


# Response to consultation on Telecommunications and Other Legislation Amendment (Assistance and Access) Bill 2018

Chris Culnane and Vanessa Teague  
University of Melbourne



September 2018

There are some admirable efforts in the bill:

- the separation between an “access notice,” for information that the company already has, and a “capability notice,” for building new capacity to access additional information,
- the attempt to exclude the introduction of “systemic weaknesses,”
- the prohibition against demanding that a systemic weakness not be rectified.

However, as it stands the bill could have serious negative consequences for the cybersecurity of Australians. There is no definition of “systemic weakness” and no good method for assessing the unintended consequences for the security of other users.

There are two important misconceptions:

- that tech companies represent the best interests of their users and
- that tech companies and some Australian authorities can adequately assess the unintended security consequences of technical changes.

We support the efforts of law enforcement in catching cybercriminals and also conventional criminals who use the Internet. If the legislation had better provisions for transparency, judicial oversight and review, and if the notice came from a police officer with an appropriate warrant, we would support the hand-over of information that the company already had—this seems to correspond to a Technical Assistance Notice. However, we have serious concerns about the unintended weakening of cybersecurity as a result of a Technical Capability Notice or Technical Assistance Request. It would be a serious mistake if a well-intentioned government effort to make it easier to catch criminals also made it far easier to commit cybercrime.

Our research is in applied cryptography. This submission focuses on the cybersecurity implications of the proposed legislation, not the implications for due process, international relations or human rights. We concentrate on the technical notices in Schedule 1.

## 1 The balance between security and security

Weak cybersecurity threatens national security and the security of individuals. Ordinary Australian people, business and government depend on the security of their devices and communications for banking, health data, identity documents, elections, land titles and other matters central to our national, personal and financial security. There are numerous cases of criminals and foreign spy agencies using cybersecurity weaknesses to commit crimes or conduct espionage against targets in Australia and other democracies. The important balance in this discussion has always been between weakening security for catching criminals, and weakening security for everyone else as well. Every instance we know of government-mandated weakening of cryptographic protections has eventually been shown to be exploitable by bad actors too.

The major tech companies have been extensively consulted during the preparation of this bill, but they are not the ones who are likely to be most harmed if a newly installed capability undermines user security and privacy. Both Google<sup>1</sup> and Facebook<sup>2</sup> have recently been fined for anticompetitive or privacy-invading behaviours against their users' best interests. Their real customers are the advertisers. The draft bill's indemnity provisions, and the secrecy that binds both corporations and law enforcement, serve the tech companies' interests against those of their users. This is particularly so for Technical Assistance Requests, which are covered by neither transparency requirements nor the prohibition against introducing a "systemic weakness."

Ordinary users should have the opportunity to walk away based on their understanding of their risks, even if the corporation consents to the risks they are being asked to put their users' data to. Public awareness of the extent or usage of surveillance tools is critical to allowing ordinary consumers to make appropriate risk-management decisions about the trust they place in technology.

There is no way for a mathematical tool (whether for offence or defence) to behave differently depending on the morality of the person using it. The main risk of this legislative program is, by focusing solely on the law enforcement aspect, to underestimate the consequences of undermining cybersecurity for the millions of ordinary Australians who are much more likely to be the target of cybercriminals than of a police investigation. For example, if we think only about the police investigation, then it might seem like a positive step to make it easier for police officers to take control of other people's cameras, in order to

---

<sup>1</sup><http://www.abc.net.au/news/2018-07-19/eu-fines-google-a-record-6.8-billion-over-android-mobile-system/10010510>

<sup>2</sup><https://www.theguardian.com/technology/2018/jul/11/facebook-fined-for-data-breaches-in-cambridge-analytica-scandal>

observe their behaviour and gather evidence on crimes such as child sex offences. However, it would be a serious mistake, because malware can be (and already has been) used by criminals to control other people’s cameras in the course of committing sex offences and extortion against women and girls.<sup>3</sup> Any decision about any improvement in law enforcement access needs to take into account the likelihood that criminals will use the same access vector.

