Weaver Spring 2021 — CS 161 Computer Security — Final

For questions with **circular bubbles**, you may select exactly *one* choice on the Exam Tool.

◯ Unselected option

● Only one selected option

For questions with **square checkboxes**, you may select *one* or more choices on the Exam Tool.

■ You can select

■ multiple squares

For questions with a **large box**, you need to write your answer in the text box on the Exam Tool.

There is an appendix at the end of this exam, containing descriptions of all C functions used on this exam.

You have 170 minutes, plus a 10-minute buffer for distractions or technical difficulties, for a total of 180 minutes. There are 10 questions of varying credit (200 points total).

The exam is open note. You can use an unlimited number of handwritten cheat sheets, but you must work alone.

Clarifications will be posted on the Exam Tool.

Q1 *MANDATORY – Honor Code* (5 points)

Read the following honor code and type your name on Gradescope.

> I understand that I may not collaborate with anyone else on this exam, or cheat in any way. I am aware of the Berkeley Campus Code of Student Conduct and acknowledge that academic misconduct will be reported to the Center for Student Conduct and may further result in, at minimum, negative points on the exam and a corresponding notch on Nick's Stanley Fubar demolition tool.

Q2 *True/false* **(38 points)**

Each true/false is worth 2 points.

Q2.1 True or False: WPA2 is a protocol that translates IP addresses to MAC addresses.

◯ True ◯ False

Q2.2 True or False: Specification-based detection uses a blacklist.

◯ True ◯ False

Q2.3 True or False: If a pseudorandom number generator (pRNG) is secure, then an attacker who only sees the output of the pRNG is unable to learn its internal state.

◯ True ◯ False

Q2.4 True or False: Argon2 and PBKDF2 are appropriate algorithms to use when hashing and storing passwords in a database.

◯ True ◯ False

Q2.5 True or False: All forms of two-factor authentication (2FA) are resistant to phishing attacks.

◯ True ◯ False

Q2.6 True or False: Logging is a method of intrusion detection in which server log files are preserved so they can be asynchronously scanned to detect malicious activity.

◯ True ◯ False

Q2.7 True or False: Cryptographically secure MACs can be constructed using secure cryptographic hash functions.

◯ True ◯ False

Q2.8 True or False: When analyzing a cryptographic hashing scheme, preimage resistance (one-way) implies collision resistance.

◯ True ◯ False

Q2.9 True or False: Sending all DNS requests and responses over HTTPS can be used as an effective defense against censorship by preventing censors from knowing what websites you are visiting.

◯ True ◯ False

Q2.10 True or False: One-time pads, as long as they are used correctly, are secure against an adversary with infinite computational power.

◯ True ◯ False

Q2.11 True or False: TLS is able to prevent on-path attackers from learning metadata about your communications (e.g. request and response times, message length) by encrypting communications from a client to a server.

○ True ○ False

Q2.12 True or False: Publicly accessible stairs, walkways, and elevators can be considered part of the physical equivalent of a trusted computing base for airport security.

○ True ○ False

Q2.13 True or False: Clickjacking refers to a class of attacks where the attacker manipulates the user interface of a website to convince the user to click something that they did not intend to click on.

○ True ○ False

Q2.14 Consider two different detectors with the same false positive rate and false negative rate. Assume that false negatives and false positives are equally costly.

True or False: A website with a high volume of users but a low volume of attacks would benefit more from placing the detectors in series rather than in parallel.

○ True ○ False

Q2.15 True or False: Signature-based intrusion detection systems are good at identifying novel network attacks that have not been previously seen.

○ True ○ False

Q2.16 True or False: A primary advantage of a host-based intrusion detection system (HIDS) over a network-based intrusion detection system (NIDS) is that traffic can be analyzed in plaintext, since the host can access decrypted TLS traffic.

○ True ○ False

Q2.17 True or False: For organizations with a large number of network devices, network-based intrusion detection systems (NIDS) are easier to deploy and manage than host-based intrusion detection systems (HIDS).

○ True ○ False

Q2.18 True or False: The UDP protocol guarantees that packets are delivered to the destination server by detecting dropped packets and retransmitting them until they are acknowledged.

○ True ○ False

Q2.19 True or False: Alice decides to use Tor to protect herself from tracking and surveillance online. The Tor circuit contains three Tor nodes: an entry node, a relay node, and an exit node. Assume the nodes do not collude. The exit node knows Alice's IP address but not the domain of the website she is visiting.

○ True ○ False

Q2.20 True or False: `EvanBot is a real bot.` (0 points)

○ True ○ False

Q3 *Full Stack Security* **(17 points)**

Examtool is a test-taking website located at `https://exam.cs161.org/`. Assume that all network connections are made over HTTPS, unless otherwise specified.

Examtool uses session tokens for user authentication. Session tokens are stored as cookies with `Domain=exam.cs161.org` and no other cookie attributes (no `Secure` flag, no `HttpOnly` flag, `Path=/`).

