

Submission to PJCIS inquiry into the 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 is 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 handover of information that the company already had. 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.

This submission is very similar to our earlier one to the Department of Home Affairs' informal consultation on the exposure draft of the bill. The bill's current version has three important changes since then, each partially addressing one of our recommendations:

- assistance requests are now to be included in the Annual Reports defined in 317ZS,
- protection of public revenue has been removed as a reason for issuing notices or requests, and
- there is some improvement to the process of deciding whether a particular modification (from a technical assistance request or technical capability notice) constitutes a systemic weakness.

The above constitute steps in the right direction, but do not adequately address the problem of the possible introduction of a serious security problem that affects ordinary innocent users of the Internet. The prohibition against systemic weaknesses has still not been extended to Technical Assistance Requests, and the process for assessing a systemic weakness still does not involve any transparent or detailed examination, but instead assumes that a trustworthy and technically expert auditor will be found to take responsibility for the assessment. This seems implausible since the responsibility is huge and the key term, systemic weakness, is not defined. So although the process has been slightly improved, we remain concerned that it is not adequate for preventing Technical Assistance Requests or Technical Capability Notices from introducing or exacerbating cybersecurity risks to innocent users.

Additionally, we are concerned that the consultation process has been undermined by unnecessary haste and a lack of transparency. The invasive nature of the legislation warrants careful and thorough consideration. All submissions to the informal consultation should be made public, subject to authors approval, to ensure an accurate picture of public sentiment is available to all. Furthermore, we are concerned that by holding an informal consultation beforehand it will act as a sinkhole for submissions, with those not familiar with parliamentary inquires being unaware of the need to submit a second response.

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 Google¹ and Facebook² have recently been fined for anticompetitive or privacy-invasive 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.

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 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.³ 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

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

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

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

human and economic effect. The WannaCry ransomware incident demonstrated how vulnerable essential services like hospitals can be.”⁴

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 to attack hospitals and numerous other targets.

2 What is a systemic weakness?

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.*

New subsection 317ZG(4) “clarifies that subsections (2) and (3) are enacted for the avoidance of doubt and do not change the ordinary meaning of the terms ‘systemic weakness’ or ‘systemic vulnerability’, ” but neither of these terms has an ordinary meaning, either in the technical literature or in the international political conversation. The new Explanatory Memorandum still has no definition, but does contain more hints including: “The mere fact that a capability to selectively assist agencies with access to a target device exists will not necessarily mean that a systemic weakness has been built.” This is true. It may mean that a systemic weakness already exists.

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 building blocks to forge a digital certificate and hence make the malware appear to come from Microsoft. Efforts such as Google’s Certificate Transparency Project⁵ 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⁺17], though they are not yet widely used.

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

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

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 asked whether their cloud storage was end-to-end encrypted, they produced 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 and could hence be served with a Technical Assistance Notice.

2.1 Suggestion: Transparency from the providers about their data storage

If every corporation that operates in Australia were obliged to publish a complete and transparent account explaining what data it collects about ordinary users, this would be of great value to both users and law enforcement officers. We are not suggesting that the data itself be published, but rather the meta-data describing what sort of information is collected from which users under what circumstances. This would allow police officers to make more targeted requests, and users to make better privacy choices. It would also allow for a clearer distinction between Technical Access Notices and Technical Capability Notices, because it would be more immediately evident what data the corporation already possessed, even if it had not implemented a program to retrieve it.

Some companies store a great wealth of information about millions of Australians, but do not advertise the fact. A great deal of information may be available to police under an appropriately-directed Technical Assistance Notice, without needing to undermine protections with a Technical Capability Notice.

2.2 Removing single points of failure is good for security

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. One possible resolution of the Apple/FBI controversy is that Apple will (if they haven't already) design a phone 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.

Other suggestions 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 Signal 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.⁶ 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 surveillance⁷. 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)

⁶This 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.

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

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.

2.3 Why transparency?

The security implications of a particular proposal are incredibly difficult to understand, even for experts. There is now a consensus that cryptographic “backdoors” are counterproductive, but it wasn’t always obvious. It is the result of 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.

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⁺15, ABD⁺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⁺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 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 draft bill, which may not have sufficient expertise, incentive or time to thoroughly examine the system—see below.

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.⁸ This may not have been discovered before deployment if it hadn’t received careful, open review from other scientists.

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

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

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 the technical mechanisms 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.

2.4 Why a trusted auditor of systemic weaknesses is not enough

The bill's new version has slightly improved the process for assessing whether a Technical Capability Notice introduces a systemic weakness. Rather than by ministerial determination (as in the exposure draft), the bill now includes a possibility for appointing an auditor by agreement between the service provider and the government. Although this is a (slight) improvement, it doesn't address the fundamental problem of a lack of expertise or incentive. There is no reason to expect that the appointed auditor will have sufficient expertise, diligence or time to find subtle unintended consequences, nor is there anyone involved in the process with any incentive to protect users' interests. Since there is still no precise definition of systemic weakness, it isn't even clear what the auditor is attempting to detect.

In a series of papers over many years, we and our coauthors have discovered numerous security vulnerabilities and software bugs in election-related software that has been audited or certified by trusted authorities appointed by electoral commissions. This has included vulnerabilities in Internet voting software which exposed the system to denial of service attacks, vote privacy breaches and manipulation [HT15, CEET17]. We also found a coding error in the official NSW vote counting software which with very high probability caused the wrong person to be elected to a local council [CBNT17]. This code had been formally certified perfectly correct by trusted auditors engaged by the electoral commission. When we demonstrated the error, the code was corrected and re-certified perfectly correct again, after which we found two more bugs. We do not know whether the same auditor was engaged to re-re-certify the third version perfectly correct again. This series of weaknesses and errors in certified and audited software indicates a serious problem in the process, which lacks any incentive either for the commissions to engage skilled auditors, or for the auditors to do a thorough job.

