Department of Electrical Engineering and Computer Science

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

6.858 Spring 2020

Quiz II

You have 120 minutes to answer the questions in this quiz. In order to receive credit you must answer each question as precisely as possible.

Some questions are harder than others, and some questions earn more points than others. You may want to skim them all through first, and attack them in the order that allows you to make the most progress.

If you find a question ambiguous, be sure to write down any assumptions you make. Be neat and legible. If we can't understand your answer, we can't give you credit!

Online quiz instructions

Write down your answers in the supplied ASCII text file template. Keep the template formatting unchanged.

Upload the answer file through the submission web site at the end of the quiz time. The filename should be `quiz2.txt`.

You have an additional 30 minutes after the official quiz end time to upload the quiz. You can upload your quiz answers as many times as you want; we will grade the last submission.

If you have questions during the quiz, please ask a private question through Piazza.

This is an open book, open notes, open laptop exam.
NO COMMUNICATION OR COLLABORATION DURING THE QUIZ.

I (xx/16)	II (xx/16)	III (xx/16)	IV (xx/14)	V (xx/32)	VI (xx/10)	VII (xx/4)	Total (xx/108)

I Paper Reading / Lecture Questions

1. [8 points]: Based on Max Burkhardt's guest lecture, which of the following statements are true? **(Circle True or False for each choice; we subtract points for incorrect answers.)**

A. True / False The cloud allows security engineers to scale making software secure.

B. True / False Network intrusion detection software is directly usable to detect intrusion of a web app.

C. True / False Detecting missing access-control checks in Web apps can be done by keeping statistics on Direct Object References, and monitoring those statistics.

D. True / False To develop intrusion detection for Web apps, it is important to have a strong background in machine learning techniques.

2. [8 points]: Based on Max Krohn's guest lecture about Keybase, which of the following statements are true about the Keybase chat application? **(Circle True or False for each choice; we subtract points for incorrect answers.)**

A. True / False The administrators of the servers for `keybase.com` can read the chat messages of all users.

B. True / False The administrators of the servers for `keybase.com` can modify Alice's chat message without Bob noticing it. (You can assume Bob's knows Alice's public key).

C. True / False By periodically publishing the Merkle root of keybase in the Bitcoin ledger, administrators of the servers for `keybase.com` cannot perform a permanent fork attack.

D. True / False If a user has a single keybase device and an attacker steals that device, then the user can revoke that device by logging into `keybase.com` and remove the stolen device.

II SSL/TLS

The SSL 3.0 handshake that is evaluated by Wagner and Schneier in the “Analysis of the SSL 3.0 protocol” is as follows:

```
 1. C -> S: Hello: client version, randomC, session_id, ciphers
 2. S -> C: Hello: server version, RandomS, session_id, ciphers
 3. S -> C: ServerCertificate: cert list
 4. S -> C: HelloDone
 5. C -> S: ClientKeyExchange: encrypt (pre_master_secret, PK_S)
 6. C -> S: ChangeCipherSpec
 7. C -> S: Finished, MAC({master_secret ++ msg 1,2,3,4,5}, C_key) (and encrypted)
 8. S -> C: ChangeCipherSpec
 9. S -> C: Finished MAC({master_secret ++ msg 1,2,3,4,5,7}, S_Key)  (and encrypted)
10. C -> S: encrypt(sign(data, MAC_secret), C_key)
```

3. [8 points]: The paper observes that a weakness in the protocol is that message 6 isn’t included in the MAC in messages 7 and 9, and on an authenticated-only SSL connection an adversary can launch a MITM attack: delete the ChangeChipherSpec message and strip off record-layer authentication fields from Finished message and session data. The authors argue that as long as the server doesn’t accept a Finished message before receiving ChangeChipherSpec message this MITM attack isn’t possible. Explain briefly why this small implementation change avoids the exploit?

4. [8 points]: Assume that the SSL implementation is modified as described in the previous question. Consider now a client that performs the handshake for an encrypted and authenticated SSL session with strong ciphers. Suppose the attacker doesn’t delete the ChangeCipherSpec message but replaces it with a ChangeCipherSpec listing a weak cipher. Would this attack allow the attacker to get the server to accept a weak cipher and learn the content of the SSL session by breaking the weak cipher? (Briefly explain your answer.)

