Law, Technology, and Cybersecurity

ISACA – Ottawa Valley Chapter & AEA Ottawa-Gatineau Chapter

Ottawa, ONT – June 2018

Paul Rosenzweig

Good morning. Thank You very much for the invitation to speak with you today. I especially want to thank Robert Weismann for reaching out and asking me to join you. I also want to apologize – as a lecturer I know that as IT professionals you love power point, but I have none for you. You'll have to just listen I guess.

My topic today is the coming cyber security liability revolution. And let me start by quoting something a colleague said just the other day: "The cyber security of the Internet of things is a national security issue. It is long past time for the law to impose liability on those write insecure code."

Those aren't my words, they are from a speech given a last year by Melissa Hathaway to a small group of cybersecurity policy experts in Washington DC. You'll be forgiven if you don't know who Melissa is (especially since she is an American), but a quick Google search will tell you that she worked as a senior cyber policy analyst in the White House for both President Bush and President Obama. When you want to know what official Washington is thinking, Melissa is a good leading indicator.

Think of what the implications are of what she said. For you, in this room, the security of the systems you are deploying in the information systems you are designing is a matter of good programming, good engineering and good enterprise architecture. You are worried about systems security; about effectiveness; about cost; and about efficiency.

In Washington and, no doubt also here in the halls of government in Ottawa, we are worried about war and national security. And the plan is to enlist you as soldiers in the fight, whether you want to or not.

I hope I have your attention. To be fair, I'm being a bit overly apocalyptic for dramatic effect. You don't really need to worry about being drafted anytime soon.

But I do think that it is long past time for anyone in the cybersecurity industry, especially those in fields like the information systems design and architecture that bear on the life and death of consumers, to realize that the easy times of code development and system design are coming to an end. You will be enlisted in fixing that problem – not as draftees working for the government, but as private sector actors who, by hook or by crook, will be obliged to modify your business practices in light of the perceived threat.

Please note that, like many in Washington, I share the general assessment of the threat environment. Sitting here in Ottawa you may disagree and think that the threat is wildly overstated. But I don't think that matters because I doubt you can

convince Washington or the Canadian government down the street that it is wrong – so expect it to act.

And since governments are governments, the way they will act is through what they do best – law. In my judgment, we are on the cusp of a liability revolution that will change the way you do business. My goal in this talk is to tell you a little bit about what I think that revolution will look like – all with the usual caveats that predicting the future is not a terribly easy business, so don't bet too much money on me!

Let me first set the scene by talking about where we are today, as a matter of law, in the area of cybersecurity liability.

First, a caveat. I am an American lawyer talking to information systems designers and enterprise architects in Canada. I do not pretend to know the details of Canadian law. But I do know enough about our shared common law tradition to be reasonably confident that what I'm suggesting is the course of the future is a realistic prospect in both our countries. And, of course, to the extent that you continue to economically interact with the United States – despite the best efforts of my current President to sever those ties – the law and policy of the United States will be relevant to you.

Second, let me offer a couple of definitions.

When I talk about liability for information security failures I am talking about responsibility for "bad code" and bad system design. What I mean by bad code is code which can be manipulated to perform in a way contrary to that anticipated by its writers. This can range from the common, almost trite example like a buffer overflow or SQL injection, to sophisticated, unrecognized errors like the one that the WannaCry ransomware exploited which allowed remote code execution in Windows Vista and Windows 7 operating systems. I am also talking about bad system design – that is enterprise architecture that does not provide for adequate controls or audits of system performance and capabilities.

Importantly, then, I am not talking about broader categories like errors in autonomous decision making or privacy concerns regarding large collections of data. Those are big and significant issues and some of what I am saying will apply to them as well … but for now they aren't my topic.

And when I talk about "liability" though I mostly mean legal civil liability for damages, I also want to talk about it in the broader context of responsibility and obligation – which may mean things like regulation, for example, or social liability.

So, let me start with a summary of the law today – and I apologize in advance to any lawyers in the room who will recognize that I am summarizing in three minutes a field of law that requires three years of study.

Today, we live in a world where for the most part liability for cybersecurity failures is non-existent. The standard shrink-wrap contract that comes with any

software package disclaims liability and for the most part courts in both countries have upheld those contracts. Likewise systems design contracts often disclaim liability. Responsibility (and liability) lies with the end user – say someone like Target who uses a software package and loses the personally identifiable information of its customers in a data breach. They pay damages, but the software writer who drafted the buggy code they used and the enterprise architect who allowed systems access for an HVAC contractor do not. Indeed, I am aware of only two American cases to date where liability beyond data breach has been imposed – and neither of those involved the software developers or enterprise architects.

In short, the cybersecurity liability system is, today, almost exactly where the rules for automotive liability were in the mid-1960s in America. Bear with me because even though the story is about cars in America and not information systems in Canada, it seems to me pretty relevant.

