Cisco's Response to "Strengthening Australia's Cyber Security Regulations and Incentives – a call for views"

https://www.homeaffairs.gov.au/reports-and-pubs/files/strengthening-australia-cyber-security-regulations-discussion-paper.pdf

27 August 2021

Why should the government take action?

1. What are the factors preventing the adoption of cyber security best practice in Australia?
2. Do negative externalities and information asymmetries create a need for Government action on cyber security? Why or why not?

The major factors affecting cyber security best practice uptake in Australia are common around the world. Cybersecurity is a complex technical risk domain and for many can be a complex business risk domain. Whilst general awareness of cyber security threats may have grown in the last 12-24 months, knowledge of what to do and where to invest may not have. Whilst the ASD's Essential 8 provides a great resource for guidance, especially where immaturity lies, they are a subset of ACSC prioritised mitigations for cyber security incidents. However, the adoption of the Essential 8 or any similar best practices (or bare minimums) is not driven by the Essential 8 themselves.

In the Strategies to Mitigate Cyber Security Incidents, ACSC suggest the "pre-requisites" to following the strategies – pre-requites which accurately summarise the factors preventing adoption (bolding added for emphasis)

> *"Prior to implementing any of the mitigation strategies, organisations need to* ***identify their assets*** *and* ***perform a risk assessment*** *to identify the level of protection required from various cyber threats. Furthermore,* ***organisations require motivation*** *to improve their cyber security posture,* ***supportive executives****, access to* ***skilled cyber security professionals*** *and* ***adequate financial resources****. Motivators can include a significant cyber security incident, a penetration test, mandatory data breach reporting, mandatory compliance,* ***and evidence of a lower cyber security posture or higher threat exposure than previously realised****.[1]"*

Similarly, the current review of the security of Critical Infrastructure and Systems of National Importance started with a focus on Governance Rules and a Risk Management Program. Many critical infrastructure operators and large enterprises already capture cyber security risk as part of their risk management programs. However, given the diversity the types of organisations and the size of those organisations from the micro-business to the

[1] Strategies to Mitigate Cyber Security Incidents | Cyber.gov.au

multi-national, there is no single risk management framework that is suitable for all. Adding to voluntary or mandatory cyber security best practices will not maximise the desired outcome if organisations have not acknowledged the risks that the best practices nominally treat. Investigation and guidance on the most appropriate risk management framework for different businesses would be a starting point to addressing barriers to adoption.

Successful examples exist where businesses have had to implement cybersecurity capabilities as a condition of doing business. The cost of a compliance-based focus creates barriers of entry and disproportionate costs to smaller business. Guidance based approaches, incentive and assistance frameworks and centralised government funded solutions should also be considered. An example is the Centre for Defence Industry Capability (CDIC) provides advisory services and potential Capability Improvement Grants to assist SME in the Defence sector meet DISP requirements.

The current regulatory framework

3. What are the strengths and limitations of Australia's current regulatory framework for cyber security?
4. How could Australia's current regulatory environment evolve to improve clarity, coverage and enforcement of cyber security requirements?

There is a risk that in advocating for cyber security best practices, the outcome will always be zero cyber security incidents. Defenders must be right 100% of time the time and attackers only once; therefore, greater emphasis should be given to the concept of cyber resiliency as part of normal business continuity planning. Examples of strengths of Australia's regulatory framework include the shift to more risk-based assessments rather than compliance such as the ACSC's Cloud Assessment and Authorisation Framework and the Department of Home Affairs adopting a principles based rather than prescriptive approach in the proposed changes to the security of critical infrastructure.

APRA's recent announcements on focusing on "operational resilience" including cyber-resilience for the banking and finance sector is an example of the type of guidance that could be provided to other sectors which ultimately underpins the need for and adoption of cyber-security best practice.

Regulation has the capability to focus on one part of the cyber security landscape, commonly observed as a problem across three areas of business operations: people, process and technology. Business cannot solve all these issues, which requires both looking outside of business regulation and how people are interacting with services.

The largest gap we have in effectively identifying regulatory gaps is the lack of data showing where vulnerabilities have been exploited, or in other words, breach data. Data breaches notification rules currently only addresses PII data and contains several excluded entities. A good area of investigation would be to understand what is being breached and

how it is being breached, and then map government mandatory breach notification policy to that information to gain further visibility.

