Speech to the Privacy Forum: 7 May 2014

Ian Fletcher
Director, Government Communications Security Bureau

Privacy & Security: Identity, society & the state in the internet age

The internet changes everything. Well, not everything: it doesn't change human nature, and it has not yet really changed our sense of how the world ought to work for each of us.

Today I want to reflect on recent debates around the world concerning the internet, privacy and security, and the role of the state. It will not be a commentary on intelligence work. But later I will have something to say about the expectations people in a liberal democracy ought to have of their government's ability to operate covertly on the internet.

But first things first. The internet is big, getting bigger very quickly, and changing very rapidly. Almost any generalisation about the internet is likely to be wrong, or out of date, or both. So, instead, some facts (I am indebted to UK colleagues for these):

Next year, it will take 5 years to watch all of the video crossing the global Internet every second;

Today, every 60 seconds, 204 million e-mails are sent. 47,000 apps are downloaded; 1.3 million individual views try to keep up with the 30 hours of video material uploaded in each 60 second period.

But bad things happen too. Each minute on the Internet today, 20 new victims of identity theft occur, 135 botnet infections take place.

The Internet is big. It's also changing very rapidly: today, the number of networked devices is about the same as the global population. By the end of next year, the number of networked devices will equal approximately double the global population.

The Internet is becoming mobile. 1300 new mobile users join the Internet every 60 seconds. Increasingly, social, commercial and public services are delivered and consumed mainly over digital networks.

Generally, this has proved to be a real social and economic benefit. Access to information shifts economic power from producers to consumers, as it reduces asymmetric control of and access to information. It brings producers and service providers closer to their customers, sharpening competition and benefiting consumers and society generally. As governments get better at delivering public services digitally, costs go down and convenience goes up. Public services should be as economically delivered as possible, as simply as possible, and as conveniently as possible. Proper use of digital technologies helps all three.

Yet, as we know, things can go wrong. Mishandling of private data, privacy breaches, and the consequences of cyber crime are all too common. I said at the outset that the Internet changed everything, but not human nature. This is particularly true in respect of cyber crime, where criminals have been particularly inventive adopters of digital technology. The Internet does not change human nature but it can industrialise human urges, including the bad ones. Cyber-delivered espionage is part of this continuum too: the internet doesn't change the reasons for espionage, but it does make it easier, cheaper and much less risky.

There is also the question of surveillance, which has become a controversial issue in Western countries. Here, we need to remember that a good deal of domestic electronic surveillance in Western countries is undertaken by police forces investigating crime. Contemporary policing relies heavily on electronic records.

And, importantly for the argument I am about to set out, everyone in our society and in societies like ours thinks that there ought to be an efficient police force, linked to a system of public justice, both giving effect to rules we generally accept. In our daily lives, we look for order, and we rely on the state to provide the framework.

But a lot of debate in Western countries over the past year [following the unauthorised disclosures attributed to Edward Snowden] has concentrated on claims about state surveillance for security and intelligence purposes. At the heart of this debate there lie three assumptions which I want to set out:

The first is that privacy (seen as a desirable attribute on the internet) is in some way inimical to the value of security, which is linked to the state. Both are caricatures, I will argue.

Secondly, the internet is and should be "free" in the sense that it should not be controlled. And if rules are needed, then they should be in the hands of the private sector, not governments.

Finally, there's the conspiracy theory: this argues that if governments have technical capacity they will use it even if it's not lawful or proper, and the solution to this is

greater transparency. The fundamental claim here is that transparency would solve a trust gap. I will argue that this is wrong.

Privacy & security

But let us start with privacy and security. It seems to me to be both right and reasonable to assert that people (individually, but also at a social, societal and economic level) ought to be able to do their business on the internet with a reasonable expectation of privacy. I would add that to get the benefit of privacy, people will also expect to be able to have an expectation of authenticity: to be able to assert their identity, and have confidence in the identity of others. Privacy without authenticity is a cyber Kingdom of the Blind. Identity matters in all human interactions.

