Interview with Nick Weaver

RIK FARROW

Nick Weaver received a BA in astrophysics and computer science in 1995, and his PhD in computer science in 2003 from the University of California at Berkeley. Although his dissertation was on novel FPGA architectures, he was also highly interested in computer security, including postulating the possibility of very fast computer worms in 2001. In 2003, Nick joined ICSI, first as a postdoc and then as a staff researcher. His primary research focus is network security—notably, worms, botnets, and other Internet-scale attacks—and network measurement. Other interest areas have included both hardware acceleration and software parallelization of network intrusion detection, defenses for DNS resolvers, and tools for detecting ISP-introduced manipulations of a user's network connection. nweaver@icsi.berkeley.edu

Rik is the editor of *;login:* rik@usenix.org

I attended the first Enigma conference in January 2016 and was pleased by the quality of the talks as well as by the depth managed by speakers, who had just 20 minutes to make their points. While I had lots of favorite talks, such as those by Stefan Savage and Ron Rivest [1], I found myself wanting to dig a bit deeper into Nick Weaver's talk.

Nick talked about "The Golden Age of Bulk Surveillance," but in my mind his talk was really about what can be done with captured data. Nick walked the audience through the process of de-anonymization and just how easy it is when you have most traffic, particularly the metadata for the traffic [2].

Rik: Your PhD dissertation involved FPGAs, but somehow you got interested in TCP/IP and security at the network layer. Can you explain how that happened?

Nick: Well, during my dissertation I participated in the NIST Advanced Encryption Standard (AES) process, evaluating how easy the various five final candidates would map to high performance hardware.

But overall it was Code Red. When Code Red hit, the worldwide reaction was "13 hours, god that's fast." A friend of mine, Michael Constant, and I were sucking down sodas and going "13 hours, god that's slow. We're computer people, we should be able to do it faster than that."

So we started sketching out concepts that became the "Warhol Worm" concept: how to efficiently infect all the vulnerable machines in 15 minutes [3]. Once that happened, then you have to think about various automated defenses. So I came into security during the dawn of the "Worm Era," where high speed, broad malcode became a thing.

Rik: There really was something like the Warhol Worm, a worm that abused some Windows RPC over UDP, wasn't there?

Nick: SQL Slammer which targeted MS-SQL. This was how the UCSD/ICSI collaboration formed: Slammer hit, spread worldwide in ~10 minutes, and it was a rush analysis between Vern Paxson, Stuart Staniford, myself, Stefan Savage, Colleen Shannon, and David Moore at CAIDA to figure out just Whisky Tango Foxtrot happened [4].

It actually used a much simpler scheme than the one I proposed; it just sent infectious UDP packets at line rate. We've since seen a similar one with the "Witty" worm [5].

Rik: In your Enigma talk, you mentioned Bro [6]. Is that something you use in your work at ICSI?

Nick: I'm mostly just a Bro user. I've had a part in some high level and researchy stuff with it (e.g., hardware acceleration and parallelization), but for the most part I'm just a user who happens to work down the hall from the developers.

My primary focus at ICSI is a catch-all of network measurement and network security, including, in the past, malcode, packet injection, network monitoring, and network mapping. Additionally, "security@ICSI" is composed of not just system administrators but myself and other researchers, so we end up having to do our own incident response.

When the Snowden slides came out, I looked at them and went "Well, they pretty much do what we do."

Rik: So your Enigma presentation really comes out of a combination of the work you do and the Snowden slides?

Nick: Yep. I looked at the slides and saw basically what I do with more money.

If you discount my civil liberties streak, my abysmal management skills, and my blanket refusal to get a security clearance, I would make a great technical director for the NSA. If anything, what often frustrates me the most is where the NSA is inefficient or inelegant. If the NSA is going to try to spy on the rest of the world and annoy everyone else in the process, at least they should do a good (and less expensive) job!

Rik: It sounds like the way you work at ICSI has taught you how to put together metadata. You also put together a monitoring device that collects data but fits inside a lunchbox. Tell us about that.

Nick: I don't personally work at LBNL (Lawrence Berkeley, not Lawrence Livermore), but Vern does, and I'm fairly familiar with how they operate; I've been involved in multiple studies where I or someone else have some analysis that should be performed, and Vern then takes the analysis and runs it on LBNL's data.

Thus, for example, in "The Matter of Heartbleed," the LBNL bulk recording was used to verify that Heartbleed wasn't exploited by someone against LBNL before public disclosure.

The lunchbox is largely taking advantage of the IDS flow's scalability not just up (to 100 Gbps+ installations) but down: you can run on slower links with cheaper hardware. So the bulk of the work in making the lunchbox was simply deciding what additional analyses to run. One of my minor to-dos is to see how well I can get things to run on the latest Raspberry Pi 3.

In looking at the Snowden documentation, the big difference between what the NSA does on the network is focus on users, not just machines. Thus they have some specific metadata they want to extract to identify "who" rather than "what" is on the network, and most of the work in the lunchbox was focusing on adding these analyses and creating a snazzy Web interface: it was more work for me to figure out enough Bootstrap to make it look good than it was to figure out how, through passive analysis, to identify people in the network traffic.

Rik: Give us a paragraph about the hardware in your lunchbox.

Nick: The hardware is simply an Intel NUC (now two generations old) combined with a really nice DualComm Ethernet Tap/Switch. I admit I gilded the NUC (16 GB RAM and a 120 GB SSD), and you can buy cheaper taps (e.g., SharkTap) that could substantially reduce the cost.

Everything else is simply stuff needed to create an access point for people to connect to since I don't want to tap people who don't consent to be monitored.

References

[1] Enigma YouTube channel: https://www.youtube.com/channel/UCIdV7bE97mSPTH1mOi_yUrw.

[2] Nicholas Weaver, "The Golden Age of Bulk Surveillance": https://www.youtube.com/watch?v=zqnKdGnzoh0.

[3] Stuart Staniford, Vern Paxson, and Nicholas Weaver, "How to 0wn the Internet in Your Spare Time," in *Proceedings of the 11th USENIX Security Symposium (USENIX Security '02)*, 2002: http://www.icir.org/vern/papers/cdc-usenix-sec02/.

[4] Inside the Slammer Worm: http://www.icsi.berkeley.edu/pubs/networking/insidetheslammerworm03.pdf.

[5] Abhishek Kumar, Vern Paxson, and Nicholas Weaver, "Exploiting Underlying Structure for Detailed Reconstruction of an Internet-Scale Event," in *Proceedings of the Internet Measurement Conference 2005 (IMC 2005)*: http://www.icir.org/vern/papers/witty-imc05.pdf.

[6] The Bro Network Security Monitor: https://www.bro.org/.