Bill Cheswick on Firewalls

An Interview

RIK FARROW

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

Ches is an early innovator in Internet security. He is known for his work in firewalls, proxies, and Internet mapping at Bell Labs and Lumeta Corp. He is best known for the book he co-authored with Steve Bellovin and now Avi Rubin, *Firewalls and Internet Security: Repelling the Wily Hacker.* Ches is now looking for the next interesting thing to do, and is open to suggestions.
ches@cheswick.com

Like many USENIX members, I first met Bill Cheswick at a conference. Ches tended to stand out from the crowd, whether he was encouraging people to learn juggling using beanbags or making a presentation.

Ches later became famous after he co-authored the first book about firewalls with Steve Bellovin in 1994. This was not Ches's first adventure in the world of security by far, as he had been working with a firewall as early as 1987 while at Bell Labs.

Rik: Tell us about how you got involved with firewalls.

Ches: My first day at the Labs turned out to be the Christmas party in 1987. I walked up to Dave Presotto, the author of upas, and told him I wanted to be postmaster. I told him that networks were the wave of the future, and I wanted to learn a bit more about them.

He was delighted, and showed me the ropes. He also showed me around his application-level firewall, a VAX 11/750 running 4.3 BSD. Sendmail was replaced by upas, and IP forwarding was disabled. Insiders could use some software he and Howard Trickey had written to access the Internet from inside the company. I later modified this software, changed its name to "proxyd" because "gated" was already taken. I realized a couple decades later that this was the first use of the word "proxy" as it is now used. I need to drop a note to the editor of the OED.

In many ways postmaster was a much more interesting job at the time. There was no spam, which casts a gloomy gray pall over today's email. What we did have were a sea of networks and email addresses. In those days, my two main addresses were research!ches and ches@att.arpa.

Rik: How did you learn about TCP/IP? Back in the late '80s, this was a difficult topic to learn, with no books out yet.

Ches: Steve Bellovin and Dave Presotto had taught me the basics of TCP/IP, and I had *plenty* of lab work. I was using research UNIX, soon to be called the tenth (and last?) edition, and messing around with Plan 9, a lovely and cleaner rethink of UNIX designed from the ground up.

Then, one morning in November 1988, Marion Harris called, saying there was something bad happening on the ARPANET. NPR confirmed this.

I had a sinking feeling: would our firewall hold up to whatever was happening? That feeling has dominated my thoughts about security ever since. Besides, there would be no end to the complaints, ribbing, and whining if the attack got through. Working in the UNIX room toughened the skin. I rushed into work, and the whole place was abuzz.

The short answer was yes, the firewall had held. Peter Weinberger was on the phone basically saying "neener neener" to a variety of sites, especially Bellcore. Those folks had completely rejected the idea of a firewall, and were completely bogged down with the Morris worm. Exponential growth is very hard to control, and Robert Morris had gotten a bit of the worm wrong. This is actually an enduring lesson: the pros behind Stuxnet had a similar problem.

Back to that sinking feeling. How did we really do?

There was no direct IP connectivity to the intranet through our gateway, so the inside was safe there. But if the worm captured the gateway machine, the kernel allowed incoming connections to SMTP servers, most of which were running Sendmail. I ran a scan of the burgeoning Bell Labs/AT&T intranet and found more than 1,300 susceptible hosts.

Rik: What did you use to scan for port 25 in November 1988?

Ches: I probably used sweep, a shell script written by Mike Muuse. It tested machines for five different vulnerabilities.

One of the worm's attacks had been through Sendmail, a common source of security problems at the time. Dave had replaced it, so it wasn't an issue on the gateway.

"Best block is no be there" —Mr. Miyagi, *Karate Kid 2*.

We did not run any of the r* services on the gateway. No trust, no access, so that worked. The worm also broke in via a hole in the finger daemon. About six months before, toward the end of a day's hard work, I felt lazy and decided to wander around looking at the various hosts I administered.

I checked things out on our gateway, and noticed a number of services in /etc/inetd.conf that I was not familiar with. They ran as root, and that's not a good thing. Instead of diligently checking each one out, I simply disabled the lot of them, including fingerd. There were no users on the machine besides administrators. If somebody didn't like it, I could revisit the decision.

