Submission to the inquiry into the impact of new and emerging information and communications technology

Dr Vanessa Teague Melbourne School of Engineering
The University of Melbourne, Victoria 3010 Australia

This submission discusses the role and use of encryption, encryption services and encrypted devices, and the possible unintended consequences of legislation intended to facilitate law enforcement access to encrypted data. This is a complex and controversial topic. My intention is to raise some important questions, not to suggest that I know the right answers.

What is the discussion about?

Some scenarios apparently under consideration are:

- end-to-end encrypted communications between two or more devices, in which the decryption key is available only to the endpoint devices.
- encrypted cloud storage, in which the decryption key is available only to a device/person who generated the encrypted files.
- encrypted on-device storage, in which files on the device are available only after the user enters a password/PIN etc.

It is important to distinguish three different kinds of possible requirements. The UK's Investigatory Powers Act seems, unfortunately, to be highly ambiguous. It doesn't seem clear whether companies have to provide, on the spot, all data that it is reasonable for them to extract (1), whether they will be served with warrants for data extraction that is actually not feasible (2), or whether they have to make all reasonable efforts to redesign their systems to facilitate data extraction (3).

1. **Available Data** A company could be asked to furnish unencrypted data that they have, or (equivalently) decryptions of encrypted data for which they have the key. This would make sense for many services in which the provider already has access to the information, including many ordinary cloud services (Google Drive, Microsoft Cloud) and social media posts (most of Facebook) and communications (Skype, Google Hangouts). Even many end-to-end encrypted services (such as Signal) expose some metadata, such as who communicated with whom or who accessed what file. Compliance would be entirely feasible because the company already has the data.
2. **End-to-end Encrypted Data** A company could be asked to furnish the decryptions of encrypted data for which only the end-user has the decryption key or password/PIN, such as Signal, WhatsApp, and some iPhone storage. Assuming the encryption had been implemented properly, this would not be possible (barring implementation errors, unknown backdoors, etc – see below). When Facebook says that they *cannot* decrypt WhatsApp messages, this is true.

3. **Re-engineering to sidestep end-to-end encryption** All companies that trade in Australia could be forced to engineer their systems so that case (2) never arises, so all encrypted data can be decrypted either by the company or by Australian law enforcement, or some combination, on request. This is the controversial case, which I discuss in the rest of this note. This might be achieved by a key escrow system, a cryptographic backdoor, a weakening of encryption parameters, an option to push a signed update onto the device to extract the data (as the FBI requested of Apple), or in various other ways. When Apple was asked to do this, they replied that they *wouldn't*, not that they *couldn't.*

Is it feasible to re-engineer to sidestep end-to-end encryption?

Yes and no. Yes, it is feasible for any company to re-engineer the system so that any encryption implemented by that company can be decrypted. However, there are two important limitations:

[Evasion] Any compliance applies only to encryption implemented by the company, not to other encryption software downloaded independently by the user. For example, Microsoft might promise to comply, but a user might install some encryption software from elsewhere and use it to encrypt files – Microsoft would not be able to decrypt them, even if the encrypted file was present on a Microsoft machine. (This is how my research group worked on our draft paper about the full implications of the MBS-PBS data breach, before it was delivered to the government. We encrypted the file on a non-Internet-connected device using a freely available implementation of gpg, then sent the encrypted file as an attachment using ordinary email. There would have been no way for the email provider to read it.)

Similarly, a device manufacturer (e.g. Dell) might promise to comply, but a user might wipe the default operating system and install a user-controlled operating system such as Linux, which comes with various options for secure encryption.

I suspect that clever criminals are (unfortunately) already learning how to communicate in a way that evades metadata retention. These same people would, probably more easily, learn to use encryption independent of the big service providers and device manufacturers. They could download encryption code from overseas, from open source projects, and from companies that do not operate in Australia.

[Weakening of security] The real question is whether it is feasible to facilitate decryption by legitimate law enforcement, without also making it easier for bad actors such as criminals and foreign spy agencies to access the data too. "No," is the general opinion of the security community, articulated in an influential paper entitled "Keys under doormats"[1] and also by cybersecurity experts who filed amicus briefs in the San Bernadino case[2]. The reason is simply that the legitimate law

[1] https://dspace.mit.edu/handle/1721.1/97690

[2] https://cyberlaw.stanford.edu/files/blogs/CIS%20Technologists%20Apple%20Brief%20Final.pdf

https://www.eff.org/files/2016/03/03/16cm10sp_eff_apple_v_fbi_amicus_court_stamped.pdf enforcement operatives are doing (for good reasons) exactly what criminals and other bad actors do: exposing someone else's data without their consent. Any change that makes this easier is likely, unfortunately, to make malicious hacking easier too. There are numerous examples of tools or weaknesses that were employed first for legitimate law enforcement and intelligence purposes, but were later shown to be exploitable by everyone (FREAK/Logjam, Dual-EC-DRBG, Wannacry).

Why do some companies deliberately design their systems so that the company can't read the data, while others seem to be able to read everything?

These are two different strategies for making money, not one noble group of cooperators and an opposing group of companies who are deliberately being difficult. All other things being equal, a device or service that doesn't collect or expose information without the user's explicit consent is keeping that user more secure by making it harder for criminals to extract the data. Most of the people trying to extract data from others' devices are criminals or scammers – the same defences work against both criminals and government. Companies that gather more data are taking a calculated risk, offering their users richer and more targeted services, but knowing that *if* their privacy is breached the results could be catastrophic. The very technically savvy companies in this space are among the most powerful entities in the world (Google, Facebook, Microsoft); those who gather lots of data but don't preserve its security precipitate disasters (Yahoo, Ashley Madison).

