ENCRYPTION

Warning Signs: A Checklist for Recognizing Flaws of Proposed "Exceptional Access" Systems

By **Daniel J. Weitzner** Wednesday, May 11, 2016, 12:00 PM

Author's note: Despite appearing under my byline, this post actually represents the work of a larger group. The Keys Under Doormats group includes Harold Abelson, Ross Anderson, Steven M. Bellovin, Josh Benaloh, Matt Blaze,Whitfield Diffie, John Gilmore, Matthew Green, Susan Landau, Peter G. Neumann, Ronald L. Rivest, Jeffrey I. Schiller, Bruce Schneier, Michael A. Specter, Daniel J.Weitzner,who jointly authored the report "Keys Under Doormats: Mandating Insecurity" last year.The following is a follow-up in light of recent events.

In the eighteen months since FBI Director James Comey raised alarm bells about encryption and surveillance, there have been many calls for the technology community to solve the problem. Director Comey's call to action was a genuine statement of law enforcement concern but sparse on operational details. However, technical security analysis of any proposal necessarily relies on such details. Some technologists have begun to offer ideas on how to solve the exceptional access problem; indeed last week, Matt Tait proposed a scheme for providing exceptional access to encrypted data on devices such as smartphones. At a recent House Energy and Commerce hearing on the matter, Captain Charles Cohen, Commander of the Office of Intelligence and Investigative Technologies at the Indiana State Police made a similar suggestion in his testimony.

Tait's proposal seems straightforward: give each law enforcement agency a public key and have the user's private key stored on the device encrypted with that public key together with the public key of the manufacturer. Because you need both the agency's private key and the manufacturer's private key to decrypt the user's private key, Tait believes this would keep the user's private key secure. It would also guarantee that law enforcement could access plaintext by showing the manufacturer a warrant.

This particular proposal has many of the general risks that we warned about in our paper. As it is hard to evaluate every design idea that comes along, we've distilled some of the general problems associated with exceptional access systems into a short list of warning signs to look out for in any new proposal. We are not attempting to be complete; rather, we are providing a condensed checklist of crucial vulnerabilities to watch for.

1. Watch for systems that rely on a single powerful key or a small set of them.

2. Watch for systems using high-value keys over and over and still claiming not to increase risk.

3. Watch for the claim that the abstract algorithm alone is the measure of system security.

4. Watch for the assumption that scaling anything on the global Internet is easy.

5. Watch for the assumption that national borders are not a factor.

6. Watch for the assumption that human rights and the rule of law prevail throughout the world.

Any proposal for exceptional access must be able to answer these questions, so policy-makers or systems designers should consider these six warning signs when looking at new technical proposals.

1. Watch for systems that rely on a single powerful key or a small set of them.

Proposals that require both law enforcement and platform vendors to keep keys for providing access to encrypted data make those keys into very valuable targets for hackers. The keys created for law enforcement access become sources of system-wide vulnerability. Systems that rely on a "master" key that protects the security of millions of people or billions of dollars are huge targets. Tait's proposal depends on each platform provider and each law enforcement agency having a master key but fails to consider the risk of key compromise and key loss. In Tait's design, an attacker would have to get access to both a key from law enforcement and from the smartphone company, but this is not farfetched. From a US viewpoint, a smartphone made by a Chinese company has its password encrypted only once relative to Chinese intelligence. From a US-LE viewpoint, the key is dubiously accessible; what are our warrant agreements with China? Would the key have to be held by a Chinese subsidiary in the US.

History shows that even keys from governments and major companies can be stolen. Recent examples of private keys being stolen and abused exist in the DigiNotar hack, StuxNet's use of a stolen RealTek certificate, and any number of thefts of bitcoin wallets. The lesson is that any organization with a valuable private key becomes a high-value intelligence and/or organized crime target. Today, companies are attempting to avoid the risk of stolen keys by moving to specialized hardware security modules. Tait's proposal contains no explanation of how to do this in a way that manages the new and vulnerable keys he proposes to add to systems all around the world. Keys might conceivably be lost too. The loss of a key would

either make some class of phones more desirable as unescrowed or force the users to get new ones. Would governments then keep the loss secret or force people to update their phones? What would a company's liability be if it lost its key?