When a student fills out or changes an answer, their browser makes a POST request to `https://exam.cs161.org/submit_question` with the student's updated answers.

Q3.1 (5 points) Which of the following attacks could allow an adversary to read the session token cookie? Select all that apply.

☐ (A) Reflected XSS attack at `https://exam.cs161.org/`

☐ (B) Stored XSS attack at `https://exam.cs161.org/`

☐ (C) Exploitable buffer overflow vulnerability in the student's browser

☐ (D) Root access to another device on the same Wi-Fi network that the student is using

☐ (E) Root access to the Wi-Fi access point that the student is using

☐ (F) None of the above

Q3.2 (4 points) For a question on an exam, Alice first submits "A" and then later changes her answer and submits "B". What could a MITM attacker between Alice's computer and the `exam.cs161.org` server do? Select all that apply.

☐ (G) Perform a DoS attack to prevent Alice from submitting an answer choice

☐ (H) Perform a replay attack to restore Alice's saved answer to "A"

☐ (I) Modify Alice's submitted answer choice to "C"

☐ (J) Run JavaScript in Alice's browser

☐ (K) None of the above

☐ (L) ——

Q3.3 (4 points) Suppose the MITM attacker has identified a vulnerability in HTTPS that allows them to arbitrarily read and modify data in transit without detection. Alice submits another answer. What could a MITM attacker between Alice's computer and the `exam.cs161.org` server do? Select all that apply.

☐ (A) Set cookie values for the page at `https://exam.cs161.org/`

☐ (B) Redirect Alice's browser to `https://evil.com/`

☐ (C) Access any file on Alice's computer

☐ (D) Change Alice's answer choice without detection

☐ (E) None of the above

☐ (F) ——

The following subparts are independent of the previous subparts.

An instructor uploads an exam to Examtool by applying some cryptography to the exam and sending it over an insecure channel.

Assumptions:

- m is the message to encrypt (i.e. the exam).
- $\|$ is concatenation.
- k_1 and k_2 are two different secret keys known only to the Examtool server and the instructor.
- $E(k, m)$ is the encryption function of an IND-CPA secure symmetric encryption scheme.
- $\text{MAC}(k, m)$ is a secure MAC function.

For each pair of cryptographic schemes, select the scheme with fewer potential vulnerabilities.

Q3.4 (2 points) Select the more secure scheme:

◯ (G) $C = C_1 \| C_2$, where $C_1 = E(k_1, m)$ and $C_2 = \text{MAC}(k_1, C_1)$

◯ (H) $C = C_1 \| C_2$, where $C_1 = E(k_1, m)$ and $C_2 = \text{MAC}(k_2, C_1)$

◯ (I) ——

◯ (J) ——

◯ (K) ——

◯ (L) ——

Q3.5 (2 points) Select the more secure scheme:

◯ (A) $C = C_1 \| C_2$, where $C_1 = E(k_1, m)$ and $C_2 = \text{MAC}(k_2, C_1)$

◯ (B) $C = E(k_1, m || \text{MAC}(k_2, m))$

◯ (C) ——

◯ (D) ——

◯ (E) ——

◯ (F) ——

Q4 *"Bank-Grade" Security* **(28 points)**

Bear Bank is using a third-party analytics service called ABtesters. To use it, the bank website includes a tag to load the ABtesters JavaScript library.

Bear Bank's website is located at `https://bearbank.com` and contains the following HTML:

```
<script src="https://cdn.abtesters.com/lib.js"></script>
<form name="login" action="/login" method="POST">
  <input type="text" name="username" />
  <input type="password" name="password" />
</form>
```

Q4.1 (5 points) In the same-origin policy, which of the following are used in determining the origin of an HTTP webpage? Select all that apply.

☐ (A) IP ☐ (C) Protocol ☐ (E) Request path

☐ (B) Port ☐ (D) Domain name ☐ (F) None of the above

Q4.2 (3 points) Bear Bank is concerned that the ABtesters JavaScript library could steal customer passwords from the login form if the JavaScript library were compromised. Is this a valid concern?

◯ (G) Yes, because the ABtesters JavaScript library executes with the origin of ABtester's webpage.

◯ (H) Yes, because the ABtesters JavaScript library executes with the origin of Bear Bank's webpage.

◯ (I) No, because `https://cdn.abtesters.com` uses a certificate that is signed for different domain name.

◯ (J) No, because the ABtesters JavaScript library can only execute specific JavaScript functions required for its basic functionality.

◯ (K) ——

◯ (L) ——

Q4.3 (3 points) Bear Bank decides to move the login form to `https://auth.bearbank.com` and embed it on the homepage (`https://bearbank.com/`) in an iframe.

Can the ABtesters JavaScript library running on Bear Bank's homepage steal customer passwords from the login form in the iframe?

◯ (A) Yes, because the ABtesters JavaScript library is running on the same page as the iframe.

