Institute of Corporate Directors

Podcast Transcript: Be It Resolved that Boards Need to Fundamentally Re-think Cyber Risk in the Age of AI.

Rahul Bhardwaj (00:04): Welcome to Be It Resolved, the podcast where bold ideas meet courageous leadership. I'm Rahul Bhardwaj, president and CEO of the Institute of Corporate Directors in Canada. In each episode, I speak with experts to delve into pressing issues, impacting directors and decision making in the boardroom.

My guest today is Stephen Burns, partner at the law firm Bennett Jones, LLP and co-head of the ITB Practice. His work focuses on technology law including data governance, cybersecurity, privacy, and intellectual property. He regularly appears before Canada's Information and Privacy Commissioners and publishes on key topics like AI and cyber risk.

Now, today, Stephen is going to share his insights on why cybersecurity must be seen as a strategic priority for directors, and along the way, we're gonna learn about the Watermelon Effect. Looking forward to that now.

Today's resolution: be it resolved that boards need to fundamentally rethink cyber risk in the age of AI.

Welcome Stephen. It's great to have you with us today.

Stephen Burns (01:14): Thank you. It's great to be here.

Rahul Bhardwaj (01:16): So here we are talking about the shift in cyber risks. For those who might not be as well schooled in this, when we're talking about cyber risk, what exactly do we mean, and particularly in the boardroom?

Stephen Burns (01:28): It's a great place to start, and I think it's very important when we start talking about cyber risk. We can easily go down the rabbit hole of all the different ways that an organization can be attacked from a cyber point of view, all the different tools and methodologies. What I think for the purposes of the board conversation we're going to have here, it's better to stand back and think of it in terms of what are we trying to protect; which is we're trying to protect our information, what the organization collects, uses, and learns so that it's able to then drive its competitive advantage and deliver its goods and services to market. That information is often unique to the organization and very valuable.

The second thing we're trying to protect is the operational assets of the organization, where there's some electronic or computer element to it that runs that operation. So, is it the control systems running a pipeline? Is it the control system running a particular machine that is punching out a particular product?

When you start thinking in terms of what we're trying to protect is the information of the organization and its operational technology, then you can start looking at, well, what are the factors around how that is protected? Do we have adequate physical procedural and technical controls in place, and do we have - I'll add a fourth - adequate management controls as the board to be comfortable that the physical, technical and procedural protections are being considered in light of the risk of the organization of that particular information or operational control asset and the potential that the changing technology will have an impact on it? I think that creates the right sort of structure for the conversation going forward.

From there, you're able to then look at what is the design of that particular system, what is the value of that information and how is it that we, the organization, are going to change that information or operational system so that we can take advantage of the new artificial intelligence capabilities, the new tools that let us better understand our data, better control our operations and ultimately drive value out to our organization.

Rahul Bhardwaj (03:50): So, today's resolution is about potentially a change in the face of AI, but this is not to say the directors weren't already doing a lot of good work in this area. Let's touch on that a little bit. What's evolved to give directors comfort that they've got oversight in this space?

Stephen Burns (04:06): So, a lot of work has gone on to educate boards and bring up to the board-level the skills and information they need to ask the right questions. They want to ask, “Where is my data? Is it secure? Can I rely on it? How are my operations being secured? Is there a material risk that a particular piece of my operation could be turned off or have a catastrophic effect if I'm unable to use it?” That's really been driven based on the architecture of how those information and operational systems have been put together.

Typically, we've been looking at protecting systems by not having them connected to anything, having an air gap, as it's sometimes called and having it just as a standalone system. Then we've had systems where they haven't been very interconnected. So, my data on human resources was not in the same system as my financial data or in my customer-facing or client data. That's all started to change as we've evolved quickly here into starting to understand the value of data and bringing it together into large data lakes so that we can start analyzing and understanding the totality of our operations, not individual systems.

By bringing those systems together, bringing that data together, we've lost the protection of having small discreet systems that might not have been interesting and instead have started creating very large pools of data which attract the interest of either the criminal element or the activist element or potentially a state-based actor. As we look to be able to do more with our own stuff, we've made our own stuff more valuable and thus of more interest to the threat actor.

