
Photo by Martin Gundersen

Cybersecurity for the Public Interest

Bruce Schneier
Harvard University

The Crypto Wars have been raging off and on for a quarter century. On one side is law enforcement, which wants to be able to break encryption to access devices and communications of terrorists and criminals. On the other is almost every cryptographer and computer-security expert, repeatedly explaining that there's no way to provide this capability without also weakening the security of every user of those devices and communications systems.

It's an impassioned debate, acrimonious at times, but there are real technologies that can be brought to bear on the problem: key-escrow technologies, code obfuscation technologies, and backdoors with different properties. Pervasive surveillance capitalism—as practiced by the Internet companies that are already spying on everyone—matters. So do society's underlying security needs. There is a security benefit to giving access to law enforcement, even though it would inevitably and invariably also give access to others. However, there is also a security benefit of having these systems protected from all attackers, including law enforcement. These benefits are mutually exclusive. Which is more important, and to what degree?

The problem is that few policy makers are discussing this policy issue from a technologically informed perspective, and very few technologists truly understand the policy contours of the debate. The result is that both sides consistently talk past each other and policy proposals—which occasionally become law— that are technological disasters.

This isn't sustainable, for either this issue or any of the other policy issues surrounding Internet security. We need policy makers who understand technology, but we also need cybersecurity technologists who understand—and are involved in—policy. We need *public-interest technologists*.

Let's pause at that term. The Ford Foundation defines *public-interest technologists* as "technology practitioners who focus on social justice, the common good, and/or the public interest." A group of academics recently wrote that public-interest technologists are people who "study the application of technology expertise to advance the public interest, generate public benefits, or promote the public good." Tim Berners-Lee has called them *philosophical engineers*. I think of public-interest technologists as people who combine their technological expertise with a public-interest focus: by working on tech policies, by working on a tech project with a public benefit, or by working as a traditional technologist for an organization with a public benefit. Maybe it's not the best term—and I know not everyone likes it—but it's a decent umbrella term that can encompass all these roles.

We need public-interest technologists in policy discussions. We need them on congressional staff, in federal agencies, at nongovernmental organizations (NGOs), in academia, inside companies, and as part of the press. In our field, we need them to get involved not only in the Crypto Wars but everywhere cybersecurity and policy touch each other: the vulnerability equities debate, election security, cryptocurrency policy, Internet of Things safety and security, big data, algorithmic fairness, adversarial machine learning, critical infrastructure, and national security. When you broaden the definition of *Internet security*, many other areas fall within the intersection of cybersecurity and policy. Our particular expertise and way of looking at the world are critical for understanding a great many technological issues, such as net neutrality and the regulation of critical infrastructure. I wouldn't want to formulate public policy about artificial intelligence and robotics without a security technologist involved.

Digital Object Identifier 10.1109/MSEC.2018.2889891
Date of publication: 20 March 2019

continued on p. 82

 Copublished by the IEEE Computer and Reliability Societies

a witness that enables the deterministic reproduction of the bug. Sanitization, the process of instrumenting code with additional software guards, helps in discovering bugs closer to their source. Overall, security testing remains challenging, especially for libraries or complex code, such as kernels or large software systems. As fuzzers become more domain specific, an interesting challenge will be to make comparisons across different domains (e.g., comparing a grey-box kernel fuzzer for use-after-free vulnerabilities with a black-box protocol fuzzer). Given the significant recent improvements in fuzzing, exciting new results can be expected. Fuzzing will help make our systems more secure by finding bugs during the development of code before they can cause harm during deployment.

Fuzzing is a hot research area with researchers striving to improve input generation, reduce the impact of each execution on performance, better detect security violations, and push fuzzing to new domains, such as kernel fuzzing or hardware fuzzing. These efforts bring excitement to the field. ■

References

1. M. Zalewski, "American fuzzy lop (AFL)," 2013. [Online]. Available: http://lcamtuf.coredump.cx/afl/technical_details.txt
2. R. Swiecki, "Honggfuzz," 2010. [Online]. Available: https://github.com/google/honggfuzz
3. K. Serebryany, D. Bruening, A. Potapenko, and D. Vyukov, "AddressSanitizer: A fast address sanity checker," presented at the 2012 USENIX Annual Technical Conference, Boston, MA. [Online]. Available: https://www.usenix.org/conference/atc12/technical-sessions/presentation/serebryany
4. Y. Jeon, P. Biswas, S. A. Carr, B. Lee, and M. Payer, "HexType: Efficient detection of type confusion errors for C++," in *Proc. 2017 ACM SIGSAC Conf. Computer and Communications Security*, pp. 2373–2387. doi: 10.1145/3133956.3134062.
5. N. Stephens et al., "Driller: Augmenting fuzzing through selective symbolic execution," in *Proc. ISOC Network and Security System Symp.*, 2016. doi: 10.14722/ndss.2016.23368.
6. S. Raway, V. Jain, A. Kumar, L. Cojocar, C. Giuffrida, and H. Bos, "VUzzer: Application-aware evolutionary fuzzing," in *Proc. ISOC Network and Security System Symp.*, 2017. doi: 10.14722/ndss.2017.23404.
7. H. Peng, Y. Shoshitaishvili, and M. Payer, "T-Fuzz: Fuzzing by program transformation," in *Proc. 2018 IEEE Symp. Security and Privacy*. doi: 10.1109/SP.2018.00056.
8. I. Yun, S. Lee, M. Xu, Y. Jang, and T. Kim, "QSYM: A practical concolic execution engine tailored for hybrid fuzzing," presented at the 27th USENIX Security Symp., Baltimore, MD, 2018.
9. G. Klees, A. Ruef, B. Cooper, S. Wei, and M. Hicks, "Evaluating fuzz testing," in *Proc. ACM Conf. Computer and Communications Security (CCS)*, 2018. doi: 10.1145/3243734.3243804.