Cybercrime is an increasingly important threat to Australia’s wellbeing. Home Affairs Minister Peter Dutton recently said,

“Cybercriminals are mounting increasingly sophisticated and discreet attacks employing credential-harvesting, ransomware, and social engineering. On conservative estimates, cybercrime currently costs Australians upwards of \$1 billion per year [...] A successful attack on critical infrastructure could have a potentially catastrophic human and economic effect. The WannaCry ransomware incident demonstrated how vulnerable essential services like hospitals can be.”<sup>4</sup>

The WannaCry Ransomware is credibly attributed to a leaked NSA toolkit designed for allowing law enforcement and intelligence operatives to catch criminals and spy on foreign threats. Unfortunately, exactly the same tools were easily redeployed by criminals for attacking hospitals and numerous other targets.

## 2 What is a systemic weakness?

We are glad that nobody intends to mandate the introduction of systemic weaknesses, but we doubt whether a systemic weakness will be recognised before it is too late, and even if it is, are concerned that it is permissible to be requested. The paragraphs included in the draft bill (317ZG) “for the avoidance of doubt,” rely entirely on a term—systemic weakness—that is not standard and not defined anywhere in the bill. *The tremendous difficulty of understanding the unintended consequences and unforeseen security problems caused by a particular modification make Technical Capability Notices (and their voluntary equivalents) dangerous.*

There is no inherent reason why the manufacturer, designer or supplier of a software or hardware system should continue to be able to attack the system successfully after the user takes control. The fact that they often can in practice (for example via targeted software updates) is *already a systemic weakness* in the sense that it represents a single point of trust which, if compromised, could be used by bad actors to break into innocent people’s devices or communications. The Flame malware took advantage of weaknesses in some cryptographic

---

<sup>3</sup><https://www.justice.gov/usao-cdca/pr/temecula-student-sentenced-federal-prison-sextortion-case>

<sup>4</sup><https://www.smh.com.au/politics/federal/increasing-cyber-crime-attacks-costing-up-to-1b-a-year-20180410-p4z8ui.html>

building blocks to forge a digital certificate and hence make the malware appear to come from Microsoft. Efforts such as Google’s Certificate Transparency Project<sup>5</sup> aim to mitigate the effects of one bad digital certificate, by allowing devices to check in real time whether the certificate they are being asked to trust is properly installed on a public ledger. It isn’t perfect, and it isn’t yet very widely used, but it should make it much harder for malware like Flame to infect properly-configured devices, even if the malware has a valid digital signature apparently from the software provider. These ideas are already being extended to individual software updates [NKKJ<sup>+</sup>17], though they are not yet widely used. These efforts make ordinary users more secure by limiting the attacks that can be effectively performed even if a trusted supplier of their software or device is compromised or spoofed.

Of course, they also make it less effective for legitimate law enforcement operations to compel a certificate authority to issue a certificate for their use in deceiving criminals, or compel a provider to issue a targeted update to a particular user’s device.

The draft bill frequently refers to “removing protections,” but it isn’t clear whether this means removing encryption (which might often be impossible after it has been properly applied), or updating the software so that it doesn’t encrypt any more. The explanatory note makes a clear and important distinction between assistance that the organisation is already capable of providing (with an Assistance Notice) and re-engineering the system to expand capability (a Capability Notice). It is important to make this clearer in legislation—Assistance Notices should include information for which the provider has all necessary data but may not have written a retrieval program, such as encrypted information for which the provider has the decryption key (regardless of whether they applied the encryption). Dr Teague has listened to one well-known multinational software company pretending that it could “push back” on government data access warrants—when Dr Teague asked them whether their cloud storage was end-to-end encrypted, she received a long and contrived story about how the company themselves couldn’t access the data, which was simply not true. If the company has the decryption keys (which they do unless the system is securely designed around user control of their own keys), then they have the information necessary to decrypt.

We already see that end-to-end encryption services frustrate interception efforts, because even the provider of the software does not see the data or hold the keys to decrypt it. This isn’t a special case: it’s a general principle of good security design to avoid a single point of failure that can compromise the whole system. The rise of end-to-end encryption has greatly improved the security of ordinary people’s data against malicious actors. Other technologies for mitigating the single point of failure (*i.e.* the supplier) in a device or protocol will likewise make users more secure, though they will unfortunately also frustrate legitimate law enforcement. Dr Teague’s prediction for the resolution of the Apple/FBI controversy is that Apple will (if they haven’t already) design a phone

---

<sup>5</sup><https://www.certificate-transparency.org/>

that does not accept firmware updates without both the user’s pincode and online evidence that the same update is being sent to all users. Other device manufacturers will quickly follow. This will defend users against sophisticated targeted malware (unless it is sent to everyone) and will also have the side effect of rendering any court order against Apple for a targeted firmware update moot.

