;login:
DECEMBER 2014 VOL. 39, NO. 6
ELECTRONIC SUPPLEMENT

Conference Reports

23rd USENIX Security Symposium

August 20–22, 2014, San Diego

Summarized by Rik Farrow, Zhigong Li, Johanna Ullrich

Tor Panel and Lightning Talks

Summarized by Rik Farrow (rik@usenix.org)

Kevin Fu introduced Roger Dingledine and Paul Syverson, pointing out that their paper about Tor in 2004 did not win an award. The two other papers in the same session won Best Paper awards that year. When Tor started, there were just 32 relays, and now there are over 6000. In 1995, Syverson had been working on an earlier version that the Navy used, and had been able to make that system public because anonymity networks only work when there is enough background traffic to hide what online search the Navy wanted to accomplish.

Kevin then asked what were the big issues when they started working on the version of Tor that came out in the paper. Syverson said that getting funding was hard, and they really needed to work on making Tor more robust, with better crypto, but all in a form that people would readily understand. Dingledine said that they wanted to have wide acceptance, and that meant they couldn't just publish a paper and move on. They kept working on Tor, but they also traveled around the world talking about it. Syverson said it was also important to talk about the hard problems, as that helped to get other researchers to work on Tor. They also needed good documentation and metrics, so research was feasible.

Kevin than asked them to talk about risk-taking, the risks they took on as young researchers. Syverson said, "You should do it. Dare to be stupid, perhaps you will be dumb, but you could be brilliant." Syverson had tried to get a patent on onion routing and find VCs to commercialize it, but they discovered that the system worked much better as a volunteer system. He also said that you want to think about your problems long and hard. Their own first thoughts were not successful.

Kevin asked if their paper was accepted the first time it was submitted, and Syverson said it failed to get into IEEE Security and Privacy (Oakland). Dingledine said there were lots of papers on mixer systems, but they actually had a distributed system with real users. Syverson said that another reason they got into USENIX Security was because Niels Provos said that he used Tor and it worked.

Kevin asked them to summarize their experience, and Syverson said that you want to produce something that works, that people can actually use, and make it as secure as you can, as well as be as usable as possible. Dingledine added that you want to build a community that cares about your topic. They spent a lot of time talking to communities about problems they wanted to solve. Syverson ended by telling the audience not to be afraid of failure.

After the panel finished, the lightning talks relied on short videos provided by authors of papers. Some were really good, with music and, occasionally, animation; in some the author provided a summary into a video camera; and two just showed a summary slide. Hopefully, if another chair tries to have a lightning talks session, presenters will have a better idea of what makes an interesting video. I found the videos useful, as there were two tracks of papers, and having more than an abstract, and a suggestion of what the presenter's presentation style might be like, helped me to decide which track to sit in on.

Work-in-Progress Reports

Summarized by Johanna Ullrich (jullrich@sba-research.org)

USENIX Security's Work-in-Progress session contained eight presentations in total and presented a wide range of work covering privacy, security, and funding. Abe Singer (Laser Interferometer Gravitational Wave Observatory) began by discussing the failures of user-selected passwords. Even in the case where passwords are not based on ordinary dictionaries, more sophisticated dictionaries can be built from existing passwords. He claims that these dictionaries allow attackers to guess 80 percent of user-selected passwords, and the only mitigation is their prohibition.

Adrian Dabrowski (SBA Research) presented digital self-defense in mobile networks to detect IMSI catchers. IMSI catchers temporarily fake a network cell to control your mobile. Mobile IMSI catcher catchers are mobile applications that learn about your network environment and compare the current network to the learned network to detect temporary cells from IMSI catchers. Additionally, a number of fixed stations have been installed over the city of Vienna to permanently sense the network's behavior.

Giselle Font (NIC Chile Research Lab) works on guaranteeing location privacy for users of mobile devices participating in a monitoring system to measure access quality. Passive monitoring should be enhanced by active monitoring. Clients' privacy should be protected while still allowing researchers to compute aggregate data, which is tackled by homomorphic encryption.

Jeremy Epstein (National Science Foundation) insistently asked the audience to take his money and outlined the US flagship program "Secure and Trustworthy Cyberspace." That program funded 128 new projects in 2014 and targets a wide variety of topics, including recent issues, e.g., cyber-physical system security. He strongly advised listeners to join the newsletter (listserv@listserv.nsf.gov).

David Mazières (Stanford University) presented cryptographically enforced flow control. His goal is the implementation of cryptographic instructions into CPUs to prevent attacks from overwriting return pointers. An implementation already exists and a draft of the paper can be found at the homepage.