There is nothing in the draft bill that suggests that the audit process for systemic weaknesses will be any better than the equivalent processes for election

software. In both cases, the overwhelming incentive is to reassure the public rather than to find and correct vulnerabilities or errors. If anything, the issues surrounding Technical Capability Notices will be even more complex, and the opportunity for independent parties to examine the results will be even more limited, than they are for elections. Remember that many of the serious consequences of prior efforts to undermine cybersecurity for facilitating law enforcement were not apparent at the time.

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.

2.5 Whose responsibility?

The draft bill, Section 317ZJ, provides communications providers with immunity from civil liability associated with data breaches resulting from cooperation with requests or notices. However, it does not provide a way for the providers to prove which breaches were, or were not, the result of such cooperation. Data breaches happen all the time—how could a communications provider prove that the breach was caused by cooperation under the act? How can an affected user prove that it wasn't?

Immunity is not extended to the auditors who certify that a given Capability will not increase the risk to other users. Although we are not suggesting that the auditors should be given immunity, this omission does seem to decrease the set of qualified people who could rationally take on the job for a reasonable fee.

There does not seem to be any other opportunity for redress for ordinary users harmed by a data breach that is inadvertently (directly or indirectly) caused by a Technical Capability Notice. This seems to be an oversight, which we are not sure how to rectify—perhaps some sort of compulsory insurance scheme would be appropriate.

3 Privacy, transparency and human rights

Although we have focused primarily on cybersecurity, not on privacy and human rights, we generally agree with Dr Monique Mann and other human rights law experts, who have noted⁹ the bill's "limited oversight and accountability structures," are a serious concern given Australia's "limited human rights and privacy protections." We have concentrated here on the security implications of Schedule 1, but we agree with the numerous concerns that have been raised about the human rights implications of that and the other schedules too.

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

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

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.

There are undoubtedly times when successful interception is good for national security, and there are sometimes legitimate invasions of privacy or limits on transparency in the hope of improving security. However, all three of those concepts (privacy, transparency and security) are important. Security would not trump all other values even if the bill were good for security overall. This would lead to the formation of a surveillance state, and ultimately a police state. We do not live in a risk-free world, with freedom comes the risk that some will abuse those very freedoms to cause harm and attack the open and free society we aspire to be. Whilst we should endeavour to prevent such attacks, we cannot sacrifice the very freedoms that are under threat—to do so would hand a victory to the attackers.

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.

3.1 Requests vs. Notices

Whilst there are limitations 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. If an act is not appropriate to be mandated, why is it acceptable to request that same act to be performed voluntarily?

4 Surveillance Capitalism

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 [Zub15]. 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). In a timely example, Facebook has recently admitted to a security breach in which access tokens for 50 million user accounts were stolen.¹⁰ The concentration of so much valuable information in so few hands is itself a serious weakness—when Facebook has a security problem, it affects millions or billions of people. 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. This is in stark contrast with the EU’s “Data protection by design and default,”¹¹ which aims to protect its ordinary citizens by ensuring that it is difficult for others to access their data without their consent.

5 Summary and 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.”¹²

We do not know of any way to allow the improvements in Case 1 to continue, while avoiding the introduction of a systemic weakness (by any definition) and for which failure to gain access would be rare.

¹⁰<https://www.theguardian.com/technology/2018/sep/28/facebook-50-million-user-accounts-security-berach>

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

¹²<https://www.homeaffairs.gov.au/about/national-security/five-country-ministerial-2018/access-evidence-encryption>

5.1 Recommendations

1. Remove Technical Capability Notices. 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.
2. Restrict the scope of (voluntary or involuntary) re-engineering of a system (to extract more data) to only the most serious of crimes or threats to national security.
3. Limit Technical Assistance Requests as described in Division 7, *i.e.* they should not be able to request systemic weaknesses, nor develop new techniques for removing electronic protection.
4. Mandate transparency reports, including all requests and notices, from all recipients.
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. Add a definition of Systemic Weakness. Without it, the promise not to introduce any seems very hard to keep.

References

- [ABD⁺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⁺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.
- [CBNT17] Andrew Conway, Michelle Blom, Lee Naish, and Vanessa Teague. An analysis of new south wales electronic vote counting. In *Proceedings of the Australasian Computer Science Week Multiconference*, page 24. ACM, 2017. <https://arxiv.org/abs/1611.02015>.
- [CEET17] Chris Culnane, Mark Eldridge, Aleksander Essex, and Vanessa Teague. Trust implications of ddos protection in online elections. In *International Joint Conference on Electronic Voting*, pages 127–145. Springer, 2017. <https://arxiv.org/pdf/1504.05646>.
- [CMG⁺16] Stephen Checkoway, Jacob Maskiewicz, Christina Garman, Joshua Fried, Shaanan Cohneney, 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.
- [HT15] J Alex Halderman and Vanessa Teague. The new south wales ivote system: Security failures and verification flaws in a live online election. In *International Conference on E-voting and Identity*, pages 35–53. Springer, 2015. <https://arxiv.org/pdf/1504.05646>.
- [NKKJ⁺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.
- [Zub15] Shoshana Zuboff. Big other: surveillance capitalism and the prospects of an information civilization. *Journal of Information Technology*, 30(1):75–89, 2015.