Rahul Bhardwaj (05:47): So, it sounds like a double-edged sword, somewhat. And before we dive into the AI, under the regime that people had put together who had responsibility for oversight of cyber risk, and I mean specifically at the board level, was this the audit committee largely?

Stephen Burns (06:04): Ultimately, it's always the board and then they've tended to delegate down either to the audit committee or actually just handle it directly through their executive. A lot of organizations have let their CEO hire a CIO - chief information officer - and then a few years ago, we started seeing chief information security officers be hired. So, they've put a member on the executive team who's responsible for understanding where the data is, understanding where the threat is, understanding where they are most likely to be at risk, and then putting their resources in the right place.

And I think it's important, just as a sort of a side comment, cybersecurity is expensive and you don't try and secure everything at the same standard. There's a risk assessment that you go through to understand, you know, where is my most valuable information, where is my most important operations and how are they secured so that you're investing money wisely?

That's typically where you've seen the board interested is make sure that I've got a strategy that I want to be secure, but I wanna do so in an economically reasonable manner and I want to have trusted people in the organization responsible for that task. That's sort of where most organizations were, trying to do a good job but not overinvest or overspend on a risk that in many ways was thought of as being theoretical. Cyber attacks weren't very common or certainly were not the same volume as they were say 15 years ago against today.

Rahul Bhardwaj (07:38): So, we've spoken about this before, that boards have got the duty to act in the best interest of the corporation. They've got oversight of culture, strategy and risk. And when we start to look at data as the crown jewels of many corporations, they're taking a risk lens to this, and something tells me that your view that the change in a new AI world is shifting from that risk perspective to a little bit more of a strategic lens. Let's go there. Can you explain that a little bit?

Stephen Burns (08:08): So artificial intelligence has been around in various forms for decades and many organizations have used various forms of machine learning to be able to do credit card adjudication, determining whether a loan should be granted. Many automated processes have been there using information wisely.

What's happened and why we're seeing a sudden increase in threat and the sudden increase in opportunity is that we've broken a number of barriers. One of those barriers, a couple of years ago, was that we broke the language barrier. So a large language model can now understand language better than a human can, not by much yet, but still enough that it lets them start creating a whole new range of tools that we haven't seen previously. Tools that open up the ability to analyze very large pools of data and create some very interesting insights.

Artificial intelligence isn't magic, it's not this sort of single tool that does everything. What it is is a whole series of different tools that have come together under one title or grouping that lets you look at your data and start looking for deep insights; lets you look at your operations and start looking at how can I drive out costs? It starts letting you look at how do I best understand how the business is operating?

That process is driving or these tools are enabling us to look at data in ways we've not looked at it previously. That process is driving us in the cyber risk conversation. We're having to pool lots of data together to move a lot of information into bigger and bigger lakes or pools, which lets us then do the sort of broader, deeper analysis and see what we can find and understand. But it also means then that we've got a bigger, more evident crown jewel for someone to go after.

Rahul Bhardwaj (10:02): You said before that there were three tsunamis, I think it's the word you use, that you were thinking about describing the impact of AI on the cyber risk environment.

Stephen Burns (10:14): There are really three macroeconomic impacts that each board needs to understand. When we think about technology, we really are talking about business process automation. We're really talking about how do I undertake the tasks that I need to do to deliver what my organization does in the marketplace? How do I use my people and the technology to undertake each of the things that we need to do to deliver to our customers what it is they want to buy from us?

That suggests that we know both how our business operates and what our technology can do for us. Those two elements are being impacted by these three tsunamis that we're seeing.

The first one is the silver tsunami, which is that the Baby Boomers are starting to retire, and we're seeing large amounts of institutional knowledge exiting the organization due to retirement. And what we're finding is organizations typically don't have a really strong understanding of how they deliver on the activities they need to deliver to market their goods and services. They know that they do it, but they're not really certain how they do it, and a bunch of that knowledge is starting to go out the door.

The second thing that you have is that technology is rapidly changing. So where previously we've used computers on the computational basis, they've done two plus two equals four. They haven't really been able to handle huge amounts of data and come out with new insights. What we have with these new artificial intelligence models is the ability to take massive amounts of data and begin to predict outcomes in a way we've never been able to do before. So, we're able to significantly improve the efficiency and the quality of the work and the undertakings that we're doing.