Though there were solid principles involved, to this day I consider this to be Security by Luck. Luck is a handy but unreliable tool, and does not help mitigate that sinking feeling I mentioned before.

There was one other stroke of luck. Steve Bellovin had an unprotected 56 Kb link from our intranet to Bellcore's network, where the worm was seething. The worm spread by using /etc/hosts as a list of target machines. Steve's machine at the Labs was at the end of the list and the worm always started attacking targets from the top of /etc/hosts. Any infected machine bogged down completely before reaching the last entries in the /etc/hosts, so it never reached us.

Security by Luck indeed! I had to fix this. I want security without that sinking feeling (there's a book title). Not confidence in security due to hubris, but due to "no being there." This has guided my approach since then.

Rik: I can see how that incident could inspire you. What did you do next?

Ches: Fresh from the uneasy victory over the Morris worm, I decided it was time to create a new firewall. The old one was hopelessly overloaded, and I wanted a design I could rely on. I created a belt-and-suspenders two machine solution out of a couple of MIPS machines. This was a very robust and high-performance design. Steve pointed out that both machines ran upas, which gave it the possibility of a common mode failure, but I had high confidence in upas. It never did let me down.

At about this time (late 1980s) Mark Horton obtained a class A address for AT&T from the powers-that-be by simply asking. That was a different era. The address block lay fallow for a while. We wanted to use it inside the growing AT&T intranet, but the routers of the day had subnetting problems and we couldn't deploy it. Oddly enough, our Cray computer seemed to *require* a class A network: I never did figure that one out.

So I took 12.0.0.0/8 and announced it to the Net, feeding the packets to a non-existent Ethernet address and running tcpdump on the traffic, which came to about 12 to 25 MB/day. Steve analyzed that traffic and wrote a fine paper. Basically, we were watching the death screams of attacked hosts that used IP address-based authentication. One of the steps was to flood a client machine with traffic so it couldn't complain about the attack on the associated server machine. Apparently the author of the code thought it would be fun to use AT&T's network address for the spoofed packets.

This is the first packet telescope I can remember, and I think I might even have coined the term "packet telescope," but my memory is fuzzy on that. I do know that monitoring unused IP addresses remains a very useful tool, and kc [Claffy] from CAIDA [1] gave a nice talk on current uses of the technology just recently.

Rik: I seem to recall that you started publishing about this time.

Ches: There were a lot of PhDs writing papers at the Labs. I had done some new work for the fancy firewall, so I wrote one and showed it to Fred Grampp. I was delighted when he nodded his head and said, "This is a nice paper." So I submitted it and gave the talk at Winter Anaheim USENIX in 1990. It was my first paper, my zeroth being a Permuted Index for TeX and LaTeX commands, something I still use occasionally.

Rik: How about your paper analyzing interactions with a honeypot you designed?

Ches: Chasing down the "attacks" was interesting. Was the attack casual, accidental, or evil? Discerning intent is important in security, which is where those little sticks that the border guards used to use in *Mission Impossible* came from. Who was willing to crash through a stick? The guys with the real barrier down the road wanted to know, and with just a little lead time.

I wanted to catch an actual bad guy and watch what he did. We eventually did, using an early honeypot, and wrote it up in An Evening With Berferd [2]. I almost didn't write it; it was like an optional English paper. I don't like to write, but I do love to have written.

By 1993, it was clear that firewalls were important, but the only real coverage was in a chapter in a book from Gene Spafford and Simson Garfinkel. I mentioned to Steve that we could probably find a dozen papers to staple together into a decent book. He suggested this to John Wait at Addison-Wesley, who had been bugging Steve for a book for about a decade. John said the book was a great idea, but we had to write it from scratch.

Thirteen English papers assigned! Steve and I worked great together; I had never been allowed to write an English paper with a co-author. We settled on a table of contents pretty quickly. Chapters would bounce back and forth for a couple of days and be nearly completed. Some sections demanded information I hadn't thought about. This provided incentive to fill in the knowledge gaps.