◯ (B) Yes, because the ABtesters JavaScript library can execute any JavaScript it wants on the Bear Bank's homepage.

◯ (C) No, because the ABtesters JavaScript library is not developed by Bear Bank itself.

◯ (D) No, because the ABtesters JavaScript library has a different origin than the login form.

○ (E) ——

○ (F) ——

After a user successfully logs into their account, Bear Bank's website sets a `session_token` cookie to track the user's logged in status and allows users to transfer funds by making a GET request to `https://bearbank.com/transfer`.

Q4.4 (3 points) Which of the following cookie attributes would cause the `session_token` cookie to be sent in a request to `https://bearbank.com/transfer`? Select all that apply.

☐ (G) Domain=bearbank.com; Path=/transactions

☐ (H) Domain=bearbank.com; Path=/transfer; Secure

☐ (I) Domain=auth.bearbank.com; Path=/login; HttpOnly; Secure

☐ (J) None of the above

☐ (K) ——

☐ (L) ——

Q4.5 (3 points) Bear Bank realizes that there are no CSRF protections on the transfer form, which means attackers can steal money from users' accounts.

Which of the following are reliable defenses against CSRF attacks? Select all that apply.

Clarification during exam: Everyone will receive credit for this question because we did not specify what it means for a defense to be "reliable."

☐ (A) Add a random CSRF token to the transfer form each time the page loads

☐ (B) Check the referrer header on the server when processing the transfer form submission

☐ (C) Move the transfer form to an iframe hosted at `https://transfer.bearbank.com`

☐ (D) None of the above

☐ (E) ——

☐ (F) ——

The following subparts are independent of the previous subparts.

Tree Bank is different bank considering alternative security methods. Once a user is logged in, they can send HTTP requests to Tree Bank to make transactions. Each request contains a session token set by the server when the user first logged in. The requests do not contain any counters or timestamps. The requests are sent over HTTP (not HTTPS).

Eve is an on-path attacker.

Q4.6 (4 points) Eve observes a single request from EvanBot to Tree Bank, which contains a transaction. What can Eve do? Select all that apply.

☐ (G) Learn EvanBot's session token

☐ (H) Learn the contents of EvanBot's transaction

☐ (I) Learn EvanBot's password

☐ (J) Repeat EvanBot's transaction

☐ (K) None of the above

☐ (L) ——

Q4.7 (4 points) Assume that the user knows Tree Bank's public key, and Tree Bank's corresponding private has not been compromised. Suppose that Tree Bank requires that the user encrypt the entire HTTP request (including the transaction and token) with the ElGamal scheme from lecture before sending it to the bank.

Eve observes a single encrypted request from EvanBot to Tree Bank, which contains a transaction. What can Eve do? Select all that apply.

☐ (A) Learn EvanBot's session token

☐ (B) Learn the contents of EvanBot's transaction

☐ (C) Learn EvanBot's password

☐ (D) Repeat EvanBot's transaction

☐ (E) None of the above

☐ (F) ——

Q4.8 (3 points) What is the best way for the bank to defend against Eve's attacks, and what concept best describes the design flaw that allowed Eve to compromise EvanBot's requests?

◯ (G) Use DNSSEC. Don't build your own crypto.

◯ (H) Use TLS. Don't build your own crypto.

◯ (I) Use DNSSEC. Security is economics.

◯ (J) Use TLS. Security is economics.

◯ (K) Use DNSSEC. Least privilege.

◯ (L) Use TLS. Least privilege.

Q5 *TC sPeedy* **(15 points)**

To improve the speed of TCP, Alice suggests modifying the TCP protocol to allow data to be sent in the SYN and SYN-ACK packets during the 3-way handshake. The data in the SYN packet is immediately accepted by the server during the initial handshake (before the 3-way handshake finishes).

Clarification during exam: In sub-parts 4 and 5, in subsequent connections, the token is sent only in the SYN packet.

Q5.1 (3 points) Which of the following attacks are possible on this modified scheme? Select all that apply.

Clarification during exam: "Reliably"means that the attacker doesn't have to guess any values.

☐ (A) An off-path attacker can reliably inject packets after a connection has been established.

☐ (B) An off-path attacker can reliably execute a RST injection attack.

☐ (C) An off-path attacker can fool the server into accepting some spoofed data.

☐ (D) None of the above

☐ (E) ——

☐ (F) ——

Q5.2 (2 points) Alice notices that her modified scheme may be vulnerable to a DoS attack where the attacker sends a large data payload in the SYN packet without completing the TCP handshake. She proposes including SYN cookies as part of her modification.

TRUE or FALSE: SYN cookies provide a valid defense against the proposed DoS attack.

◯ (G) True ◯ (H) False ◯ (I) —— ◯ (J) —— ◯ (K) —— ◯ (L) ——

Q5.3 (4 points) Alice uses her modified 3-way handshake to form a TCP connection with a server. Assume that source port randomization is not in use.

