Lockstep Consulting
11 Minnesota Ave
Five Dock NSW 2046

8 January 2013

Angelene Falk
Acting Assistant Commissioner -- Compliance
Office of the Australian Information Commissioner
Level 3, 175 Pitt Street
Sydney NSW 2000
Email: Redacted

Dear Assistant Commissioner,

Comments on Consultation Draft Guide to information security

Lockstep applauds the OAIC's efforts to develop guidance for ICT practitioners seeking to interpret the Privacy Act. In my work over the past 15 years or more, I have continually encountered examples of project managers and architects misjudging what privacy means for their work. So this guide addresses a very real need.

Introduction: technologists' privacy blind spot

I have come to believe that a systemic conceptual shortfall affects typical technologists' thinking about privacy. It may be that engineers tend to take literally the well-meaning slogan that "privacy is not a technology issue". I say this in all seriousness. Only last month, during the course of conducting a PIA, I spent time with the development team for a new government register. These were good, senior people, but they harboured typically restrictive views about privacy. For instance, I noticed they tended to refer to "private" information rather than *Personal Information*. After explaining that Personal Information is the operable term and reviewing its definition from the Privacy Act, I found that the team had failed to appreciate the extent of the PI in their system. They overlooked that most of their audit logs *collect* PI albeit indirectly and automatically. Further, they seemed not to understand that information about clients in their register provided by third parties was also PI, a virtual blind spot I attributed to their weak informal frame of "private" information.

About Lockstep

Lockstep Consulting is a specialist advisory firm dedicated to identity security, digital risk management and privacy. Established in 2004, Lockstep provides research, analysis and advice to private and public sector clients across the Asia Pacific. Sister company Lockstep Technologies undertakes innovative, award winning R&D in PETs and digital identity protection. Lockstep has performed Threat & Risk Analyses and PIAs for clients including the Victorian smart meter program, DFAT, Queensland Health, Vicroads, NEHTA and

Australia Post. Our unique statistical security ROI calculator developed for the NSW Government was adopted and adapted by the US Department of Defence.

I have long been involved in public policy development for security and privacy. I was a member of the Australian Law Reform Commission's Developing Technology Advisory Sub-committee (2007-08), the National Electronic Authentication Council (1998-2001), the Federal Privacy Commissioner's PKI Reference Group (2000), and the APEC TEL eSecurity Task Group (1998-2001). Recently I have been selected by the Identity Ecosystem Steering Group (IDESG) to lead the privacy evaluation of the US National Strategy for Trusted Identities in Cyberspace (NSTIC).

The state of the draft

I find that the draft *Guide to information security* is certainly a worthy document, but it seems to fall short on two broad fronts. The guide catalogues a great many worthwhile security measures, yet it doesn't often explain the privacy rationale for any specific measure, nor does it much discuss how to gauge the appropriate strength to be deployed. Instead the guide for the most part simply asks if such-and-such is in place, or not. This is the first broad problem.

The second problem probably explains the first: there is no obvious unifying framework or theme to knit all the security advice together, and to allow readers to extend the guidance to their own circumstances. There are conventional ways for information security measures to be analysed and selected—in particular, Threat & Risk Assessment (TRA)—and I suggest that the guide be based on these methods. I discuss TRAs in further detail below.

Scoping the information security guide

As mentioned, many information architects and software designers seem to misjudge what privacy means for their work. I have noticed many recurring misconceptions including:

- *Personal Information* is sometimes taken to mean especially sensitive information such as payment card details, rather than *any* information pertaining to an identifiable individual[1]
- the act of collecting PI is sometimes regarded only in relation to direct collection from the individual concerned; technologists can overlook that PI provided by a third party to a data custodian is nevertheless being *collected* by the custodian, and they can fail to appreciate that generating PI internally, through event logging for instance, can also represent collection
- even if they are aware of the Access and Correction Principle, database administrators can be unaware that, technically, individuals requesting a copy of

[1] An exchange between US data breach analyst Jake Kouns and me over the Epsilon incident in 2011 is revealing of a technologists' systemically narrow idea of PII; see http://datalossdb.org/incident_highlights/52-epsilon-bingo.

information held about them should also be provided with pertinent event logs (a non-trivial case where individuals can have a genuine interest in reviewing event logs is when they want to know if an organisation's staff have been accessing their records; the ability to furnish this type of event log is critical in some systems but can be overlooked by designers).

These instances, among many others in my experience working across both information security and privacy, show that ICT practitioners suffer important gaps in their understanding. Security professionals in particular may be forgiven for thinking that nine of the 10 NPPs are legal niceties with only NPP 4 being relevant to them. Yet every one of the privacy principles is impacted by information technology and security practices [1]. I believe the gaps in their privacy knowledge are not random but are systemic, probably resulting from privacy training for non-privacy professionals being ad hoc and not properly integrated with their particular world views.