Back in the 1916, Justice Benjamin Cardozo first advanced the idea that anyone who manufactures an "inherently dangerous" product was responsible for making sure that it was safe. Here is how it put it in a famous case involving Buick motor company, back at a time when cars were brand new inventions and might be thought of as far more dangerous than the horse and buggy:

> "If the nature of a thing is such that it is reasonably certain to place life and limb in peril when negligently made, it is then a thing of danger. . . . If to the element of danger there is added knowledge

> that the thing will be used by persons other than the purchaser, and used without new tests, then, irrespective of contract, the manufacturer of this thing of danger is under a duty to make it carefully."

So here we have the first idea – that if you aren't careful when you make a dangerous consumer good – that is, in the language of the law, if you are "negligent" in how you build something – you might be liable for the consequences. Let that sink in for a minute and think about what it means to be careful in, say, designing the architecture for an enterprise involving, say, the transmission of electricity, which would, I think, be in Cardozo's view, a thing of danger.

The auto industry's response to that decision was instructive – it started putting disclaimers of liability into its contracts. For a while those disclaimers had legal effect – and American courts told car buyers that the rule was "buyer beware" or "caveat emptor." In other words, they got the car that they got. The contracts were the historical equivalent of today's shrink-wrap or click-through contracts and liability disclaimers.

Then, in the 1960s came the product liability revolution. The idea was that manufacturers of products were liable for any defect in their products. Here is how Justice Roger Traynor of the California Supreme Court put it in a famous case called *Greenman*:

> "To establish the manufacturer's liability, it was sufficient that plaintiff proved he was injured while using the [product] in a way it was intended to be used as a result of a defect in the design and manufacture which the plaintiff was not aware that made the [product] unsafe for its intended use."

Note especially that this change in law does not speak about taking care, or about negligence. It speaks of design defects and was the precursor of what we have come to think of as strict product liability – if you design it or build it and it malfunctions, you are responsible. Period. Put another way, the courts said that contracts that disclaimed liability were unconscionable – what we call contracts of adhesion where one side has all the power – and could not be enforced. To continue the metaphor, imagine that disclaimers, and click-through and shrink-wrap limitations were outlawed.

The change was driven by many sociological factors – the growing industrialization of North America and its dependence on the new technology; a sense of unfairness and injustice that the small guy should pay; and also some of the safety issues that were beginning to plague the auto industry and that led to so many of the safety standards that are now deeply embedded in that industry.

That description should sound awfully familiar. Today, in my view, we are one major disaster away from product liability for bad code or bad systems architecture in consumer products. I don't know if that disaster will be in the medical field (say a massive outage of heart monitors) or in, say, the automotive

field. But imagine if an autonomous bus were to drive off the road killing 30 school children. Or imagine if ransomware akin to WannaCry exploited a vulnerability in a system of sewage treatment facilities, making hundreds of them inert. You really don't need me to tell you what the public uproar would be like. If you want just a taste, consider this headline from just two weeks ago "It's not just the elections: Russia hacked the electric grid too." And official Washington and Ottawa =will= think of it as a national security problem.

So … what will this look like? Here's the short answer: Where life and death are at issue, responsibility and liability cannot be far behind. This will have impacts on the insurance industry and may result in regulatory intervention if the industry is not proactive in its approach. In addition to best practices for cybersecurity and standards of software development, this will also require the development of, as yet non-existent, audit and grading mechanisms to support insurance risk rating. Let's take each of these in turn.

First, there is the question of liability. As I've said there are two potential types of liability the might be adopted – one is a form of strict liability for any defect in a product and the other is a negligence/reasonable care standard. Under a strict liability standard manufacturers of any driverless car might be liable for any defect in the code that is executed to make their system function. Since, as we both know, it is impossible to write bugless code – that is code without a defect in it – and since it is equally impossible to design an information system that is completely secure -- strict liability would amount to pervasive, persistent, and absolute liability.

That way, I think, lies disaster. The costs to manufacturers, designers and developers would be so high and so persistent that innovation in this domain would quickly grind to a halt. This would, of course, be true of almost all other products, like medical devices, or other common goods, like wastewater treatment where safety is at risk. With the wrong liability rule we could, as a society, effectively stifle technological development – at great cost to our overall well-being.

Because the costs to innovation would be so severe, I feel fairly comfortable in predicting that the liability standard that will eventually be developed by the courts (or possibly by legislation) is one that will focus on reasonableness and best efforts. In other words, the expectation will be a requirement that those designing information security audit and control system do their best to make their systems secure against cyber intrusions, but there will not be an expectation of perfection. As to what that reasonableness standard will mean in practice – if you'll permit me I will save that for the end of this talk as a series of recommendations for your consideration.

