CEO Leadership Series: Vol 28

Clearwater

Understand Cyber Risks within the Healthcare System

September 7, 2023

Jon Moore

Chief Risk Officer,
Head of Consulting Services
and Customer Success,
Clearwater

Key Takeaways

Digesting the Staggering Scope of Cyber Risk

- When you start to look at the size and scale and cost of recovering from a cybersecurity incident, particularly in the U.S. healthcare industry, it can be terminal.
- A breach in the U.S. in healthcare is the most expensive breach you can have.
    - It's now an average of $10.9 million and can climb to over $100 million for recovery from a ransomware attack, which are extremely common in the industry.
    - Beyond the direct cost of responding to and recovering from the attack, there are likely to be one or more class action lawsuits following a breach.
- There are anywhere between 100 and 120 nation-state and cyber-criminal gangs targeting the U.S. healthcare industry at any time. Particularly in the case of cyber criminals, they are often going to go after the most vulnerable. Making sure that you're not one of the most vulnerable is certainly a step in the right direction.
- If you're in certain healthcare segments - e.g. pharmaceuticals, medical technology - they tend to be more common targets for the advanced persistent threat. Those are the nation-state types of actors who are often engaging in some sort of espionage, where they're more interested in stealing intellectual property than they are in stealing patient data or shutting down systems and extorting you.
- The cyber risk from third-party service providers is one of the greatest risks out there for healthcare organizations. If you look back over the last few years, the largest breaches in healthcare have been as a result of a breach of third-party service providers - usually technology providers – which has a cascading effect across the industry because their clients or customers are impacted as a result.
- Several liability insurance carriers are becoming more careful about the coverages they provide. They're asking much more detailed questions. They have increasing expectations as to the controls that an organization has in place before they issue a policy. And then they'll oftentimes put limitations or restrictions on those policies.

Understanding Components of Effective Cyber Preparedness

A lot of organizations invest heavily in security tools, many of which go unused. Oftentimes because these tools are very sophisticated, they take experts to use them appropriately - resources that understand and can deploy them and know what to do with them. Oftentimes, that hasn't been factored into the equation. As a result, there's a lot of cybersecurity investment, a lot of tools that are out there that are not really being used.

There is an established body of guidelines for determining reasonable and appropriate cyber programs

- The Section 405D task group created in response to the cybersecurity act of 2015, provided the Health Industry Cybersecurity Practices (HICP) to address the most common threats to healthcare. The task group broke down those practices into those that are most appropriate for small, medium and large organizations.
- National Institute of Standards & Technology ("NIST") cybersecurity guidance and in particular the NIST Cybersecurity Framework.
- There are also active rule making at the SEC for greater management of cyber risk within publicly traded organizations.

Robust cyber preparedness involves proactive and recurring completion of exercises such as:

- **Tabletop exercises** with executive leadership and IT teams to prepare for and practice cyber incident response.
- **Attack simulation**, which involves deploying sensors and monitoring response without actually hacking into the systems. This helps understand how well an organizations existing security investments are working.
- **Penetration testing**, which involves a technical testing team carrying out a mock attack to see where vulnerabilities exist.
- **Assumed breach penetration testing**, which assumes a bad actor has received low level credentials access and seeks to determine how pervasive the breach can escalate to.

Despite Robust Preparedness, Incidents Can Still Happen – Measures Can be Taken Proactively To Reduce the Cost of Cyber Incidents

- **Reduce Downtime:** Appropriate procedures, technology and secure backups can facilitate reduced downtime and, therefore, reduce a lot of the cost associated with ransomware.
- **Limit Spread:** The learnings from proactive exercises can help leadership teams in incident response by highlighting controls that can be put in place to limit the spread of the ransomware, such as endpoint detection and response, as well as other response capabilities or tools that can limit the impact of a ransomware attack.
- **Avoid Signs of Negligence in Class Action Lawsuits:** To the extent that a company can demonstrate that it had a reasonable and appropriate cybersecurity program in place, that it was doing the things that should have been done, companies can typically limit their culpability or liability as part of a class action lawsuit.