Specific comments

In this section I offer specific comments on the draft, many of which will substantiate my concern that many of the suggestions are lacking a rationale and fall short of quantifiable.

- Page 6: *Privacy and your business*, para 2: "It is important for entities to integrate privacy compliance into their risk management strategies" (underline added). This is true of course but as a matter of emphasis, I suggest "compliance" strikes the wrong note at this point in the document. The previous paragraph is written in a more inclusive way, showing that privacy is central to customer relations, so it seems a little heavy handed to then switch to the harsher language of compliance. I would simply delete "compliance".[2]

- Page 7: *Privacy by design and privacy impact assessments*, para 1: "For example, entities should consider the security of personal information before they purchase, build or update information technology (IT) and electronic records management (ERM) systems." This example seems ad hoc, coming at the end of a passage dealing with much larger issues. Why would ERM be germane here, and why is it appended to "IT" systems? More importantly, what exactly should entities do in 'considering the security of personal information'? This example does little to help the reader if it doesn't include specifics.

[2] I'd like to make a larger point here, albeit one that is tangential to the information security guide. At the recent iappANZ annual conference, there was abundant and earnest talk of Privacy by Design, but frankly very little practical detail on how to go about it. Moreover, many of the presentations focused on compliance and law reforms. I do not wish to diminish the importance of compliance, but if the profession is serious about Privacy by Design, it needs to re-balance the discourse. Privacy professionals' comfort zone historically is in regulations, compliance and audit, which is not the natural frame for ICT design.

— Page 8: *Circumstances that affect assessment of reasonable steps*: the title itself strikes an odd note. If the guide is addressed at security professionals doing design, then the emphasis should not be legalistic. I would suggest a title more along the lines of *How do you decide which security controls are reasonable?*

— Page 8: *Circumstances that affect assessment of reasonable steps*, bullet 3: "… if the information is not secured" would be better written "… if the information is breached" (harm is only done after a breach; there is no actual harm in being "not secured" until a breach occurs).

— Page 8: *Nature of the entity*, para 1: "Factors include the size of the entity and the business model on which the entity operates". I would have expected many more factors than these two. ICT practitioners will need a more complete account of what to think about here.

— Page 8: *Nature of the entity*, para 1: "if an entity operates through franchises or dealerships, or gives database and network access to contractors, the steps that it is reasonable for it to take may differ from the steps that it is reasonable for a centralised entity to take". This point needs to be expanded. What difference do these factors really make?

— Page 8: *Nature of the entity*, boxed case study: This is an excellent case (as are all the others in the draft). However, it is not apparent to me how the Vodafone Hutchison misadventure illuminates the particular topic of *Nature of the entity.*

— Page 9: *Nature of Personal Information held*: overall, this section really says little about the *nature* of PI held. Most of the section is taken up by the legal definitions. The last para says that "as the quantity, extent and sensitivity of personal information held increases, so too will the steps that it is reasonable for that entity to take to protect that information" but there is no concrete guidance. The boxed case that follows is wholly concerned with processes (such as reviews and Protective Security Assessment); it doesn't describe any steps as such that a security designer might adopt. Finally, given that the section is meant to guide security people handling different types of PI, it would be good to have more than one case study, to provide some sort of contrast.

— Page 11: *Risk of harm*, para 3: "The Commissioner concluded that Medvet did not have reasonable steps in place to protect the personal information". Perhaps there is more detail in the linked investigation report, but it would be better if the guide described the sorts of concrete steps that the Commissioner found wanting.

— Page 11: *Data handling practices*, para 1: "Relevant factors include whether those [data handling] suppliers are subject to the Privacy Act." I suggest this is exactly the sort of technicality where readers would benefit from more detail. For starters, the guidance appears to put the cart before the horse: if the entity is itself subject to the Privacy Act, then my understanding is that any data handling it outsources to another party is by

extension also subject to the Privacy Act. Users of the guide may also find it a little frustrating to read "relevant factors include …" when only one item is then provided. These matters may be too important to be treated open-endedly.

— Page 12: *Ease of implementation and proportionality*, para 1: "It may not be reasonable to implement a measure if doing so will be impracticable or unduly expensive when balanced against the risks" (underline added). This is another instance where the guide makes a generalisation without drilling down another level or two to suggest how exactly the "balance" can be reached.

— Page 14: *Steps and strategies which may be reasonable to take*: In general, this section recounts a great many controls that would *routinely* be considered by security designers in any organisation. There is not much here to indicate how privacy considerations would alter the way security professionals approach their work.

— Page 14: *IT security*: Itemising "IT Security" alongside "data breaches", "personnel security" and so on seems arbitrary. I wonder if the eight bullet point structure at the top of p14 is adopted from some external reference?