III Spectre

Spectre V2 as described in the paper "Spectre attacks: Exploiting Speculative Execution" uses the following gadget to learn a secret v:

```
    if (off < sz) {
        v = array1[off]
        (*f)(v)
    }
```

The variable f is a pointer to a function. The idea is that the attacker trains the branch predictor so that on a miss on "sz" the processor will speculatively execute a function f of the attacker's choosing. (You may assume the attacker can reverse engineer how the branch predictor works.)

Assume there is a large array2 that the attacker can read and a function g:

```
void g(unsigned char v) {
  v1 = array2[v*v]
}
```

5. [8 points]: Briefly explain if (*f)() invoked g, is g suitable for learning the value v? (You may assume that the attacker has set up everything up correctly to handle alignment, cache lines, etc. correctly and that array2 is large enough.)

Consider the following function h:

```
uint32 z(unsigned char v) {
  // Efficiently return a pretty strong
  // cryptographic hash of v.
}

void h(unsigned char v) {
  v1 = array2[z(v)]
}
```

6. [8 points]: Briefly explain if h is suitable for learning the value v, and, if so, how the attacker would learn the value of v? (You can make the same assumptions as in the previous question.)

IV Messaging Security

Consider the following messaging protocol. If sender A (with private key SK_A) wants to send message m to recipient B (with public key PK_B), A generates a fresh symmetric key K, and transmits the following to B:

- Encrypt(PK_B, K)
- Sign(SK_A, K)
- MAC(K, m)
- m

7. [7 points]: Does the protocol provide authenticity (that is, B should be convinced that the message m came from A, assuming A is honest)? Explain why or why not.

8. [7 points]: Does the protocol provide deniability (that is, there is no way for B to provide cryptographic evidence that A must have sent m)? Explain why or why not.

V Web Security

This question is based on the Zoobar lab assignments, the web security lecture, and the “SoK: SSL and HTTPS: Revisiting past challenges and evaluating certificate trust model enhancements” paper.

Ben Bitdiddle runs a restaurant called Ben’s Bites in downtown Boston. Due to the coronavirus lockdown, he decides to set up a website for business continuity, where customers can make takeout and delivery orders. He needs to set up the website quickly, so he uses the entirety of the 6.858 Zoobar application as a template, and parks it under the domain name `bensbites.com`.

To start off, Ben patched all cross-site scripting vulnerabilities in the Zoobar application. In addition, the online login process has been modified to include the following "Secret Image" security control:

- At Registration: The user has to choose ONE (1) image out of a large library of preset images, which will be associated with user’s username. The choice of image is kept secret.
- At Login: Initially, the user is required to type in ONLY their username. The website then shows the password field and an image. The user is instructed to enter their password IF AND ONLY IF they see the image they chose at the time of registration. (See Figure 1)

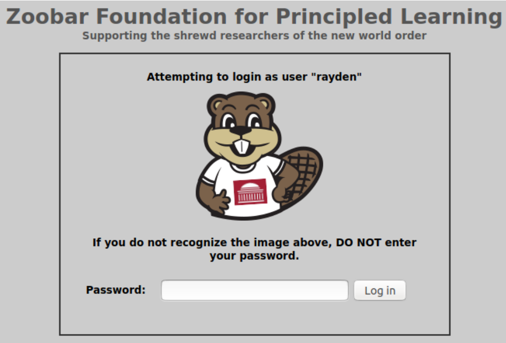

Zoobar Foundation for Principled Learning

Supporting the shrewd researchers of the new world order

Attempting to login as user "rayden"

If you do not recognize the image above, DO NOT enter your password.

Password:

Figure 1: Example `bensbites.com` login page after entering username.

Feature	Details
Public/Private Key	3072-bit RSA
Protocol Support	$\geq$ TLS 1.2
Revocation Information	CRL, OCSP
OCSP Stapling	NOT enabled
Public Key Pinning (HPKP)	NOT enabled
HTTPS-only Pinning (HSTS)	NOT enabled
`TLS_FALLBACK_SCSV`	ENABLED
TLS Compression	NOT enabled

Figure 2: Some TLS configuration details about `bensbites.com`.

Ben is very concerned about security, so he enforces TLS on all connections to the site as shown in Figure 2. Assume that the TLS cipher suites supported by `bensbites.com` are **cryptographically secure**, and that the certificate was issued by a well-known and trusted CA.