Background

I started my professional life as a practicing attorney, clerking for a trial court judge and serving as a public defender. We started to have more and more technology in our law practice, which really interested me. I went back to graduate school at Carnegie Mellon and got a joint degree from their school of Computer Science as well as their Business School.

While there, I was recruited by a federal contractor, who was doing cybersecurity work, primarily related to Department of Defense security architecture. I ended up at PricewaterhouseCoopers, who at the time, was restarting their consulting practice and, in particular, their federal consulting practice. They had a partner in their healthcare practice, federal healthcare practice, who needed assistance and she asked me to help with an engagement, which went particularly well and I was asked to stay on the Healthcare team. I spent about seven years working with customers in the federal healthcare space. From there, it's really been that combination of my background in cybersecurity and privacy with my legal background, and healthcare background that has made me an ideal candidate for the position here at Clearwater. I've been here for about five years now.

The term Chief Risk Officer is almost a novel concept to a lot of otherwise seemingly large organizations in healthcare services. What does the Chief Risk Officer provide that these organizations lack by not filling this position?

The Chief Risk Officer is a pretty broad role across my organization in terms of the typical business risk compliance issues that we may face as an organization in the context of working with our clients. The space that I am more focused on is enterprise cyber risk management, which is becoming an increasingly prominent issue here and around the world. There are active conversations at the SEC for greater management of cyber risk within publicly traded organizations. But, that doesn't mean that those same risks don't exist across all sizes of organizations. And, because of the nature of cyber risk in particular, not managing that risk well, or not understanding that risk can mean the end of those organizations. When you start to look at the size and scale and cost of recovering from a cybersecurity incident, particularly in the US healthcare industry, it can be terminal.

The nature of cyber risk, as some of us have understood, might be too narrow given the rapidly evolving world that we're living in. I'm thinking about the influence of artificial intelligence ("AI"). We're now seeing the power of virtual misrepresentation - individuals that don't exist, or individuals that do exist but that are being misrepresented with incorrect identities.

First, let me kind of address the initial point you made around cyber risk and what that even means. It's frankly not a hundred percent clear to me sometimes because historically cyber risk referred to the risk to those information systems that were exposed to cyber incidents. Often those are targeted with hacking type events. Then we had IT security risk, which was a bit broader and included things like environmental risks. Typically, we are talking about risks to the confidentiality, integrity and availability of our information. And as these information systems have become embedded into everything that we do, it's rare that you see any process or function within businesses that our data systems don't touch in some way, shape or form. We've become dependent on our information systems - so, there's the complexity of that which in and of itself poses a risk. There's risk where systems are misconfigured and information is exposed that shouldn't be exposed, or where folks are allowed access to information that shouldn't be accessible, often resulting in confidentiality breaches, and high levels of corresponding disruption.

We tend to underestimate the degree of disruption that comes from a breach of a system that is at the core of everything.

I wish I knew with certainty what AI risk represents. Unfortunately, there's a group of actors out there - nation state type of actors, primarily coming out of China, North Korea and the former Soviet block – as well as sophisticated cyber-criminal gangs. There's anywhere between 100 and 120 of these groups targeting the U.S. healthcare industry at any given time. And they are investing, reinvesting in their businesses just like any business would. And when they reinvest, they look for more sophisticated ways to take advantage, to gather the information they're trying to gather to exploit businesses. We already know of situations where Chat GPT is being leveraged to write more effective phishing emails, as well as where it's being leveraged to actually write malicious code. They'll inject a piece of malware, but that malware itself doesn't necessarily have the code. What it does is call back to Chat GPT and have Chat GPT write the malicious code that then is deployed within the environment that's been hacked. Those types of things are occurring right now. We also have the deepfakes where AI is used in social engineering, to create a voicemail that sounds like the CEO or a voicemail that sounds like a message from the CFO. Usually in those cases, they're trying to have funds redirected into bank accounts of their choosing. We have a situation right now with web beacons and pixel trackers, where these tools that are typically used for web marketing purposes, are used in conjunction with AI and other sources of information to essentially build profiles of people that include protected health information, which is a breach of HIPAA regulations.