Eric Eide (University of Utah) offered a thrilling possibility for cloud security research—CloudLab (www.cloudlab.us). As a joint project involving three universities, a metacloud has been built as a scientific infrastructure with more than 15,000 cores for cloud-related topics. CloudLab makes it possible to set up your personal cloud in a few steps. Although being further extended in the next year, it is available now and is free to use for research and educational purposes.

J. Bonneau (Princeton) presented a Goldfinger attack against Bitcoin, including rental of computing power, where all participants lose and bring the system to an end. But his theoretical model implied that the attack's cost was rather low. Thus, he invited you firmly to contribute to the model's enhancement to tackle reality. His colleague, Steven Goldfeder, questioned how individuals are able to store their bitcoins in a secure matter. He proposed a two-factor security solution using threshold signatures that are stored on your laptop as well as on your mobile device. An implementation is currently underway.

Poster Session

Summarized by Zhigong Li (lizhigong1991@gmail.com)

Hiding Shellcodes in Korean Texts

Ji-Hyeon Yoon and Hae Young Lee, Seoul Women's University

Ji-Hyeon Yoon and Hae Young Lee presented this work for hiding shellcodes, which are small pieces of malicious codes, in Korean texts. Shellcode can be transformed to English-like text in ASCII with complex and time-consuming processes. However, hiding shellcodes in Korean text does not require complicated encoders and decoders. In addition, they may be undetected by manual and automated inspections. The transformation method may also be applied to some other Asian languages, such as Chinese and Japanese.

Secrets in Public Repositories

Kyle McGuire, Hao Bai, and David Evans, University of Virginia

Secrets are everywhere in today's programs, including passwords, API keys, server logins, etc. But these secrets may be visible in public repositories such as GitHub. Kyle McGuire presented this work to find out how often secrets end up in publicly visible code and how to mitigate this risk. By scanning four large repositories, they found that lots of secrets exist. More than 60% of owners of these repositories fixed the problem after receiving the warning emails from them.

SystemLeakalyzer: Systematically Detecting System Side-Channels

Qi Alfred Chen and Yunhan Jia, University of Michigan; Zhiyun Qian, NEC Labs America, Inc.; Z. Morley Mao, University of Michigan

Qi Alfred Chen presented this work to detect system side channels. However, exhaustively discovering them is hard. Therefore, they modeled the side-channel attacks as an information leakage problem and used the techniques of program analysis to detect them. In the initial results for off-path TCP attacks, they found 25 potential leakages. Six of them were potential vulnerabilities and 19 were false positives. They wanted to improve the method through refining static taint analysis implementation and adding more information sinks.

Measuring Privacy Disclosures in URL Query Strings

Andrew G. West, Verisign Labs (Verisign, Inc.); Adam J. Aviv, U.S. Naval Academy

Session and form data are often included in the query strings of URLs. This project looks at how often these sensitive data appear in the URLs. By analyzing 892 million user-submitted URLs, they found that 55% of URLs have one or more key-value pairs. The keys include geo-location, network, identity, etc. Even cleartext passwords appear. In addition, about 40% of the problem URLs come from mobile devices.

Exploring Movement-Pattern Based Authentication for Mobile Platforms

Dustyn James Tubbs and Khandaker Abir Rahman, Saginaw Valley State University

This project presents a new method of user authentication for mobile platforms, movement-based authentication. This method recorded values of four sensors: accelerometer, linear accelerometer, gyroscope, and tilt sensor. After 29,450 authentication attempts, they achieved an Equal Error Rate (EER) of 20.22%. The authors considered that this alternate authentication method can eliminate shoulder surfing and smudge attacks to some extent. In addition, more optimizations may be done to remove outliers and achieve high-level features.

Panel

The Future of Crypto: Getting from Here to Guarantees

Summarized by Rik Farrow (rik@usenix.org)

Moderator: Sandy Clark, University of Pennsylvania
Panelists: Daniel J. Bernstein, Technische Universiteit Eindhoven and University of Illinois at Chicago; Matt Blaze, University of Pennsylvania; and Tanja Lange, Technische Universiteit Eindhoven

Sandy Clark started by asking each panelist for an opening statement. Bernstein started off funny, but quickly proceeded into a clever diatribe dripping in sarcasm. Bernstein said that while cryptography could be a barrier to an attacker, we want to remove those guarantees so that cryptography always fails. If people interested in cryptography don't go to work for the public sphere (NSA), we can hope they go to grad school and spend their time researching fully homomorphic encryption and side-channel leakage, or focus on attacks. And if there is the risk of people actually getting helpful crypto, we can make things as complex as possible, put it on vulnerable devices, include lots of knobs and switches, and be sure it is complicated enough that it will fail.

