CSci 5271: Introduction to Computer Security

Exercise Set 4 **due: Wednesday, April 10th, 2019**

Ground Rules. You may choose to complete these exercises in a group of up to two students. Each group should turn in **one** copy with the names of all group members on it. You may use any source you can find to help with this assignment but you **must** explicitly reference any source you use besides the lecture notes or textbook. An electronic (plain text or PDF) copy of your solution should be submitted on the Canvas by 11:59pm on Wednesday, April 10th.

1. Protocol (an)droids. (20 pts) Two robots Artoo and C3-2-0 often fly on different starships and need to alert each other to their presence when their ships come in contact—otherwise they might accidentally blow each other up! They agree on a shared key K and a MAC algorithm that outputs 256-bit tags to use in the following protocol.

1. $A \longrightarrow C$: a random 256-bit string N_A and $\mathsf{MAC}_K(N_A)$.
2. C: on message n, t check that $\mathsf{MAC}_K(n) = t$, and if so, accept A, otherwise blow up the other party.
3. $C \longrightarrow A$: $\mathsf{MAC}_K(t)$.
4. A: on message t' check that $t' = \mathsf{MAC}_K(\mathsf{MAC}_K(N_A))$. If so, accept C, otherwise blow C up.

The idea here is that A proves it is A by correctly MACing N_A (which, if the key is secret, only A or C could do) and C proves it is C by MACing the MAC. But...

(a) A and C use this protocol for a while and then discover, to their dismay, that sometimes the evil galactic robo-emperor, E, has been successfully fooling C into believing it is A. Even supposing that robot-in-the-middle attacks are prevented by speed-of-light limitations or some other plot contrivance, what is a simple way for E to do this?

(b) A and C decide that one way to prevent the attack is for C to remember every value of N_A used in a previous challenge and reject if one is ever reused. Suppose E sees one authentication between A and C. How can it fool C into believing it is A as many times afterwards as it wants?

2. Random numbers with limited entropy. (30 pts) Alice, Bob, and Carol are employees of a company (in a small island nation) setting up an online casino website based on card games like blackjack. They realize that if users could predict the sequence of pseudorandom numbers used to deal cards, they could win reliably and hurt the company's bottom line. They've found a good cryptographically-strong pseudorandom number generation algorithm to use in the shuffling process, but they're having trouble deciding what to use as the seed when they initialize the generator at the start of each user's session.

(Following the usual good security design principles, they don't want the security of the games to depend on the choice of the pseudorandom generator or the shuffling algorithm being secret; they might also want to franchise their casino out in the future. But practically speaking, reverse-engineering those algorithms would be a significant effort, so attacks that worked without the attacker needing to do so would be particularly damaging.)

(a) Alice suggests seeding the PRNG with the time: specifically the date and time as returned by the Unix `time` system call, equal to the number of seconds since midnight, January 1st 1970 UTC. Explain why this is a bad idea by describing an easy attack.

(b) Bob suggests seeding the PRNG with the process ID of the login CGI script. Assuming this script runs once each time a player logs in, and process ID numbers are assigned sequentially in the range of 2 to 65535, describe an attack against this scheme.

(c) Carol suggests combining Alice and Bob's ideas by taking the time and the PID and XORing them together. But Alice points out a problem with this scheme that involves a user logging in once every second. Explain the details of her attack and why it's a problem.

(d) After the problems with their previous schemes, Alice, Bob, and Carol have called you in as a consultant. Suppose that because of the architecture of the system, the seed is required to be a deterministic function of the time in seconds and the PID. Propose a better combining function that takes these two pieces of information as input and produces a bit string (of any length) than can be used as a seed. Would it help if the function could also take another input that was like a key, fixed per-site but secret? Evaluate the security of your approach.

3. Cross-site scripting variations. (20 pts) There are a lot of different kinds of cross-site scripting vulnerabilities, but for space reasons we only covered one of them in hands-on assignment 2. This question covers another. Here's an excerpt from some Java code in the 2014 implementation of question 6 from hands-on assignment 2:

```
public class MACCookieServlet extends GroupServlet {
    @Override
    protected void doGet(HttpServletRequest req, HttpServletResponse resp)
                    throws ServletException, IOException {
        String username = req.getParameter("username");
        if (username == null)
             username = "";
        String digest_hex = ...;
        resp.setStatus(HttpServletResponse.SC_OK);
        resp.setContentType("text/html;charset=utf-8");
        resp.getWriter().print("User \"");
        resp.getWriter().print(username);
        resp.getWriter().print("\" is identified with the MAC ");
        resp.getWriter().print(digest_hex);
        resp.getWriter().println("\".");
    }
}
```

This code suffers a reflected XSS vulnerability: the `username` parameter is under the control of the untrusted user, and it is copied directly into the HTML output. So if it contained JavaScript, that code would run with the site's permissions. There's no similar problem with `digest_hex`, because the omitted code ensures that it contains only hexadecimal digits.

One way to fix this vulnerability would be to sanitize the contents of the `username` string using HTML entities; for instance, translating each "`<`" into "`<`". This is what we did for the newer version of the question (in PHP, we used `htmlspecialchars`). But suppose the programmer didn't know what library would contain a good implementation of that translation or was too lazy to implement it him or herself. What other simple change could you make to this code to avoid the cross-site-scripting danger?

4. TCP-Unfriendly. (30 pts) TCP's "congestion control" mechanism relies on *end-hosts* (i.e., users) to respond appropriately to network congestion by backing off their sending rate. One potential problem with this mechanism is what's called by economists the "tragedy of the commons." Suppose Alice knows that everyone else obeys TCP's congestion control mechanism. Then if she continues sending at the same rate, everyone else will slow down a little bit more and she will get better service from the network. So Alice has no motivation to obey TCP congestion control (other than the fact that not doing so involves finding or writing her own TCP stack—details, details) and in fact neither does anyone else. But if no one obeys the mechanism, the network (commons) becomes useless, which is the tragedy.

(a) Bob the Network Builder has an idea about how to solve this problem. He reasons that congested routers can see the *exact* state of a TCP connection. So if a particular connection does not slow down in response to dropped packets, the router can send a `RST` packet to each end of the connection. This will cause both ends of the connection to drop the connection, much more painful than just dropping an odd packet or two. From a *security* standpoint, what's the problem with Bob's idea—that is, if I'm an unscrupulous user intent on communicating at a high rate, can I circumvent this mechanism?

(b) When Bob realizes that reset packets aren't sufficient, he proposes a more direct approach: *blacklisting*. Under this idea, routers that notice TCP senders that don't respond to dropped packets appropriately will just stop routing packets for that sender. List several ways in which this is both ineffective against adversaries and a generally bad idea if adversaries got wind of it.