Are there data analytics tools out there that provide executives with comprehensive, ongoing real time assessments, or is it still more analog manual expert assessments? Where are we in terms of providing solutions?

For any organization, what you should be seeking to achieve is putting in place a reasonable and appropriate enterprise cyber risk management program for your business. So what does that mean? Every business is different. They may vary in size, they may vary in the nature of the services they deliver, geographic diversity, systems that they use. All of these components may vary. From a regulatory perspective, as well as from a best practice perspective, we want to understand what's reasonable and appropriate for our particular business. There was a lot of work done by the health and public health sectors joint cybersecurity work group, 405 D task group, which was tasked with fulfilling the expectations of section 405D of the cybersecurity Act of 2015 that provided cybersecurity practices to address the most common threats to healthcare. They broke down those practices into those that are most appropriate for small, medium and large organizations. For example, if you're a managed services organization in healthcare, the recommendation is to follow the medium practices of 405D Health Industry Cybersecurity Practices. So that's one of the starting points that we recommend for organizations. We typically would use the NIST cybersecurity framework in conjunction with the health industry cybersecurity practices to start to build out an appropriate target profile, as well as a reasonable and appropriate cybersecurity program for that particular organization. There are a number of things that drive the considerations, including your business goals and objectives, the results of your risk analysis, different compliance requirements you might have. These variables drive the design of what's reasonable and appropriate, and the controls and safeguards that you should have in place as an organization. So that's a starting point for building a program.

So that brings us to your second question, which is, "how do I know, what's happening with that program and whether or not it's effective?" There are different things that you might put into place to understand that. When we work with our customers, we use several tools. For example, we use a program management tool to help folks we're working with understand where we are with the design, implementation and operation of that reasonable and appropriate program and the safeguards

that need to go into place, as well as ongoing tracking of those activities. Then, there are other activities that you need to be doing on an ongoing basis, both from a best practices perspective as well as from a regulatory perspective, including risk analysis and risk management, risk response types of activities. Over a decade ago now, Clearwater became the first developer of a cloud-based SaaS product for the conducting of risk analysis and risk management in compliance with, with HIPAA, but also from a best practices security perspective. And within that system, we also provide dashboards, though perhaps as dynamic as what you were referencing.

When you get into the area of security operations itself, now we're talking the day-to-day, minute by minute type of monitoring and responding to threats as they occur within your environment. And similarly, when you get into managed detection response, we have a customer portal where you can see where all of the vulnerabilities are from the last scan. So, it's more dynamic at that operational level - you can get a bit more insight into what's happening minute by minute within your IT environment. For organizations that are more cloud-based, it's easier to start to develop dashboards that provide additional insight into the client's environment and what's happening in that environment, and where the organization is relative to control requirements, different compliance requirements, as well as just from an operational perspective.

Do you run cybersecurity crisis simulations? Do you pretend to be an outside nefarious influence to see how strong your client's infrastructure is at withstanding and responding? Or is that something that you haven't yet engaged in?

One of the ways that simulation can manifest is in tabletop exercises. And those are quite common for helping leadership teams understand where they are from a maturity perspective in their preparations and ability to respond to a cybersecurity incident. We'll do tabletop exercises both with the leadership team itself, CEO etc., and we'll also do that with the technology teams to help them understand how they would respond to and how prepared they are for any sort of cybersecurity incident. That's one way that we can do simulations.

There are also software tools that are available that will simulate attacks on your environment. You put sensors on your environment, and they send traffic over the environment to simulate that attack without hacking into your systems. Those are becoming more common. And they can simulate some of the most common attacks. You can see how your security controls respond to those simulated attacks. We call that your controls validation assessment - looking at how well your controls would respond to different types of attacks like ransomware, different flavors of ransomware.

Then there is penetration testing itself, which is more of the historic way that organizations have tested their controls. In those scenarios, we use our technical testing team. Most of those folks have come out of a government background, very sophisticated in their abilities to simulate the same types of attacks that would be used by nefarious actors out there in the wild. You can have a basic network external penetration test where we're going try to get in from the outside, and there may be rules of engagement limiting what they do. For example, they may not phish to gain credentials. They may limit that to look for specific vulnerabilities in the network that they can exploit.

