Internet Policy Research Initiative
Massachusetts Institute of Technology

Australian Government
Department of Home Affairs
By email to: AssistanceBill.Consultation@homeaffairs.gov.au

Regarding the TELECOMMUNICATIONS AND OTHER LEGISLATION AMENDMENT (ASSISTANCE AND ACCESS) BILL 2018

I. Background and Statement of Interest

The MIT Internet Policy Research Initiative is pleased to offer these comments on the proposed TELECOMMUNICATIONS AND OTHER LEGISLATION AMENDMENT (ASSISTANCE AND ACCESS) BILL 2018 pending before the Parliament of Australia. These comments address the technical challenges recognized in the Bill, particularly the importance of avoiding the introduction of systemic weaknesses or vulnerabilities that can result from requiring technology providers to implement exceptional access capabilities. We have written extensively about the technical security risks associated with exceptional access requirements[1], so we are pleased to see that the drafters acknowledge these risks and look forward to the chance to help identify approaches that work to this end.

The mission of the Internet Policy Research Initiative (IPRI) is to work with policy makers and technologists to increase the trustworthiness and effectiveness of interconnected digital systems through engineering and public policy research, education and engagement. There is a pressing need to bridge the gap between the technical and policy communities, and we are doing this with our fully interdisciplinary research approach that pulls together expertise from across MIT and beyond. IPRI is led by faculty researchers from engineering, social science, and management labs at MIT and is located at the MIT Computer Science and Artificial Intelligence Lab (CSAIL).

IPRI has been actively engaged with governments, industry, civil society and fellow academics from around the world on the pressing questions of surveillance and encryption. MIT IPRI's

[1] Keys under doormats: mandating insecurity by requiring government access to all data and communications. 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. Journal of Cybersecurity Nov 2015. https://academic.oup.com/cybersecurity/article/1/1/69/2367066. (See attached copy)

Founding Director Daniel Weitzner recently chaired, together with Professor Joan Feigenbaum (Yale University) and Timothy Edgar (Brown University), a workshop at the Crypto 2018 conference on Encryption and Surveillance. This workshop brought together senior government officials, academic cryptographers, systems security researchers, and human rights advocates to consider legal, policy, and technical aspects of the exceptional access debate. We were pleased to have a presentation from Mr. Adam Ingle of the Australian Department of Home Affairs on the Bill and welcome continued engagement with the Australian government on this issue going forward. Our comments here are informed by discussions at that workshop but represent only the views of the MIT Internet Policy Research Initiative.

II. Summary

- It is still an open question whether it is possible to design a secure EA system, and in the course of our work at MIT, we have yet to find an EA design that would satisfy the requirement of avoiding the introduction of systemic weaknesses or vulnerabilities.
- Among existing legislation or proposals, both the UK's Investigatory Powers Act 2016 and Australia's proposed Bill need explicit transparency provisions covering the underlying design protocols, cryptographic algorithms, and software that allow security researchers and the public to evaluate TCN requests for systemic weaknesses and vulnerabilities.
- Communication providers subject to TCNs should be allowed to publicly disclose what they think is necessary about how their systems implement the TCNs they receive.
- Any decision to mandate and implement an EA system should take into account the global ramifications of such a decision. This includes the impact on local firms serving a global clientele and the potential that a system designed for domestic use could also threaten human rights in other countries.

III. Transparency

The security of digital systems, especially Internet-wide systems, depends critically upon transparency of the underlying design protocols, cryptographic algorithms, and, in many cases, the software's code itself. As the security research community has demonstrated over and over again, design flaws and implementation vulnerabilities in critical code is often discovered by third parties, not the engineers who design and implement the systems themselves. Hence, it is vital that the Bill encourage, not penalize, transparency of relevant details of any technical requirements that might be imposed.

The importance of design transparency for large scale systems, especially security systems, can be seen in recent experiences with Internet-scale vulnerabilities. Consider the well-known and very serious Heartbleed vulnerability. Heartbleed was discovered in OpenSSL, an open-source software library used by millions of websites to encrypt information sent over the Internet[2]. Attackers can send maliciously-crafted heartbeat messages that trick a server running OpenSSL into divulging the contents of RAM, the memory of the computer running the web server. An attacker could then leverage Heartbleed to force any server using the affected versions of OpenSSL to give up information. Web servers with this vulnerability can be tricked into divulging user names, passwords, secret keys used to encrypt data, and other security credentials, enabling the attacker to hijack a user's account, access any amount of private information, or take over the target server.

The harm averted by patching Heartbleed was enormous[3], and these concerns aren't just academic; Mandiant noted that the vulnerability was being exploited in the wild, and many of the world's most popular websites were vulnerable at some point. The security analysis company Hacklabs determined that 10% of the top 200 websites in Australia were vulnerable to this attack at the time that it was discovered.