2. Watch for systems using high-value keys over and over and still claim not to increase risk.

Any key required for exceptional access will need to be used frequently by law enforcement agencies with varying degrees of technical skill, from all around the world. Software and hardware companies do have experience protecting high value keys for tasks such as code-signing, those keys are used relatively infrequently, perhaps a dozen or so times per year, and are under the exclusive control of one company. The existing keys that come closest to the proposed master keys are the keys banks use to link customer personal identification numbers (PINs) to bank account numbers. These keys are kept in tamper-responding hardware security modules and used only in very restricted, automated, and audited ways. Even so there are regular security failures.

Master keys used for exceptional access would not only be used frequently – probably several times per day – but would be used by the law enforcement and intelligence agencies of numerous countries. Even if these keys are stored on hardware security modules, the systems that drive the modules will need to be implemented on computers, which would immediately be among the highest-priority targets for the world's intelligence agencies, drug cartels, and other well funded criminal syndicates. If these computers were connected to the Internet, managed by people who are not security experts, and running software systems with zero-day vulnerabilities, there would be a very high probability of key compromise. If they were classified systems produced in government labs, which other government would trust them? In many cases, key compromises may not even be detected. The public should be made aware that systems such as these have built-in insecurities and are likely to be rapidly compromised.

3. Watch for the claim that the abstract algorithm alone is the measure of system security.

A secure system depends on two distinct elements. First, a secure system will use strong cryptographic algorithms, the complex math that is the basis for encrypting (scrambling) and decrypted (unscrambling) valuable data. Second, the engineering of the entire system--including everything from the underlying hardware, operating system, and programming

languages--must be resistant to tampering, and the abstract algorithm must be implemented correctly. Both the math and the engineering must be done properly in order to protect user's data.

Tait's proposal is based on a well-known cryptographic algorithm. This basic concept—nested layers of encryption—is not new, having been part of many schemes dating back to the Clipper Chip, and used by a variety of other systems, including Tor. Cryptographers can analyze that algorithm and prove that it is secure against attack in the ideal world. But when it was engineered in the real world, errors emerged. Governments proposed other designs, but these also turned out to have vulnerabilities. Can we reasonably expect current implementations to be better than those developed by the agencies twenty years ago? Many of the practical advances in computer security since then have come from learning what *not* to do. Specifically, keeping many keys around can increase the potential of loss and abuse. Anyone proposing an abstract design for exceptional access has to show how it is possible to actually deploy it securely. Cryptographic systems are subtle; the errors generally appear in the detailed engineering process, even when the underlying mathematical algorithm is proven to be secure. So don't expect new wine to come out of old bottles.

4. Watch for the assumption that scaling anything on the global Internet is easy.

Designs on paper may seem simple, but security flaws tend to appear when they are implemented at large scale across the global Internet. That one large company can implement a simple system for its individual customers in one country doesn't explain how to get competing providers with different hardware, software, and business models to agree on a single design. And even if we might have confidence that one or two big tech companies could implement such a system on their own, the challenge of replicating that accomplishment in myriad hardware and software contexts is much harder.

Turning algorithms into protocols, turning protocols into code, and integrating code into products are each challenging and expensive steps in designing, building and deploying any large system. Tait does not explain how this system would actually be reduced to a set of technical standards implementable by thousands of device manufacturers and deployable by millions of firms in over a hundred mutually suspicious countries. Global scale systems are never perfect when first deployed. The successful ones, such as the Internet and the World Wide Web, evolve progressively as problems are found and fixed; the fixes introduce more complexity that must be dealt with in turn. Then these systems have to be tested for security vulnerabilities and fixed worldwide whenever they fail. We also know from the Heartbleed

SSL flaw that even widely-used code can harbor security risks for years without being detected. Tait offers no account of which capable actors would be sufficiently motivated to guard and update the security of his proposed exceptional access system.

