Hacked Hardware Could Cause The Next Big Security Breach

Microchips govern our homes, cities, infrastructure, and military. What happens when they're turned against us?

By P.W. Singer Posted February 17, 2015

In late summer of 2006, the Japanese division of McDonald's decided to run a new promotion. When customers ordered a Coca-Cola soft drink, they would receive a cup with a code. If they entered that code on a designated website and were among 10,000 lucky winners, they would receive an MP3 player pre-loaded with 10 songs.

Cleverly constructed, the promotion seemed destined for success. Who doesn't like a Coke and a free MP3 player? But there was one problem the marketers at McDonald's could not anticipate: In addition to 10 free songs, the music players contained QQPass malware. The moment winners plugged their players into a computer, the Trojan horse slipped undetected into their system and began logging keystrokes, collecting passwords, and gathering personal data for later transmission.

McDonald's eventually recalled the devices and issued an apology, but not before an unknown number of users had fallen prey to the malware. In the annals of fast food promotions, the incident is still regarded as one of the worst of all time (even beating the ill-conceived McAfrika burger—an African–inspired sandwich released at the height of a famine). For security professionals, it was notable too, but for entirely different reasons: It offered a terrifying glimpse at how hackers could build a cyberattack directly into the very systems we depend on.

In the past year, cybercrime has blossomed into a pandemic, consuming more than $445 billion in lost time, jobs, and intellectual property. Hackers compromised 233 million personal records from eBay; they intimidated Sony into scuttling the theatrical release of *The Interview*; they even commandeered the Pentagon's Twitter account. But as varied as those assaults were, they shared a trait: Somone hacked software to penetrate a network or account. What set the McDonald's incident apart—and what strikes fear into cybersecurity professionals everywhere—is that the perpetrator hacked hardware instead.

In computing terminology, hardware boils down to microchips, the integrated circuits that run our devices. They are in our phones, refrigerators, electric grids, planes, and missiles. And many more are on the way. Cisco estimates that more than 50 billion Internet-connected devices will come online by 2020, all communicating ceaselessly with the world around them.

Microchips are the bedrock upon which our digital world is based, and they are almost entirely unsecured. Whereas software security is on pace to become a $156 billion industry in the next five years, hardware security gets relatively little mention. Yet the challenges hardware presents are in many ways more extensive, more dangerous, and more difficult to combat. When the

marketers at McDonald's ordered their MP3 players, they simply chose a device from a catalog. It just happened that someone at a production line in Hong Kong decided to load it with malware. We'll likely never know why that person chose those particular MP3 players, and that's not really the point. This kind of attack could have hit anywhere hardware exists, from coffeemakers to fighter jets, and the consequences could have been much, much worse.

Problem From Hell

When Jack Kilby of Texas Instruments invented the first integrated circuit in 1958 (for which he later won a Nobel Prize), the age of the microchip was born. These early processors cost $450 and consisted of a few transistors, diodes, resistors, and capacitors placed onto a slice of germanium and linked by gold wires. The unit was about 10 millimeters across.

Today's microchips follow the same principles but are exponentially more complex. They consist of billions of transistors and are divided into multiple sub-units (called "blocks," as Kilby first labeled them), each of which carries out a specific function. A smartphone's processor, for example, may have some blocks whose purpose is to store frames of video and others to convert them so they can be sent over an antenna.

As the nature and complexity of chips has changed in five and a half decades, so too has their design and manufacture. In the 1970s and '80s, there were just a handful of known and trusted chip designers; now there are a huge number of companies creating more than 5,000 new designs each year, spread from the U.S. to Asia. These teams, in turn, involve hundreds or thousands of people at multiple locations—each working on different blocks. Chips have become so intricate that no one person can see, let alone understand, every detail of their architecture.

These developments have, by and large, been positive. The more powerful our microchips, the more capabilities we have. But when such complexity is paired with massive scale—$333 billion–worth of chips were sold in 2014 alone—it also creates significant vulnerabilities, and an ever-more irresistible opportunity for hardware hackers. In a recent report for the Brookings Institution, John Villasenor, a professor of electrical engineering and public policy at University of California at Los Angeles, wrote, "The laws of statistics guarantee that there are people with the skills, access, and motivation to intentionally compromise a chip design." In other words, more frequent and large-scale hardware attacks are just a matter of time. And when they come, whether from a nation state, a crime syndicate, or a rogue employee, they will arrive in one of two forms: overt or covert.

