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

Feinstein-Burr 2.0: The Crypto Backdoor Bill Lives On

Julian Sanchez

September 9, 2016

When it was first released back in April, a "discussion draft" of the Compliance With Court Orders Act sponsored by Sens. Dianne Feinstein (D-CA) and Richard Burr (R-NC) met with near universal derision from privacy advocates and security experts. (Your humble author was among the critics.) In the wake of that chilly reception, press reports were declaring the bill effectively dead just weeks later, even as law enforcement and intelligence officials insisted they would continue pressing for a solution to the putative "going dark" problem that encryption creates for government eavesdroppers. Feinstein and Burr, however, appear not to have given up on their baby: Their offices have been circulating a series of proposed changes to the law, presumably in hopes of making it more palatable to stakeholders. I recently got a look at some of those proposed changes. (NB: I rather imprecisely referred to these changes collectively as a "revised draft" in an earlier version of this post. I've edited the post to more accurately characterize these as*proposed* revisions to the previously circulated draft that are currently under consideration.)

To protect my source's anonymity, I won't post any documents, but it's easy enough to summarize the four main changes I saw (though I'm told others are being considered):

(1) Narrower scope

The original discussion draft required a "covered entity" to render encrypted data "intelligible" to government agents bearing a court order if the data had been rendered unintelligible "by a feature, product, or service owned, controlled, created, or provided, by the covered entity or by a third party on behalf of the covered entity." This revision would delete "owned," "created," and "provided"—so the primary mandate now applies only to a person or company that "controls" the encryption process.

(2) Limitation to law enforcement

A second change would eliminate section (B) under the bill's definition of "court order," which obligated recipients to comply with decryption orders issued for investigations related to "foreign intelligence, espionage, and terrorism." The bill would then be strictly about *law enforcement* investigations into a variety of serious crimes, including federal drug crimes and their state equivalents.

(3) Exclusion of critical infrastructure

A new subsection in the definition of the "covered entities" to whom the bill applies would specifically exclude "critical infrastructure," adopting the definition of that term from 42 USC §5195c.

(4) Limitation on "technical assistance" obligations

The phrase "reasonable efforts" would be added to the definition of the "technical assistance" recipients can be required to provide. The original draft's obligation to provide whatever technical assistance is needed to isolate requested data, decrypt it, and deliver it to law enforcement would be replaced by an obligation to make "reasonable efforts" to do these things.

It's worth noting that I *haven't* seen any suggestion they're considering modifying the problematic mandate that distributors of software licenses, like app stores, ensure that the software they distribute is "capable of complying" with the law. (As I've argued previously, it is very hard to imagine how open-source code repositories like Github could effectively satisfy this requirement.) So what would these proposed changes amount to? Let's take them in order.

The first change would, on face, be the most significant one by a wide margin, but it's also the one I'm least confident I understand clearly. If we interpret "control" of an encryption process in the ordinary-language sense—and in particular as something conceptually distinct from "ownership," "provision," or "creation"—then the law becomes radically narrower in scope, but also fails to cover most of the types of cases that are cited in discussions of the "going dark" problem. When a user employs a device or application to encrypt data with a user-generated key, that process is not normally under the "control" of the entity that "created" the hardware or software in any intuitive sense. On the other hand, when a company *is* in direct control of an encryption process—as when a cloud provider applies its own encryption to data uploaded by a user—then it would typically (though by no means *necessarily*) retain both the ability to decrypt and an obligation to do so under existing law. So what's going on here?

One obvious possibility, assuming that narrow reading of "controlled," is that the idea is to very specifically target companies like Apple that are seeking to combine the strong security of end-to-end encryption with the convenience of cloud services. At the recent Blackhat security conference, Apple introduced their "Cloud Key Vault" system. The critical innovation there was finding a way to let users users back up and synchronize across devices some of their most sensitive data—the passwords and authentication tokens that safeguard all their *other* sensitive data—without giving Apple itself access to the information. The details are complex, but the basic idea, oversimplifying quite a bit, is that Apple's backup systems will act a like a giant iPhone: User data is protected with a combination of the user's password and a strong encryption key that's physically locked into a hardware module and can't be easily extracted. Like the iPhone, it will defend against "brute force" attacks to guess the user passcode component of the decryption key by limiting the number of permissible guesses. The critical difference is that Apple has essentially destroyed their own ability to change or eliminate that guess limit.

