

Handout 1 (Security Engineering)

Much of the material and inspiration in this module is taken from the works of Bruce Schneier, Ross Anderson and Alex Halderman. I think they are the world experts in the area of security engineering. I especially like that they argue that a security engineer requires a certain *security mindset*. Bruce Schneier for example writes:

“Security engineers — at least the good ones — see the world differently. They can’t walk into a store without noticing how they might shoplift. They can’t use a computer without wondering about the security vulnerabilities. They can’t vote without trying to figure out how to vote twice. They just can’t help it.”

“Security engineering...requires you to think differently. You need to figure out not how something works, but how something can be made to not work. You have to imagine an intelligent and malicious adversary inside your system ..., constantly trying new ways to subvert it. You have to consider all the ways your system can fail, most of them having nothing to do with the design itself. You have to look at everything backwards, upside down, and sideways. You have to think like an alien.”

In this module I like to teach you this security mindset. This might be a mindset that you think is very foreign to you—after all we are all good citizens and not hack into things. I beg to differ: You have this mindset already when in school you were thinking, at least hypothetically, about ways in which you can cheat in an exam (whether it is about hiding notes or looking over the shoulders of your fellow pupils). Right? To defend a system, you need to have this kind mindset and be able to think like an attacker. This will include understanding techniques that can be used to compromise security and privacy in systems. This will many times result in insights where well-intended security mechanisms made a system actually less secure.

Warning! However, don’t be evil! Using those techniques in the real world may violate the law or King’s rules, and it may be unethical. Under some circumstances, even probing for weaknesses of a system may result in severe penalties, up to and including expulsion, fines and jail time. Acting lawfully and ethically is your responsibility. Ethics requires you to refrain from doing harm. Always respect privacy and rights of others. Do not tamper with any of King’s systems. If you try out a technique, always make doubly sure you are working in a safe environment so that you cannot cause any harm, not even accidentally. Don’t be evil. Be an ethical hacker.

In this lecture I want to make you familiar with the security mindset and dispel the myth that encryption is the answer to all security problems (it is certainly often part of an answer, but almost always never a sufficient one). This is actually an important thread going through the whole course: We will assume that encryption works perfectly, but still attack “things”. By “works perfectly” we

mean that we will assume encryption is a black box and, for example, will not look at the underlying mathematics and break the algorithms.¹

For a secure system, it seems, four requirements need to come together: First a security policy (what is supposed to be achieved?); second a mechanism (cipher, access controls, tamper resistance etc); third the assurance we obtain from the mechanism (the amount of reliance we can put on the mechanism) and finally the incentives (the motive that the people guarding and maintaining the system have to do their job properly, and also the motive that the attackers have to try to defeat your policy). The last point is often overlooked, but plays an important role. To illustrate this lets look at an example.

The questions is whether the Chip-and-PIN system with credit cards is more secure than the older method of signing receipts at the till. On first glance Chip-and-PIN seems obviously more secure and improved security was also the central plank in the “marketing speak” of the banks behind Chip-and-PIN. The earlier system was based on a magnetic stripe or a mechanical imprint on the card and required customers to sign receipts at the till whenever they bought something. This signature authorised the transactions. Although in use for a long time, this system had some crucial security flaws, including making clones of credit cards and forging signatures.

Chip-and-PIN, as the name suggests, relies on data being stored on a chip on the card and a PIN number for authorisation. Even though the banks involved trumpeted their system as being absolutely secure and indeed fraud rates initially went down, security researchers were not convinced (especially the group around Ross Anderson). To begin with, the Chip-and-PIN system introduced a “new player” that needed to be trusted: the PIN terminals and their manufacturers. It was claimed that these terminals are tamper-resistant, but needless to say this was a weak link in the system, which criminals successfully attacked. Some terminals were even so skilfully manipulated that they transmitted skimmed PIN numbers via built-in mobile phone connections. To mitigate this flaw in the security of Chip-and-PIN, you need to vet quite closely the supply chain of such terminals.

Later on Ross Anderson and his group managed to launch a man-in-the-middle attacks against Chip-and-PIN. Essentially they made the terminal think the correct PIN was entered and the card think that a signature was used. This was a more serious security problem. The flaw was mitigated by requiring that a link between the card and the bank is established at every time the card is used. Even later this group found another problem with Chip-and-PIN and ATMs which do not generate random enough numbers (nonces) on which the security of the underlying protocols relies.

The problem with all this is that the banks who introduced Chip-and-PIN managed with the new system to shift the liability for any fraud and the burden of proof onto the customer. In the old system, the banks had to prove that the customer used the card, which they often did not bother with. In effect, if fraud occurred the customers were either refunded fully or lost only a small

¹Though fascinating this might be.

amount of money. This taking-responsibility-of-potential-fraud was part of the “business plan” of the banks and did not reduce their profits too much.

