CS 491/591 Security and Privacy Spring 2010 Lab 2b

50 points (lab 2a was the other 50)

Assigned: 1 April 2010

Due: Thursday, 15 April 2010, at 11:59pm

The purpose of this lab is to explore information flow and side-channel inference attacks in a systematic way.

You will turn in a single writeup for your group, with the names of the group members in alphabetical order. You may discuss the assignment at a high-level with other groups and trade snippets of code, but each group is expected to develop their own idle scan. You will also turn in any source code you wrote as a separate tar ball.

Remember that your purpose is to do all reconnaissance from the `192.168.33.0/24` network for all of lab 2. You should not attempt to use your lab 1 exploit on any machines at this point.

In lab 2a, you might have discovered that ports `8000-9000` for the `192.168.111.0/24` subnet were filtered by shasta, which sends an ICMP host prohibited errror and drops such packets rather than forwarding them. This rule is based on the network interface, not on the source IP address. The first thing you should discuss as a group is why the idle scans currently at your disposal (rstscan.pl and nmap's builtin idle scan) won't work and what the requirements for a new idle scan should be.

On one machine on the `192.168.111.0/24` subnet, somewhere between ports `8000` and `9000`, there are 50 ports that the capitalization service is listening on. Your job is to identify the machine and the fifty ports. Educated guesses about what type of machine would be offering such a service can help to narrow down your search.

Note that developing an idle scan and answering this question is only one part of the lab. You also need to do a thorough, systematic study of information flow in this setting. You can use Kemmerer's matrix methodology, Wray's clock methodology, covert flow trees, non-interference, or something new that you create. You must either find a new idle scan technique that fits the requirements or reasonably prove that in the scenario given the SYN cache idle scan is the only one possible. You must include any network features in your model that shasta and typical networks include (ARP, TCP/IP, ICMP, etc.). You may include other features (packet fragmentation, stateful firewall rules, etc.) and I'll add them to the shasta environment if you find something you think will be fruitful.

If you find a new idle scan technique, you can implement it with perl, python, or an nmap script. If not, you can implement the SYN cache technique that is in the Ensafi et al. technical report.

Your writeup should thoroughly document what you did and show me all the steps. For example, if you choose to use Kemmerer's matrix methodology you should show me the list of all subjects and objects and the initial rights you assigned to them, and then show me the transitive closure and highlight any new possibilities for idle scans. The presentation of your work is as important as the work itself, so assign a group member for this that will do a very neat job of it. You should also include a discussion of what technique you used to study information flow in this context and why you chose that technique over others. What limitations of the technique did you identify and how did you deal with them?

At the end of the writeup, you should state which group members showed up to which of the three in-class group meetings (April 2nd, 9th, and 14th), other meetings you had (not required, but encouraged), and a brief statement of how you divided up the work. Specialization is encouraged, but the discussion and process of carrying out the thorough, systematic study of information flow should involve all members of the group. A reasonable way to divide up the work might be, *e.g.*, "Hector did the implementation of the idle scan in perl, Jane collected all of the subjects, objects, and rights to build our matrix and wrote the code to calculate a transitive closure, and Jon did the writeup. All three group members discussed and contributed to our matrix model of information flow." Remember that the quality of the writeup is very important and that everyone should contribute to the information flow analysis part.

I plan to assign one grade for the whole group, but if a particular member is delinquent or only contributed something trivial then I may give that member a lower score or even a zero.