Such visibility would inform and help prioritise other efforts. For example, a rule forcing a security risk reporting framework on all business may have significantly reduced value if the major factor contributing to breaches is in fact email security. One in six Australians fell victim to cybercrime during lockdown in 2020[2] which may warrant a different focus of efforts to address these threats, such as end user protection and company digital communication standards.

A need for lightweight reporting regulation will best deliver the desired policy outcome. For example, bodies such as ASIC could require cyber breach notification requirements for Australian business. The objective should also be to broaden the coverage of notifications to all incidents that affect the availability, integrity or confidentiality of business or the data it collects. Such measures should not overlap with other regulations, resulting in multiple reporting obligations.

Lastly, small business and consumer level cyber security needs can be very similar. It is possible to increase the availability of effective cyber solution through existing channels. Technical capabilities (discussed below) may be incentivised, encouraged or mandated through existing regulators or centrally funded by government. These can be more cost effective as a whole and even cover consumers potentially.

Exclusions from cyber law can also be seen as a weakness. An example are exclusions from the Privacy Act for State and local government agencies. Whilst there will always be valid exclusions, these need to be rationalised and minimised. We note that the Privacy Act Review process is still ongoing and there has been no public release of findings from the Issues Paper shared with industry and citizens in late 2020. The outcome of this review needs to be understood to inform possible evolution of cyber security requirements that could be linked to the Act. Government should be an exemplar for cybersecurity and cybersecurity measures should apply to government entities by default.

Governance standards for large businesses

5. What is the best approach to strengthening corporate governance of cyber security risk? Why?
6. What cyber security support, if any, should be provided to directors of small and medium companies?
7. Are additional education and awareness raising initiatives for senior business leaders required? What should this look like?

Cisco has strong corporate governance of cyber security and recommend a combination of the options presented – a mandatory requirement to implement corporate governance of cyber security risks, possibly through strengthened wording of APP11's "reasonable

[2] NortonLifeLock Digital Transformation Report 2020

steps", combined with a voluntary governance standard that can be adopted or an equivalence demonstrated. Flexibility in the proposed standards and frameworks will result in better outcomes across all sectors rather than a one size fits all approach. Critical to all discussions regarding selection or development of standards is to adopt a principle of "standards equivalency" and "alignment with the intent of a standard". Often standards are rooted in the history of a particular industry sector and are difficult to transpose to other sectors or organisations of differing size or orientation.

Compliance with standards can be costly and generally more costly the more comprehensive or complicated the standard is (e.g. ISO27001). This can disproportionately impact small to medium businesses; however this is the sector that remains the largest area of concern, both from a cyber maturity and visibility (or lack thereof). While there is still work to be done, Australian large businesses are generally mature and/or well regulated.

When looking at large business in Australia, attempts to take a guidance based framework such as the ASD's Essential 8 as a 'tick the box' compliance framework is not suitable for some sectors and large enterprises. Point controls developed and prioritized for a homogeneous business environment do not translate into complex, heterogeneous and more dynamic business environments. While the Essential 8 and other related or similar recommendations provides good guidance, it has not been developed as a compliance framework or tool and should not be treated as such.

The recommendation of a more comprehensive, risk based approach is part of the E8 Maturity model itself. There is discussion on the treatment of exceptions, yet still meeting the requirements of a given maturity model using compensating security controls. Where controls are difficult to apply to a particular technology domain – ACSC already provide specific guidance for Linux systems as an example – alternate treatments are valid.

The World Economic Forum's recommendations[3] offer a principles based approach on corporate cyber risk governance. Whilst directed at business, the guidance outlined provides good focus areas for efforts in Australia.

Minimum standards for personal information

8. Would a cyber security code under the Privacy Act be an effective way to promote the uptake of cyber security standards in Australia? If not, what other approach could be taken?
9. What cost effective and achievable technical controls could be included as part of a code under the Privacy Act (including any specific standards)?
10. What technologies, sectors or types of data should be covered by a code under the Privacy to achieve the best cyber security outcomes?

