

Conference on Formulating a New Approach to Privacy and Trust in the Information Age

***Parliament of New South Wales
Chamber of the Legislative Assembly***

4 July 2007

Closing Remarks from the Chair

***An edited version of remarks made by
Malcolm Crompton
Managing Director, Information Integrity Solutions Pty Ltd***

It falls upon me to try and summarise the day and then point to some directions forward. I will give you a particular perspective which is mine and mine alone.

I'm going to summarise quickly particular aspects of what many people said at various times then make some concluding observations.

My observations start with something said quite early in the day. The basic concept of individual control became problematic the first time that systems engineers figured out how to link many data sets in a system and then to link many systems. As a result, frameworks for protecting personal information have been in breakdown mode for quite some time. Then on top of that we had the arrival of data analytics and from there probabilistic predictive analysis, all exacerbating risks to individuals, as Marty Abrams pointed out in his description about false positives.

And now we're in an era of what some people would call information governance and especially the associated processes of risk management and management of failure.

Against this background, the Conference then began to hear particular perspectives: a business perspective and a consumer perspective. I want to draw your attention to a couple of the points that were made. From a business perspective, Andrew Want drew attention to four kinds of risk faced by a business in managing personal information: reputation risk, compliance risk, (including unpredictable legal framework change risk and how to manage for it) investment risk and reticence risk.

Andrew pressed very heavily the idea of transparency as a solution to building a privacy compliance framework because it would leverage off these risks. But is it enough to simply say what you're going to do? As Carolyn Bond said, from a consumer perspective does this simply amount to "anything goes as long as you disclose"?

Carolyn Bond also made a point that maybe it's time to end the separation of consumer advocacy and privacy advocacy as things somehow different. Rather, she considered that they are actually much more conjoined than we thought before. Certainly we should look at the Productivity Commission with regard to its review of consumer protection law.

We then heard views on other areas of law that appeared to have been successful and possibly instructive. Philip Argy, for example, observed that parts of the Trade Practices Act made a company liable for the consequences of supplying unsafe goods. This included definitions as to what constitutes an "unsafe good", namely one about which a reasonable person would have expectations as to safety but it fails that test, if necessary as decided in a court of law.

We heard some other analogies. For example in addition to consumer protection laws, securities regulation, and the responsive regulation principles were put forward as instructive. Indeed, we heard that John Braithwaite highlighted responsive regulation principles in an APEC data privacy seminar in January this year. They're still resonating through the analytical process that we're going through in APEC. We used them again in Cairns last week.

We also heard a very timely reminder not to forget Generation Y and the emergence of online social networking. This in turn raised the question of whether we should be protecting this group from themselves as another vulnerable group, or whether something else is happening in the way of a generational shift in attitudes towards privacy.

Now let me make some observations about what I thought came out of the work groups and the ensuing reports back to the conference in plenary. At the risk of gross simplification, though, I won't repeat material that we've only just heard.

Let me summarise those reports this way. We heard about a large array of risks to which we're exposed if we don't act and of a different array of risks if we do act. And

we tried to capture that array of risks. And hopefully we will be able to use that for further analysis from now.

But first, as Nigel Waters said from a privacy advocacy perspective, we need to put on the table that there is not necessarily universal agreement that we have a broken system on our hands.

That said, in thinking through some principles on where we might move forward, a number of observations were put forward:

1. Whatever system is built, it should build in “no surprises” principles. People should be able to feel as though the world around them is moving in an expectable or expected way.
2. There may be a whole series of circumstances that we currently treat as binary that maybe aren't binary. For example, is it really useful to have a definition of personal information and then treat a data element quite differently depending on whether it is or isn't personal information? Is it that instead, we really have only consumer interests or business interests? Maybe the system is not as binary as we've been thinking.
3. We heard calls for simplicity combined with adaptability to respond to such things as change to law, context and to what's happening in other parts of the world. In this regard, we heard a call for greater emphasis on using standards, or some other mechanism that sets rules about what is acceptable that is not directly delivered by Parliament or other legislatures. We have seen this operate very successfully in the great world of the safe motor car, the safe driver and the highway. Australian Design Rules are not pieces of legislation passed by Parliament but they certainly have a lot of legislative banking.
5. Enforceable remedies and strong enforcement also seem to be part of any of our future.
6. Moreover, it was strongly put to us that enforcement effort should be directed strategically. In particular, regulators should focus on where harm is greatest.
7. This perspective comes with an extremely important rider: harm mustn't simply be dollar based. There are all sorts of harms that are very difficult to convert into dollars but just because they can't be converted into dollars don't mean that they should be ignored.

8. Whatever we do, we should be helping people to understand if a company or a new offering or different channel is trustworthy, in order to help them form a view on what can be trusted. This is a particular issue in the online world. In order to understand what is trustworthy, what is the evidence: how do we replace the historical cues that our species evolved over millions of years and that we've become extraordinarily subtle at using, when they are simply not available to us when we're transacting remotely or on line?

In other words, how do we put back the cues? Some speakers thought that trust marks or other evidentiary processes like that might help us.

9. However, there's no point in having evidence we can't use or are less than fully able to use. Thus we must include significantly better education in our package of principles, again just as we have done in the package of measures that has been required over a long period of time to improve road safety.