Since banks managed to successfully claim that their Chip-and-PIN system is secure, they were under the new system able to point the finger at the customer when fraud occurred: they must have been negligent loosing their PIN. The customer had almost no means to defend themselves in such situations. That is why the work of *ethical* hackers like Ross Anderson’s group was so important, because they and others established that the bank’s claim that their system is secure and it must have been the customer’s fault, was bogus. In 2009 for example the law changed and the burden of proof went back to the banks. They need to prove whether it was really the customer who used a card or not.

This is a classic example where a security design principle was violated: Namely, the one who is in the position to improve security, also needs to bear the financial losses if things go wrong. Otherwise, you end up with an insecure system. In case of the Chip-and-PIN system, no good security engineer would claim that it is secure beyond reproach: the specification of the EMV protocol (underlying Chip-and-PIN) is some 700 pages long, but still leaves out many things (like how to implement a good random number generator). No human being is able to scrutinise such a specification and ensure it contains no flaws. Moreover, banks can add their own sub-protocols to EMV. With all the experience we already have, it is as clear as day that criminals were eventually able to poke holes into it and measures need to be taken to address them. However, with how the system was set up, the banks had no real incentive to come up with a system that is really secure. Getting the incentives right in favour of security is often a tricky business.

Of Cookies and Salts

Lets look at another example which should helps with understanding how passwords should be verified and stored. Imagine you need to develop a web-application that has the feature of recording how many times a customer visits a page. For example to give a discount whenever the customer visited a webpage some x number of times (say x equal 5). There is one more constraint: we want to store the information about the number of times a customer has visited inside a cookie. I think, for a number of years the webpage of the New York Times operated in this way: it allowed you to read ten articles per months for free; if you wanted to read more, you had to pay. My guess is it used cookies for recording how many times their pages was visited, because if you switched browsers you could easily circumvent the restriction about ten articles.

To implement our web-application it is good to look under the hood what happens when a webpage is requested. A typical web-application works as follows: The browser sends a GET request for a particular page to a server. The server answers this request. A simple JavaScript program that realises a “hello world” webpage is as follows:

```

1 var express = require('express');
2 var app = express();
3
4 app.get('/', function(request, response){
5     response.write('Hello World');
6     response.end()
7 });
8
9 // starting the server
10 app.listen(8000);

```

The interesting lines are 4 to 7 where the answer to the GET request is generated...in this case it is just a simple string. This program is run on the server and will be executed whenever a browser initiates such a GET request.

For our web-application of interest is the feature that the server when answering the request can store some information at the client's side. This information is called a *cookie*. The next time the browser makes another GET request to the same webpage, this cookie can be read by the server. We can use cookies in order to store a counter that records the number of times our webpage has been visited. This can be realised with the following small program

```

1 var express = require('express');
2 var cookie = require('cookie-parser')
3
4 var app = express();
5 app.use(cookie());
6
7 app.get('/', function(req, res){
8     var counter = parseInt(req.cookies.counter) || 0;
9     res.cookie('counter', counter + 1);
10    if (counter >= 5) {
11        res.write('You are a valued customer ' +
12                'visting the site ' + counter + ' times. ');
13    } else {
14        res.write('This is visit number '+ counter + '!');
15    }
16    res.end();
17 });
18
19 app.listen(8000);

```

The overall structure of this program is the same as the earlier one: Lines 7 to 17 generate the answer to a GET-request. The new part is in Line 8 where we read the cookie called counter. If present, this cookie will be send together with the GET-request from the client. The value of this counter will come in form

of a string, therefore we use the function `parseInt` in order to transform it into an integer. In case the cookie is not present, we default the counter to zero. The odd looking construction `... || 0` is realising this defaulting in JavaScript. In Line 9 we increase the counter by one and store it back to the client (under the name `counter`, since potentially more than one value could be stored). In Lines 10 to 15 we test whether this counter is greater or equal than 5 and send accordingly a specially grafted message back to the client.

Let us step back and analyse this program from a security point of view. We store a counter in plain text on the client's browser (which is not under our control). Depending on this value we want to unlock a resource (like a discount) when it reaches a threshold. If the client deletes the cookie, then the counter will just be reset to zero. This does not bother us, because the purported discount will just not be granted. In this way we do not lose us any (hypothetical) money. What we need to be concerned about is, however, when a client artificially increases this counter without having visited our web-page. This is actually a trivial task for a knowledgeable person, since there are convenient tools that allow one to set a cookie to an arbitrary value, for example above our threshold for the discount.

There seems to be no real way to prevent this kind of tampering with cookies, because the whole purpose of cookies is that they are stored on the client's side, which from the the server's perspective is a potentially hostile environment. What we need to ensure is the integrity of this counter in this hostile environment. We could think of encrypting the counter. But this has two drawbacks to do with the key for encryption. If you use a single, global key for all the clients that visit our site, then we risk that our whole "business" might collapse in the event this key gets known to the outside world. Then all cookies we might have set in the past, can now be decrypted and manipulated. If, on the other hand, we use many "private" keys for the clients, then we have to solve the problem of having to securely store this key on our server side (obviously we cannot store the key with the client because then the client again has all data to tamper with the counter; and obviously we also cannot encrypt the key, lest we can solve a chicken-and-egg problem). So encryption seems to not solve the problem we face with the integrity of our counter.