Whatever the Australian government eventually decides to do, it is very important not to tilt the market to the advantage of those companies that gather vast data about their users. Services that gather less data, or gather only encrypted data, and devices that facilitate user control, are architected to preserve the security and privacy of their users. The overwhelming majority of these are ordinary citizens on whom the police will never want to serve a warrant.

GCHQ/ASD or a security contractor says they can break into a device or protocol anyway. Is this true?

I do not have a security clearance, so there is probably a lot about that side of the operation that I do not know. Certainly various agencies and their contractors seem to know how to extract data from some devices or communications, because of occasional discoveries of errors that undermine the encryption and other security protections. It would not surprise me if there are ways to eavesdrop on end-to-end encrypted messaging apps that a diligent, lucky and well-funded agency can employ, but which the civilian security community does not currently know about. There is a constant process of vulnerabilities being discovered, exploited, and patched, by intelligence and law enforcement services, civilian researchers, and criminals. There are also simple ways to bypass encrypted apps, such as putting false (weak) substitute apps into the appstore instead of the real thing, and substituting people's public keys and hoping their friends don't notice (WhatsApp had a bug related to this). If the reports from the FBI/San Bernadino case are to be believed, there was a way of breaking into iPhone storage that Apple and the FBI didn't know about until a private contractor disclosed it. Everything I have said above about the infeasibility of decrypting messages excluded the use of these kinds of tools and tricks.

The trustworthiness of software updates

Many of the proposals for intercepting encrypted communications without a traditional "backdoor" involve a specific software update for a particular device or person. This is what the FBI asked of Apple in the San Bernadino case. This plan rests on the assumption that it is even possible to make such an update – in fact, many open source software projects already provide ways for users (and their devices) to defend themselves. *Repeatable builds* (or *verifiable builds*) allow users to check that they have an accurate compilation of some open source software. *Update transparency*[3] allows users to verify that they are getting a genuine update, without having to trust a single software provider's digital signature. These techniques are not perfect, but they make this sort of forced individual update substantially harder. They might also be adopted (at least partially) by commercial software providers, because they defend against criminals using the update system as an attack vector.

Some other points for discussion

- If we force a company, *e.g.* Apple, to be able to turn over data, what happens if other governments (perhaps ones we don't like) insist on Apple turning over data on visiting Australians' devices?
- Data breaches and security problems happen all the time. In the last few weeks there are reports of a serious vulnerability in most of the world's computer processors[4], a ridiculous error that undermined Outlook encrypted mail[5] and many, many more. Insisting on feasible interception is likely to increase the (already high) rate of these sorts of problems. It also allows the Australian government to be blamed, even for problems that may have happened anyway.
- Australian companies trying to sell trusted, secure software already face a highly competitive international environment. Many of these companies enhance users' security by designing systems that give users full control over their decryption keys, so the services are end-to-end encrypted and the provider does not see the data. Forcing them to acquire and hand over data would undermine their customers' trust in two ways: it would greatly increase the complexity of the system, hence increasing the likelihood of accidental errors, and it would raise the possibility that the Australian (or some other) government could abuse its power and demand the data. Since there are already documented cases of Australian journalists' metadata being read without the necessary warrant,[6] this concern would be justified, especially for export customers but also for Australians.

[3] https://eprint.iacr.org/2017/648.pdf

[4] https://spectreattack.com/

[5] https://www.theregister.co.uk/2017/10/11/outlook_smime_bug/

[6] http://www.abc.net.au/news/2017-04-28/afp-officer-accessed-journalists-call-records-in-metadata-breach/8480804

- The UK's Investigatory Powers Act lacks sufficient oversight and transparency to allow the public to understand how and where it is being used. At an absolute minimum, Australian legislation should mandate an opportunity for open public scrutiny and review. Older implementations of key escrow, such as the Clipper Chip, were discontinued when security researchers demonstrated that they introduced serious weaknesses into ordinary encrypted communications. If they hadn't been able to access the chip, those weaknesses might have gone unnoticed. There must be some process for allowing independent external scrutiny of the wider implications of whatever is eventually mandated in Australia.
- Precision and detail are good for public discourse, rational decisionmaking, fairness, and accuracy. The more detail that the government provides, and the clearer and more specific the draft legislation, the higher the chances of a rational discussion about what the implications are, and in particular whether the risks are worth it.

Conclusion: A comparison with other Australian efforts at cybersecurity-critical tasks

I led the team that re-identified suppliers, and later patients, in the MBS-PBS open dataset.

The MBS-PBS dataset release, like compulsory interceptions of data, was motivated by a genuine and worthwhile need. Everyone would like legitimate medical researchers to access health records for research, just as everyone would like police officers and intelligence services to catch terrorists and other serious criminals. Unfortunately the high-level instructions given for the release of sensitive unit-record-level data did not adequately consider whether the release *with adequate protection of privacy* was feasible. The result was a serious data breach.

The consequences of technology-related laws can be very different from their authors’ intentions. It is perfectly reasonable to insist on the handover of data that is already available to a service provider, if a proper warrant is produced. But ordinary Australians would be less secure if companies had to re-engineer their products for easier data access.