The other suggestions Dr Teague has heard for law enforcement access, to data the company doesn’t already have, generally involve exploiting something that could be accurately described as a systemic weakness that already exists. For example, one popular suggestion is to add a surreptitious participant to an end-to-end encrypted group communication. Many end-to-end encrypted services such as Skype or Zoom allow groups to communicate together—the security of this process relies heavily on a non-cryptographic user interface that shows participants who has joined in their chat. The software could easily be tweaked to suppress some participants, so that the members of the group didn’t even know that their encrypted communications were also being sent to another party.<sup>6</sup> This is an entirely legitimate thing for law enforcement to do (with a proper warrant). However, it also represents a weakness that could be exploited by bad actors against innocent targets—imagine the opportunities for corporate espionage against high-level online meetings, or for political surveillance<sup>7</sup>. Methods for circumventing and detecting this (and a hamfisted effort would be easy to detect) could therefore be used either by criminals against police surveillance, or by innocent people against criminal hacking. It is likely that group end-to-end encrypted messaging services will start introducing cryptographic means for participants to verify who is participating in the group, just as Signal (and other end-to-end encrypted services) already provide a way to check that a one-to-one connection isn’t being intercepted. In other words, this systemic weakness will probably be removed. This will defend ordinary users against criminal interception, and unfortunately also impede police efforts to invisibly join criminal groups.

The draft bill’s penalties for counselling circumvention of a notice (317ZA) might accidentally catch people who explain to people how to keep their data secure. Methods for circumventing or detecting police access are going to be exactly the same as methods for circumventing or detecting criminal interception.

## 2.1 Why transparency?

The security implications of a particular proposal are incredibly difficult to understand, even for experts. The main reason there is now a consensus that “backdoors” are counterproductive is a long history of independent security analysis demonstrating that efforts to allow law enforcement access have exposed ordinary users to compromise. This has happened for a number of mechanisms.

---

<sup>6</sup>This should not be confused with police officers posing as paedophiles or terrorists openly in online groups, which uses overt social methods rather than covert cryptographic methods to learn what the group is doing.

<sup>7</sup><http://www.abc.net.au/news/2018-08-26/barrie-cassidy-reveals-liberal-whatsapp-messages/10166050>

1. The Clipper Chip’s key escrow mechanism was shown to have a weakness in its authentication system that allowed a target to substitute an innocent person’s key to be decrypted by law enforcement [FY95].
2. US key-length restrictions for export-grade cryptography created widespread vulnerabilities in many TLS implementations [BBDL<sup>+</sup>15, ABD<sup>+</sup>15], decades after the rules passed into abeyance.
3. The dual-EC-DRBG pseudorandom number generator, widely believed to have been deliberately chosen as an easy method for NSA surveillance, was found in Juniper Networks’ code with its “backdoor” rekeyed, presumably by someone else [CMG<sup>+</sup>16].

None of these systemic weaknesses was intended by, or even known to, those deploying the capability or insisting on the rule at the time. Indeed, the key-length restrictions probably didn’t even present a systemic weakness at the time—the problem only arose after decades of speedups in computing power. It was only noticed when multiple large teams of independent researchers communicated together about the theory and practice of TLS. Open, independent review isn’t a perfect or immediate way to achieve a completely accurate assessment of the risks of a proposal, but it is better than the limited process proposed in the exposure draft, which does not seem to involve any technical expertise or anyone representing users’ interests.

There is no reason to be secretive about the potential use of particular mechanisms, especially when those mechanisms are already in the public domain. A new proposal for secure law enforcement access by Ray Ozzie was recently shown to allow criminals to misuse it to expose another (innocent) person’s data.<sup>8</sup>

The example of a corporation (such as Apple) being asked to issue a signed firmware update to bypass user authentication (as they were by the FBI in the San Bernadino case) is already available for public discussion including public amicus briefs, whistleblowing from within the FBI, and extensive public analysis. Australia doesn’t need to make a decision right now about whether Apple’s argument about the increased risk to its other customers was valid, but we do need to design a good, open process for assessing such concerns.

Any proposal for exceptional access should mandate the release of enough public detail about technical mechanisms being required to allow independent analysis and user choice based on as accurate as possible an understanding of the consequences for the security of ordinary users.