What fields would an **on-path attacker** have to guess in order to inject some data from Alice's client to the server?

☐ (A) Client IP address and port

☐ (B) Server IP address and port

☐ (C) Server sequence number

☐ (D) Client sequence number

☐ (E) None of the above

☐ (F) ——

Alice modifies her protocol to use a cryptographic token. When a client and server connect for the first time:

1. The client sends a SYN packet with a token request.

2. The server generates a token using a MAC function with a key known only to the server and responds with a SYN-ACK packet to the client containing the token. The client and server both store the token.
3. The client responds with an ACK packet, as in normal TCP.

In subsequent connections, the client skips the 3-way handshake by sending the SYN packet with both the token and data (similar to Alice's modification from previous parts). The server verifies the value of the token and acknowledges both the SYN and the data. The server may begin sending data to the client before receiving the client's ACK as part of the handshake. The server rejects the SYN and data if the token is invalid.

Here are diagrams detailing the protocol:

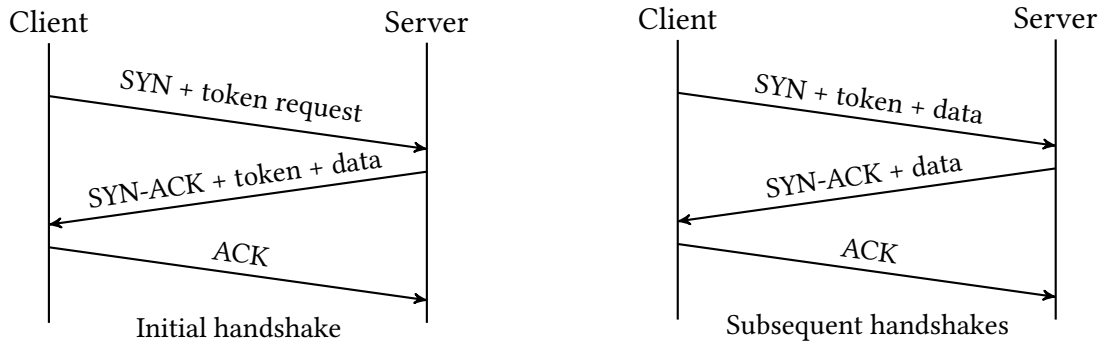

Q5.4 (3 points) Which of the following attacks on TCP becomes more difficult with the addition of the token? Select all that apply.

☐ (G) RST injection

☐ (H) Blind hijacking

☐ (I) MITM hijacking

☐ (J) None of the above

☐ (K) ——

☐ (L) ——

Q5.5 (3 points) A major issue with this protocol is that it is vulnerable to replay attacks, as an adversary can spoof a connection by replaying the token. A potential workaround is to modify the TTL (time to live) of the token. Name **one** benefit and **one** drawback of using a shorter TTL rather than a longer TTL.

Enter your answer in the text box on Exam Tool.

Q6 *UnicornBox v2* **(17 points)**

UnicornBox decides to implement 2-factor authentication (2FA).

The server stores a table of active codes with the following schema:

```
CREATE TABLE IF NOT EXISTS users (
    username TEXT,
    code TEXT,
    -- Additional fields not shown.
);
```

When a user wants to log in:

1. The user logs in by making a POST request with their username and password.
2. The server randomly generates a 10-digit numerical code and stores it in the `users` table.
3. The server sets a cookie with name = `auth_user` and value = the user's username in the user's browser. The server also sends a text to the user's phone with the code.
4. The user makes a GET request to `https://unicornbox.com/confirm?code=$code`, where `$code` is the code that was entered.
5. The server runs the SQL query `SELECT username FROM users WHERE code = '$code'`, where `$code` is the value submitted by the user.
6. The server checks that the value returned by the SQL query matches the username sent in the `auth_user` cookie in the request submitted by the user.

Clarification during exam: For all sub-parts, the user has an entry in the table.

Clarification during exam: "CalCentral" should be "UnicornBox" in the question text.

Clarification during exam: In step 1, the server verifies the password and will not proceed if the password is wrong.

Q6.1 (5 points) Assume that `evan` is the name of an account in UnicornBox with an entry in the `users` table.

Construct an input for `$code` that would cause the SQL query in step 5 to return `evan`.

Enter your answer in the text box on Exam Tool.

Q6.2 (4 points) How can you log in as `evan` without knowing their password? You may use `PAYLOAD` to reference your answer from the previous part.

Hint: You will need 2 steps. List both.

Enter your answer in the text box on Exam Tool.

Q6.3 (4 points) Which of these defenses would stop your exploit from above? Select all that apply.

☐ (A) Using SQL prepared statements

☐ (B) Rate limiting requests to the UnicornBox server

☐ (C) Putting the hash of the username in the cookie instead of the username

☐ (D) Using a 20-digit code instead of a 10-digit code

☐ (E) None of the above

☐ (F) ——

Q6.4 (4 points) Consider a modification to Steps 5 and 6. If there are any rows returned by the SQL query, then the verification succeeds without checking the value of the returned username. However, the server returns an error without executing the query if the format of the code is not exactly 10 numerical digits.