— Page 16: *Access*, para 1: the little tutorial about multiple authentication factors seems disconnected, and contains nothing that is actionable with respect to privacy. It is important to not advocate generically for any particular number of factors. There should be no implication than two factors is inherently better than one; for example, using a well-chosen passphrase alone (one factor) is probably better than a cheap fingerprint scanner plus a lazy password (two factors). Specific authentication solutions must always be guided by a TRA tailored to the organisation's circumstances.

— Page 16: *Access*, major bullet 6: audit trails for access to databases are indeed a vital privacy enhancing measure. I would actually like to see auditing discussed more prominently and in more detail in the guide. In general, I find that information systems tend to *over-log* the interactions of database subjects (allowing extraneous PI to build up in logs for no real purpose), and at the same time *under-log* the activities of administrators (making it too difficult to trace possible illicit access to records).

— Page 17: *Encryption*, major bullet 1: The guide does nothing here to explain the importance of encryption, or gauge any of the many design variables that go with it. For example, who administers encryption keys? How long should the keys be? What algorithms are best? Key recovery mechanisms are essential but how are they administered so privacy is preserved? And finally, what privacy protecting options are available if encryption is not deployed?

— Page 17: *Encryption*, minor bullet 2: It is important to not over-simplify advice around encrypted email. There are in fact very few easy-to-use cross-platform email solutions in existence. In my experience, well-meaning ambitions to deploy secure

email almost always collapse, which presents a policy dilemma to authorities that sometimes mandate encryption without due regard to its practicality.

— Page 18: *Network security*, para 2: "What sorts of firewalls are employed …?" How does this matter to privacy?

— Page 18: *Testing*: This section does not seem to me to be well integrated with privacy. There is only a vague link made in the guide between software quality and privacy, with some breaches being made possible by bugs or failures. So for that reason, the reader may presume that testing is important, but the connection is weak.

  Having said that, I regard testing as a vital topic, with its own special privacy issues that should be fleshed out in greater detail. Testers need to take care with test data, and have formal means for suppressing PI for real subjects, especially when testing is contracted outside the organisation. Yet in some cases, using dummy data can compromise the meaningfulness of testing; if "live" data must be used for whatever reason, then a PIA should be undertaken, and the organisation's Privacy Policy reviewed in case it ought to disclose any special PI flows caused by testing. Diagnostic software features sometimes used in testing (like extra detailed event logging) can threaten privacy if they are inadvertently left in production code.[3]

— Page 19: *Physical security*, major bullet 5: "Is there a clean desk policy?" This is another example where the guide implies something is necessary but provides no calibration. The criticality of a clean desk is highly context dependent, varying according to the type of work and its location.

— Page 20: *Personnel security*, para 2: "'Spear phishing' is a personalised attack utilising personal identifiers to attempt to appear legitimate to a particular user" (underline added). I suggest changing "personal identifiers" to "personally relevant information". Identifiers are not generally relevant here; most spear phishing appeals to its targets by referencing details that are especially applicable to their workplace, such as company information or recent news.

Information security Threat & Risk Assessment

Threat and Risk Assessment (TRA) is an information security requirements analysis tool, widely practiced in the public and private sectors. There are a number of standards that guide the conduct of TRAs, including the Information Security Manual (ISM) for federal government, and ISO 31000 (which supersedes the long standing Australian AS 4360).

A TRA is used to systematically catalogue all foreseeable adverse events that threaten an organisation's information assets, identify candidate security controls (considering technologies, processes and personnel) to mitigate those threats, and—most

[3] Google explained away the illicit collection of home network Wi-Fi transmissions by StreetView cars as an inadvertent consequence of test code being released in the production software build.

importantly—determine how much should be invested in each control to bring all risks down to an acceptable level. The TRA process delivers real world management decisions, understanding that non zero risks are ever present, and that no organisation has an unlimited security budget.

Most TRAs proceed according to the schematic workflow below, beginning with an inventory of the information assets of the organisation. We then identity all realistic potential threats to those assets, and for each threat, estimate the likelihood of it occurring and the impact if it does. In general, likelihood, impact and risk are rated on an ordinal scale such as *negligible, low, medium, high* and *extreme.* TRA standards allow the relative severity of a threat to be gauged in the context of the business; for example, the loss of a computer might be counted as a *minor* incident for a large business, but could be *grave* in a small one.

