Everything Is Broken
Quinn Norton
A journalist of Hackers, Bodies, Technologies and Internets. ‘’Useless in terms of… tactical details’’ -Stratfor
==========

Once upon a time, a friend of mine accidentally took over thousands of computers. He had found a vulnerability in a piece of software and started playing with it. In the process, he figured out how to get total administration access over a network. He put it in a script, and ran it to see what would happen, then went to bed for about four hours. Next morning on the way to work he checked on it, and discovered he was now lord and master of about 50,000 computers. After nearly vomiting in fear he killed the whole thing and deleted all the files associated with it. In the end he said he threw the hard drive into a bonfire. I can’t tell you who he is because he doesn’t want to go to Federal prison, which is what could have happened if he’d told anyone that could do anything about the bug he’d found. Did that bug get fixed? Probably eventually, but not by my friend. This story isn’t extraordinary at all. Spend much time in the hacker and security scene, you’ll hear stories like this and worse.

It’s hard to explain to regular people how much technology barely works, how much the infrastructure of our lives is held together by the IT equivalent of baling wire.

Computers, and computing, are broken.

Build it badly, and they will come.

For a bunch of us, especially those who had followed security and the warrantless wiretapping cases, the revelations weren’t big surprises. We didn’t know the specifics, but people who keep an eye on software knew computer technology was sick and broken. We’ve known for years that those who want to take advantage of that fact tend to circle like buzzards. The NSA wasn’t, and isn’t, the great predator of the internet, it’s just the biggest scavenger around. It isn’t doing so well because they are all powerful math wizards of doom.

> The NSA is doing so well because software is bullshit.

Eight months before Snowden’s first revelation I tweeted this:

> Sec spoiler alert: Everything's got 0days, everyone's tracked, all the data leaks, all the things are vulnerable. It's all fucking pwned.
> — @quinnnorton

It was my exasperated acknowledgement that looking for good software to count on has been a losing battle. Written by people with either no time or no money, most software gets shipped the moment it works well enough to let someone go home and see their family. What we get is mostly terrible.

Software is so bad because it’s so complex, and because it’s trying to talk to other programs on the same computer, or over connections to other computers. Even your computer is kind of more than one computer, boxes within boxes, and each one of those computers is full of little programs trying to coordinate their actions and talk to each other. Computers have gotten incredibly complex, while people have remained the same gray mud with pretensions of godhood.

Your average piece-of-shit Windows desktop is so complex that no one person on Earth really knows what all of it is doing, or how.

Now imagine billions of little unknowable boxes within boxes constantly trying to talk and coordinate tasks at around the same time, sharing bits of data and passing commands around from the smallest little program to something huge, like a browser — that’s the internet. All of that has to happen nearly simultaneously and smoothly, or you throw a hissy fit because the shopping cart forgot about your movie tickets.

We often point out that the phone you mostly play casual games on and keep dropping in the toilet at bars is more powerful than all the computing we used to go to space for decades.

> NASA had a huge staff of geniuses to understand and care for their software. Your phone has you.

Plus a system of automatic updates you keep putting off because you’re in the middle of Candy Crush Saga every time it asks.

Because of all this, security is terrible. Besides being riddled with annoying bugs and impossible dialogs, programs often have a special kind of hackable flaw called 0days by the security scene. No one can protect themselves from 0days. It’s their defining feature — 0 is the number of days you’ve

had to deal with this form of attack. There are meh, not-so-terrible 0days, there are very bad 0days, and there are catastrophic 0days that hand the keys to the house to whomever strolls by. I promise that right now you are reading this on a device with all three types of 0days. “But, Quinn,” I can hear you say, “If no one knows about them how do you know I have them?” Because even okay software has to work with terrible software. The number of people whose job it is to make software secure can practically fit in a large bar, and I’ve watched them drink. It’s not comforting. It isn’t a matter of if you get owned, only a matter of when.

Look at it this way – every time you get a security update (seems almost daily on my Linux box), whatever is getting updated has been broken, lying there vulnerable, for who-knows-how-long. Sometimes days, sometimes years. Nobody really advertises that part of updates. People say “You should apply this, it’s a critical patch!” and leave off the “…because the developers fucked up so badly your children’s identities are probably being sold to the Estonian Mafia by smack addicted script kiddies right now.”