[3] http://www3.weforum.org/docs/WEF_Cyber_Risk_Corporate_Governance_2021.pdf

The Privacy Act Review process is still ongoing and there has been no public release of findings from the Issues Paper shared with industry and citizens in late 2020. Some of the issues discussed such as the current exclusion of small business and other entities directly relate to the applicability of APP11 and hence uptake of cyber security standards. Additionally, Section 11.8[4] already lists ICT security domains where "reasonable steps" should be taken. What is missing is further advice on reasonable steps that is suitable for the wide range (and possibly expanding) of organisations covered by the Privacy Act. Existing referred guidance is to the PSPF and the ISM – which is government focused.

We do not see a need for a separate cybersecurity code for protection of personal information. The code should be about protection of information in general that address personally identifiable information (PII)issues, however a "one-size-fits-all" approach to cybersecurity protection of personal information may not be ideal given the broad definition of PII and the associated risks to different categories. For instance, the same level of protections may not be justified for a product or service which only handles name and email address, or a MAC address, versus one which manages more sensitive personal information fields such as government service identifiers. The code in general should adopt a risk-based, flexible (and not prescriptive), interoperable and voluntary approach to help manage the cybersecurity risk in processing of personal information. Any such Code should also correspond to and recognize globally accepted cybersecurity standards, encourage accountable practices such as security & privacy by design and take into consideration the size, complexity and type of information being processed. This will be particularly useful for businesses in ensuring that their security compliance and efforts are not only effective but aligned globally and help reduce compliance costs (of having to meet different security requirements in different jurisdictions). The Australian government could take into consideration and make references to frameworks such as NIST cybersecurity framework, ISO 27001 etc.

Mandatory product standard for smart devices

11. What is the best approach to strengthening the cyber security of smart devices in Australia? Why?
12. Would ESTI EN 303 645 be an appropriate international standard for Australia to adopt as a standard for smart devices?
13. Would you be willing to voluntarily remove smart products from your marketplace that do not comply with a security standard?
14. What would the costs of a mandatory standard for smart devices be for consumers, manufacturers, retailers, wholesalers and online marketplaces? Are they different from the international data presented in this paper?
15. Is a standard for smart devices likely to have unintended consequences on the Australian market? Are they different from the international data presented in this paper?

[4] Chapter 11: APP 11 — Security of personal information — OAIC

Cisco's position outlined in our Cisco 2020 IoT Code of Practice Submission has not changed. Smart devices and the IoT represents a proliferation of endpoints that are difficult to secure, manage and update. In the SMB and consumer space, cost pays a very prominent role in buying decision. Vendors with lower security, and hence generally lower cost, are often rewarded with more business.

Recognising that the prevalence of parallel importation and direct manufacturer to consumer models, there is significant online ordering from overseas. Online marketplaces are difficult to regulate, and any additional steps local manufacturers and retailers are required to take has the ability to further challenge their competitiveness. Care must be taken as small cost differential increases can shift a large portion of buying decisions.

Much of what is in the ESTI EN 303 645 is very endpoint focussed. The scale and scope of the endpoint related problem in this space, as well as conventional security thinking, namely 'defence in depth', dictates that we look at other layers. The internet gateway and network layers are well positioned to assist with endpoint problems[5] more generically and is a largely unexplored and under-represented area in security controls in SMB and consumer space. More comprehensive IoT security frameworks such as IoTopia[6] looked beyond the endpoint problem for this reason.

Labelling for smart devices

16. What is the best approach to encouraging consumers to purchase secure smart devices? Why?
17. Would a combination of labelling and standards for smart devices be a practical and effective approach? Why or why not?
18. Is there likely to be sufficient industry uptake of a voluntary label for smart devices? Why or why not?
19. Would a security expiry date label be most appropriate for a mandatory labelling scheme for smart devices? Why or why not?
20. Should a mandatory labelling scheme cover mobile phones, as well as other smart devices? Why or why not?
21. Would it be beneficial for manufacturers to label smart devices both digitally and physically? Why or why not?

The concept behind labelling, notably the ability to provide relevant information to the consumer at time of purchase has merit, at first glance. There are a number of underlying issues with this approach that relates to smart devices and security that lead us to advocate priority of digital labelling concepts over physical.