Overt actions are perhaps the simpler of the two: They make it apparent that the system isn't working properly. The best example would be a so-called kill switch, in which an enemy or criminal could selectively turn off chips at will. Doing this is easier than one might think. For example, the different blocks in a chip can communicate and coordinate via a "system bus," which they take turns using so as not to create interference. If one block was corrupted so it would not give up access to the system bus—something well within reach of many mid-level chip designers—it would prevent the other blocks from getting data, effectively disabling, or bricking, the system.

Just one small corruption can have grave consequences. In 2011, faulty transistors were found in an electromagnetic interference filter destined for a U.S. Navy helicopter (an SH-60 deployed to a destroyer in the Pacific Fleet). Though never installed, that defective part would have compromised the SH-60's ability to fire its Hellfire missiles, making it practically useless in combat. The manufacturer of the filter, Raytheon, and the U.S. Senate Committee on Armed Services had to trace the transistors through five companies before finding their origin in China.

An investigation later proved the flaws were an honest production error. But had someone intentionally pursued this sort of hack, the result could have been different. More than three-quarters of the field-programmable gate arrays in the F-35 strike fighter are made in China and Taiwan. So are the majority of chips in automobiles and wireless medical devices, such as pacemakers and dialysis machines. If that hardware was modified ever so slightly, a kill code could selectively disable the chip and the systems that depend on it. And that code could come from any number of sources. A command could originate in a text or email message. It could be delivered by radio signal to a micro-antenna hidden on the chip. It could even be a simple internal time bomb, programmed at the chip's inception, to trigger a coordinated shutdown on a certain time and date, as in the first episode of *Battlestar Galactica*.

If an overt action is the equivalent of dropping a bomb, a covert one is like laying a landmine. A compromised chip may appear to function normally while secretly collecting and transmitting information, launching malware from inside the system, or even coordinating with other corrupted chips to carry out a larger attack. In 2007, for example, the Taiwanese Ministry of Justice discovered that a number of Seagate hard drives had two separate Trojans built into them by someone in the design or manufacturing process. The malware would phone home to a pair of websites hosted out of Beijing, which would then cause the hard drive to upload all its data. More recently, the Star N9500, a knockoff of the Galaxy S4 smartphone, shipped from a factory in China preloaded with a Trojan masquerading as the Google Play Store. It allowed the attackers to record phone calls, read emails, intercept financial information, and remotely watch and listen in via the phone's camera and microphone.

Even hardware generally considered innocuous could be exploited by hackers and used for covert acts. Modified third-party phone chargers have served as vehicles for malware, as have game consoles. In the world of hardware hacking, any smart device—a refrigerator, clock, even a wearable fitness monitor—could be weaponized.

Such covert actions could inflict even greater harm were they to work their way into the backbone of the Internet: the servers and other networking equipment that comprise the infrastructure of the IT world. Instead of gathering embarrassing emails from a handful of executives, hackers with compromised servers could monitor most of the world's Internet messages. As companies such as Huawei Technologies and ZTE Corporation—both of which supply telecommunication equipment and have ties to the Chinese military—continue to grow, so too will concerns about network security. Add to that, the revelations by Edward Snowden indicate the National Security Agency (NSA) has moved from hacking individual computers to network hardware.

Perhaps the most devastating form of covert attack would be one that turns kinetic. Imagine a single employee at a microchip foundry hellbent on engineering an international crisis. Knowing the foundry's chips go into drone systems, that employee could embed a malfunction into the hardware that would activate only at a certain GPS point. When the drone reaches the designated position, say in northwest Pakistan, it would fire a missile at a school or dam instead of a militant camp.

The example is a worst-case scenario but hardly inconceivable. At a cybersecurity panel at the Aspen Institute in 2011, General Michael Hayden, a retired Air Force four-star general who headed both the CIA and NSA, was asked about hardware hacking, and his response was simple: "It's the problem from hell."