The really bad bugs (and who knows which ones those are when they click the “Restart Later” button?) can get swept up by hackers, governments, and other horrors of the net that are scanning for versions of software they know they can exploit. Any computer that shows up in a scan saying “Hey! Me! I’m vulnerable!” can become part of a botnet, along with thousands, or hundreds of thousands of other computers. Often zombied computers get owned again and become part of yet another botnet. Some botnets patch computers to throw out the other botnets so they don’t have to share you with other hackers. How can you tell if this is happening? You can’t! Have fun wondering if you’re getting your online life rented out by the hour!

____
Next time you think your grandma is uncool, give her credit for her time helping dangerous Russian criminals extort money from offshore casinos with DDoS attacks.
____

Recently an anonymous hacker wrote a script that took over embedded Linux devices [http://internetcensus2012.bitbucket.org/paper.html]. These owned computers scanned the whole rest of the internet and created a survey that told us more than we’d ever known about the shape of the internet. The little hacked boxes reported their data back (a full 10 TBs) and quietly deactivated the hack. It was a sweet and useful example of someone who hacked the planet to shit. If that malware had actually been malicious, we would have been so fucked.

This is because all computers are reliably this bad: the ones in hospitals and governments and banks, the ones in your phone, the ones that control light switches and smart meters and air traffic control systems. Industrial computers that maintain infrastructure and manufacturing are even worse. I don’t know all the details, but those who do are the most alcoholic and nihilistic people in computer security. Another friend of mine accidentally shut down a factory with a malformed ping at the beginning of a pen test. For those of you who don’t know, a ping is just about the smallest request you can send to another computer on the network. It took them a day to turn everything back on.

Computer experts like to pretend they use a whole different, more awesome class of software that they understand, that is made of shiny mathematical perfection and whose interfaces happen to have been shat out of the business end of a choleric donkey. This is a lie. The main form of security this offers is through obscurity – so few people can use this software that there’s no point in building tools to attack it. Unless, like the NSA, you want to take over sysadmins.

A well written encrypted chat, what could go wrong?

Let’s take an example computer experts like to stare down their noses at normal people for not using: OTR. OTR, or Off The Record messaging, sneaks a layer of encryption inside normal plain text instant messaging. It’s like you got on AIM or Jabber or whatever and talked in code, except the computer is making the code for you. OTR is clever and solid, it’s been examined carefully, and we’re fairly sure it hasn’t got any of those nasty 0days.

____
Except, OTR isn’t a program you use, as such.
____

There is a standard for OTR software, and a library, but it doesn’t do anything on its own. It gets implemented in software for normal human shlubs to use by other normal human shlubs. By now, you know this ends in tears.

The main thing that uses OTR is another piece of software that uses a library called libpurple. If you want to see infosec snobs look as distressed as the donkeys that shit out their interfaces, bring up libpurple. Libpurple [https://developer.pidgin.im/wiki/WhatIsLibpurple] was written in a programming language called C.

____

C is good for two things: being beautiful and creating catastrophic 0days in memory management.

———

Heartbleed [http://heartbleed.com/], the bug that affected the world over, leaking password and encryption keys and who knows what? Classic gorgeous C.

Libpurple was written by people who wanted their open source chat client to talk to every kind of instant messaging system in the world, and didn't give a shit about security or encryption. Security people who have examined the code have said there are so many possible ways to exploit libpurple there is probably no point in patching it. It needs to be thrown out and rewritten from scratch. These aren't bugs that let someone read your encrypted messages, they are bugs that let someone take over your whole computer, see everything you type or read and probably watch you pick your nose on your webcam.

This lovely tool, OTR, sits on top of libpurple on most systems that use it. Let me make something clear, because even some geeks don't get this: it doesn't matter how good your encryption is if your attacker can just read your data off the screen with you, and I promise they can. They may or may not know how to yet, but they can. There are a hundred libpurples on your computer: little pieces of software written on a budget with unrealistic deadlines by people who didn't know or didn't care about keeping the rest of your system secure.

Any one of these little bugs will do when it comes to taking over everything else on your computer. So we update and update, and maybe that throws any intruders out, and maybe it doesn't. No one knows!

———

When we tell you to apply updates we are not telling you to mend your ship. We are telling you to keep bailing before the water gets to your neck.

———

To step back a bit from this scene of horror and mayhem, let me say that things are better than they used to be. We have tools that we didn't in the 1990s, like sandboxing, that keep the idiotically written programs where they can't do as much harm. (Sandboxing keeps a program in an artificially small part of the computer, cutting it off from all the other little programs, or cleaning up anything it tries to do before anything else sees it.)