So, organizations are very excited. They want to use these new tools because they know they can drive improvement in quality and improvement in efficiencies or a reduction in cost and they can be more competitive in the marketplace. That's driving the knowledge we need to figure out how to use those tools is retiring exiting the organization and the tools are rapidly changing so that we have these opportunities to use them in new ways, creating a real void. And that's the third tsunami.

The third tsunami is this transformation tsunami that organizations are being forced either because of the retirements or because of the new technology to change how they're doing business. We're in a time of significant change. That is this transformational tsunami which is creating what I think is actually the real risk to the board at the time, or at this time.

We know that we need to do more and we need to use this technology more, but we don't know how we're going to do more or how we're gonna successfully change. And unfortunately on the transformation side you have very low success rate.

There are some good studies that show that while day-to-day change, organizations are quite good at transformational change, actually restructuring and being able to take a strategy. [Organizations] think about how they're going to deliver on that strategy and then transform the organization with these new tools to meet that strategy. You're talking about success rates that are under 20%, and then the 80% that are unsuccessful, you're talking about maybe 30-40% of that. So, half of that is having those projects go forward, but they're doing it at two x to five x cost.

That's this sort of pressure cooker that the boards are under. They are looking out at new tools, big piles of information they wanna do new things with. They know that the understanding of how their business operates is starting to retire, and so they are, and their competitors are going through the same process. They don't know what to do.

They don't have the skills and the focus at the board-level to yet really know what the questions they should be asking of their management team of how are we handling the lost knowledge of the silver tsunami, how are we handling all these new opportunities coming from the technology and how are we able as board members to know that we are ready in making our best possible investments in the transformations that we're being told we need to do to be able to handle these changes. Still, we don't really know how we're going to measure our success along that journey?

Rahul Bhardwaj (14:42): So, this is a lot to unpack, but we've got these three tsunamis that are changing everything and from a board perspective, it's valuable, it's alluring, competition's doing it, but it's really raising the risk as well. So, let's just focus for a second on exactly how is AI in your mind reshaping the cyber threats?

Stephen Burns (15:04): So, let's go back to where we were and where we are today.

Where we were is a series of separate systems secured individually and often separated so that if one system was attacked and compromised, the other systems would not be at risk. That's an old model based on servers under our control, data centers under our control, and systems that are under our control internally; that's moved. Cyber is expensive, and the protection of systems is expensive. It's also expensive to maintain those different systems.

There's been this big movement to the cloud. Lots of organizations have looked to bring and simplify their systems into cloud-based services and look to outsource what used to be in-house skills and understanding to their outsourced providers, to their cloud providers. And that was a natural evolution. We can take advantage of scale by going and putting our stuff in the same cloud as other people were putting their stuff in the same cloud and then we could all cover the cost of securing it.

What's happened is the same tools that are creating opportunities for the organization to go and do new things with data and do better understandings of their business and better delivery of their services is exactly the same tools that are enabling the threat actors to better understand how to get after the data they want, better understand how to get around the securities around it, better enable all the social engineering and attacks they need to get to a password, they're able to do it.

We're seeing the governments and the various security organizations in the world putting their hands up, identifying that in the right hands, the new tools enable a whole range of attacks that have never been done before. We all saw the very successful use of an artificial intelligence-created avatar to persuade a person and organization to start making wire transfers because they believed that the individual instructing them was, in fact, the executive. So they're able to marry voice and video and information to convince an individual in the organization to wire money out to the cyber criminals.

That was impossible to do five years ago, and now it is something that the AI enables to be done right now.

Rahul Bhardwaj (17:28): So boards hear this scary story and they say, wait a minute, we've got competent management. They'll put in the training and the systems internally to minimize that. But then the board looks at itself and they say, "How do we know we've got the best practices in place? And we thought we had it on the right track under cyber risk 1.0, now we've got AI-enabled cyber risk." Let's talk a little bit about that, and you were starting to describe the Watermelon Effect as maybe a place for the directors to start their inquiry.

Stephen Burns (18:01): So let's just talk about what the Watermelon Effect is first.