The discovery of Heartbleed depended on unimpeded access to the underlying software and would have been far more difficult without such access. Heartbleed was discovered by security researchers at Codenomicon and Google auditing the source code of OpenSSL. Expeditious discovery and remediation of such vulnerabilities is vital to the security of the Internet environment, both in Australia and globally. Without guarantees of transparency, both of the details of the SSL protocol and the underlying code, the harm to the global infrastructure would have been far more severe.

It is worth noting that failed cryptographic protocols can cause outsized damage in unexpected ways that last far beyond when they were discovered to be faulty. For example, the FREAK and DROWN exploits were only possible because earlier regulatory mandates to weaken encryption on products exported from the United States left critical systems perpetually vulnerable as Internet servers continued to support out-of-date software exported under the regulation.

---

[2] Sydney Morning Herald, Revealed: How Google engineer Neel Mehta uncovered the Heartbleed security bug. 9 Oct 2014.
[3] OpenSSL Usage Statistics

Similar to Heartbleed, these so-called “export grade encryption” cipher suites resulted in a class of vulnerabilities that caused colossal damage to the internet infrastructure. At one point, roughly 12% of the top million most visited websites were completely interceptable, allowing attackers to gain user credentials, passwords, and other private data.

Transparency will be particularly important for technical requirements that arise from the Bill as such requirements would address the very security features upon which all users (law-abiding citizens, businesses, as well as potential suspects) depend. As Heartbleed, FREAK and DROWN illustrate, system components that are designed to provide security features are critical pieces of code that can become important sources of security vulnerabilities. Beyond that, vulnerabilities may lurk in software, hardware, cryptographic protocols, and other places for years before discovery, and many suspect that vulnerabilities live longer in less-transparent systems.

Yet as we understand the Bill, there would be substantial penalties for disclosing information about required changes to system design and implementation, whether through technical assistance notices or technical capacity notices. Such penalties would thwart the increasingly vital process of subjecting widely-used software to maximum public scrutiny so that third-party security researchers can have the best chance of discovering vulnerabilities.

Enabling third-party, adversarial scrutiny of features associated with TAN and TCNs anticipated by the bill requires transparency in certain respects, but can be implemented in a way to avoid operational risks to law enforcement or national security investigations. Given the substantial concern that EA features can cause systemic security vulnerabilities, it is vital that the BIll provide adequate transparency in two dimensions: 1) the ability for the public to have access to the technical details of the TCNs, and 2) the ability for providers subject to TCNs to disclose what they think is necessary about how their systems implement the TCNs.

First, technical details of any mandate through the Bill should be made publicly available in order to enable the technical community to scrutinize such requirements for potential security risks. Although well-intentioned engineers may make their best effort to carefully craft a requested system, designing any cryptographic code is a difficult and error-prone process. Vulnerabilities are still very likely to be discovered by independent third parties, and as the examples above illustrate, security risk emerges not only due to carelessness, but because it is

often difficult to anticipate all of the risks inherent in the design. Systems can appear secure when they are first designed, but unexpected vulnerabilities become more visible when the operation of those systems is subject to public scrutiny.

Second, vendors of large-scale systems should not be forced to hide security features from their users. Large providers including Apple, Microsoft, Google, WhatsApp, Signal and others have published more and more technical details about their security architectures, building trust with users and helping users and third-party designers build and operate systems in a more secure fashion. It would be a real and dangerous step backward for the security of the global Internet environment if vendors were forced to hide relevant security design details from their users and customers.

IV. Standards and methods to assess systemic weakness

As the drafters of the Bill are well aware, there is ongoing concern in the technical community about risks associated with exceptional access systems. Some in the law enforcement community have called on systems designers to propose their own security EA designs, reasoning that this is a technical problem that can be solved with sufficient technical effort[4]. The drafters of the Bill should be commended for moving beyond this unilateral approach. We offer two observations about the current state of the technical debate.

First, it is still an open question if it is possible to design a secure EA system, and it has been resolved that such a system cannot be reasoned about without the context of a specific set of functional requirements and implementation parameters. In other words, even given a useful specification, a full understanding of the security risks of any given EA design are still far off in the future.

This is not solely the opinion of the authors of this document, but appears to be widely shared in the research community. The workshop mentioned in the introduction included presentations of

[4] "After all, America leads the world in innovation. We have the brightest minds doing and creating fantastic things. If we can develop driverless cars that safely give the blind and disabled the independence to transport themselves; if we can establish entire computer-generated virtual worlds to safely take entertainment and education to the next level, surely we should be able to design devices that both provide data security and permit lawful access with a court order," FBI Director Christopher Wray, "Raising Our Game: Cyber Security in an Age of Digital Transformation." (Fordham University, January 9, 2018)

several EA designs, all of which were presented as very early design sketches of systems that would require considerable elaboration and analysis before any judgement could be made about their risk profiles. Indeed, vulnerabilities are often found during the implementation stage of such systems, and, to our knowledge, no such system has been so developed.