Mathias Payer is a security researcher and an assistant professor at the EPFL School of Computer and Communication Sciences, leading the HexHive group. His research focuses on protecting applications in the presence of vulnerabilities, with a focus on memory corruption and type violations. Contact him at mathias.payer@nebelwelt.net.

Last Word *continued from p. 84*

Public-interest technology isn't new. Many organizations are working in this area, from older organizations, such as EFF and EPIC, to newer ones, such as Verified Voting and Access Now. Many academic classes and programs combine technology and public policy. My cybersecurity policy class at the Harvard Kennedy School is just one example. Media startups like The Markup are doing technology-driven journalism. There are even programs and initiatives related to public-interest technology inside for-profit corporations.

This might all seem like a lot, but it's really not. There aren't enough people doing it, there aren't enough people who know it needs to be done, and there aren't enough places to do it. We need to build a world where there is a viable career path for public-interest technologists.

There are many barriers. A report titled "A Pivotal Moment" (https://www.netgainpartnership.org/s/pivotalmoment.pdf) includes this quote:

> While we cite individual instances of visionary leadership and successful deployment of technology skill for the public interest, there was

> a consensus that a stubborn cycle of inadequate supply, misarticulated demand, and an inefficient marketplace stymie progress.

That quote underscores the three places for intervention. One: the supply side. There just isn't enough talent to meet the eventual demand. This is especially acute in cybersecurity, which has a talent problem across the field. Public-interest technologists are a diverse and multidisciplinary group of people. Their backgrounds come from technology, policy, and law. We also need to foster diversity within public-interest technology; the populations using the technology must be represented in the groups that shape the technology. We need a variety of ways for people to engage in this sphere: ways people can do it on the side for a couple of years, between more traditional technology jobs, or as a full-time rewarding career. We need public-interest technology to be part of every core computer science curriculum, with "clinics" at universities, where students can get a taste of public-interest work. We need technology companies to give people sabbaticals to do this work and then value what they've learned and done.

Two: the demand side. This is our biggest problem right now: not enough organizations understand that they need technologists doing public-interest work. We need jobs to be funded across a wide variety of NGOs. We need staff positions throughout the government: executive, legislative, and judiciary. President Obama's U.S. Digital Service should be expanded and replicated; so should Code for America. We need more press organizations that perform this kind of work.

Three: the marketplace. We need job boards, conferences, and skills exchanges—places where people on the supply side can learn about demand. Major foundations are starting to provide funding in this space—the Ford and MacArthur Foundations, in particular, but there are others as well.

This problem in our field has an interesting parallel with the field of public-interest law. In the 1960s, there was no such thing as public-interest law. The field was deliberately created, funded by organizations like the Ford Foundation. They financed legal-aid clinics at universities, so students could learn housing law or discrimination, or immigration law. They funded fellowships at such organizations as the ACLU and NAACP. They created a world where public-interest law is valued, where all the partners in major law firms are expected to have done some public-interest work. Today, when the ACLU advertises for a staff attorney, paying one-third to one-tenth normal salary, it gets hundreds of applicants. Today, 20% of Harvard Law School graduates go into public-interest law, and the school has soul-searching seminars because that percentage is so low. Meanwhile, the percentage of computer-science graduates going into public-interest work is basically zero.

This is bigger than computer security. Technology now permeates society in a way it didn't just a couple of decades ago, and governments move too slowly to take this into account. That means technologists are now relevant to all sorts of areas with which they traditionally had no connection: climate change, food safety, future of work, public health, bioengineering.

More generally, technologists need to understand the policy ramifications of their work. There's a pervasive myth in Silicon Valley that technology is politically neutral. It's not, and I hope most people reading this today know that. We built a world where programmers felt they had an inherent right to code the world as they saw fit. We were allowed to do this because, until recently, it didn't matter. Now, too many issues are being decided in an unregulated capitalist environment where significant social costs are too often not taken into account.

This is where the core issues of society lie. The defining political question of the 20th century was "What should be governed by the state, and what should be governed by the market?" This defined the difference between East and West and the difference between political parties within countries. The defining political question of the first half of the 21st century is "How much of our lives should be governed by technology and under what terms?" In the last century, economists drove public policy. In this century, it will be technologists.

The future is coming faster than our current set of policy tools can deal with. The only way to fix this is to develop a new set of policy tools with the help of technologists. We need to be in all aspects of public-interest work, from informing policy to creating tools for building the future. The world needs our help. ■

Bruce Schneier is a lecturer and fellow at the Harvard Kennedy School and special advisor to IBM Security. His new book is *Click Here to Kill Everybody: Security and Survival in a Hyper-Connected World*. He organized a one-day track on public-interest technology at the RSA Conference this March in San Francisco. He maintains a resources page at www.public-interest-tech.com. Contact him via www.schneier.com.