Lines Of Defense

At this point, hardware hacking is still in its infancy, and so too are solutions to it. Chip designers primarily rely on protocols that have not appreciably changed for years. For that reason, Villasenor wrote in 2010, "Defensive strategies have not yet been fully developed, much less put into practice."

And so protection for consumers at this point comes down to common sense: If you don't know where something is from, it's generally not a great idea to plug it into your network. The advice sounds obvious, but it bears stating that the worst hack in U.S. military history occurred when someone found a corrupted memory stick outside a base in the Middle East and plugged it into a classified network.

Beyond simple schoolyard rules, creating defenses becomes much more difficult. To stop hardware hacking at the design and manufacture stage, the Pentagon has launched its "Trusted Foundry" program. To qualify, foundries that build integrated circuits must pass a rigorous accreditation process. It's a good first step, but it affects only a small fraction of the chips the U.S. military needs, let alone the rest of us. The next step would be to expand the network of trusted chipmakers and punish companies found to be untrustworthy. But given the layers of buyers and sellers involved, that will be difficult. The researchers that detected the hack in the Star N9500 smartphone spent more than a week trying to find the source of the malicious chip, to no avail.

In the world of hardware hacking, any smart device—a refrigerator, clock, even a wearable fitness monitor—could be weaponized.

As foundries strive to improve their security, some researchers are investigating the development of digital watermarks, such as holograms or bits of plant DNA, that could be authenticated at key points in the supply chain. Other researchers are looking upstream to secure the microchip design process. More robust encryption programs could track design changes, making it harder for someone to initiate a hack in the first place.

Testing, too, requires an overhaul. Tests today are "usually designed to weed out accidental defects and design flaws, not identify parts that counterfeiters have specifically altered to masquerade as something they are not," Villasenor wrote in an article with co-author

Mohammad Tehranipoor. And only a small percentage of the millions of chips produced each year are tested anyway. To fortify this vulnerability, DARPA created the Integrity and Reliability of Integrated Circuits program. Its projects include an advanced scanning optical microscope that will use an infrared laser to probe “microelectronic circuits at nanometer levels, revealing information about chip construction as well as the function of circuits at the transistor level.”

The agency also launched the Supply Chain Hardware Integrity for Electronics Defense program. It aims to develop a dielet, a 100-micron-by-100-micron component that could be attached to chips at less than a penny per unit. It would carry an encryption engine to help secure data and sensors to detect any tampering.

Each program holds a lot of promise, but to truly safeguard hardware vulnerabilities chip designers need to rethink chips themselves. That means building defenses directly into integrated circuits. One example could be to install input and output monitors that stop chips from communicating with unauthorized connections and memory gatekeepers that prohibit access to off-limits areas. Another would be to incorporate a “no execute” bit, which cordons off certain areas of memory and prevents the processor from executing any code from there. The appetite for such solutions, however, is still very limited.

Chronic Condition

A few years ago, Cody Brocious, a 24-year-old researcher at Mozilla, began to investigate the security of the electronic room-lock systems used at many hotels, most of which can be programmed to accept master keys. At the 2012 Black Hat security conference, he showed off how to spoof a master key with little more than $50 worth of homebrewed hardware. The lock manufacturer developed a defense against this attack, but it involves replacing the hardware in more than four million locks.

In the end, that’s truly what makes hardware hacking the “problem from hell”: The potential avenues of attack are so numerous and insidious, they can be hard to contemplate. Addressing them will be neither easy nor fast—but it can be done. The challenge of software security appeared equally insurmountable at one time, but now cybersecurity professionals are doing a better job of understanding and confronting those risks than ever before. As with software, the decision to pursue hardware security will ultimately come down to cost-benefit analysis. Added defenses often come with tradeoffs, namely lower performance, increased cost, or both. Until now, the decision to adopt them has been pretty easy—don’t bother. Going forward, the thought process will change. As James Hayward, the CEO of Applied DNA Sciences, said in an interview, “A $100 microchip might keep a $100 million dollar helicopter on the ground.”

That new calculus will hopefully spur governments and companies to attack hardware vulnerabilties before criminals do. “Frankly, it’s not a problem that can be solved,” General Hayden said of hardware hacking in Aspen. “This is a condition that you have to manage.”

This article was originally published in the March 2015 issue of Popular Science, *under the title, "Nowhere To Hide."*