Who wants to intrude on our privacy? The threats are clear: cyber-delivered crime, including identity-based crime. It includes large scale theft of and dealing in credit-card and other identity-related data. The scale is breathtaking; we need only recall the US retailer Target's loss of tens of millions of individuals' records last year. And that's just an example.

Of course, the Police may have reason to intrude. But their work on the internet (child protection excepted) is largely backward looking – investigating crimes that have occurred, rather than preventing or deterring those that have yet to occur.

Big data. The phenomenon is well known: companies collecting the electronic trail we leave behind, and using it to gain commercially valuable insights into our behavior, so as to better market goods and services. The issue with big data is not aggregation (anonymised insights through aggregate data date back to biblical census stories). Rather, the issue is disaggregation: pulling out insights about me, you, your family, and monetising that. It's not the census that threatens privacy; rather, it's the permanent Domesday Book of disaggregated insights. But it pays, and there seem to be strong economies of scale, and these combined with trusted brands help create large, profitable players who actively seek to defend their market positions.

And the state. The state has legitimate concerns with the prevention of terrorism, with the protection of its own information from espionage, and perhaps too with organised crime, and non-proliferation. A surprisingly limited list, I would argue. And a list which directly affects very few people (all of whom know who they are because they're doing really bad stuff), which is good because the scale of the internet (and its potential for anonymity) means that almost all data, and almost all people, are simply out of scope for even the best resourced.

The contrast is the analogue world, where the State regulates a lot more.

In the analogue world there is the Hobbesian bargain, where we give up any private rights to use coercive force, in return for a framework of rules and a means of enforcement. Thomas Hobbes published his book (popularly known as *Leviathan*) in 1651. Written against the background of the English civil war (100000 had died from a

population smaller than New Zealand today; the cost of war was horribly at hand), Hobbes made the case for Government rule making and enforcement, contrasting it to the 'war of all against all' that was the alternative. In fact, he argued for what we would see as small, authoritarian government.

Hobbes wrote Leviathan in a time when religious confession was a divisive and defining issue, and the question of the nature and legitimacy of Kingship was seen through that lens. But more recent research shows a continuing relevance: the existence of a government has a significant positive effect on the way social and economic interactions are likely to work out. It means there is likely to be more advantage in cooperation (and in sticking to our promises) than in competition, or aggression. It resolves the so-called Prisoners' Dilemma problem consistently, the right way. It also means that the costs of security (Police, economic rules, and the justice system) are efficiently allocated, and consume a smaller proportion of output than would otherwise be the case.

Hobbes' rather pessimistic insights were often unpopular. Others, like Rousseau and others, argued that man was inherently good. But, as Ian Morris sets out in a recent book (*War: what is it good for?*), the historical record is on Hobbes' side: order (via the existence of the state) leads to prosperity (even if the state itself is, at least initially, the product of conflict).

Contrast the internet. Privacy here comes without anything like as much supporting security as the rest of our lives. Some regulatory safeguards cross over from the analogue world – especially consumer rules, company law, and a lot of contract law. In one area, child protection, the state makes particular efforts. But mostly, we are comparatively on our own.

The result is that we live an increasing part of our lives in an environment which is much less well governed than we are used to. Trust online is an issue for adults as well as children. Identity online is shaky too, because the non-rivalrous nature of data (it can be in two places at the same time) combines with a ready market in stolen data to create a really significant cyber crime economy. It also means that we support a private security industry on the internet which is probably larger than it needs to be, imposing a significant efficiency cost.

My fundamental contention is that privacy without some level of security is likely to be sub-optimal for many people, and we should think about that. The sort of security I am referring to is not censorship, nor really anything to do with content. Rather, it is the kind of framework of law and order, supporting our ability to go about a lawful business, which we have built up so painstakingly and painfully in the analogue world, really since civilisations began. I am not advocating any particular solution; I am arguing that it's an important issue we should think about, and caricaturing privacy and security as conflicting values is not an adequate response.

Which leads to the question of trust. If lives online were to benefit from any kind of framework, who can we trust to make the rules, and to see them through?