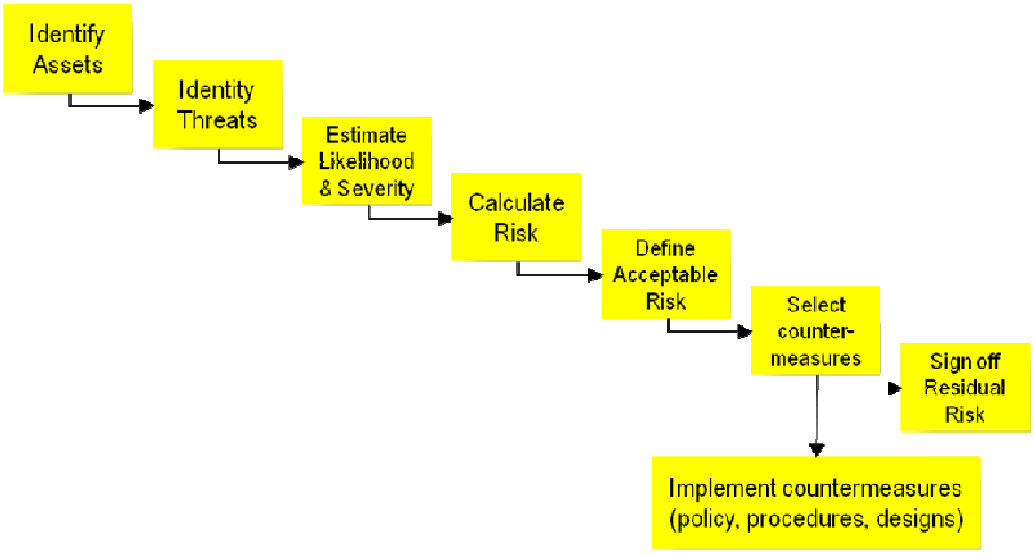

Each organisation needs to define what risk is acceptable for each identified threat. One might presume that *negligible* would generally be the target risk level, but in reality the cost of reducing the likelihood and/or severity of actual threats all the way down to *negligible* can be prohibitive, and so it is often acceptable for target risks to merely be *low*.

Once all identified threats are assigned a risk rating, security designers then work on controls to tackle likelihood or severity or both. Security controls can take substantial time and money, and may need to be prioritised for implementation according to risk. The final step is to re-calculate the residual risks expected to remain once the agreed security controls are in place, so that management can sign off.

Further reading

More material on the intersection of privacy, security and ICT practice is available at our website. The following may be useful:

- http://lockstep.com.au/library/privacy/googles-wifi-misadventure-and
- http://lockstep.com.au/blog/2011/01/26/public-yet-still-private
- http://lockstep.com.au/blog/2012/10/22/types-of-PI
- http://lockstep.com.au/blog/2011/10/24/marrying-privacy-and-infosec

Conclusions and submissions

I have two reasons for raising awareness of Threat & Risk Assessment in the context of the *Guide to information security*. Firstly and most immediately, I suggest that the methodology of TRA is the best way to unify the disparate security measures that are provided in the draft.

More strategically, Lockstep Consulting has found that in practice, the TRA exercise is readily extensible as an aid to Privacy by Design. A TRA can expressly incorporate *privacy* as an attribute of information assets worth protecting, alongside the conventional security qualities of confidentiality, integrity and availability (sometimes collectively dubbed "C.I.A."). A crucial subtlety here is that privacy is not the same as confidentiality, yet as discussed, ICT practitioners frequently confuse the two. A fuller understanding of privacy leads designers to consider the Collection, Use, Disclosure and Access & Correction principles, over and above confidentiality when they analyse information assets.

Lockstep Consulting is actively researching the closer integration of security and privacy practices. We would welcome the opportunity to discuss our findings and our new approaches.

Summarising the concerns set out above, I respectfully submit the following suggestions:

- The guide should start out with a fresh account of privacy framed in line with how security professionals view the world, with a focus on controlling information flows rather than confidentiality or secrecy.
- The guide should be framed around conventional security Threat & Risk Assessment. The reasonableness of any particular security step is usually gauged through a TRA. That would allow adverse impacts on privacy to be analysed alongside other types of information asset, and the reasonableness of candidate privacy enhancing measures to be gauged in ways familiar to ICT practitioners.
- Security professionals might be more comfortable with privacy if it were positioned as another aspect of the value of the information assets that they work to secure.

— The guide should show that security professionals ought to concern themselves with all the legislated privacy principles, and not just IPP 4 and NPP 4. Privacy is more than confidentiality: it has just as much to do with an individual's right to know how and why their information is being used, when and where it is used and by whom.

In closing, I wish you well with further refinement of the *Guide to information security*. This is important work and I expect it will be well received by ICT practitioners, especially if it is more closely aligned to the way they practice and think about security.

Yours sincerely,

Stephen Wilson
Managing Director

By e-mail.

References

[1]. *Mapping Privacy requirements onto the IT function*, Stephen Wilson, Privacy Law & Policy Reporter, Vol. 10.1& 10.2, May-June 2003 (archived at http://lockstep.com.au/library/privacy/mapping_privacy_requirements_).