5. Watch for the assumption that national borders are not a factor.

Perhaps the most intractable problem with Tait's proposal is that of managing keys across national borders. What really happens when a phone crosses a border? Do the encryption keys automatically change by some undescribed mechanism? Does an owner hand over her phone at customs to have new keys put in and the old ones removed?

Country-specific wiretapping was relatively easy when all phones were wired and immobile. For mobile phones without encryption, it's not much harder; a mobile call between countries can be tapped at both ends. But as the encryption mechanisms move from the network to the phones themselves, and to the servers with which they communicate, security becomes far more complex. Country-specific access keys inserted at the point of sale are not sufficient. Some countries won't have the necessary infrastructure or may reject the very concept of a lawful access feature (the Netherlands comes to mind). The scenario of foreign-trained terrorists entering the US is high on the FBI threat list. The Paris terrorists brought their own phones with them (along with their own encryption programs), and such phones can be bought in countries that won't cooperate.

6. Watch for the assumption that human rights and the rule of law prevail throughout the world.

Tait acknowledges that putting exceptional access in the hands of tyrants and autocrats can be dangerous, but he waves off the seriousness of this challenge by saying that "If the device manufacturer has ethical questions about servicing the decryption, it can simply refuse to decrypt its layer." But that is a clear violation of national sovereignty, and a manufacturer who refuses to cooperate will be sanctioned. Tait mentions the "Chinese democracy activist." But iPhones are manufactured in China, and China is Apple's largest phone market. The company is in a much stronger position if it cannot comply than if it refuses to. In addition to that, the Snowden revelations showed that even developed countries' agencies were engaged in practices now recognized are contrary to the protection of human rights.

In the end, we must recognize that security is a systems property. Not only must all aspects of an engineering design be correct, their interactions must be as well. Without clear answers to the questions we've posed here, it is not possible to be confident about the

security risk associated with any exceptional access system. Nevertheless, the watch list we have outlined is should be used only as an initial guideline; if analysis seems to indicate that everything seems (barely) adequate with respect to these six items, there are still many details to consider in order to assess full system security. We don't relish the role of naysayers. We do think that applying this six-part watch list to new exceptional access proposals will help policymakers understand when they have an approach that can help law enforcement without putting vital communications and information services at risk.

Topics: Encryption, Privacy Paradox, Cybersecurity, Privacy: Technology

Tags: going dark, Encryption, Privacy

Daniel J. Weitzner is Director of the MIT Internet Policy Research Initiative and Principal Research Scientist at the MIT Computer Science and Artificial Intelligence Lab. From 2011-2012, Weitzner was United States Deputy Chief Technology Officer for Internet Policy in the White House. Weitzner's computer science research has pioneered the development of Accountable Systems architecture to enable computational treatment of legal rules and automated compliance auditing. He teaches Internet public policy in MIT's Electrical Engineering and Computer Science Department. Before joining MIT in 1998, Weitzner was founder and Deputy Director of the Center for Democracy and Technology, and Deputy Policy Director of the Electronic Frontier Foundation. Weitzner has law degree from Buffalo Law School, and a B.A. in Philosophy from Swarthmore College.

djweitzner

MORE ARTICLES >

RELATED ARTICLES

Annie Hall’s Prescient Insight on ODNI and the Don’t Panic report

Herb Lin Sat, May 14, 2016, 10:29 PM

Setting up a Straw Man: ODNI's Letter in Response to "Don't Panic"

Susan Landau Thu, May 12, 2016, 9:00 AM

The IC Thinks Harvard Is Wrong About Encryption

Paul Rosenzweig Sun, May 8, 2016, 8:47 AM

The Apple Vulnerability Disclosure Question -- Looks Like I Will Win

Paul Rosenzweig Thu, Apr 28, 2016, 10:39 AM

An Approach to James Comey's Technical Challenge

Matt Tait Wed, Apr 27, 2016, 7:00 AM