Rayden is a loyal customer of Ben’s Bites and has created an account on `bensbites.com`. One day, he decides to order lunch on the website. Unfortunately for Rayden, Noah carried out a man-in-the-middle (MITM) attack and stole Rayden’s password.

9. [8 points]: Briefly describe how and why Noah is able to carry out an MITM attack to steal Rayden's password when Rayden browses to `bensbites.com`.

You tell Ben Bitdiddle about the attack in Question 9, and he patches it so that it no longer works. Not to be outplayed, Noah creates a malicious site `evil.com` and manages to steal Rayden's password for a second time via a phishing attack!

10. [8 points]: How is Noah able to carry out a phishing attack on Rayden using `evil.com`, despite the fact that Rayden entered his password only when he saw the secret image he chose at the time of registration?

In response to Noah's phishing attack, Ben modifies the "Secret Image" security control to be more secure. Here is the modified control with the changes in **bold**:

- At Registration: The user has to choose ONE (1) image out of a large library of preset images, which will be associated with user's username. The choice of image is kept secret. **A browser cookie named `_secimg_nonce` will be set on the user's browser for `bensbites.com` with a 64-byte string value initialized from a CSPRNG, and the cookie value is bound to the user's username in the database.**
- At Login: Initially, the user is required to type in ONLY their username. **The user's browser then sends the `_secimg_nonce` cookie to the server which verifies that the cookie value matches the bound value in the database. If the cookie is not set or it does not match the bound value, then an error is raised and the login process halts.** Otherwise, the website shows the password field and an image. The user is instructed to enter their password IF AND ONLY IF they see the image they chose at the time of registration. (See Figure 1)

11. [8 points]: Does Ben's modified security control mitigate Noah's phishing attack in Question 10? **Explain why.** Be sure to explicitly state any assumptions for your security model.

In order to better utilize his server network bandwidth, Ben decides to enable TLS compression using the DEFLATE compression algorithm, which replaces repeated byte sequences with a pointer to the first instance of that sequence. The longer the repeated sequences, the higher the compression ratio. For example, the string "Oooooooooompa Looooooooompa" compresses to "Oooooooooompa L(-14,12)" since the repeated "ooooooooompa" sequence is 12 characters long and located at relative offset -14.

When Rayden successfully logs in to bensbites.com, Rayden's browser sends requests in the following format (shown before any encryption:

```
GET / HTTP/1.1
Host:  bensbites.com/zoobar/index.cgi
Cookie:  PyZoobar=c30e8a3b8ed2323bccfaf2224a35716d
...(other headers, HTML body)...
```

Here, the PyZoobar cookie is a session token that uniquely identifies a user to bensbites.com.

12. [8 points]: Describe how Noah, who is passively eavesdropping on encrypted TLS traffic between Rayden and bensbites.com, can steal Rayden's PyZoobar session cookie when Rayden visits evil.com. Assume that Rayden logged on to bensbites.com minutes before visiting evil.com.

VI SUNDR

Ben Bitdiddle is trying to optimize the serialized SUNDR protocol, as described in section 3.3 of the SUNDR paper. Ben thinks he can leverage the fact that the serialized SUNDR protocol allows only one user to make a change at a time. Instead of storing a full version vector in the version structure (see Figure 3 in the SUNDR paper), Ben's version structure stores just two elements of the version vector: the version counter for the user that created the version structure, and the version counter for the previous user that made a change. The rules for validating version vectors remain the same.

For example, if Alice made two changes, and then Bob made a change, Bob's version vector would contain 2 for Alice and 1 for Bob. If Bob made three more changes in a row, he would still keep Alice in his version vector, and Bob's version structure would contain 2 for Alice and 4 for Bob.

13. [10 points]: Does Ben Bitdiddle's variation provide fork consistency? Explain why or why not. Assume that only one user performs an operation at a time.

VII 6.858

We’d like to hear your opinions about 6.858. Any answer, except no answer, will receive full credit.

14. [2 points]: Are there any papers in the second part of the semester that you think we should definitely remove next year? If not, feel free to say that.

15. [2 points]: Are there topics that we didn’t cover this semester that you think 6.858 should cover in future years?

End of Quiz