TRUE or FALSE: The modified scheme is no longer exploitable using SQL injection. Briefly justify (1 sentence) your answer.

◯ (G) True ◯ (H) False ◯ (I) —— ◯ (J) —— ◯ (K) —— ◯ (L) ——

Enter your answer in the text box on Exam Tool.

Q7 *Plaintext Feedback* **(15 points)**

Consider the "plaintext feedback" (PFB) mode where the encryption formula for ciphertext block C_i is given as follows:

$$C_0 = IV$$
$$C_i = E(K, P_i) \oplus C_{i-1}$$

E is AES encryption and D is AES decryption.

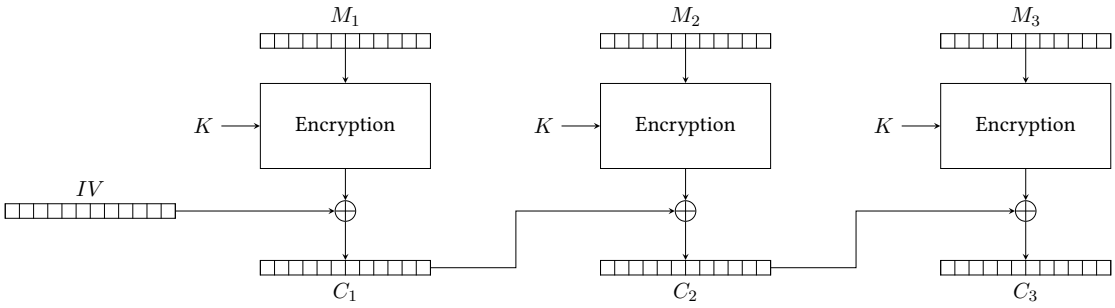

Clarification during exam: IVs are always randomly generated and never reused in this question.

Clarification during exam: M_i in the encryption diagram refers to plaintext block P_i.

Q7.1 (3 points) Which of these is the corresponding decryption equation?

- ◯ (A) $P_i = D(K, C_i \oplus C_{i-1})$
- ◯ (B) $P_i = D(K, C_i \oplus P_{i-1})$
- ◯ (C) $P_i = D(K, P_i \oplus P_{i-1})$
- ◯ (D) $P_i = E(K, C_i \oplus C_{i-1})$
- ◯ (E) $P_i = E(K, C_i \oplus P_{i-1})$
- ◯ (F) $P_i = E(K, P_i \oplus P_{i-1})$

Q7.2 (3 points) Alice and Bob are communicating using PFB mode. Alice encrypts and sends a 10-block message encrypted using PFB. Bob receives the message, but the 6th ciphertext block C_6 is lost in transmission. Which blocks of plaintext can Bob recover? Assume Bob is aware that C_6 was lost in transmission.

- ◯ (G) Bob can recover all blocks of the message.
- ◯ (H) Bob can recover all blocks up to and including P_6, but no block after that.
- ◯ (I) Bob can recover all blocks up to and including P_5, but no block after that.
- ◯ (J) Bob can recover all blocks except for P_6 and P_7.
- ◯ (K) Bob can recover all blocks except for P_6.
- ◯ (L) Bob cannot recover any block of the message.

Q7.3 (3 points) PFB mode is not IND-CPA secure. To prove this, the adversary will win the IND-CPA game against the challenger as follows:

First, the adversary sends two messages, P and P'. The first message P is 3 unique, randomly generated blocks, $P = P_1 \| P_2 \| P_3$. Which of the following values of P' would allow the adversary to win the IND-CPA game?

◯ (A) $P' = P'_1 \| P'_1 \| P'_1$, where P'_1 is a randomly generated block

◯ (B) $P' = P'_1 \| P'_2 \| P'_3$, where P'_1, P'_2, and P'_3 are unique, randomly generated blocks

◯ (C) $P' = P_1 \| P_2 \| P_3$

◯ (D) $P' = P'_1 \| P'_2 \| P'_3$, where P'_i is the same as P_i, but with the last bit flipped

◯ (E) $P' = P'_1 \| P'_2 \| P'_3$, where P'_i is the same as P_i, but every bit flipped

◯ (F) ——

Q7.4 (3 points) The challenger sends back a ciphertext $C = C_0 \| C_1 \| C_2 \| C_3$, which is an encryption of either P or P'. Describe a strategy that the adversary should use to deduce whether P or P' was encrypted that would allow them to win the IND-CPA game with probability greater than $\frac{1}{2}$.

Enter your answer in the text box on Exam Tool.

Q7.5 (3 points) Which of the following are true about PFB mode? Select all that apply.