[5] https://csrc.nist.gov/publications/detail/white-paper/2021/05/20/trusted-iot-device-network-layer-onboarding-and-lcm/final
[6] https://globalplatform.org/iotopia/

A label must provide meaningful value. In order to do this, a label has to not only be understood, but also be something that is cared about by the purchaser. Research[7] has shown that a label can provide a false sense of security to consumers, ultimately resulting in a reduced security posture when the consumer is led to believe that a standard outlined on a label provides protection.

As threats evolve, the standard attached to the label must also evolve. A simple example of this evolving threat is what we must do to enable the standard for a post-quantum world. This creates an administrative burden in addition the existing efforts to ensure compliance, requiring management of the evolution of the labelling that is still understood by the public. It will raise questions such as "Is Labelling v2 better than v1? Does my Labelling v1 product comply with v2?".

Consumer awareness and understanding of cyber security concepts, or even why they should care about elements such as product support lifetimes, vary across the population. Cost most likely remains the single biggest factor in buying decisions and any security comes at increased cost.

Aspects that would ordinarily be considered standard of any product, need to apply to smart devices. "Fit for purpose" can and probably should apply to cybersecurity principles. There is opportunity for consumer law provided by ACCC to enforce concepts like software updates during the device warranty period, including frequency and timelines for divergent CVSS ranges.

However, there will always be vulnerable devices at any given point in time and we must extend autonomous capabilities to protect against exploit. The concept of a digital label becomes something that can inform autonomous functions, explicit management capabilities, as well as the end consumers, and hence a compelling area to focus on.

The Digital Label

A digital label is a set of claims that can be received and processed by automation. It should be signed by someone who has certified the product as having met a standard. It may be renewable. It may contain information about whether a product continues to be certified. In addition, a product supplier can apply a new digital label, should a product be qualified for additional certification. The label can indicate when a product is outdated and should be replaced. The label can also indicate other information, such as what sort of network protections the product needs, and what sort of vulnerabilities it may have.

Most importantly, digital information in standardized, machine-readable formats can enable intelligent, intuitive network services to automatically identify, provision, and protect devices by enabling only those permissions necessary for the device to operate as intended by the device manufacturer, the purchaser, and the network operator. Digital labels can be built on previous work of the National Cybersecurity Centers of Excellence (NCCoE) such as NIST SP 1800-15, and existing standards such as Manufacturer Usage Descriptions (MUD) [RFC 8520]. MUD provides a framework for information exchange

[7] Adelman, B., Adverse Selection in Online "Trust" Certifications, Harvard University, 2006. https://www.benedelman.org/publications/advsel-trust-draft.pdf

about a device that can easily be extended to include certifications that can include software bills of materials and how to find security advisory information.

This information can be further leveraged by controllers that can be embedded in gateway devices such as home routers[8]. They can then automate functions such as network based controls and IPS ruleset for attacks against the known vulnerabilities. Such functions do not require human input or can be used to inform human workflows better. As such layers of automated protection do not require consumer input, this would lift the bar of cybersecurity across Australia.

Responsible disclosure policies

22. Would voluntary guidance encourage Australian businesses to implement responsible disclosure policies? If not, what alternative approaches should be considered?

Cisco are advocates of transparency and vulnerability disclosure. However, we recognise that this is not necessarily the case across all vendors, and practises vary widely. The industry needs to normalize vulnerability research as a standard cost of doing business. Australia would do well to lead the way internationally on this activity. Responsible disclosure is very important and any new regulations should focus on responsible disclosure. Such efforts must consider the vulnerability management part in conjunction with disclosure practices of researchers, without which the disclosure becomes meaningless.

Key items need to be incorporated in disclosure, allowing researchers to contact the company to give them time to validate and patch the said vulnerability. Contact details and appropriate policy are things that can be required by regulation. The vulnerability management piece is more difficult.

Limitations on disclosure period need to be considered. Examples exist such as Google Project Zero which has a 90 day disclosure deadline from the notification to public sharing. In many cases, fixes for hardware related issues (e.g. Spectre) can take much longer and circumstances may warrant non-disclosure to protect industry. Policy mechanisms to allow for this are difficult implement and mediation by government entities would be problematic. Some initiatives have a longer mandatory period and start form the time of reporting, which appears a more reasonable approach.