Assuming that we have liability for poor system design, that means we are going to need insurance. Perhaps as IT specialists and engineers that isn't obvious to you, but as lawyers the need for insurance after liability is as obvious as the premise that day follows night.

Liability, after all, means money. It means paying money to those who are injured by the alleged flaw in the design of the operating system for a chemical

production facility. It means paying the family of those 30 kids on the bus that went over the cliff. And nobody wants to bear that cost by themselves. So that means that insurance companies are going to be asked to write insurance policies to insure against the risk that enterprises will have to make liability payments for cybersecurity flaws in their systems.

Right now, most general liability or products liability policies – that is the non-cyber specific policies – have exclusions that preclude coverage for cyber incidents.

And there are precious few cyber-specific products. As I stand here today, I am aware of only one company that writes cybersecurity liability insurance policies that cover damage or injury to third parties, like the consumers injured by an outage at the local water treatment facility. In short, the idea of insurance for cyber systems is in its barest infancy.

To be honest, as insurance products I don't think much of them – there value has never been tested because it seems as though no claims have been paid out. Right now, the premiums being charged and the coverage being offered seem more like guesswork to me than like sound actuarial science.

One other point bears noting about insurance – everyone needs it. Everyone. When a wastewater treatment facility or a manufacturing plant gets sued, thier insurers will likely make claims against any third-party software or systems designers (that's most of you in the room) for their part in the accident or

incident. So you too will need insurance. There is a market for that, but it is limited, and they are also just making up numbers for premium without a database. A widely-based, mature insurance market will require meaningful risk rating tools.

And that, really, is the nub of the problem for insurers. When they write insurance policies insurers engage in a practice known as rating a risk. That is, they evaluate a risk and assign it a numerical score that is a reflection of the probability of the risk occurring and their best estimate of how much the damages will be (that is how much they will have to pay out) if the risk does occur.

A perfectly good example of this is the insurance against hurricane damage in the Southeastern United States. Insurers, and climatologists, have nearly 150 years of data on the frequency and intensity of hurricanes that strike America's shores. As an aside (and this is really just a fun fact that I hope you will enjoy) did you know that no hurricane has ever made landfall in Georgia in recorded US history. Florida yes. South Carolina of course. But never Georgia.

The insurers also have robust, extensive data sets about the value of construction in potentially affected areas and about the nature of that construction and its resistance to wind and water damage. With that history they can make a pretty good estimate of how much they are going to have to pay out for hurricane damage – on a short time scale like this year, the estimate may be inaccurate, but on a longer time scale they can have great confidence in their risk ratings.

And they worry about the fidelity of their data. Insurers worry about global climate change not so much for ecological reasons as because of their concern that it might change their hurricane models.

The more data you have, the better your predictive models. To take another example, American insurers can predict with stunning accuracy how much they are going to pay out on automotive accident claims this year. Their frequency; the repetitiveness of their causes; and the systematic distribution of damage numbers makes it actuarially a pretty precise calculation.

Of course, that is the exact opposite of where we are today for cybersecurity risks. We have very little good data on frequency, on causality, or on potential damages. If I ask you to assess the comparative cybersecurity of one product over another, or one enterprise over another, or make a judgement of security on some absolute scale, you, rightly respond that you don't have a good metric for measuring it. We don't even have a good public repository for the existence of cyber incidents that have actually occurred – in the US a Department of Homeland Security initiative that began in 2014 is still in the "proof of concept" stage (at least as of May). And I know of no comparable initiative here in Canada.

So we are in a conundrum. Liability means there will be a demand for insurance. But insurance for cybersecurity cannot be written if we can't rate the risk. I am therefore also quite confident in predicting that a small industry will soon grow up in rating cybersecurity risks.

Some of you may be familiar with the hacker who goes by the handle Mudge. His real name is Pieter Zatko and he is a legend in the hacker world. He started a non-profit a couple of years ago with one goal – to develop a way of rating or grading cybersecurity risks. That effort is still a work in progress – late last year he announced the development of an algorithm to measure software security.

The idea is that one day, perhaps relatively soon, programs developed in this room or systems designed by you could be sent to Mudge's lab for testing and come back with a seal of approval, in much the same way that electrical outlets get an Underwriters Laboratory stamp before they are offered in the market place. Underwriters was a private sector group set up by the insurance industry, but we can imagine this as a government activity – just recently Germany began to consider setting up a government office for auditing algorithms and IT control and audit systems for efficacy and security.

Candidly, I have some skepticism that a testing methodology will work – and I'll outline my proposed alternative in a second. But I am certain of one thing – some form of grading for cybersecurity risk is essential if a functioning insurance market is going to be developed.