☐ (A) Decryption is parallelizable

☐ (B) PFB provides integrity

☐ (C) The plaintext must be padded to a multiple of the block length

☐ (D) None of the above

☐ (E) ——

☐ (F) ——

Q8 *Caltopia DNS* **(21 points)**

EvanBot is trying to determine the IP address of `caltopia.com` with DNS. However, some attackers on the network want to provide EvanBot with the wrong answer.

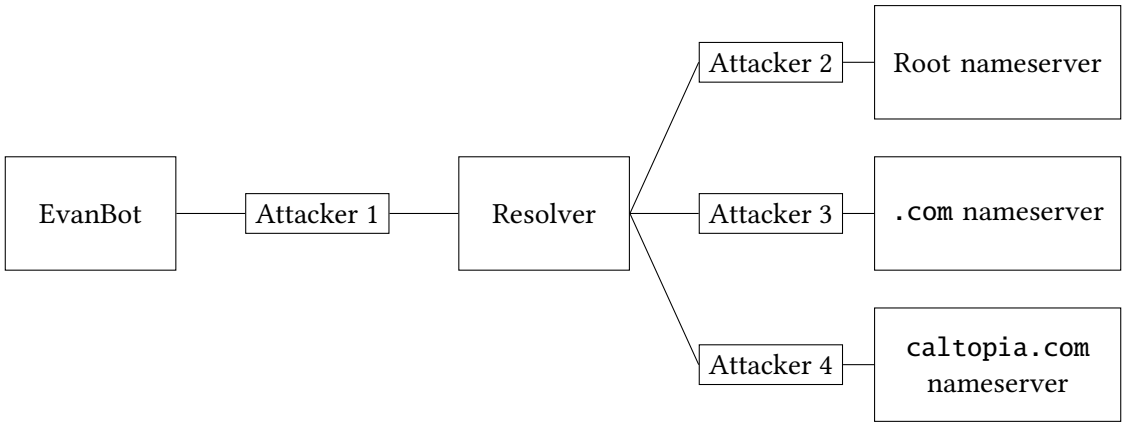

Assumptions:

- Each attacker is a man-in-the-middle (MITM) attacker between their two neighbors on the diagram above.
- No attackers can perform a Kaminsky attack.
- Standard DNS (not DNSSEC) is used unless otherwise stated.
- No private keys have been compromised unless otherwise stated.
- In each subpart, both EvanBot's cache and the local resolver's cache start empty.
- Each subpart is independent.

Clarification during exam: Assume that bailiwick checking is in use for this entire question.

In each subpart, EvanBot performs a DNS query for the address of `caltopia.com`.

Q8.1 (4 points) In this subpart only, assume the attackers only passively observe messages.

Which of the attackers would observe an `A` record with the IP address of `caltopia.com` as a result of EvanBot's query? Select all that apply.

☐ (A) Attacker 1 ☐ (C) Attacker 3 ☐ (E) None of the above

☐ (B) Attacker 2 ☐ (D) Attacker 4 ☐ (F) ——

Q8.2 (3 points) Which of the attackers can poison the local resolver's cached record for `cs161.org` by injecting a record into the additional section of the DNS response? Select all that apply.

Note: Attacker 1 has intentionally been left out as an answer choice.

☐ (G) Attacker 2 ☐ (I) Attacker 4 ☐ (K) ——

☐ (H) Attacker 3 ☐ (J) None of the above ☐ (L) ——

Q8.3 (4 points) Assume that the resolver and the name servers all validate DNSSEC, but EvanBot does not validate DNSSEC. Which of the attackers can poison EvanBot's cached record for `caltopia.com` by modifying the DNS response? Select all that apply.

☐ (A) Attacker 1 ☐ (C) Attacker 3 ☐ (E) None of the above

☐ (B) Attacker 2 ☐ (D) Attacker 4 ☐ (F) ——

Q8.4 (5 points) In this subpart only, assume the attackers only passively observe messages.

Assume that everyone validates DNSSEC. Which of the following records would Attacker 3 observe as a result of EvanBot's query? Select all that apply.

☐ (G) `DS` record with hash of the `.com` name server's public KSK

☐ (H) `DS` record with hash of the `caltopia.com` name server's public KSK

☐ (I) `A` record with the IP address of `caltopia.com`

☐ (J) `A` record with the IP address of the `caltopia.com` name server

☐ (K) `DNSKEY` record with the `.com` name server's public KSK

☐ (L) None of the above

Q8.5 (3 points) Assume that everyone validates DNSSEC, and the `caltopia.com` name server's private KSK has been compromised (i.e. all attackers know the `caltopia.com` name server's private KSK). No other private keys have been compromised.

Can EvanBot trust that they received the correct IP address of `caltopia.com`?

◯ (A) Yes, because the ZSK that signs the `A` record has not been compromised

◯ (B) Yes, because the trust anchor (the root's KSK) has not been compromised

◯ (C) No, because the compromised KSK can be used to sign a malicious `A` record

◯ (D) No, because the compromised KSK can be used to sign a fake ZSK that is used to sign a malicious `A` record

◯ (E) ——

◯ (F) ——

Q8.6 (2 points) TRUE or FALSE: DNSSEC prevents Attacker 4 from learning the IP address of `caltopia.com`.