Matt Blaze argued that we don't need guarantees that crypto will fail. Instead, he can think of a giant, government conspiracy, where the government managed to brainwash security researchers into fighting over the export of PGP instead of spending our time making security usable. We should have spent time working on the human interface, since by 1990 we had crypto that worked well enough, perhaps with a few tweaks for Moore's law, but

instead we were exhausted after fighting for the use of PGP. Blaze said he was as guilty as anyone, glad the battle was won, but still regrets the opportunity lost while fighting the wrong battle.

Lange took a slightly more positive view on things. She pointed out that even if we have something that works well and is secure, people will complain that it's not certified. Lange, along with Bernstein and Peter Schwabe (Radboud Universiteit Nijmegen) built NaCl, an easy-to-use, high-speed crypto library without all types of knobs, just one way of doing each task right (http://nacl.cr.yp.to/). Lange said that NaCl is actually better than 1990s crypto because we didn't know about time-based attacks on the key or message back then. NaCl is designed with this in mind, and will take a constant amount of time for different keys and sizes, within reason.

Sandy then opened up the floor for questions.

Jeffrey Goldberg (AgileBits) wanted to follow up on usability, saying that's where research should be, and asked how do you teach people tough crypto concepts? Goldberg said he had spent years trying to teach people how to use PGP. Bernstein countered by saying that you need to start by making an insecure connection to an untrusted key server, part of the basis for PGP, and not a very good one. Blaze pointed out that some things have been done better, such as encryption of over-the-air cell calls, which Blaze said is a great victory because it was completely invisible to most people. And while he knows how to use encryption, he still doesn't know how to send Bernstein an encrypted email.

Lange said she would never recommend NaCl for users but for programmers. And while she might think something is easy, her students will remind her that it's not. Bernstein then shifted, saying he would take off his attacker's hat, and suggest that there are usable ways of managing keys, such as using a domain system. Lange objected, asking why we should trust government-issued keys.

Sandy then posed the question: what does the community have that it can rely on, and Bernstein reverted to style: nothing. Blaze pointed to the elephant in the room (Bernstein is really average size), the one person on the planet who can write secure software, but he can't explain it to anyone else, so we have crypto in an unusable package. Bernstein countered that we need to focus on simplicity in libraries. He called OpenSSL a collection of cryptographic functions, with multiple versions of elliptic curves, each with the same code but different seeds, essentially providing programmers with too many options. Bernstein said you can find thousands of examples in cryptographic software. Lange countered that there are some differences, such as one is constant time and another isn't. Bernstein maintained that we just need to implement strong algorithm X in the simplest way.

David LeBlanc said that his job at Microsoft is to help people put the building blocks in the right order: for example, MAC then encrypt or encrypt then MAC. Lange countered that for the research community, that's a solved problem, so you just ignore the wrong way of doing things. Bernstein stated that the crypto community needs to provide just the one thing that the user needs, not options.

Matt Blaze decided that he should say that while Bernstein's comments suggest that there is no value in fundamental and theoretical research in crypto, Blaze believes that there is a middle ground. Bernstein commented on the Security Symposium, saying that he believes the people here are more concerned about physical reality than they are at the annual CRYPTO conference occurring in Santa Barbara at the same time.

Brian Warner (Mozilla) pointed out that it was difficult to use libraries that aren't 20 years old. Blaze commented that two years ago, you would be using a trusted standard, like the FIPS library that contained the Dual Elliptic Curve Random Number Generator with the back door. Lange pointed out that we are in a better position today because of the revelations and can push for going beyond NIST EC.

Someone who had driven from CRYPTO to Security asked if cryptographers should be telling people to do what we tell them, or just let them do whatever they want. Lange said that the question is long, but the answer is simple: use NaCl, and she can explain it if they are interested. Bernstein stated that the situation is already dire: 99% of all traffic is not encrypted, which is completely disastrous. There is a real tension between just getting the basics done, but instead we have a general purpose solution, OpenSSL, that's not specific. It's like you can teach lots of people to write kernel code, and you will have a fancy kernel with more lines of code than you can support, and loads of features, but it will never be secure. And you can do that with crypto, too.

Blaze asked for a show of hands: Should Bernstein spend the next *N* months of his life picking apart an NIST standard or developing secure email? Secure email was the clear winner.

Rachel Greenwood (Drexel) asked how to get standards going forward. Lange said that she felt the IETF was moving in the right direction in their crypto group, that the NSA guy is leaving at the end of the year. Blaze agreed, saying that the IETF tries to be fairly lightweight.

Bill Simpson pointed out that he had added a hash to the IETF standard, which even though it used MD5 still worked since his hash included adding a secret. Simpson said that we do read the papers. Bernstein replied that it doesn't matter who creates the standards, as long as they are well designed and survive a serious review.

Sandy Clark then asked for closing statements, and all three panelists were brief. Lange apologized for being a cryptographer and unpleasant to deal with. Blaze said that he looks forward to using Bernstein's secure email system, and Bernstein stated, keep it simple, stupid.