackers from using precomputed hash lookup tables to crack passwords.
    (d) To reduce the storage space required for hashed passwords.

(A) Incorrect: Salting does not change the actual password, only how it is hashed.

(B) Incorrect: Salting does not help with password resets; reset mechanisms are separate.

(C) **Correct**: To prevent attackers from using precomputed hash lookup tables (rainbow tables) to crack passwords. Salting adds a unique, random value to each password before hashing, ensuring that even if two users have the same password, their stored hashes are different.

(D) Incorrect: Salting slightly increases storage size by adding extra data (the salt).

24. What is the main reason *credential stuffing* attacks succeed?
    (a) Many websites do not require passwords for authentication.
    (b) Hackers can break even relatively strong passwords easily using brute force.
    (c) Some modern password hashing algorithms are weak.
    (d) Many users reuse the same password across multiple sites.

(A) Incorrect: Websites do require passwords, but users often reuse them across multiple accounts.

(B) Incorrect: Brute-force attacks take time, but credential stuffing avoids brute-force methods by using known passwords.

(C) Incorrect: Modern password hashing is strong, but credential stuffing bypasses hashing by using real, stolen passwords.

(D) **Correct**: Many users reuse the same password across multiple sites. If a password is exposed in a data breach, attackers can use it to access multiple accounts belonging to the same user (e.g., if someone reuses the same password for banking, email, and social media).

25. Why is TOTP (time-based one-time passwords) generally considered more secure than HOTP?
    (a) TOTP passwords are encrypted, while HOTP passwords are not.
    (b) HOTP passwords can be used multiple times, while TOTP passwords are for one-time use only.
    (c) TOTP-generated passwords expire after a short time, reducing opportunities for a replay attack.
    (d) TOTP does not rely on the security of a cryptographic hash function.

(A) Incorrect: Neither TOTP nor HOTP encrypts OTPs; they use hashing (HMAC).

(B) Incorrect: Both TOTP and HOTP generate one-time-use passwords, but HOTP passwords remain valid until used.

(C) **Correct**: TOTP-generated passwords expire after a short time, preventing replay attacks. Since TOTP passwords change every few seconds, attackers cannot reuse an intercepted OTP, reducing the risk of replay attacks.

(D) Incorrect: TOTP and HOTP both use hash functions.

The end.