Department of Electrical Engineering and Computer Science

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

6.858 Fall 2014

Quiz II

You have 80 minutes to answer the questions in this quiz. In order to receive credit you must answer the 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!

Write your name and submission website email address on this cover sheet.

This is an open book, open notes, open laptop exam.
NO INTERNET ACCESS OR OTHER COMMUNICATION.

This quiz is printed double-sided.

Please do not write in the boxes below.

I (xx/16)	II (xx/24)	III (xx/10)	IV (xx/16)	V (xx/8)	VI (xx/12)	VII (xx/8)	VIII (xx/6)	Total (xx/100)

Name:

Submission website email address:

I User authentication

1. [8 points]: A company named Vault wants to offer a secure, cloud-based backup system. When the user updates a local file, her Vault client opens a TCP connection to a Vault server, and uses the Diffie-Helman protocol to establish a secret symmetric key K with the server. Then, the client generates this string s:

$$s = \langle \text{ documentName, documentContent, userName, userPassword, randomNumber } \rangle$$

and sends the following message to the Vault server:

$$E_K(s, \text{HMAC}_K(s))$$

where $E_K(m)$ denotes encrypting message m using key K, and $\text{HMAC}_K(m)$ denotes computing an HMAC message authentication code of message m using key K.

The server decrypts the message, verifies the user's password, and verifies the integrity of the message using the HMAC. If all of the checks succeed, the server stores the document. If the server sees more than 10 messages with the wrong password, all future accesses to that account are blocked.

How can a network attacker reliably obtain the user's password?

2. [8 points]:

Suppose that a user wants to verify that a server knows the user's password. The user and the server engage in the following challenge/response:

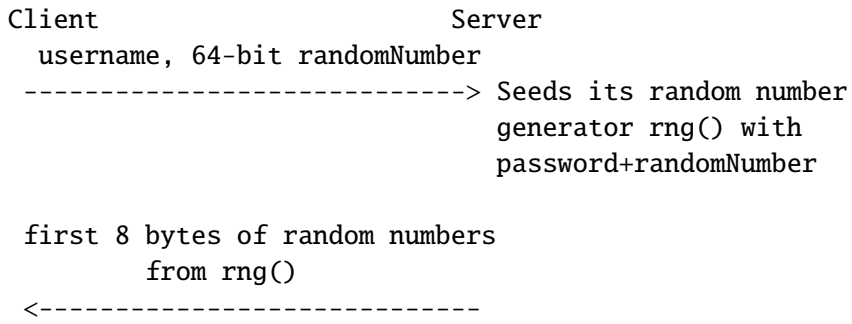

Everyone knows which algorithm the server uses for `rng()`, so the client can validate that the server's response contains the expected bytes.

Suppose that many different servers use this protocol. Further suppose that the attacker has a list of popular usernames, and a separate list of popular passwords. Explain how an attacker can abuse the protocol (without otherwise compromising any servers) to build a rainbow table of passwords and easily determine the passwords of many users.

II Tor and Private browsing

Imagine that the website `https://foo.com` embeds a JavaScript file from `http://attacker.com`. Suppose that a user's browser allows an HTTPS page to run JavaScript code fetched from an HTTP URL; however, the code cannot read or write any persistent client-side state. For example, the JavaScript code cannot read or write cookies, nor can it read or write DOM storage. The browser also ensures that the `attacker.com` web server cannot set client-side cookies using the `Set-Cookie` header, or receive client-side cookies in the HTTP request for the JavaScript file.

Suppose that the user visits `https://foo.com` once in private browsing mode, closes the tab, and then visits the site again in regular browsing mode; in both cases, the user's web traffic goes through Tor. The `attacker.com` web server would like to determine with high likelihood that the same user has visited the site twice. However, the attacker does not control any Tor nodes.

3. [8 points]: Why is it unlikely that the `attacker.com` server can use TCP fingerprinting to identify the user? Recall that TCP fingerprinting involves looking at TCP connection parameters, such as the initial TCP window size, TCP options, TCP flags, etc.

4. [8 points]: Describe a way that the attacker can fingerprint the user with a much higher likelihood of success.

5. [8 points]: Browsing the web through Tor can be slow. This is because user traffic is forwarded between volunteer computers that are scattered across the world, overburdened with traffic, and potentially situated behind slow network connections.