The draft bill’s 5 years imprisonment for exposing information (317ZF) could prevent valuable security analysis. Although it is fair to criminalize deliberately undermining a police investigation, it is important not to criminalize legitimate research that could lead to the identification and removal of weaknesses. It is also important to provide an opportunity for legitimate whistleblowing in cases such as misuse of a capability or data breaches affecting ordinary users.

---

<sup>8</sup><https://www.cs.columbia.edu/~smb/blog/2018-05/2018-05-02.html>

### 3 Broad application and limited oversight

Although we have focused primarily on cybersecurity, not on legal aspects, we generally agree with Dr Monique Mann and other human rights law experts, who have noted<sup>9</sup> the bill’s “limited oversight and accountability structures,” are a serious concern given Australia’s “limited human rights and privacy protections.”

The interception of private communication is a serious invasion of privacy and should be reserved for only the most serious incidents. Likewise, secrecy provisions covering the application of justice should only be used in the rarest of cases, and only with judicial oversight. To do otherwise fundamentally undermines the principle of open justice on which Australian legislation is built. It is therefore particularly concerning to see such invasive powers reside in the hands of agencies and not the justice system, and for those powers to be protected by overly broad and punitive secrecy measures. Application of those powers to matters of protecting the public revenue is massively disproportionate, particularly given the possible financial risks to the rest of the economy from impaired cybersecurity.

Capability notices are not operational notices and are not being applied in the context of an active threat. They are about building capability. Any time the state is building secret capabilities is a cause for concern. In the past these capabilities were generally outwardly focussed, *i.e.* defence capabilities that were never intended to be targeted on the population itself. In this case it is different, the state is building secret capabilities that are specifically being targeted at Australians. That presents a dangerous precedent, potentially shifting power and sovereignty away from the population. Whilst there is justification for not revealing active operations, keeping capabilities secret risks preventing public oversight, and is likely to lead to abuse.

Keeping these methods secret does not make the system more secure, or reduce the likelihood of accidentally introducing a systemic weakness. It reduces the likelihood that such weaknesses would be widely understood and identified by the scientific community before the insecure capability was widely deployed.

#### 3.1 Requests vs. Notices

Whilst there are both limitations and reporting requirements associated with the assistance and capability notices, no such restrictions apply to the requests. Requests can ask for the implementation of systemic weaknesses, whatever they turn out to be, and can ask for new capabilities to be implemented to remove protections from electronic communication. Requests may not be legally enforceable. However, the government wields enormous soft power—to suggest that a “request” from the Australian Government can be ignored is ridiculous. As it stands the requests are probably the most powerful aspect of the legislation. They have few limitations, and are inexplicably excluded from the annual

---

<sup>9</sup><http://theconversation.com/the-devil-is-in-the-detail-of-government-bill-to-enable-access-to-communications-data-96909>

reporting requirements imposed on the Government in Section 317ZS. It is inexplicable that what limited public oversight is provided for in the legislation excludes one of the most powerful components of that legislation.

## 4 Summary & Recommendations

Consider the following three efforts to ensure different kinds of security.

1. Increasingly many device manufacturers and software providers improve users' security by ensuring that encrypted data cannot be accessed by anyone other than its owner, not even the company that provided the device or service.
2. Everyone agrees that no "systemic weakness" should be introduced that undermines this improvement in security, though there is no clear definition of "systemic weakness."
3. A recent statement from the five eyes security alliance included, "Governments should recognize that the nature of encryption is such that that there will be situations where access to information is not possible, although such situations should be rare."<sup>10</sup>

We do not know of any way to allow the improvements in case (i) to continue, while avoiding the introduction of a systemic weakness (by any definition) and for which failure to gain access would be rare.

Though data minimisation and removing single points of failure are good security designs, many large tech companies do exactly the opposite. Massive data gathering for the purposes of targeted advertising drives much of the Internet economy. The company sits at the centre of the network, with the ability to read and control all user communications. This structure puts users at risk. Individuals risk their particular data being exposed or stolen (consider Equifax and Ashley Madison). Democracy itself is threatened by our dependence on a small number of providers for political advertising and news.

It is vitally important that any Australian legislation discourage—or at least not further encourage—massive data gathering about Australians, whether for primarily commercial purposes or for helping law enforcement, because such data can easily be used for purposes detrimental to our society and our democracy, even if it is occasionally also useful for catching criminals.