The Watermelon Effect is a reference that your dashboard, your management information system, is showing a whole bunch of green lights in that, you know, either a three-light system or a five-light system. Still, often it's green, yellow, red for a risk profile on a management report that says, here's what we're measuring, here's our key performance indicator, here's our management performance indicator, and we're showing green all the way down. And it's showing green, green, green until you're hacked and you realize that you weren't measuring the right thing; that your board and management systems were showing that you were paying attention to and looking at risk from a historic point of view when you haven't updated your management systems to start thinking about risk in the more modern current point of view.

If your management reporting system is showing you the same dashboard that you had five years ago to today as it relates to information and technology risk, then that's a really good indication that it hasn't actually updated in a timely manner; that it's not showing you the right information and you might be walking yourself into a Watermelon Effect problem.

Rahul Bhardwaj (19:18): So green on the outside, red on the inside, you've got a problem.

So, if a board now is wise to this and a director goes back to their board and shares the stories that I heard Stephen Burns talking about the Watermelon Effect, how do we start the path to ensure that we don't fall into that trap?

Stephen Burns (19:36): Well, I think it comes down to the fundamental question [is]: where is value in the organization? What are we trying to protect? And the value in the organization is typically twofold. Either it's the buckets of data that we have that we wanna make sure that we are using and protecting, and that others can't take to their advantage. And then our infrastructure, our operations, how are we actually performing? What are the machines doing, and are they separate and protected?

I think the question becomes, okay, once we understand where our valuable things we want to protect are, then we go to look how would they be attacked? And we do that with a modern lens of how would someone look to attack them today? How would someone look to get control and access to that system? And from there, we start looking at our protections. What are our physical procedural, electronic protections around those things? But also, really start to look at our management controls as well.

It's interesting to think this way. A lot of generative AI tools are free. A lot of systems that your staff want to use are free or low cost and they can use them with a credit card if they have to buy them. So where historically the organization has used dollars as a proxy for risk and have looked to have more control, the more expensive something is, that really doesn't work in the cyberspace because their most valuable data might be exposed because someone wants to use it with a free generative AI tool they found on the internet that's gonna do something cool and make them look great to their bosses. But because it's free, no one will have caught and done the right level of diligence or management thought to say, “Is this something we want to have the organization use?”

So, it's a good example of where a historic control using who can sign what for how many dollars, and having a hierarchy of authority based on dollars doesn't work in the modern world of technology, because so many of them are given away for free or for low dollars.

Rahul Bhardwaj (21:54): You're gonna hear some directors say, well, look, we've got this risk under control. We've got the policies and procedures and everybody's been following them. We haven't had a big cyber risk incident to worry about. I'm guessing you'd say something along the lines of that's not gonna cut it in this new environment. And they're gonna say, well what do we need to do to change? What would be a couple of the top things in your mind? And also let's touch on board competencies. Do you need to change the folks that are on your board to really upper up or oversight?

Stephen Burns (22:22): So I'm not gonna advocate that you need to change your boards out. That's a big bold statement. But I am going to say that there is an advantage to bringing technology expertise into the board, certainly into the board discussions.

So, what we're seeing is a movement towards having technology expertise as being one of the members of the board. The value of the board is the difference in experiences and backgrounds that each brings into that conversation. And so adding that strong technology voice to the conversation is going to help.

The second thing that the boards are looking for is the technology competency at the management level. Is their senior leadership team aware of and thinking about both the strategic risks, which is what we're talking about, but also the strategic importance of looking at the new technologies, seeing where they actually drive value in the organization and finding a way to advance their position competitively by making these investments.

Sometimes it doesn't make sense to make the investment because it's $4 million of it to make $100,000 savings in a business process, that's not a good choice. But if you can make a couple hundred thousand dollars investment into a tool and you can get a 20% improvement in productivity and as a result you can get better quality and faster outputs from your teams, then that begins to make a bunch of sense.

The second thing that the board needs to do is break down what its controls are around risk to make sure that they are adapted to the new environment. And a good example of that is organizations are adopting data schemes where they're looking at their data and they're breaking it into five or six categories and identifying, of those five or six categories, what level of management authority is required before data sitting in those categories gets to interact with the new system.