10. Do we need identifiable data as often as it is claimed we need, or can more be achieved with non identifiable data or other mechanisms such as the one I call "pinging"? This is where an inquirer goes to where the data is already housed and asks of data already held by the inquirer: "Is this true or false?" The only answer that the inquirer gets back is "true" or "false". The inquirer doesn't get any other information. That's how Australia has reengineered some of the protections around the Electoral Roll, for example.

11. Even if identifiable data is needed for other purposes such as analytics and other probabilistic mechanisms, what is needed to ensure that such data is not used beyond those original purposes? Is stronger law or some other protection mechanism needed to make sure that's the only thing for which the identifiable data can be used?

Now let me try and draw together these observations into a couple of final thoughts. They also bring me right back to the first couple of speeches that we heard today.

One of them is the Attorney-General's observation: There does appear to be a growing interest and importance in a thing called privacy. We're still not sure what it is but there certainly seems to be a re-emerging interest as seen in an array of forums and for a wide array of reasons. It is certainly coming out of the national security debate, the identity theft debate, the review by the Australian Law Reform Commission of privacy law, APEC and elsewhere.

Second, we have the very interesting proposition that Steve York posed as a question to the Attorney, which in my view probably needs further development. He asked the Attorney, “Don’t we need a chief risk officer for the country?”

This is a really significant insight and deserves very serious consideration. The additional point I would ask is “whose risk?” Unless we ensure the question is framed properly, the “risk” question too often focuses on only one perspective. I would suggest that when we hear that question asked, it usually leads to the chief risk officer analysing risk only from the entity perspective: what risks is the entity facing and how can it manage them, preferably minimise them. If the entity is a government agency, government department, business, or other organisation, it can be an extremely important perspective.

Rarely, though, does the risk question get asked in such a way as to include all the risks faced by all the stakeholders, including those faced by the individual.

Sometimes it is asked sufficiently broadly to include malevolent risks faced by the individual. By that I mean risks such as those arising from data being stolen, misused, lost, or misused in some other way where malevolence by another party is involved.

Rather, the element that is most rarely included in this kind of analysis is the last component and possibly the most subtle. Yet it is the one that will be as important as anything else in generating trustworthiness as an essential precursor to trust.

And this is the risk generated by use of personal information in a way that other parties consider beneficent or beneficial but is not necessarily seen this way by the individuals affected.

Unless the chief risk officer is taking in to account all of these perspectives, and especially the last one, the analysis won’t go right and won’t deliver what the community wants.

Let me give an example, and I pick on it in particular: identity management systems, whether they’re access cards, whether they’re banking processes, or whether they’re simply to get in to the building in the morning, are often not actually reducing risk.

They are transferring risk.

If the organisation is putting in place an identity management system and seeking to reduce its own risks it’s an extraordinarily rational thing to do. But if that’s done at the

expense of somebody else's risks going up, and that party is very poorly positioned either to bear that risk or to mitigate it, have we got either a sensible social policy outcome or a sensible outcome in terms of building trustworthiness in a system around us? Because most identity management systems are aimed at locking people out rather than facilitating service provision, individuals are treated as guilty until proven innocent in any misalignment or incorrect inference and left to carry the consequences.

This is an important point. Think about when things go wrong, who carries the can? Because risk management is too often focused solely on the entity perspective, many have failure management strategies such as business recovery plans, but when things go wrong for you and me, too often it's you and me that end up carrying the can. We are the ones who end up listening to muzak for an awful lot of time while some stupid voice recognition technology doesn't work, or some press-button mechanism won't work or if you do eventually get through to a real person, that individual has no power to do anything except read from a script that is irrelevant. For far too many of us far too often, that is a significant daily risk in using new technologies.

So: when we take the idea of the chief risk officer, and we combine it with some of the harms based thinking that Phil Argy and Marty Abrams were putting forward, have we got the glimmer of how we might rebuild our privacy and trust framework so that is more "future proofed"?

Such an approach would need to be based on:

1. a very expansive view of risk that ensures we're bringing to the table all risks to which all parties might be exposed; and
2. a conscious allocation of all those risks between the parties, based on who should be bearing them, which is often the party also most capable of managing them down.

Now let me add in another risk that is often not factored into "risk" thinking. This is the risk of failure.

With that rather gloomy thought I will cease talking.

To conclude, what have we done today? We had three objectives for the day.

First of all, to foster the capacity for robust, sustained conversation amongst four stakeholder groups, and by that we meant consumer, business, regulator and government.

Second, to begin to identify barriers to value and trust creation in the information economy.

Thirdly, to identify current and future signs of stress or use-by-date in the current privacy framework.

I think we did the third of these quite well. The way I viewed the conversation being conducted today suggest that we at least got a pass on the first test as well: we fostered the capacity for robust, sustained, conversation amongst four stakeholder groups. We have more work to do on the second, although we did get some insights.

But you will note that today one of the things we did not set out to do, and which we certainly didn't achieve, was building through from saying have we got issues on our hands, to a solution. And that's because the Privacy and Trust Partnership is about a journey taking at least the next six months, not just today. It is a journey comprising a conversation between the four groupings that begins to put on the table some possible solutions. But not today.

I would like to think that today is the beginning of that conversation and where we've laid the groundwork, we've asked some questions, but that is all we have done. We're going to take away the materials that we've learned today, the people, the friendships, the networks, and hopefully some of the trust, that we've created today. We will use all this to continue a conversation over the next six months in a way that will be structured but hopefully leads to a second White Paper that goes a lot further than the first White Paper that was given to you to read for today's session.

Thank you all.