

An Internet X-Ray Machine for the Masses

Author : Paul Ohm

Date : June 12, 2015

Aldo Cortesi, et al., [mitmproxy](#).

Thank you to the Jotwell editors for indulging me as I stretch their [mission statement](#) (and quite possibly their patience) by highlighting not an article nor even a conventional work of scholarship but rather a piece of software as the “thing I like (lots)”: [mitmproxy](#), a tool created by [Aldo Cortesi](#) who shares [authorship credit](#) with Maximilian Hills and a larger “[mitmproxy community](#).”

mitmproxy does just what it says on the tin (assuming you know how to read this particular kind of tin). It’s a Man-In-The-Middle Proxy server for the web. In English, this means that this tool allows you to reveal, with finely wrought control, exactly what your browser is saying and to whom. It is an X-ray machine for the web, one which lays many of the Internet’s secrets bare. Let me extol the many virtues of this well-designed piece of software, and after I do that, let me explain why I think this strikes me as an important contribution to *legal* scholarship.

There are [many other tools](#) that do what mitmproxy does. Where mitmproxy shines relative to everything I have tried is the way it embraces both usability and power without compromising either.

Take usability first. Especially for Mac OS X users, mitmproxy is the single easiest tool of its kind I have encountered. Here is [what you need to do](#) to begin wiretapping yourself as you browse the web:

- Step 1: Install the OSX binary available at <https://mitmproxy.org/>
- Step 2: [Open a terminal window](#) and extract, find, and start the binary.¹
- Step 3: Open a browser and configure it to use the IP address and port of the computer running mitmproxy (probably 127.0.0.1 and 8080) as its web proxy.
- Step 4: Surf the web.

At this point, the mitmproxy display will fill with lots of http requests down the screen. The controls to navigate these requests are so intuitive they [require little documentation](#): arrows scroll up and down, enter reveals more detail about the current request, escape returns to the previous screen, etc.

By performing the steps above, the student or scholar of technology law or policy who has never operated a packet sniffer above can more deeply understand some of the secrets of web surveillance. mitmproxy is, most importantly, packet sniffing for the masses. For the first time, we are given a tool which is simple to understand, relatively easy to operate, free to download, and available to people lacking root access to their computers. These qualities make this a powerfully democratizing tool.

All of this makes mitmproxy also a wonderful tool for teaching. For three years, I have taught a course on “[The Technology of Privacy](#),” in which the students have spent an hour or two sniffing packets. Until this year, my students toiled with [Wireshark](#)—an old tool, but still the industry standard for packet sniffing. To say that Wireshark confused my students is an understatement. The semester-end reviews were replete with comments like, “Great class, but I have no idea what was going on with Wireshark.”

This year, I taught the same unit using mitmproxy. The experience could not have been more different. After walking through the steps above and watching a demo for two minutes, my students started monitoring their own web traffic,

needing no further guidance. My only instruction was “find something interesting,” and within five minutes, that’s exactly what they did.

Perhaps the most astonishing thing the tool makes easy is the sniffing of encrypted web traffic. Techies might scoff at my being impressed by this, because it’s almost tautological; that’s what a MITM proxy permits. But look again at how simply this has been implemented. Here are the steps required to permit the monitoring of encrypted traffic:

- Step 5: From your browser, visit mitm.it. (This won’t send you to an Italian webpage; mitmproxy intercepts the request and sends you its own content instead.)
- Step 6: Follow the simple instructions at that page.
- Step 7: Surf the encrypted web.

If all mitmproxy did was bring packet sniffing to the masses, it would still do plenty. But mitmproxy is not only easy-to-use, it is also so powerful and robust that it has become a serious tool of web-based forensics.

Take the work of Ashkan Soltani, who introduced me to mitmproxy. Ashkan is well-known in the privacy law community as the current Chief Technologist of the FTC. He made his first big splash as the technological brains behind many of the groundbreaking studies conducted by Julia Angwin and her fellow journalists at the Wall Street Journal, in the [“What They Know” series](#). The great impact of those studies—and what qualify them in my mind as scholarly research as much as investigative journalism—stems from the rigorously obtained and compellingly presented data revealing [third-party tracking of the web](#) and [invasive tracking of mobile apps](#). It is my understanding that at least some of these important results were obtained using mitmproxy.

Others have used mitmproxy to “slay dragons.” It is [credited with revealing privacy violations](#) in mobile apps. It has allowed researchers to [peer into opaque private APIs](#) to learn how companies are protecting their users’ secrets (spoiler: not always well).

There is too much more to praise in full about mitmproxy, so let me summarize the rest. It is released under a [GPL open source license](#) and [distributed via github](#), so anybody can tinker under the hood. It is [written in python](#), so you’re likelier to understand what you’re looking at under that hood. It allows you to “replay” web requests and responses from the past, giving you fine-tuned controls for testing. It lets you monitor the activity of mobile apps as seamlessly as web browsing. You can easily automate it.

All of this power can be used for evil as well as good, of course. If I trick your browser into using my mitmproxy, then [with a few lines of code](#), I can flip all of the images sent to the browser upside down or [replace all images with photos of kittens](#), or do something even more evil.

Finally, back to the question I started with: why does mitmproxy belong on a website dedicated to celebrating *scholarship*? mitmproxy is a scholarly tool or methodology, akin to R or logistic regression, something that too few legal scholars use and many more should embrace. That alone is probably enough to justify this review.

But in some sense, a packet sniffer is the key to my personal origin story as a scholar of Internet privacy. In my first job after college—helping develop and defend the networks of the RAND Corporation—in what I think was my first week on the job, I ran a packet sniffer—one much clunkier to use than mitmproxy—on our local network segment. Entirely by happenstance, the first screenful of packets [intercepted contained a packet revealing the RAND vice president’s username and password](#) in plaintext, right on my screen. I don’t think I ever closed an application as quickly as I did at that moment, and my manager (who was standing behind me) said, with a smile on his face, “we shall never speak of this again.” I can draw a direct line from that moment to many thoughts I have had and things I have written about Internet privacy.

We scholars of internet policy spend most of their time focused on the abstract and intangible. The things we

investigate flit through the aether (or ethernet) near the speed of light. There is value in finding ways to reify these abstractions into something closer to the tangible and concrete, the way sniffing tools like mitmproxy do. It is one thing to write about, say, privacy as an abstraction, it is another altogether to capture a password or set up a proxy server. Doing little things like this will remind us that what we are investigating is real and within our reach.

1. For the truly uninitiated, this step might require a bit more elaboration.
 1. On the mitmproxy website, click the link next to the big apple logo, which reads, “OSX (Mountain Lion and later)”.
 2. It should drop a file into your “Downloads” directory, which is probably the icon on your desktop dock next to your trash can. Click that icon.
 3. Click the file you just downloaded. It’ll be called something like “osx-mitmproxy-0.11.3.tar.gz”, although the version numbers may vary.
 4. Now, open the terminal (see link in step three above).
 5. Type: “cd Downloads/osx; ./mitmproxy” without the quotes and then enter.
 6. You should be running mitmproxy. (It should look like the screen depicted at <https://mitmproxy.org/>.)
 7. Macs running older versions of OS X might encounter errors at this point.

If you run linux, you should be able to figure out the install for yourself. If you run windows, you’re probably out-of-luck, although I hear good things about (but cannot vouch personally for) [Fiddler](#). [?]

Cite as: Paul Ohm, *An Internet X-Ray Machine for the Masses*, JOTWELL (June 12, 2015) (reviewing Aldo Cortesi, et al., mitmproxy), <http://cyber.jotwell.com/an-internet-x-ray-machine-for-the-masses/>.