JUST SECURITY
A FORUM ON LAW, RIGHTS, AND U.S. NATIONAL SECURITY

One from the Vault 7: Wikileaks and the CIA's Hacking Arsenal

Julian Sanchez

March 8, 2017

It's a cliche of political scandals that "the coverup is worse than the crime": Attempts to conceal misconduct, because they're easier to prove and provide otherwise elusive evidence of a guilty mind, often end up being more politically damaging than the underlying misconduct would have been. In the case of the latest Wikileaks document dump, the first in a planned series from a cache the site has dubbed "Vault 7," we have an apparent reversal of the formula: The un-coverup—the fact of the leak itself—is probably more significant than the substance of what has thus far been revealed.

There are, of course, some points of real interest in the archive of documents, mostly concerning an array of hacking tools and software exploits developed or used by the Central Intelligence Agency's Engineering Development Group—and it's likely more will emerge as reporters and analysts churn through more than 8,000 files and documents. We've confirmed that the CIA has hung onto and exploited at least a handful of undisclosed "zero day" vulnerabilities in widely-used software platforms, including Apple's iOS and Google's Android, the operating systems on which nearly all modern smartphones run.

We also learn that—as many of us expected—the obstacles to conventional wiretapping posed by the growing prevalence of encryption have spurred intelligence agencies to hunt for alternative means of collection, which include not only compromising communications endpoints such as smartphones, but also seeking to repurpose networked appliances on the Internet of Things as surveillance devices. The latter goal has even spawned its own research department, the Embedded Development Branch.

Still, in light of what we already knew about the National Security Agency's own efforts along similar lines, thanks to Edward Snowden's disclosures about the agency's Tailored Access Operations division, this is—at least from a policy perspective—not so much revelation as confirmation. Moreover, there's little here to suggest surveillance that's either aimed at Americans or indiscriminate, the features that made Snowden's leaks about NSA surveillance so politically explosive. One of the more widely-reported projects in Vault 7, for instance, has been the Doctor Who–referencing "Weeping Angel" implant, which can turn Samsung televisions into surveillance microphones even when they *appear* to be turned off. Yet, at least at the time the

documentation in the Wikileaks release was written, Weeping Angel appeared to require physical access to be installed—which makes it essentially a fancy and less detectable method of bugging a particular room once a CIA agent has managed to get inside. This is all fascinating to surveillance nerds, to be sure, but without evidence that these tools have been deployed either against inappropriate targets or on a mass scale, it's not intrinsically all that controversial. Finding clever ways to spy on people is what spy agencies are *supposed* to do.

What is genuinely embarrassing for the intelligence community, however, is the fact of the leak itself—a leak encompassing not only thousands of pages of documentation but, according to Wikileaks, the actual source code of the hacking tools those documents describe. While Wikileaks has not yet published that source code, they claim that the contents of Vault 7 have been circulating "among former U.S. government hackers and contractors in an unauthorized manner," which if true would make it far more likely that other parties—such as foreign intelligence services—had been able to obtain the same information. Worse, this comes just months after the even more disastrous Shadow Brokers leak, which published a suite of exploits purportedly used by the NSA-linked Equation Group to compromise the routers and firewalls relied upon by many of the world's largest companies to secure their corporate networks.

That's of great significance for the ongoing debate over how intelligence agencies should respond when they discover vulnerabilities in widely-used commercial software or firmware. Do they inform the vendor that they've got a security hole that could put their users at risk, or do they keep quiet and make use of the vulnerability to enable their own surveillance? If the latter, how long do they wait until disclosing? In 2014, the White House's cybersecurity czar attempted to reassure the public that the government's mechanism for making such decisions—an informal "Vulnerability Equities Process" designed to weigh the intelligence benefit of keeping an exploit against the public's interest in closing security holes—was strongly biased in favor of disclosure. The number of critical vulnerabilities we now know have remained undisclosed, sometimes for years, should cast serious doubt on that assertion. But the *means* by which we know it should strengthen the case for disclosure still further.

Prior to the Shadow Brokers leak, the primary concern of security experts had been that the longer a software vulnerability is kept secret by spy agencies, the greater the risk that some malicious actor—whether a criminal hacker or another intelligence agency—would independently discover and use it. Now, however, we need to factor in the growing evidence that the Intelligence Community cannot properly secure its own hacking tools. And breaches this sort create significantly higher risks, because they result in the wide circulation, not just of individual vulnerabilities that might be of limited use to an attacker in isolation, but whole suites of them, already in weaponized form, and conveniently chained together for easy one-stop hacking. One such breach might be shrugged off as an aberrant lapse. Two—*that the public is aware of*—in the span of eight months suggest a more systematic problem. And since foreign intelligence agencies are likely to be more interested in *using* stolen cyberweapons than gifting them to the world, it seems a reasonable inference that the two publicly known instances of large-scale exfiltration aren't the only such cases.

That ought to make the public a whole lot more skittish about the prospect that a myopic focus on maintaining intelligence accesses is making all of us significantly less secure on net. And it

ought to prompt some serious reevaluation within the government about whether their purported bias in favor of disclosure shouldn't be a whole lot stronger.

Julian Sanchez is a senior fellow at the Cato Institute and contributing editor for Reason magazine.