The book came out in time for the Spring 1994 Interop. John had estimated that we would sell 8,000 to 12,000 copies. The first printing was 10,000 copies, and sold out in a week. They rushed the second printing, not even waiting for us to correct three errors. The corrections finally made it to the third printing, not long afterward. The first edition sold more than 100,000 copies in at least a dozen languages. As Steve once said, it was a 320-page business card. For me, it was certainly the most important thing I have done in my career. A decade later I'd go to a meeting with some sharp techies in a big company, and they would come up to me later and say the book got them started in network security.

Despite a number of incentives and entreaties, we didn't come out with the second edition, mostly a rewrite, until 2003, with Avi Rubin helping out.

Rik: I imagine that, with the book completed, you started working on other things, such as dealing with "split" DNS.

Ches: After the book was published, Steve and I worked to merge DNS processing for inside and outside queries at the firewall. It ended up in two patents (switching and filtering).

A couple things of note happened in 1996. One was the Panix SYN packet attack. It started me thinking about how to trace anonymous packets back through the Internet. Hal Burch joined me the next summer and we worked on an idea: applying little denial-of-service attacks on possible incoming packet paths, and seeing if they perturb the packet flow. Then repeat, attempting to trace back to the source of the packets. The DoS attacks would be done by locating packet amplifiers on the Internet and carefully applying bursts of pain. We tried this on Lucent's intranet, and it usually worked.

The USENIX paper we wrote came back with two classes of referee comments: 1) we can't accept this paper until the technique is proven on the Internet, and 2) don't you dare try this on the Internet. This approach was certainly out-of-the-box, and much better suggestions were made by others. Today, we don't seem to much care anymore: packets come from everywhere.

Around this time, my role was changing. I remember one day my boss asked if I would do some modifications to upas to make it respond to Sendmail switches. I also received a request to come review the security of AT&T Worldnet, which was going beta in six weeks. I mentioned that I came to the Labs because I didn't like Sendmail and besides, wouldn't the AT&T consulting be a more useful pursuit for the company?

Rik: What other possibilities did the success of your book create for you?

Ches: The book's publication opened many doors. I was invited to speak at numerous conferences, public and less so. A recent count showed that I have spent non-trivial amounts of time in more than 30 countries. I had the opportunity to help a number of law enforcement groups to start coming up to speed on cybersecurity. There were a few CIO breakfasts. Insurance companies wanted to write hacking insurance, and were keen to understand the worst-case scenario (the "hurricane Andrew") of a cyberattack.

Steve and I got an invite from the folks at Renaissance Weekend [3]. Steve thought it was junk mail and discarded it. My wife called and verified that yes, it is the annual get-together that the President goes to and yes, we were invited. I have gone for most years since then, and it has been a rich source of family interactions, ideas, and new friends. Heck, I shook the President's hand the first year I was there. This has also been a wonderful place to answer my science guy questions since leaving Bell Labs. And yes, Bill Nye and I had a fine walk on the beach discussing some of his work on the avionics on the 747.

The other thing in 1996 was my first Highlands forum. I met people like Esther Dyson and Fred Cohen. We did a "day after scenario" developed by the Rand corporation. It went like this:

- Imagine it is ten years from now (2006) and a series of <buzzword compliant> attacks seem to be happening. Evaluate the attacks and advise the President.
- Now it is a week later and a series of <very nasty buzzwork compliant> attacks are seen. Analyze and advise the President.
- And, finally, we are back in 1996. What should we be doing to prepare for all this?

I had remembered the early days of the MILNET (the military's ARPANET), which was connected to the ARPANET through three 56 Kb lines at one point. When bad things happened on the

ARPANET, the military folks cut the links—the turtle pulled in his head.

If we wanted to do this in 2006, would we know where the links were? Besides, maps are cool, and someone should watch and monitor connections and their changes. Who could do that?

If the Air Force pinged Finland, is that an act of war? What about traceroute? Who could collect such data? I asked if data provided by a corporate research project would be useful. The answer was an emphatic yes: of course we want your free data!

Rik: I can see where this is heading: Lumeta and networking mapping.