Certain whole classes of terrible bugs have been sent the way of smallpox. Security is taken more seriously than ever before, and there's a network of people responding to malware around the clock. But they can't really keep up. The ecosystem of these problems is so much bigger than it was even ten years ago that it's hard to feel like we're making progress.

———

People, as well, are broken.

———

"I trust you…" was my least favorite thing to hear from my sources in Anonymous. Inevitably it was followed by some piece of information they shouldn't have been telling me. It is the most natural and human thing to share something personal with someone you are learning to trust. But in exasperation I kept trying to remind Anons they were connecting to a computer, relaying though countless servers, switches, routers, cables, wireless links, and finally to my highly targeted computer, before they were connecting to another human being. All of this was happening in the time it takes one person to draw in a deep, committal breath. It's obvious to say, but bears repeating: humans were not built to think this way.

Everyone fails to use software correctly. Absolutely everyone fucks up. OTR doesn't encrypt until after the first message, a fact that leading security professionals and hackers subject to 20-country manhunts consistently forget. Managing all the encryption and decryption keys you need to keep your data safe across multiple devices, sites, and accounts is theoretically possible, in the same way performing an appendectomy on yourself is theoretically possible. This one guy did it once in Antarctica [http://www.southpolestation.com/trivia/igy1/appendix.html], why can't you?

Every malware expert I know has lost track of what some file is, clicked on it to see, and then realized they'd executed some malware they were supposed to be examining. I know this because I did it once with a PDF I knew had something bad in it. My friends laughed at me, then all quietly confessed they'd done the same thing. If some of the best malware reversers around can't keep track of their malicious files, what hope do your parents have against that e-card that is allegedly from you?

Executable mail attachments (which includes things like Word, Excel, and PDFs) you get just about everyday could be from anyone — people can write anything they want in that From: field of emails, and any of those attachments could take over your computer as handily as an 0day. This is probably how your grandmother ended up working for Russian criminals, and why your competitors anticipate all your product plans. But if you refuse to open attachments you aren't going to be able to keep an office job in the modern world. There's your choice: constantly risk clicking on dangerous malware, or live under an overpass, leaving notes on the lawn of your former house telling your children you love them and miss them.

Security and privacy experts harangue the public about metadata and networked sharing, but keeping track of these things is about as natural as doing blood panels on yourself every morning, and about as easy. The risks on a societal level from giving up our privacy are terrible. Yet the consequences of not doing so on an individual basis are immediately crippling. The whole thing is a shitty battle of attrition between what we all want for ourselves and our families and the ways we need community to survive as humans — a Mexican stand off monetized by corporations and monitored by governments.

I live in this stuff, and I'm no better. Once I had to step through a process to verify myself to a secretive source. I had to take a series of pictures showing my location and the date. I uploaded them, and was allowed to proceed with my interview. It turns out none of my verification had come through, because I'd failed to let the upload complete before nervously shutting down my computer. "Why did you let me through?" I asked the source. "Because only you would have been that stupid," my source told me.

Touché.

But if I can't do this, as a relatively well trained adult who pays attention to these issues all the damn time, what chance do people with real jobs and real lives have?

____

In the end, it's culture that's broken.

____

A few years ago, I went to several well respected people who work in privacy and security software and asked them a question.

First, I had to explain something:

____

"Most of the world does not have install privileges on the computer they are using."

____

That is, most people using a computer in the world don't own the computer they are using. Whether it's in a cafe, or school, or work, for a huge portion of the world, installing a desktop application isn't a straightforward option. Every week or two, I was being contacted by people desperate for better security and privacy options, and I would try to help them. I'd start, "Download th…" and then we'd stop. The next thing people would tell me was that they couldn't install software on their computers. Usually this was because an IT department somewhere was limiting their rights as a part of managing a network. These people needed tools that worked with what they had access to, mostly a browser.

So the question I put to hackers, cryptographers, security experts, programmers, and so on was this: What's the best option for people who can't download new software to their machines? The answer was unanimous: nothing. They have no options. They are better off talking in plaintext I was told, "so they don't have a false sense of security." Since they don't have access to better software, I was told, they shouldn't do anything that might upset the people watching them. But, I explained, these are the activists, organizers, and journalists around the world dealing with governments and corporations and criminals that do real harm, the people in real danger. Then they should buy themselves computers, I was told.