A clear area where guidance and/or regulation can assist with vulnerability management is ensuring details of vulnerabilities fixed in patches are reported. Many vendors are not disclosing CVE details of security patches, leading to a lack of clarity and uncertainty. This can assist with patch prioritisation efforts withing organisations with affected products deployed.

[8] https://www.ripe.net/publications/docs/ripe-759

Voluntary health check for small businesses

23. Would a cyber security health check program improve Australia's cyber security? If not, what other approach could be taken to improve supply chain management for small businesses?
24. Would small businesses benefit commercially from a health check program? How else could we encourage small businesses to participate in a health check program?
25. Is there anything else we should consider in the design of a health check program?

A Trustmark underpinned by a voluntary system with some government assistance has the potential to raise both posture and awareness within the small business sector. Some of the earlier discussion around the labels is applicable, depending on the framework around the certification. The UK government chose 5 key technical control areas that are fairly simple. The Essential 8, conversely would not be an appropriate framework.

A voluntary systems needs to consider ways to drive adoption or participation to enable widespread success. This could take the form of tax or fee reductions, reduced liabilities or other financial incentives.

In the section 28 below, we discuss some other options the government should consider. The intent of these is to help facilitate more affordable security options either by default, or so actions can be taken to close gaps identified from the assessment.

Clear legal remedies for consumers

26. What issues have arisen to demonstrate any gaps in the Australian Consumer Law in terms of its application to digital products and cyber security risk?
27. Are the reforms already being considered to protect consumers online through the Privacy Act 1988 and the Australian Consumer Law sufficient for cyber security? What other action should the Government consider, if any?

In the above sections there exists for opportunity for consumer law changes as per the above sections:

- "Warranty" periods to cover security patch support, with a regulated frequency and timeliness of fix based upon severity

- Australian Consumer law already has provisions to protect against misleading representations and products being fit for purpose. These concepts apply to the cyber security of the product.
- Require digital labelling requirements for smart devices

Existing regulatory bodies may be leveraged to offer security capabilities by default that support the desired outcomes. Some of these outcomes are outlined below in section 28, for example the Telecommunications Services Regulator could be leveraged to enforce cyber capabilities in end user devices, and to help prevent widespread attacks on these devices[9].

Other Issues

28. What other policies should we consider setting clear minimum cyber security expectations, increase transparency and disclosure, and protect the rights consumers?

As mentioned earlier, the ACSC Essential Eight is widely recognised and consumable useful security advice as "baseline mitigation strategies". We would support an equivalent focus on promoting an introductory framework such as the UK NCSC 10 Steps to Cyber Security which emphasise the importance of risk management, asset management, and other domains[10].

Opportunities for strengthening cyber defences of Australian business lie outside of compliance efforts. Technology based measures are an area that can potentially be less costly, easier and faster to implement, and more effective than compliance efforts, yet are rarely explored. The concept here is ensuring "security by default" through technology based capabilities that can either be subsidised/offered by government or regulated controls through existing channels.

Examples of controls that provide security by default are:

Whole of country DNS security[11] efforts targeting SMB and consumer spaces.
This is a control that can easily be opted out of, provides immediate uplift in protection and threat visibility, and is relatively lightweight and cheap. Cisco pioneered this approach to DNS security and efforts to date in this area have been sovereign capability and public sector focused. For a consumer and business focused initiative the sovereign requirements may be less important and operational and security efficacy considerations may outweigh them. Factors such as consumer trust also come into play, whereby increase opt-out activity would occur if the solution were viewed as surveillance – in this case a non-government operator would be a perceived advantage.

[9] https://www.tomsguide.com/news/arcadyan-router-malware
[10] 10 Steps to Cyber Security - NCSC.GOV.UK
[11] https://www.gartner.com/doc/reprints?id=1-26QPQDNY&ct=210708&st=sb

Uplift in internet gateway capability.

As it stands, most Australians run only the hardware provided by their ISP, which tend to be cheap and offer little in the way of security. Whilst configuring security can be complicated, a focus on support for autonomous capabilities and visibility and protection capabilities by default would normalize enhanced security. When looking at items such as low touch secure onboarding and automated protection schemes mentioned above, it requires consumer electronics support of these standards, which are further elucidated in the IoT labelling discussion. Things like Wi-Fi security and simple firewalling and traffic visibility are areas that should be considered entry level capabilities.