Suppose that CloudCo, a large technology company with datacenters scattered across the world, offers some of its machines to the Tor community for use as entry and exit nodes. CloudCo machines have plentiful RAM and CPU; CloudCo machines also have low latency, high-bandwidth network connections to all major ISPs. By using CloudCo machines as Tor entry and exit nodes, users could ensure that Tor congestion would only exist in the middle of a circuit.

Assume that CloudCo genuinely wants to help Tor users, and that CloudCo configures its machines to faithfully execute the Tor protocol. Why is it still a bad idea for users to employ CloudCo machines as entry and exit nodes?

III TaintDroid

6. [10 points]: TaintDroid defines various sources of taint, and a unique taint flag for each of those sources (e.g., `IMEI`, `PHONE_NUM`, `CAMERA`, etc.).

For the application code below, list the set of taint flags that each variable will have *after* each line of code has executed (for example, `"IMEI, CAMERA"` or `"PHONE_NUM"`). If a variable will not have any taint flags, write an Ø.

```
int x = android.getIMEI();                  x:_____________
int y = android.getPhoneNum();              y:_____________
int z;
if(x > 5000){
    z = 0;                                  z:_____________
}else{
    z = 1;                                  z:_____________
}

int arr[] = new int[2];                     arr:___________
arr[0] = x;                                 arr:___________
arr[1] = y;                                 arr:___________

char buf[] = android.getCameraPicture();    buf:___________
int val = buf[x%2];                         val:___________

x = 42;                                     x:_____________
```

IV Timing side channels

Consider the timing attack on OpenSSL RSA keys, described in the paper by Brumley and Boneh.

7. [8 points]: Suppose that OpenSSL was modified to never use Karatsuba multiplication (and instead always use "normal" multiplication), but was still using Montgomery multiplication as described in the paper. Would you expect the attack to still work with a few million queries? Explain why or why not.

8. [8 points]: Suppose that OpenSSL was modified to never use Montgomery multiplication, but was still using both Karatsuba and normal multiplication as described in the paper. Would you expect the attack to still work with a few million queries? Explain why or why not.

V Embedded device security

9. [8 points]: Ben Bitdiddle has a smartphone with an always-on voice recognition system, which runs any commands that it hears.

Alyssa P. Hacker wants to trick Ben's phone into running a command, but Ben turns off all wireless radios on the phone as a precaution to prevent Alyssa from breaking in over the network, and also keeps his phone in a locked sound-proof room in hopes of foiling Alyssa. After hearing Kevin Fu's guest lecture, Alyssa figures out how she can get Ben's phone to run a command of her choice, without breaking into Ben's room. What is Alyssa's plan?

VI Android security

Ben Bitdiddle is still excited about his phone's voice recognition system, and decides to set up a system that allows third-party applications to handle user voice commands (e.g., his music player might want to support a command like "next song", and his Facebook application might want to support a command like "post my current location on Facebook").

Ben's design runs on Android, and involves a single voice recognizer application that runs with access to the microphone. This voice recognizer transcribes whatever sounds it hears into ASCII text messages, and hands these messages off to third-party applications.

10. [12 points]: Describe a design for allowing the voice recognizer to securely send the transcribed messages to third-party applications, and allowing the third-party applications to securely receive them. Be as specific as possible; do not worry about software bugs. To help you structure your answer, consider the following:

First, what new permission labels do you propose to create? For each new permission label needed in your design, specify (1) its name, (2) who creates the label, and (3) what the type of the permission label is.

Continued on the next page.

Second, how should the voice recognizer securely send the transcribed messages? Describe (1) what component(s) the recognizer should have, (2) what permission(s) the application should ask for, (3) what label(s) should protect the recognizer's component(s), and (4) what other security-relevant steps the recognizer's code has to take.

Third, how should the third-party applications securely receive the messages? Describe (1) what component(s) each application should have, (2) what permission(s) the application should ask for, (3) what label(s) should protect the application's component(s), and (4) what other security-relevant steps the application code has to take.

VII Labs

11. [8 points]: Ben Bitdiddle's lab 6 code rewrites the following profile code:

```
var x = y[z];
```

into:

```
var sandbox_x = sandbox_y[sandbox_z];
```

Write down an exact piece of profile code that an adversary could use to call `alert(123);`.

VIII 6.858

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

12. [6 points]: Now that you are almost done with 6.858, how would you suggest we improve the class in future years?

End of Quiz

MIT OpenCourseWare
http://ocw.mit.edu

6.858 Computer Systems Security

Fall 2014

For information about citing these materials or our Terms of Use, visit: http://ocw.mit.edu/terms.