This may not sound like a big deal, but it addresses one of the big barriers to more widespread adoption of strong end-to-end encryption: convenience. The encrypted messaging app Signal,

for example,provides robust cryptographic security with a conspicuous downside: It's tethered to a single device that holds a user's cryptographic keys. That's because any process that involves exporting those keys so they can be synced across multiple devices—especially if they're being exported into "the cloud"—represents an obvious and huge weak point in the security of the system as a whole. The user wants to be able to access their cloud-stored keys from a new device, but if those keys are only protected by a weak human-memorable password, they're highly vulnerable to brute force attacks by anyone who can obtain them from the cloud server. That may be an acceptable risk for someone who's backing up their Facebook password, but not so much for, say, authentication tokens used to control employee access to major corporate networks—the sort of stuff that's likely to be a target for corporate espionage or foreign intelligence services. Over the medium to long term, our overall cybersecurity is going to depend crucially on making security convenient and simple for ordinary users accustomed to seamlessly switching between many devices. So we should hope and expect to see solutions like Apple's more widely adopted.

For intelligence and law enforcement, of course, better security is a mixed blessing. For the time being, as my co-authors and I noted in the Berkman Center report *Don't Panic*, the "going dark" problem is substantially mitigated by the fact that users like to back stuff up, they like the convenience of syncing across devices—and so however unbreakable the disk encryption on a user's device might be, a lot of useful data is still going to be obtainable from those cloud servers. They've got to be nervous about the prospect of a world where all that cloud data is effectively off the table, because it becomes practical to encrypt it with key material that's securely syncable across devices but still inaccessible, even to an adversary who can run brute force attacks, without the user's password.

If this interpretation of idea behind the proposed narrowing is right, it's particularly politically canny. You declare you're going to saddle every developer with a backdoor mandate, or break the mechanism everyone's Web browser uses to make a secure connection, and you can expect a whole lot of pushback from both the tech community and the Internet citizenry. Tell people you're going to mess with technology their security *already* depends upon—take away something they have now—and folks get upset. But, thanks to a well-known form of cognitive bias called "loss aversion," they get a whole lot less upset if you *prevent* them from getting a benefit (here, a security improvement) most aren't yet using. And that will be true even if, in the neverending cybersecurity arms race, it's an improvement that's going to be necessary over the long run even to preserve current levels of overall security against increasingly sophisticated attacks.

That strikes me, at least for now, as the most plausible read on the proposed "controlled by" language. But another possibility (entirely compatible with the first) is that courts and law enforcement will construe "controlled by" more broadly than I am. If the FBI gives Apple custody of an iPhone, which is running gatekeeper software that Apple can modify, does it become a technology "controlled by" Apple at the time the request is made, even if it wasn't under their control at the time the data was encrypted? If the developer of an encrypted messaging app—which, let's assume, technically retains ownership of the software while "licensing" it to the end user—pushes out regular automated updates and runs a directory server that mediates connections between users, is there some sense in which the entire process is "controlled by" them even if the key generation and encryption runs on the user's device? My

instinct is "no," but I can imagine a smart lawyer persuading a magistrate judge the answer is "yes." One final note here: It's a huge question mark in my mind how the mandate on app stores to ensure compliance interacts with the narrowed scope. Can they now permit un-backdoored applications as long as the encryption process isn't "controlled by" the software developers? How do they figure out when that's the case in advance of litigation?