Another small piece of the puzzle is going to be the development of a security audit capability. After all, rating risk is only useful if the companies that are being insured actually follow through on the ground implementing the security promises they have made. In every other insurance context the risk profile of an insured gets routinely audited for compliance. Look for big audit firms like KPMG

to develop cybersecurity and systems audit capabilities – and watch when, some day, they are in your space looking over your shoulders.

One thing I can certainly offer, as a cautionary note, is that the IT systems design industry should work hard to support the development of a liability standard and a functioning insurance market. Not for its own sake, though that will be good enough reason itself, but rather because the alternative is even worse.

Sometimes when I speak with clients about this they have a "put your head in the sand" sort of attitude. They hope that they can forestall security liability for the foreseeable future. I think that is highly unlikely, and that if industry groups don't proactively set themselves the goal of helping in the development of cybersecurity liability standards then Washington or Ottawa – we're at war remember -- will intervene and set the standards for them.

We've already seen some rumblings in Washington and Ottawa in that direction. Here, in Canada, I understand you are on the cusp of mandatory data breach notification and record keeping requirements – rules that will go directly to the problem of measuring security.

In the US, in the wake of the Mirai botnet attacks, which took advantage of insecurity in small consumer devices like thermostats and DVRs, the Federal Trade Commission announced that it would look into the need for security regulations. And the National Highway Transportation Safety Administration has published a report on "best practices" in automotive cybersecurity. You ignore these sorts of

government initiatives at your own peril – doing so runs the very real risk that you will be drafted into the debate against your will. And today's "best practices" will become regulatory or judicial mandates – mandates that are rigid and change on, say, a three year cycle. Think how that will effect system design and code development.

So what, in the end, does a good system of system security look like – not from a technical perspective but from the perspective a lawyer and policy maker like me who is contemplating the question from a broader national perspective? Here I want to offer a few process oriented suggestions, drawn in part from the work of a non-profit group of security professionals known as I Am The Cavalry, with whom I've been working. To my mind, industry should have a voluntary set of standards and a self-assessment model for the cybersecurity of its systems and product that asks, broadly speaking the following questions:

- Can you explain to policy makers and insurers how it is that you design and develop your systems and software products? Do you do adversarial testing programs for your products and for critical components of your supply chain? If not, why not?
- Are you open to third party research that finds flaws in your systems? Too often developers are resistant to outside scrutiny. If you have a good-faith report of a problem, how do you respond to it?
- What are the forensics of your systems? Do they provide tamper evident, forensically-sound logging and evidence capture to facilitate breach and safety investigations?

- Are your systems capable of being securely updated in a prompt and agile way? I gave advice to one client (not, I hasten to add, one in this room) who had a system that was, effectively unpatchable. My own opinion was then, and remains today, that such a design is almost per se negligent.
- Can you point to some neutral, generally accepted set of standards that define a reasonable standard of care? In other words, if some well-recognized organization has, over time, come to define what a good process is, do you follow that process? Here, I am thinking of standards like the NIST Cybersecurity Framework and more focused standards that are embedded inside it. Ask yourself questions like:
    - Does your enterprise implement Cobit 5? If not, can you explain why not?
    - What portions of the CIS Critical Systems Controls are implemented and which are not?
    - Which ISA or ISO standards do you follow? And so on.
- Finally, how are your cyber systems incorporated in the physical products with which they interact (if they have a kinetic component)? Is there, for example, a physical and logical isolation that separate critical systems from non-critical systems? Any design that relies on a unitary system of control will I think no longer be acceptable.

So. Where does all that leave us. My take aways for are you relatively few, but fairly important.

First, I see no prospect in the long run for avoiding liability for poorly designed enterprises or for deploying insecure code. The only question is whether it will be absolute liability or liability based on a reasonable care/negligence standard. You may not have thought about it before, but now that you have, I am 100% sure that you will prefer the reasonable care idea – since the requirement to design perfect systems that would come with absolute liability is simply a requirement you can't meet.

Second, liability is all about money. And that means that inevitably we will see the development of an insurance industry with requirements to evaluate or rate security risk. One of the critical problems for you is figuring out how to do that, because if you can't the natural tendency will be to impose some form of absolute liability.

Third, you owe yourself the obligation of trying to get ahead of the curve. If you don't help to design the liability system now, someone will design it for you and I suspect you will like it a lot less than if you had built it yourselves.

And, finally, do not under any circumstances, make the mistake of thinking that this is a technical problem. It is not. It is a social and economic issue of the highest order, both for those who design information systems and for those who implement those systems in any enterprise that can have real-world impact. From the perspective of governments this really is a national security problem. And if you make the category mistake of thinking otherwise … well .. then you run the very real risk of being drafted into a conflict against your will.

On that rather grim note, I want to thank you for your attention. With the time I have left, I would be happy to take a few questions.