One of the workshop participants, who has done academic work on EA schemes in the past, Prof. Mayank Varia from Boston University, wrote:

> "It's important to recognize that, whether with driverless cars or realistic virtual reality [Director Wray's examples of Silicon Valley innovation], it took years—sometimes decades—to develop the technologies, which even now remain works in progress. Secure third-party-access systems may similarly be years away from viability. That doesn't mean we shouldn't invest in the necessary research and development, just that we should have a realistic timetable in mind."[5]

And in the course of such development we should expect that there will be serious security hurdles to address. While we cannot declare the design goals impossible to meet, neither can we presume that it is possible to meet exceptional access requirements in all cases free from unreasonable risk. One of the early designs, proposed by the former Chief Technical Officer and Chief Software Architect at Microsoft, Ray Ozzie, was recently presented to a technical audience at Columbia University. During the course of the presentation, Eran Tromer, a cryptographer, identified a vulnerability which has yet to be addressed[6]. We have every reason to believe that such work will continue, but one must not infer from the existence of exceptional access design proposals that solutions are any closer.

Second, governments which plan to mandate technical capacity requirements should begin by engaging with the technical community in developing methods and standards to evaluate the security risks of such requirements. Recognizing that there can be systemic risk associated with these requirements is an important public policy step. Next there must be both a process and a technical framework for evaluating those risks. Today, neither the United Kingdom's Investigative Powers Act nor the proposed Bill specifies clear technical or operational criteria against which Technical Capacity Notices are to be assessed. This is not a simple technical

[5] Alan Z. Rozenshtein, Mayank Varia, Charles Wright, How Congress Can De-Escalate the Second Crypto War: Fund Research and Broker a Crypto Armistice, Lawfare Blog, June 5, 2018
[6] Steve Bellovin, Ray Ozzie's Proposal: Not a Step Forward (Blog Post, April 25, 2018)

task, but is essential to assure that governments avoid mandates that could put national and even global infrastructure at risk.

In summary, in the course of our work at MIT and based on our engagement with the broader systems security and cryptography research communities, we have yet to find an EA design that would demonstrably avoid the introduction of systemic weaknesses or vulnerabilities. Again, while we cannot declare the design goals impossible to meet, neither can we presume that it is possible to meet exceptional access requirements in all cases free from unreasonable risk. In the case that governments still decide to move forward, the broad technical community should be fully engaged in developing methods and standards to evaluate the security risks of such requirements.

V. The importance of a global perspective

This proceeding calls for comments on a proposed addition to Australian national law, but the implications of this proposal and others like it can have a global reach. The marketplace of global technology users, both institutions and individuals, has become sensitized to the risk that national governments may seek to weaken the security of widely-used infrastructure[7]. In the wake of the Snowden disclosures, companies in the United States faced severe skepticism from non-US buyers and increased regulatory pressure from European governments out of a belief that the US national security agencies had compromised the security infrastructure of major US Internet providers. Separate and apart from whether this was true, the global marketplace showed that it demands assurances of trust, otherwise it will punish products suspected of being under the control of government. Australia has an opportunity to learn from these adverse regulatory and market dynamics, and identify steps that will assure users around the world that they can still trust products and services made or provided in the Australian market. Transparency and independent design scrutiny is essential to assure trust in the Internet marketplace, both for domestic and global customers.

An equally-important global challenge is to reckon with the fact that exceptional access technical capabilities required under Australian law could easily become available in national marketplaces where the rule of law and respect for human rights is absent. Does the Australian

[7] Daniel Weitzner, "Weitzner: Encryption solution in wake of Paris should come from Washington not Silicon Valley" (Washington Post, Nov. 24, 2015)

law have means of preventing the spread of powerful surveillance tools to regimes in which human rights workers and dissidents depend on strong encryption without back doors to protect their political activities? Does Australia have a plan to work with other like-minded democracies to prevent to spread of technology that might simultaneously aid Australian law enforcement and threaten human rights in other countries? None of these is an easy question, but they arise inevitably when considering proposals such as included in the Bill.

VI. Conclusion

In conclusion, there is significant ongoing debate about whether it is possible to design a secure EA system. Researchers at MIT have yet to identify a system design that would allow law enforcement the requested access without introducing systemic weaknesses or vulnerabilities.

Transparency provisions and the ability to evaluate protocols, algorithms, and code are critical in the case that governments decide to push forward with EA mandates. The UK's current Investigatory Powers Act 2016 and Australia's proposed Bill both lack explicit transparency provisions that would cover the design protocols, cryptographic algorithms, and software and allow the public to evaluate TCN requests for systemic weaknesses and vulnerabilities. This transparency is necessary because it allows researchers to probe and uncover weaknesses before they become global vulnerabilities. Communication providers also need to be able to disclose what they deem necessary about how their systems implement the TCNs they receive.

Finally, the examples of Heartbleed, FREAK and DROWN highlight how domestic policy decisions have global implications and can introduce vulnerabilities that extend well into the future. In addition, decisions to mandate EA systems will affect local firms serving a global clientele and have the dangerous potential to threaten human rights in other countries.