That was it, that was the answer: be rich enough to buy your own computer, or literally drop dead. I told people that wasn't good enough, got vilified in a few inconsequential Twitter fights, and moved on.

Not long after, I realized where the disconnect was. I went back to the same experts and explained: in the wild, in really dangerous situations — even when people are being hunted by men with guns — when encryption and security fails, no one stops talking. They just hope they don't get caught.

____

The same human impulse that has kept lotteries alive for thousands of years keeps people fighting the man against the long odds. "Maybe I'll get away with it, might as well try!"

____

As for self-censoring their conversations in the face of hostile infrastructure, non-technical activists are just as good at it as Anons are, or people told to worry about metadata, or social media sharing, or that first message before OTR encryption kicks in. They blow.

This conversation was a wake-up call for some security people who hadn't realized that people who become activists and journalists routinely do risky things. Some of them joined my side of the time-wasting inconsequential Twitter fights, realizing that something, even something imperfect, might be better than nothing. But many in the security scene are still waiting for a perfect world into which to deploy their perfect code.

Then there's the Intelligence Community, who call themselves the IC. We might like it if they stopped

spying on everyone all the time, while they would like us to stop whining about it.

After spending some time with them, I am pretty sure I understand why they don't care about the complaining. The IC are some of the most surveilled humans in history. They know everything they do is gone over with a fine-toothed comb – by their peers, their bosses, their lawyers, other agencies, the president, and sometimes Congress. They live watched, and they don't complain about it.

In all the calls for increased oversight, the basics of human nature gets neglected. You're not going to teach the spooks this is wrong by doing it to them more.

There will always be loopholes and as long as loopholes exist or can be constructed or construed, surveillance will be as prevalent as it possibly can be. Humans are mostly egocentric creatures. Spooks, being humans, are never going to know why living without privacy is bad as long as they are doing it.

Yet that's the lesser problem. The cultural catastrophe is what they're doing to make their job of spying on everyone easier. The most disturbing parts of the revelations are the 0day market, exploit hoarding, and weakening of standards. The question is who gets to be part of the "we" that are being kept allegedly safe by all this exploiting and listening and decrypting and profiling. When they attacked Natanz with Stuxnet and left all the other nuclear facilities vulnerable, we were quietly put on notice that the "we" in question began and ended with the IC itself. That's the greatest danger.

_____

When the IC or the DOD or the Executive branch are the only true Americans, and the rest of us are subordinate Americans, or worse the non-people that aren't associated with America, then we can only become lesser people as time goes on.

_____

As our desires conflict with the IC, we become less and less worthy of rights and considerations in the eyes of the IC. When the NSA hoards exploits and interferes with cryptographic protection for our infrastructure, it means using exploits against people who aren't part of the NSA just doesn't count as much. Securing us comes after securing themselves.

In theory, the reason we're so nice to soldiers, that we have customs around honoring and thanking them, is that they're supposed to be sacrificing themselves for the good of the people. In the case of the NSA, this has been reversed. Our wellbeing is sacrificed to make their job of monitoring the world easier. When this is part of the culture of power, it is well on its way to being capable of any abuse.

But the biggest of all the cultural problems still lies with the one group I haven't taken to task yet – the normal people living their lives under all this insanity.

The problem with the normals and tech is the same as the problem with the normals and politics, or society in general. People believe they are powerless and alone, but the only thing that keeps people powerless and alone is that same belief. People, working together, are immensely and terrifyingly powerful.

_____

There is certainly a limit to what an organized movement of people who share a mutual dream can do, but we haven't found it yet.

_____

Facebook and Google seem very powerful, but they live about a week from total ruin all the time. They know the cost of leaving social networks individually is high, but en masse, becomes next to nothing. Windows could be replaced with something better written. The US government would fall to a general revolt in a matter of days. It wouldn't take a total defection or a general revolt to change everything, because corporations and governments would rather bend to demands than die. These entities do everything they can get away with – but we've forgotten that we're the ones that are letting them get away with things.

Computers don't serve the needs of both privacy and coordination not because it's somehow mathematically impossible. There are plenty of schemes that could federate or safely encrypt our data, plenty of ways we could regain privacy and make our computers work better by default. It isn't happening now because we haven't demanded that it should, not because no one is clever enough to make that happen.

So yes, the geeks and the executives and the agents and the military have fucked the world. But in the end, it's the job of the people, working together, to unfuck it.