Ches: 1996 had brought on the "trivestiture" of AT&T into Lucent/Bell Labs, AT&T, and NCR. I had to choose where I wanted to go, so I made a spreadsheet using pluses and minuses, sort of the way a teenager might choose a steady date. The score came out 59 to 60, clearly a draw within the margin of error. Most of my security friends, a lot of mathematicians, etc., went to AT&T Shannon Lab. I chose to remain at Lucent with the systems folk and scientists.

Hal Burch and I started on Internet mapping in 1997. Hal was a crackerjack Olympic programmer who now works for Google. I had the idea to collect connectivity data by sending traceroute-style probes to a zillion networks, and graphing and analyzing the results.

Graphing was going to be a problem: a typical scan would hit 100,000 nodes. I decided to use brute force and lay out the data using a spring force algorithm. Ace programmer Hal managed to hack together some clever optimizations to get a layout algorithm that would produce a nice display overnight. The Lucent intranet took much less time. The Internet layouts were amazing, if not entirely useful. *Wired* published one in December 1998 [5].

I hadn't checked the graphing literature before trying this, and it was a good thing: at the time, the papers said that an 800-node graph was huge. Hal's cleverness, and Moore's Law, had made large layouts much more feasible. Still, I wish I had thought of the project in 1990.

We started collecting and saving daily traceroute data in late 1998, and continued almost uninterrupted until November 2011. I have all that data lying around, available for research use.

I also came up with a way to detect leaks in an intranet perimeter: you send a packet to an inside host, with the spoofed return address of an external "mitt" machine. Packets that make it outside may not have seen a firewall. There was a similar test for packets coming inside.

When someone from Lucent New Ventures came around in late 1999 asking if I had any ideas for businesses, I told him about the maps and the leaks. Research organizations should send such queries out on a regular basis. A company would pay money for this information.

Having a company like Lucent spin off a (hopefully) hotshot startup was a little like watching an old fat man giving birth: there is a lot of grunting, but one isn't really sure what's going to happen.

On October 1, 2000, seven of us armed with VC money spun off from Lucent to found Lumeta. The sell was difficult because we were in a new category of product. Is it a security product or a network management product? Well, both, but marketing and maybe even the customers didn't like that answer. We went to a lot of VC meetings and new product shows. I saw a lot of business ideas, many of them funded, and many of which I thought were not so hot.

Lumeta started collecting and processing intranet maps from a variety of the largest corporations. I would have loved to make the data available, but it was much too sensitive. We clearly had access to the most extensive graph data on intranets in the world. Most of the customers were quite happy with the results, and we *always* found something interesting in their networks.

The daily collection of the world's path data continued, under the name of the Internet Mapping Project. Before Lumeta, back in the spring of 1999 during the NATO/Serbian war, I had focused on collecting fine network data from the area. Steve Branigan produced a fine graph and movie [4] of the effects of aerial bombing. We also found the Yugoslavian embassy in the US.

The daily network probes did bring a slow stream of complaints, which faded away by the mid-2000s: by then there was simply too much "background radiation" of evil packets on the Internet. But after 9/11, I stopped caring. I focused some extra scans on the obvious suspects and started collecting data.

A couple years later we met with Richard Clarke in Washington. I had extracted every traceroute path collected over the previous six months that contained at least one Iranian address, and mapped it, coloring rarely used links in red. He looked it over and said he had been asking for this map for the past six months. He asked if it was classified? Well, we were just a few people from NJ, no clearance, etc. He called a three star general over in DoD, and we eventually sold our product to a number of branches of government. I am not cleared to know what they found. They told us that our product made the Republic a little bit safer. That's good enough for me.

For me, Lumeta was a fine outcome for a research project. It also was a lesson in business, and has helped me separate clever research projects that will never be used from those with business potential. I have used this skill a lot since then, particularly at AT&T Shannon Lab.

Rik: Tell us about your time at Shannon Lab.

Ches: I started at AT&T Shannon Lab in April 2007. These were the AT&T researchers spun out of Bell Labs back in 1996, processed through some bad times, and augmented by sharp new additions. The security research department had been disbanded a few years before, and most of my security colleagues had scattered to various university positions.

