Remembering Peter Eckersley

J. Alex Halderman

Remarks delivered at his memorial service
Internet Archive, San Francisco, March 4, 2023

Believe it or not, barely a decade ago, it was normal to browse the Web without any encryption. Probably the page that accepted your credit card used HTTPS, and, if you were lucky, so did the form where you entered your password, but encryption was used by only about 1 in 4 Web requests. Facebook, for instance, only began using HTTPS by default in 2013. That seems almost unbelievable today, since we know that governments, ISPs, and marketers can spy on unencrypted traffic, censor messages they disfavor, and even change the data we receive to weaponize our computers to attack others. We know the Network is *evil* and *wants to kill you*, and HTTPS is the Web's most important defense. Fortunately, today, about 90% of Web requests use HTTPS—and we have Peter Eckersley, as much as anyone else, to thank for that.

When you type archive.org into your browser, how does your computer know it's talking to the real site instead of an imposter? The answer is that archive.org sends a Certificate—a statement vouching for its identity that's digitally signed by an organization your browser trusts, called a certificate authority or "CA". Every site that wants to use HTTPS first needs to go to a CA to obtain a certificate.

Back in the bad old days, before Let's Encrypt, getting a certificate meant setting up an account with a CA, paying an annual fee, jumping through hoops to prove your identity, waiting for the certificate to be issued, and manually reconfiguring your server to install it. The process could take hours or days and was highly prone to human error. It's no wonder that HTTPS adoption was so low.

Nobody even knew how many sites used HTTPS until an earlier project of Peter's, the SSL Observatory, conducted the first worldwide census! In 2010, he and Jesse Burns scanned the entire IPv4 address space—4

billion IP addresses—and collected the certificate from every public server. It's hard to convey how audacious this was—it was only the second or third time since 1982 that researchers had comprehensively scanned the Internet. With tools available at the time, it took months and heroic effort. The SSL Observatory may be Peter's most underappreciated scientific accomplishment. It helped kick-off an entire field of using Internet-wide scanning for security research and *thousands* of subsequent scientific papers.

It was during subsequent Internet-wide scanning research that Peter and I became friends and I first witnessed his seemingly limitless energy and enthusiasm for making the Web a better place. I had been incubating an idea to do that that I just *had* to share with Peter, and the opportunity arose eleven years ago this week, on March 1, 2012, when Peter and I met at the RSA Conference here in San Francisco. As we wandered the halls of the Moscone Center, Peter practically completed my sentences: "Despite its cryptographic faults, the biggest vulnerability facing HTTPS was that *not enough sites used it*. The solution was to replace the complicated and labor-intensive process of getting a certificate with highly automated software and services, based on open protocols, managed by a not-for-profit. Automation would allow certificate authorities to offer basic certificates for free. We could drive the cost and complexity of deploying HTTPS to zero and make the entire Web encrypted by default." All we had to do was build the technology and get CAs to adopt it.

By the next morning, Peter had *dozens* of ideas to move the project forward. He also had a name—Project Chocolate—since chocolate stimulates the brain to release oxytocin, a hormone that promotes the formation of trust. Peter, as always, was a whirlwind of activity. He quickly got EFF's blessing for project Chocolate and brought in Seth Schoen, Jacob Hoffman-Andrews, and later Brad Warren and others. Together with my students—especially James Kasten, who later went on to build another CA at Google—we spent the next year developing prototypes and pitching the idea to CAs, though without any success.

The most decisive step on the road to Let's Encrypt happened the next year, in 2013, when one of Peter's contacts tipped him off about a small parallel effort taking place at Mozilla. Peter and I connected with Josh Aas, Eric Rescorla (and later Richard Barnes), who had similar ideas about automating HTTPS deployment, but with a twist: instead of partnering with existing CAs, we could jumpstart adoption by founding a new certificate authority ourselves! Our thinking was so complementary that within days of discovering each other, we merged our efforts, and the founding team of Let's Encrypt was in place.

That summer, Edward Snowden came forward with chilling revelations about NSA mass surveillance, and the need for ubiquitous encryption became clearer than ever. Yet we still had a lot of work to do—it's not easy to become a CA—and it would take two more years and contributions from many people to make Let's Encrypt a reality. To operate the CA, we set up a non-profit, ISRG, where Peter remained a board member until shortly before his passing. We recruited our first sponsors, formalized the ACME protocol for automating certificate validation; developed production software; and found an existing CA willing to delegate trust so that our certificates would be accepted by browsers from day one.

The entire team met in person only once during this frenzied time, at a working retreat Peter organized. By that point in 2014, our group was large enough that Peter rented a South Bay mansion on AirBNB to house us all. There was just one catch, as Peter explained: "How do you feel about Murder Houses?" Our intended lodgings had recently been the scene of a notorious home invasion and homicide…but so goes housing in the Bay! That lent an eerie backdrop to our discussions of certificate chains and protocol syntax, as we grimly contemplated every gouge in the furniture and carpet stain…

As roles in the project became more specialized, Peter led the development of the initial ACME client, which would eventually become Certbot. In a reflection of Peter's vision for a web that is secure by default, Certbot aims not only to obtain a certificate but to fully automate HTTPS deployment by reconfiguring existing servers. Today, Certbot is among the

most popular Let's Encrypt clients, and it is developed and maintained by Peter's former team at EFF.

After more than three years of work, Let's Encrypt issued its first browser-trusted certificate in September 2015. From the beginning, we've used Certificate Transparency, a blockchain-like public ledger that allows the world to inspect every certificate Let's Encrypt ever issued. You can look back at the first few certificates …three the first days, one the next…and you'll find among them sites Peter operated. Imagine what he was feeling as he tested his brainchild for the first time and saw the lock icon appear in the browser…

Six months after Let's Encrypt launched, we had issued a million certificates. After nine months it was 10 million. Less than a year after that, 100 million. Today, we are closing in on 4 billion certificates issued and protect more than 320 million domains—more than all other CAs combined—including NSA.gov and the Whitehouse. The Web is encrypted by default; Peter's vision has come to pass, and I can't imagine a more fitting monument to our friend.

Sometimes people accomplish great things because they started out thinking, naively, that it would be easy. Peter wasn't naïve; he had the rare audacity to start things knowing for certain they'd be very hard. Thank you, Peter. The Internet is a better place for us all thanks to you.