Here, there are two kinds of claim that I consider are just wrong. The first is the Libertarian fantasy: if only governments were kept out of things everything would be ok, and human potential better fulfilled. It's a fantasy, because the unspoken premise is that human nature can be made better, just because we all agree to be better, or because we hope the unfettered internet will allow us to return to a state of Grace. This (as the contemporary philosopher John Gray points out) is a Redemption argument (ie a Christian one) taken out of context.

Human nature is unlikely to be changed by the internet, but its expression can be made more efficient – so more good stuff can be done, and more bad stuff too.

The second is the paranoid fantasy: if there is any level of public oversight of the internet, the State will be looking at everything I do. This is, I fear, the teenage boy nightmare: will Mum find out what I get up to in my room?

There's an easy, but slightly emotionally unpalatable answer. Remember the scale of the internet? Over 200 million emails every minute? No state has the resources to really monitor that, even the most fearful. This is my earlier point about scale. Chairman Mao once said the "The guerrilla moves through the people like a fish through water". The State wants to catch fish: the handful of people doing really bad things (not just thinking mildly progressive thoughts). They know who they are. If you don't know that you're a fish, then you're water, and we don't care about you.

In a country like New Zealand, we just don't know how many emails are sent and received. We do know (from published industry data) it'd take 130000 people to just listen to New Zealanders' phone calls and read their SMS messages (not doing anything with those calls and messages, just listening and reading). If it was a task given to GCSB, our salary budget alone would need to be 100 times today's total budget. Even if you don't want to trust the State, you need to know that its arms are short and the internet is big.

And so back to Hobbes (who also argued in favour of censorship, explicitly not an argument I accept or advance). There seem to me to be three points we need to consider:-

Firstly, there's a public good in being able to transact our business and live our lives safely on the internet. Safely means with a reasonable prospect of being able to avoid loss of our identities, or becoming victims of crime, and being confident that bigger threats (like organised crime) are being managed. That means there needs to be rule enforcement. I would also argue that (as much as possible) it should mirror the standards we would expect in the analogue and commercial worlds.

Secondly, that needs to be balanced by some assurance that providing for our safety is not a substitute for censorship or onerous intrusion (the thin end of the wedge argument). Democracies are actually quite good at this sort of balance, I would argue, but the system needs a bit of tension between its components to provide assurance. It's always right to ask if that balance is right, and to expect it to change over time, and with experience.

This is another aspect of the trust point. The internet is a competitive place, where small, transient knowledge advantages can yield great gains: the whole cyber security industry monetises our need to hire folk whose knowledge is only slightly behind the inventiveness of the bad guys. It really is a version of Hobbes' battle of all against all. Anything published is lost, in advantage terms. That's why in GCSB's world we keep our methods secret, and don't admit where the limits of our knowledge lie. If we're open, we're lost, and all the effort we make and you pay for is as nought.

Which is a dilemma: our society sees transparency as the basis for trust; yet full transparency protecting society on the internet is fatal to the enterprise. So, how do we get enough trust to make any sort of government presence on the internet legitimate, without undermining the whole thing? That becomes a debate about oversight systems. The point I'd make is that we have to accept that the question is both valid, and capable of a solution.

I should also say that trust and transparency are opposite values, not complements. The direct solution to a trust deficit is not transparency, but to fix whatever the root cause might be of the trust deficit. Transparency is helpful for accountability, contestability of policy and execution, and governance generally. All of which may well make public institutions more effective, and comprehensible, and customer friendly, and therefore more trusted. But it's not a direct remedy if trust is the primary issue.

And so (to quote Lenin), what is to be done? And by whom? For most of us, and our businesses and our communities, the internet is probably under-governed. The Hobbesian bargain we have worked out on the streets needs to be matched online. Relying on luck (that I won't be a victim) and on private providers of security is both risky and inefficient.

What if we do nothing? After all, the internet is big; bigger than any state can singly expect to manage. There's little consensus in our society about more involvement for the state on the internet, and larger countries are just as conflicted. And the prospect of any sort of global consensus must seem fanciful, as I shall explain shortly.