As with Bell Labs, it was an honor and pleasure to work there.

I did do some security work there. I tried a few experiments with passwords on smartphones. I recovered my stolen iPhone with the help of Steve Branigan and the NJ State police. We created a little useful case law concerning the use of WiFi localization. It turns out that the iPhone has five radios in it, and the phone company interacts with four of them. (Go ahead, Über geeks: count those radios carefully!) I did not use the phone company's resources (other than my time) to catch him, however.

But there was something in the water at Shannon. I was generating about four patent ideas a year, plus a number of business ideas. My patent count is about a dozen, with about six more crawling through the system. Alas, AT&T is not fertile ground for small new ideas, and only a couple of my efforts bore fruit.

Aside from some authentication and security patents, I created a new kind of movie (the slow movie) and a new way to see movies (movie thumbscapes). I think the latter would make terrific (and lucrative) wallpaper that would bring in some money and interest for a major movie. Unfortunately, I was unable to reach a leading filmmaker to show off the results. I also invented a new kind of extremely soft-core pornography, but I will skip that.

Yifan Hu added a terrific new layout algorithm to graphviz, much better than our efforts of a decade before, and laid out the entire AT&T corporate org chart. I added labels, colors, and other data to his positioning data to create what must be the world's largest org chart. We could color links by employee age, patent production, union membership, etc. I created one for the CEO's office. I believe this visualization would be a valuable tool for corporate consultants.

I stayed at Shannon for five years before I was laid off in April 2012.

Rik: What do you plan on doing in the future?

Ches: I am trying to figure out what to do next. It is clear that I don't fit into the usual employment slots in the usual corporate suspects. I am hanging out at Penn as a visiting scholar, and some projects are starting up. Teaching is a strong possibility: I have always enjoyed it.

I am working on iTeX, a tool for bringing LaTeX documents to the iPad, including arXiv and Project Gutenberg texts. I am translating a number of Project Gutenberg books into LaTeX.

I am always looking to work with sharp people on interesting projects. But I am not idle: I don't understand where I ever had time to go to work for 40 hours a week. I heard of one fellow who said if he retired a second time he would have to hire an assistant to get all the work done.

Resources

[1] The UCSD Network Telescope: http://www.caida.org/projects/network_telescope/.

[2] An Evening with Berferd: http://www.cheswick.com/ches/papers/berferd.pdf.

[3] Invitation-only retreats for preeminent authorities, emerging leaders, and their families: https://www.renaissanceweekend.org/home.htm.

[4] Effects of war on the Yugoslavian network: http://cheswick.com/ches/map/yu/.

[5] Internet map, circa 1998: http://www.cheswick.com/ches/map/gallery/wired.gif.

usenix

THE ADVANCED
COMPUTING SYSTEMS
ASSOCIATION

Professors, Campus Staff, and Students— do you have a USENIX Representative on your campus? If not, USENIX is interested in having one!

The USENIX Campus Rep Program is a network of representatives at campuses around the world who provide Association information to students, and encourage student involvement in USENIX. This is a volunteer program, for which USENIX is always looking for academics to participate. The program is designed for faculty who directly interact with students. We fund one representative from a campus at a time. In return for service as a campus representative, we offer a complimentary membership and other benefits.

A campus rep's responsibilities include:

- Maintaining a library (online and in print) of USENIX publications at your university for student use
- Distributing calls for papers and upcoming event brochures, and re-distributing informational emails from USENIX
- Encouraging students to apply for travel grants to conferences
- Providing students who wish to join USENIX with information and applications
- Helping students to submit research papers to relevant USENIX conferences
- Providing USENIX with feedback and suggestions on how the organization can better serve students

In return for being our "eyes and ears" on campus, representatives receive a complimentary membership in USENIX with all membership benefits (except voting rights), and a free conference registration once a year
(after one full year of service as a campus rep).

To qualify as a campus representative, you must:

- Be full-time faculty or staff at a four year accredited university
- Have been a dues-paying member of USENIX for at least one full year in the past

For more information about our Student Programs, contact
Julie Miller, Marketing Communications Manager, julie@usenix.org

www.usenix.org/students