Let's move on to the other proposed changes, which mercifully we can deal with a lot more briefly. The exclusion of intelligence investigations from the scope of the bill seems particularly odd given that the bill's sponsors are, after all, members of their respective chambers' *intelligence* committees, with the intelligence angle providing the main jurisdictional hook for them to be taking point on the issue at all. But it makes a bit more sense if you think of it as a kind of strategic concession in a recurring jurisdictional turf war with the judiciary committees. The sponsors would effectively be saying: "Move our bill, and we'll write it in a way that makes it clear you've got primary jurisdiction." Two other alternatives: The intelligence agencies, which have both intelligence gathering and cybersecurity assurance responsibilities, have generally been a lot more lukewarm than law enforcement about the prospect of legislation mandating backdoors, so this may be a way of reducing their role in the debate over the bill. Or it may be that, given the vast amount of collection intelligence agencies engage in compared with domestic law enforcement—remember, there are nearly 95,000 foreign "targets" of electronic surveillance just under §702 of the FISA Amendments Act—technology companies are a lot more skittish about being indundated with decryption and "technical assistance" requests from those agencies, while the larger ones, at least, might expect the compliance burden to be more manageable if the obligation extends only to law enforcement.

I don't have much insight into the motive for the proposed critical infrastructure carve-out; if I had to guess, I'd hazard that some security experts were particularly worried about the security implications of mandating backdoors in software used in especially vital systems at the highest risk of coming under attack by state-level adversaries. That's an even bigger concern when you recall that the United States is contemplating bilateral agreements that would let foreign governments directly serve warrants on technology companies. We may have a "special relationship" with the British, but perhaps not so special that we want them to have a backdoor into our electrical grid. One huge and (I would have thought) obvious wrinkle here: *Telecommunications systems* are a canonical example of "critical infrastructure," which seems like a pretty big potential loophole.

The final proposed change is the easiest to understand: Tech companies don't want to be saddled with an unlimited set of obligations, and they sure don't want to be strictly liable to a court for an *outcome* they can't possibly guarantee is achievable in every instance. With that added limitation, however, it would become less obvious whether a company is subject to sanction if they've designed their products so that a successful attack always requires unreasonable effort. "We'll happily provide the required technical assistance," they might say, "as soon as the FBI can think up an attack that requires only reasonable effort on our part." It'd be a little cheeky, but they might well be able to sell that to a court as technically compliant depending on the facts in a particular case.

So those are my first pass thoughts. Short version: Incorporating these changes—above all the first one—would yield something a good deal narrower than the original version of the bill, and

therefore not subject to *all* the same objections that one met with. It would still be a pretty bad idea. This debate clearly isn't going anywhere, however, and we're likely to see a good deal more evolution before anything is formally introduced.

Update: For the lawyers who'd rather rely on something more concrete than my summaries, take the original discussion draft and make the following amendments to see what they're talking about altering:
Section 3, subsection (a)(2) would read:

(2) SCOPE OF REQUIREMENT.—A covered entity that receives a court order referred to in paragraph (1)(A) shall be responsible only for providing data in an intelligible format if such data has been made unintelligible by a feature, product, or service controlled by the covered entity or by a third party on behalf of the covered entity.

Section 4, subsection (3)(B) would be deleted.

Section 4, subsection (4) would read:

(4) COVERED ENTITY.—

(A) IN GENERAL.— Except as provided in subparagraph (B), the term "covered entity" means a device manufacturer, a software manufacturer, an electronic communication service, a remote computing service, a provider of wire or electronic communication service, a provider of a remote computing service, or any person who provides a product or method to facilitate a communication or the processing or storage of data.

(B) EXCLUSION.— The term "covered entity" does not include critical infrastructure (as defined in section 5195c of title 42, United States Code.)

(The material before the first comma in (A) above would be new, as would all of section B.)

Section 4, subsection 12, would read:

(12) TECHNICAL ASSISTANCE.— The term "technical assistance", with respect to a covered entity that receives a court order pursuant to a provision of law for information or data described in section 3(a)(1), includes reasonable efforts to—
(A) isolate such information or data;
(B) render such information or data in an intelligible format if the information or data has been made unintelligible by a feature, product, or service controlled by the covered entity or by a third party on behalf of the covered entity; and
(C) delivering such information or data—
(i) concurrently with its transmission; or
(ii) expeditiously, if stored by the covered entity or on a device.

Those are the changes I've seen floated, though again, probably not exhaustive of what's being discussed.

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