But the internet changes things, and it leads to evolution and experimentation. Interestingly, very big private companies are able to operate at effective scale online, and (in return for monetising our data) provide a measure of order, if not security. History shows that early states worked this way too – they've been called fixed bandits, taxing to pay for the armies they used to keep order, learning not to tax too much along the way. We may be at the start of the genesis of a new kind of Hobbesian bargain, as corporate states (ones we currently join voluntarily) compete with each other and with territorial ones. Are we happy with that? It's a debate we might want to have.

And finally, the day job. All of this is relevant to GCSB's work too, but not as you might think. Events over the past two years, both in New Zealand and globally as a result of the disclosures attributed to Edward Snowden, have raised important questions about

the expectations citizens in a liberal democracy should have of organisations like GCSB. They deserve an answer.

The question is narrower than the broad privacy and security question I have been considering so far. But it is an important question, not least because it begins to unpack the impact the Internet is starting to have on relations between states.

Just as the Internet means we need to rethink the Hobbesian bargain between the state and people, the Internet also forces us to review the nature of international relations–it calls into question some of the fundamental principles of sovereignty, territoriality, and coherence of national identity reflected in the principles of the Treaty of Westphalia [1648]. That is a whole topic in itself (another time, perhaps); for the remainder of my time today, I want to reflect on the narrow question what should the government and people of a liberal democracy expect from a 21st century cryptologic organisation like GCSB?

My analysis is based on two assumptions:-

Firstly, cyberspace has started to become a domain in which nations struggle for advantage and enhance their national power, and into which military operations have already started to extend. This encompasses traditional espionage, as well as a spectrum of operations from subtle influencing through to actual destruction of information or systems. Cryptologic organisations are ideally placed to undertake these activities (and to defend against them). Indeed, if other types of organisation take them on, they will become cryptologic organisations themselves.

Secondly, intelligence collection and self-protection in the information domain will continue to be a proper activity for liberal democracies protecting their own people, advancing their national interests, and fighting wars.

Against those assumptions, government, Parliament and people might therefore expect to have a cryptologic organisation that looks something like this:

Firstly, it is effective at defending the government's own critical information against sophisticated cyber espionage, and against disruption from any source.

Secondly, it is highly effective at conducting sophisticated intelligence activities against any legitimate target, no matter how hard.

Thirdly, it is a potent and effective contributor to military capability.

Given the difficulty of all this, it follows that a cryptologic organisation in the Internet age must be extraordinarily proficient in all aspects of its tradecraft. It must also operate with legality and propriety. And it must clearly repay the cost of ownership in two important ways: the national advantages conferred by spying must outweigh the disadvantages of being known to undertake that spying. And those advantages should be delivered at considerably less expense than any of the alternatives, such as losing

valuable Intellectual Property, having to deploy armed forces unnecessarily, or losing the lives of armed forces personnel in conflict unnecessarily.

And it follows from that, that in order to be effective a cryptologic organisation needs to operate in a supportive environment. It needs government customers able and prepared to use intelligence as an element of statecraft. It needs a legislative and compliance regime that achieves the right balance between the needs of the state for foreign intelligence and Internet defence capability on one hand, and the constraints on the power of the state against its own people on the other. And government and the broader community need to be realistic enough to accept the need for these capabilities, and sophisticated enough to manage the resulting moral ambiguities. A successful cryptologic organisation will recognise the need for this sort of operating environment, and work with its external constituency to help achieve it.

There is much more that could be said on any of the topics I have covered today. There are no right answers, a democratic society should work its way through these issues to reach an answer which reflects its values and its interests. But I do believe that the Internet is politically and socially as well as technically and economically disruptive. Countries will need to come to terms with it, and that process will change internal and external settlements, and challenge their understanding on the way the world is ordered. New Zealand is no exception to the challenge; paradoxically, events over the last two years may give us something of edge in that debate. I hope we take up the challenge.

ENDS