◯ (G) True ◯ (H) False ◯ (I) —— ◯ (J) —— ◯ (K) —— ◯ (L) ——

Q9 *Mutuality* **(18 points)**

Recall the TLS handshake:

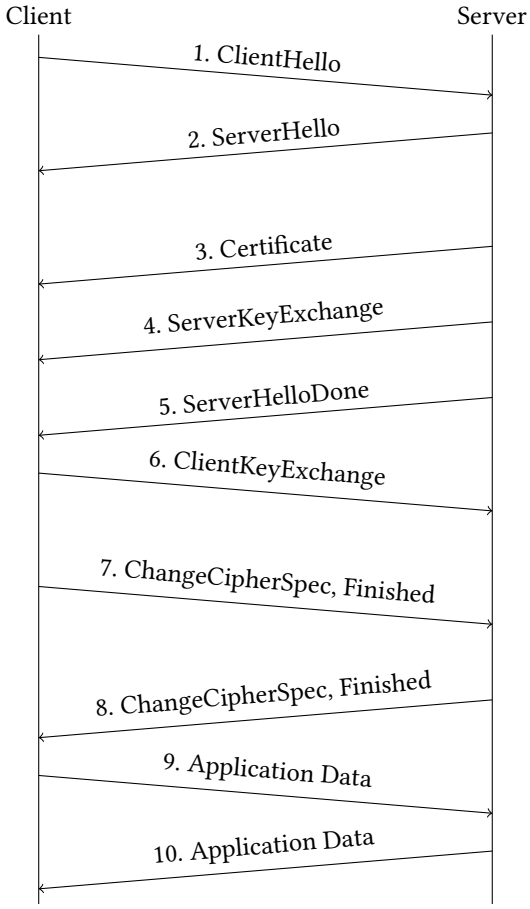

1. Client sends 256-bit random number R_b and supported ciphers
2. Server sends 256-bit random number R_s and chosen cipher
3. Server sends certificate
4. DH: Server sends $\{g, p, g^a \bmod p\}_{K_{\text{server}}^{-1}}$
5. Server signals end of handshake
6. DH: Client sends $g^b \bmod p$
   RSA: Client sends $\{PS\}_{K_{\text{server}}}$
   Client and server derive cipher keys C_b, C_s and integrity keys I_b, I_s from R_b, R_s, PS
7. Client sends MAC(dialog, I_b)
8. Server sends MAC(dialog, I_s)
9. Client data takes the form $\{M_1, \text{MAC}(M_1, I_b)\}_{C_b}$
10. Server data takes the form $\{M_2, \text{MAC}(M_2, I_s)\}_{C_s}$

In TLS, we verify the identity of the server, but not the client. How would we modify TLS to also verify the identity of the client?

Clarification during exam: All parts of this question refer to a modified TLS scheme designed to verify the identity of the client.

Q9.1 (3 points) Which of these additional values should the client send to the server?

- ◯ (A) A certificate with the client's public key, signed by the client's private key
- ◯ (B) A certificate with the client's public key, signed by the server's private key
- ◯ (C) A certificate with the client's private key, signed by a certificate authority's private key
- ◯ (D) A certificate with the client's public key, signed by a certificate authority's private key
- ◯ (E) ——
- ◯ (F) ——

Q9.2 (3 points) How should the client send the premaster secret in RSA TLS?

◯ (G) Encrypted with the server's public key, signed by the client's private key

◯ (H) Encrypted with the client's public key, signed by the server's private key

◯ (I) Encrypted with the server's public key, signed by a certificate authority's private key

◯ (J) Encrypted with the client's public key, signed by a certificate authority's private key

◯ (K) ——

◯ (L) ——

Q9.3 (3 points) EvanBot argues that the key exchange protocol in Diffie-Hellman TLS doesn't need to be changed to support client validation. Is EvanBot right?

◯ (A) Yes, because only the client knows the secret a, so the server can be sure it's talking to the legitimate client

◯ (B) Yes, because the server has already received and verified the client's certificate

◯ (C) No, the client must additionally sign their part of the Diffie-Hellman exchange with the client's private key

◯ (D) No, the client must additionally sign their part of the Diffie-Hellman exchange with the certificate authority's private key

◯ (E) ——

◯ (F) ——

Q9.4 (2 points) TRUE or FALSE: The server can be sure that they're talking to the client (and not an attacker impersonating the client) immediately after the client and server exchange certificates.

◯ (G) True ◯ (H) False ◯ (I) —— ◯ (J) —— ◯ (K) —— ◯ (L) ——

Q9.5 (3 points) At what step in the TLS handshake can both the client and server be sure that they have derived the same symmetric keys?