So if you've got a high risk data category, it's not up to just a question of, well the software's only a hundred bucks, so I don't need to get authority, I can just go do it; into, well, the software's with a hundred bucks, but that's a category five piece of information piece information so I actually need to have a vice president's signature as well before that data can proceed along. They're beginning to mirror the controls to the risks.

Rahul Bhardwaj (24:51): So, it's starting to sound a lot like what got us here won't necessarily get us there, which is the reframing and the rethinking the boards have to constantly go through to address potential new risks and strategic opportunities. It feels odd to talk about cyber risks though, without touching on that lovely topic of cyber insurance and maybe even cryptocurrency because all of these things are relatively new on the landscape too. What are your thoughts about those, Stephen?

Stephen Burns (25:18): On the insurance side, the insurance organizations are very much in the business of understanding the risk that they're insuring and making sure that they are appropriately and adequately charging the fees relative to the risk that they're taking on. And what we're seeing is where cyber insurance, it's a rapidly evolving area where you're seeing more and more exclusions, more and more due diligence, a stronger understanding of how an organization is understanding and protecting its data and its systems, and as a result, a better alignment of the insurance covers and the insurance fees to the actual risk of the organization.

Where before cyber insurance was an add-on to a policy, now you're seeing it as a separate, standalone policy with a great deal more diligence and thought into it. And it's actually a great exercise for a board to ask the very question that you just put to me, which is, so where are we on our cyber insurance and how are we doing? Because it's the third-party view as to how we're doing and managing our risks, and you can see it in the premiums they are charging.

Rahul Bhardwaj (26:28): And cryptocurrency. It seems to be out there, I would say, as it relates to a lot of this that shows up in various conversations.

Stephen Burns (26:35): It does, and it's mostly in the world of cyber risk. The blockchain is a very useful tool, and it's a tool which, when properly deployed, can make improvements in how organizations deliver their goods and services out to their market.

Separate from blockchain is this idea of blockchain used as a cryptocurrency or as a holder of value. And the idea that, because it is not coming through the banking system, it is of value to the threat actor as a means of payment.

There's a lot of discussion around whether, should a ransomware attack occur, we have to pay a ransom, we are likely to have to pay that in cryptocurrency. How are we going to fund that? Do we keep a stock of cryptocurrency around?

Whether you should or should not have an amount of cryptocurrency sitting around is a question for your treasury group, your controller, to sort out what's the right blend of assets they should have in that bucket. But I think the better place is not to get lost in the question of how we're going to pay when we're attacked, and instead focus on are we taking adequate steps to secure the assets that are most important to us and are we really understanding that hierarchy of importance within our data and our operations systems?

Rahul Bhardwaj (28:00): So, Stephen, we started this journey of a conversation saying that AI's been around in various forms for quite a while, and boards have adapted to deal with both the risk and the opportunity associated with that. And here we are landing on advanced cyber insurance and cryptocurrency, which was not in the cards when AI started, so it wasn't in the lexicon of directors, so it certainly wouldn't have been in the processes they've got.

Which brings us now to the resolution, and I'm gonna ask you, which way would you vote on, be it resolved, the boards need to fundamentally rethink cyber risk in the age of AI, which way would you vote?

Stephen Burns (28:39): So I agree with the resolution, and I do so because we've just gone through this conversation, but there's a lot of change going on. There's a lot of new functionality and new capabilities, and there's a lot more to come, right? There's a lot of research dollars going into it, and we've broken a number of problems. AI is now; some of the AI tools are now doing some incredible things, and they're likely to continue doing incredible things.

Those tools are just tools. They're focused on your data and your operations. And for every threat actor who's able to use one of those tools to get better at what they do, we need to get better at what we do and that requires us to really focus on it. So, I'm voting for, yes, we need to rethink the cyber risk.

Rahul Bhardwaj (29:24): And as you said, there are more things to come, which Stephen, that means you and I, are gonna get back together, have a conversation again when things evolve a little bit further. So, thank you for joining us today. I think this has been a really interesting wide-ranging conversation.

And to our listeners, I hope you enjoyed today's episode of Be It Resolved and that you've deepened your boardroom insights to stay ahead of these emerging trends. If you enjoyed the episode, please subscribe, rate and leave a review on your favorite streaming platform. From the Institute of Corporate Directors in Canada, I'm Rahul Bhardwaj. Until next time.