LAWFARE

CYBERSECURITY

What "Efail" Tells Us About Email Vulnerabilities and Disclosure

By **Bruce Schneier** Thursday, May 24, 2018, 7:00 AM

DayZero: Cybersecurity Law and Policy

Last week, researchers disclosed vulnerabilities in a large number of encrypted email clients: specifically, those that use OpenPGP and S/MIME, including Thunderbird and AppleMail. These are serious vulnerabilities: An attacker who can alter mail sent to a vulnerable client can trick that client into sending a copy of the plaintext to a web server controlled by that attacker. The story of these vulnerabilities and the tale of how they were disclosed illustrate some important lessons about security vulnerabilities in general and email security in particular.

But first, if you use PGP or S/MIME to encrypt email, you need to check the list on this page and see if you are vulnerable. If you are, check with the vendor to see if they've fixed the vulnerability. (Note that some early patches turned out not to fix the vulnerability.) If not, stop using the encrypted email program entirely until it's fixed. Or, if you know how to do it, turn off your email client's ability to process HTML email or—even better—stop decrypting emails from within the client. There's even more complex advice for more sophisticated users, but if you're one of those, you don't need me to explain this to you.

Consider your encrypted email insecure until this is fixed.

All software contains security vulnerabilities, and one of the primary ways we all improve our security is by researchers discovering those vulnerabilities and vendors patching them. It's a weird system: Corporate researchers are motivated by publicity, academic researchers by publication credentials, and just about everyone by individual fame and the small bug-bounties paid by some vendors.

Software vendors, on the other hand, are motivated to fix vulnerabilities by the threat of public disclosure. Without the threat of eventual publication, vendors are likely to ignore researchers and delay patching. This happened a lot in the 1990s, and even today, vendors often use legal tactics to try to block publication. It makes sense; they look bad when their products are pronounced insecure.

Over the past few years, researchers have started to choreograph vulnerability announcements to make a big press splash. Clever names—the email vulnerability is called "Efail"—websites, and cute logos are now common. Key reporters are given advance information about the vulnerabilities. Sometimes advance teasers are released. Vendors are now part of this process, trying to announce their patches at the same time the vulnerabilities are announced.

This simultaneous announcement is best for security. While it's always possible that some organization—either government or criminal—has independently discovered and is using the vulnerability before the researchers go public, use of the vulnerability is essentially guaranteed after the announcement. The time period between announcement and patching is the most dangerous, and everyone except would-be attackers wants to minimize it.

Things get much more complicated when multiple vendors are involved. In this case, Efail isn't a vulnerability in a particular product; it's a vulnerability in a standard that is used in dozens of different products. As such, the researchers had to ensure both that everyone knew about the vulnerability in time to fix it and that no one leaked the vulnerability to the public during that time. As you can imagine, that's close to impossible.

Efail was discovered sometime last year, and the researchers alerted dozens of different companies between last October and March. Some companies took the news more seriously than others. Most patched. Amazingly, news about the vulnerability didn't leak until the day before the scheduled announcement date. Two days before the scheduled release, the researchers unveiled a teaser—honestly, a really bad idea—which resulted in details leaking.

After the leak, the Electronic Frontier Foundation posted a notice about the vulnerability without details. The organization has been criticized for its announcement, but I am hard-pressed to find fault with its advice. (Note: I am a board member at EFF.) Then, the researchers published—and lots of press followed.

All of this speaks to the difficulty of coordinating vulnerability disclosure when it involves a large number of companies or—even more problematic—communities without clear ownership. And that's what we have with OpenPGP. It's even worse when the bug involves the interaction between different parts of a system. In this case, there's nothing wrong with PGP or S/MIME in and of themselves. Rather, the vulnerability occurs because of the way many email programs handle encrypted email. GnuPG, an implementation of OpenPGP, decided that the bug wasn't its fault and did nothing about it. This is arguably true, but irrelevant. They should fix it.

Expect more of these kinds of problems in the future. The internet is shifting from a set of systems we deliberately use—our phones and computers—to a fully immersive internet-of-things world that we live in 24/7. And like this email vulnerability, vulnerabilities will emerge through the interactions of different systems. Sometimes it will be obvious who should fix the problem. Sometimes it won't be. Sometimes it'll be two secure systems that, when interact in a particular way, cause an insecurity. In April, I wrote about a vulnerability that arose because Google and Netflix make different assumptions about email addresses. I don't even know who to blame for that one.

It gets even worse. Our system of disclosure and patching assumes that vendors have the expertise and ability to patch their systems, but that simply isn't true for many of the embedded and low-cost internet of things software packages. They're designed at a much lower cost, often by offshore teams that come together, create the software, and then disband; as a result, there simply isn't anyone left around to receive vulnerability alerts from researchers and write patches. Even worse, many of these devices aren't patchable at all. Right now, if you own a digital video recorder that's vulnerable to being recruited for a botnet—remember Mirai from 2016?—the only way to patch it is to throw it away and buy a new one.

Patching is starting to fail, which means that we're losing the best mechanism we have for improving software security at exactly the same time that software is gaining autonomy and physical agency. Many researchers and organizations, including myself, have proposed government regulations enforcing minimal security-standards for internet-of-things devices, including standards around vulnerability disclosure and patching. This would be expensive, but it's hard to see any other viable alternative.

Getting back to email, the truth is that it's incredibly difficult to secure well. Not because the cryptography is hard, but because we expect email to do so many things. We use it for correspondence, for conversations, for scheduling, and for recordkeeping. I regularly search my 20-year email archive. The PGP and S/MIME security protocols are outdated, needlessly complicated and have been difficult to properly use the whole time. If we could start again, we would design something better and more user friendly—but the huge number of legacy applications that use the existing standards mean that we can’t. I tell people that if they want to communicate securely with someone, to use one of the secure messaging systems: Signal, Off-the-Record, or—if having one of those two on your system is itself suspicious—WhatsApp. Of course they're not perfect, as last week's announcement of a vulnerability (patched within hours) in Signal illustrates. And they're not as flexible a email—but that makes them easier to secure.

Topics: Cybersecurity

Bruce Schneier is an internationally renowned security technologist, called a “security guru” by The Economist. He is the author of 12 books — including "Data and Goliath: The Hidden Battles to Collect Your Data and Control Your World" — as well as hundreds of articles, essays, and academic papers. His influential newsletter “Crypto-Gram” and blog “Schneier on Security” are read by over 250,000 people. Schneier is a fellow at the Berkman

Center for Internet and Society at Harvard Law School, a program fellow at the New America Foundation’s Open Technology Institute, a board member of the Electronic Frontier Foundation, and an Advisory Board member of the Electronic Privacy Information Center. He is also the Chief Technology Officer of Resilient Systems, Inc.