◯ (A) Immediately after the TCP handshake, before the TLS handshake starts

◯ (B) Immediately after the ClientHello and ServerHello are sent

◯ (C) Immediately after the client and server exchange certificates

◯ (D) Immediately after the client and server verify signatures

◯ (E) Immediately after the MACs are exchanged and verified

◯ (F) ——

Q9.6 (4 points) Which of these keys, if stolen individually, would allow the attacker to impersonate the client? Select all that apply.

☐ (G) Private key of a certificate authority

☐ (H) Private key of the client

☐ (I) Private key of the server

☐ (J) Public key of a certificate authority

☐ (K) None of the above

☐ (L) ——

Q10 *Storefront* (26 points)

Consider the following vulnerable C code:

```
1  void copy_string(char *dst, const char *src, size_t n) {
2      for (size_t i = 0; i < n + 1; i++) {
3          dst[i] = src[i];
4          if (src[i] == '\0') {
5              break;
6          }
7      }
8  }
9
10 void add_to_store(char *lst) {
11     char listing[256];
12
13     copy_string(listing, lst, 256);
14
15     printf("Contacting server to add: %s...\n", listing);
16     contact_server_and_wait(listing); // Implementation not shown.
17 }
18
19 void invoke(char *lst) {
20     add_to_store(lst);
21 }
22
23 int main(void) {
24     char buf[4096];
25     do {
26         fgets(stdin, buf, 4096);
27         invoke(buf);
28     } while (strcmp(buf, "exit") != 0);
29     return 0;
30 }
```

Definitions of relevant C functions may be found on the last page of this exam.

Assume you are on a little-endian 32-bit x86 system. Assume that there is no compiler padding or saved additional registers in all questions. For the first four parts, assume that **no memory safety defenses** are enabled.

Clarification during exam: The strcmp function is identical to strncmp, except that it doesn't take an argument n.

Clarification during exam: There are no vulnerabilities present outside of the provided source code (so there are no vulnerabilities in `contact_server_and_wait`).

Clarification during exam: Line 26 should be `fgets(buf, 4096, stdin)`.

Q10.1 (3 points) Which of the following memory safety vulnerabilities is present in this code?

◯ (A) Format string vulnerability

◯ (B) Signed/unsigned vulnerability

◯ (C) Off-by-one

◯ (D) None of the above

◯ (E) ——

◯ (F) ——

Q10.2 (3 points) Which of the following values on the stack can be partially or completely overwritten by the call to `copy_string` at line 13? Select all that apply.

Hint: Draw a stack diagram.

☐ (G) `listing`

☐ (H) SFP of `add_to_store`

☐ (I) RIP of `add_to_store`

☐ (J) None of the above

☐ (K) ——

☐ (L) ——

Q10.3 (6 points) Assume that the address of `listing` is `0xcffb5030`. Construct an input at Line 26 that would allow an attacker to execute malicious shellcode. You may reference the variable `SHELLCODE` as a 28-byte shellcode in your answer. Write your answer in Python 2 syntax (just like in Project 1).

Enter your answer in the text box on Exam Tool.

Q10.4 (3 points) Your exploit from above may not necessarily work with all possible addresses of `listing`. Provide **one** such address that would prevent your exploit from working. Write your answer in a format like `0xdeadbeef`.

Enter your answer in the text box on Exam Tool.

Q10.5 (3 points) Which of the following techniques could an attacker use to execute malicious shellcode if W^X and no other defenses are enabled? Select all that apply.

☐ (A) Return-oriented programming

☐ (B) ret2esp

☐ (C) Server-side request forgery

☐ (D) None of the above

☐ (E) ——

☐ (F) ——

Q10.6 (4 points) TRUE or FALSE: Stack canaries with no other defenses would prevent an attacker from executing malicious shellcode in this code (not necessarily using your exploit from above). Assume that all 4 bytes of the stack canary are randomized. Justify your answer.

○ TRUE ○ FALSE

Enter your answer in the text box on Exam Tool.

Q10.7 (4 points) TRUE or FALSE: ASLR with no other defenses would prevent an attacker from executing malicious shellcode in this code (not necessarily using your exploit from above). Justify your answer.

○ TRUE ○ FALSE

Enter your answer in the text box on Exam Tool.

Q11 *Cat* **(0 points)**

What is the name of Nick's gray cat?


Enter your answer in the text box on Exam Tool.

C Function Definitions

```
int strncmp(const char *s1, const char *s2, size_t n);
```

> The strncmp() function compares the first (at most) n bytes of two strings s1 and s2. It returns an integer less than, equal to, or greater than zero if s1 is found, respectively, to be less than, to match, or be greater than s2.

```
char *fgets(char *s, int size, FILE *stream);
```

> fgets() reads in at most one less than size characters from stream and stores them into the buffer pointed to by s. Reading stops after an EOF or a newline. If a newline is read, it is stored into the buffer. A terminating null byte ('\0') is stored after the last character in the buffer