Fortunately, *hash functions* seem to be more suitable for our purpose. Like encryption, hash functions scramble data in such a way that it is easy to calculate the output of a has function from the input. But it is hard (i.e. practically impossible) to calculate the input from knowing the output. Therefore hash functions are often called *one-way functions*. There are several such hashing function. For example SHA-1 would hash the string "hello world" to produce

2aae6c35c94fcfb415dbe95f408b9ce91ee846ed

Another handy feature of hash functions is that if the input changes only a little, the output changes drastically. For example "iello world" produces under SHA-1 the output

d2b1402d84e8bcef5ae18f828e43e7065b841ff1

That means it is not predictable what the output will be from just looking at input that is “close by”.

We can use hashes in our web-application and store in the cookie the value of the counter in plain text but together with its hash. We need to store both pieces of data such we can extract both components (below I will just separate them using a “-”). If we now read back the cookie when the client visits our webpage, we can extract the counter, hash it again and compare the result to the stored hash value inside the cookie. If these hashes disagree, then we can deduce that the cookie has been tampered with. Unfortunately, if they agree, we can still not be entirely sure that not a clever hacker has tampered with the cookie. The reason is that the hacker can see the clear text part of the cookie, say 3, and also its hash. It does not take much trial and error to find out that we used the SHA-1 hashing functions and then graft a cookie accordingly. This is eased by the fact that for SHA-1 many strings and corresponding hashvalues are precalculated. Type, for example, into Google the hash value for “hello world” and you will actually pretty quickly find that it was generated by input string “hello wo1rd”. This defeats the purpose of a hashing functions and thus would not help us for our web-applications.

There is one ingredient missing, which happens to be called *salts*. Salts are random keys, which are added to the counter before the hash is calculated. In our case we need to keep the salt secret. As can be see in Figure 1, we now need to extract from the cookie the counter value and the hash (Lines 19 and 20). But before has the counter again (Line 22) we need to add the secret salt. Similarly, when we set the new increased counter, we will need to add the salt before hashing (this is done in Line 15). Our web-application will now store cookies like

```
1 + salt - 8189effef4d4f7411f4153b13ff72546dd682c69
2 + salt - 1528375d5ceb7d71597053e6877cc570067a738f
3 + salt - d646e213d4f87e3971d9dd6d9f435840eb6a1c06
4 + salt - 5b9e85269e4461de0238a6bf463ed3f25778cbba
...
```

These hashes allow us to read and set the value of the counter and give us confidence that the counter has not been tampered with. This of course depends on being able to keep the salt secret.

There is an interesting point to note with respect to the New York Times’ way of checking the number visits. Essentially they have their ‘resource’ unlocked at the beginning and lock it only when the data in the cookie states the allowed free number of visits are up. This can be easily circumvented by just deleting the cookie or by switching the browser. This would mean the New York Times will loose revenue whenever this kind of tampering occurs. In contrast, our web-application has the resource (discount) locked at the beginning and only unlocks it if the cookie data says so. If the cookie is deleted, well then the resource just does not get unlocked. No mayor harm will result.

```

1 var express = require('express');
2 var cookie = require('cookie-parser')
3 var crypto = require('crypto');
4
5 var app = express();
6 app.use(cookie());
7
8 var salt = 'secret key'
9
10 function mk_hash(s) {
11     return crypto.createHash('sha1').update(s).digest('hex')
12 }
13
14 function mk_cookie(c) {
15     return c.toString() + '-' + mk_hash(c.toString() + salt)
16 }
17
18 function gt_cookie(s) {
19     var splits = s.split("-", 2);
20     var counter = parseInt(splits[0])
21     var hash = splits[1]
22     if (mk_hash(counter.toString() + salt) == hash) {
23         return counter
24     } else {
25         return 0
26     }
27 }
28
29 app.get('/', function(req, res){
30     var counter = gt_cookie(req.cookies.counter) || 0;
31     res.cookie('counter', mk_cookie(counter + 1));
32     if (counter >= 5) {
33         res.write('You are a valued customer ' +
34             'visting the site ' + counter + ' times.');
```

Figure 1:

How to Store Passwords

While admittedly silly, the simple web-application in the previous section should help with the more important question of how passwords should be verified and stored. It is unbelievable that nowadays some systems still do this with passwords in plain text. The idea behind such passwords is of course that if the user typed in *foobar* as password, we need to verify it matches with the password that is stored for this user in the system. Doing this verification in plain text is really a bad idea, but evidence suggests it is still a widespread practice. I leave you to it to think about why verifying passwords in plain text is a bad idea.