The bill would introduce an assumption of personal data availability by design and default, as indicated in the department of Home Affairs' recent statement of principles (above). This is in stark contrast with the EU's "Data protection by design and default,"<sup>11</sup> which aims to protect its ordinary citizens by ensuring that it is difficult for others to access their data without their consent.

---

<sup>10</sup><https://www.homeaffairs.gov.au/about/national-security/five-country-ministerial-2018/access-evidence-encryption>

<sup>11</sup><http://www.privacy-regulation.eu/en/article-25-data-protection-by-design-and-by-default-GDPR.htm>



## 4.1 Specific recommendations

1. We do not know of any proposal for adding a new “Technical Capability” that does not increase the risk for other users more than it benefits law enforcement efforts. We would therefore strongly argue for the removal of Technical Capability Notices.
2. The scope of (voluntary or involuntary) re-engineering of a system to extract more data should be restricted to only the most serious of crimes or threats to national security.
3. Technical Assistance Requests should be covered by the same limitations as described in Division 7, namely they should not be able to request systemic weaknesses, nor develop new techniques for removing electronic protection.
4. Recipients of notices or requests should be mandated to provide transparency reports, including all requests and notices. Technical Assistance Requests should be included in the annual report mandated in Section 317ZS.
5. Insist on full transparency of the *methods*, while acknowledging that details of particular targets and operations may need to be secret for a while. This allows for a better assessment of the unintended consequences for weakening the security of other users.
6. Remove the blanket criminal penalty for all disclosures or explanations of circumvention. Instead, make sure that criminal penalties apply only to deliberate exposure or undermining of police operations, rather than to security analysis of the unintended consequences (such as finding flaws in the Ozzie proposal) or generic counselling about improving cybersecurity (which might often have the consequence of circumventing a particular access mechanism). Ensure appropriate channels for legitimate whistleblowing in the case of improper police behaviour or undisclosed data breaches as a result of a notice/request.
7. Provide for appropriate redress for innocent parties affected by a data breach as a consequence of a notice/request. This could perhaps be achieved by a compulsory insurance program. If cooperating tech companies are to be indemnified, then there will need to be some way for them to demonstrate publicly that a particular data breach was a direct consequence of a notice/request—otherwise almost any data breach by any company that has any Australian users could potentially be blamed on this program.
8. A more precise definition and description of Systemic Weakness is also required. Without it, the promise not to introduce any seems very hard to keep.

## References

- [ABD<sup>+</sup>15] David Adrian, Karthikeyan Bhargavan, Zakir Durumeric, Pierrick Gaudry, Matthew Green, J Alex Halderman, Nadia Heninger, Drew Springall, Emmanuel Thomé, Luke Valenta, et al. Imperfect forward secrecy: How diffie-hellman fails in practice. In *Proceedings of the 22nd ACM SIGSAC Conference on Computer and Communications Security*, pages 5–17. ACM, 2015.
- [BBDL<sup>+</sup>15] Benjamin Beurdouche, Karthikeyan Bhargavan, Antoine Delignat-Lavaud, Cédric Fournet, Markulf Kohlweiss, Alfredo Pironti, Pierre-Yves Strub, and Jean Karim Zinzindohoue. A messy state of the union: Taming the composite state machines of tls. In *Security and Privacy (SP), 2015 IEEE Symposium on*, pages 535–552. IEEE, 2015.
- [CMG<sup>+</sup>16] Stephen Checkoway, Jacob Maskiewicz, Christina Garman, Joshua Fried, Shaanan Cohny, Matthew Green, Nadia Heninger, Ralf-Philipp Weinmann, Eric Rescorla, and Hovav Shacham. A systematic analysis of the juniper dual ec incident. In *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security*, pages 468–479. ACM, 2016.
- [FY95] Yair Frankel and Moti Yung. Escrow encryption systems visited: attacks, analysis and designs. In *Annual International Cryptology Conference*, pages 222–235. Springer, 1995.
- [NKKJ<sup>+</sup>17] Kirill Nikitin, Eleftherios Kokoris-Kogias, Philipp Jovanovic, Nicolas Gailly, Linus Gasser, Ismail Khoffi, Justin Cappos, and Bryan Ford. Chainiac: Proactive software-update transparency via collectively signed skipchains and verified builds. In *26th USENIX Security Symposium (USENIX Security 17)*, pages 1271–1287, 2017.