Oftentimes we'll also do what we call assumed breach penetration testing. My top hackers will tell you, "if we want to get in, we're going to get in - we'll probably phish someone in your organization and that will allow us to get in." The question then becomes, "what can we do once we're in there?" The assumed breach attack assumes that the bad actors have some low level of credentials. We look to see what the bad actors are able to do moving throughout the network, escalating their privileges and what they're able to access from data perspective once they're inside. So that's another flavor of a simulated attack, and there are a few variations depending on what it is you're trying to assess, whether that's a web application, pen test, mobile applications.

All of these things are subtly different and because of those differences, we rely on subject matter experts that specialize to a degree in these different types of attacks - they tend to need to specialize because of the complexity of the IT environments we live in today.

Simplifying for a healthcare leader who doesn't spend their life day in day out thinking about cybersecurity - if we could get to a point where every leadership team either does or does not engage in simulations on an annual basis, that binary test tells you whether you're on the right side of risk management here or not. Would you agree?

As an industry, healthcare is behind if we're thinking about where we are relative to the threats that are out there. What we have seen though, and I think this is a positive trend, is that as result of discussions around expectations for public companies, we're seeing more and more interest in understanding the risk at the board/investor level. Private equity firms we work with are trying to understand what their cyber risk management exposure is across their portfolios, asking relevant questions as to where the organization is relative to cyber risk because they're seeing the impact of some of these attacks. The Ponemon Institute, in

conjunction with IBM, came out with their most recent study on the cost of a breach. And again, a breach in the U.S. in healthcare is the most expensive breach you can have. It's now an average of $10.9 million (and can climb over $100 million) for recovery from a breach, ransomware attacks, which are extremely common in the industry. As I said, that's getting people's interest.

The other thing I think that really changed in the last couple years is there was this notion that we could just insure away this risk. That is not as easy as it used to be, because of the size of some of these claims that have resulted from breaches. Several liability insurance carriers are becoming more careful about the coverages they provide. They're asking much more detailed questions. They have increasing expectations as to the controls that an organization has in place before they will issue a policy. And then they'll oftentimes put limitations or restrictions on those policies as well. Because of that, we're seeing more activity around organizations looking to improve their cybersecurity programs than what HIPAA has done in the last decade.

Is there any good news in terms of better tools that are now available that allow for better protection, better tracking? It seems to me historically that the hackers have been sort of five steps ahead of the protectors.

Just like the attackers are using AI, the AI is being used in all sorts of security tools as well. We implemented artificial intelligence into our risk analysis software to help facilitate the identification of the likelihood and impact of a particular risk. There is increasing investment in cybersecurity tools, and those tools are getting better. I would caution, however, that this desire to just buy something to solve this is unrealistic. A lot of organizations invest heavily in security tools, many of which go unused. Oftentimes because these tools are very sophisticated and they take experts to use them appropriately, resources that understand and can deploy them and know what to do with them. And oftentimes that hasn't been factored into the equation. As a result, there's a lot of cybersecurity investment, a lot of tools that are out there that are not really being used.

What we try to encourage folks to do is to think about this programmatically and to build a program based on recognized standards. It's the steps that seem straightforward that cause the problem.

So, for example, and this is one we see all the time, one of the first things that you want to do is understand what it is you're protecting. What are those systems and associated components out there that are being used to process our information, whether that's protected health information or other confidential information, or just business information generally that we rely on for, for conducting business? You would be shocked at the number of organizations that cannot answer that question. They don't know what systems they have. They don't have a good handle on their software inventory. It's very likely that there's going to be something out there that's going to be vulnerable to attack. I don't know whether I even want to use this reference, but there's the old joke - two guys are in the woods and there's a bear. One guy says, "do you think we can outrun that?" And the other guy says, "I only need to outrun you." To a certain extent, that's the world that we live in. Oftentimes, the threat actors, particularly the cyber criminals, are going to go after the most vulnerable. So, making sure that you're not one of the most vulnerable is certainly a step in the right direction. It's fundamentals that I would recommend organizations focus on. And once you get those fundamentals in place, then you can start to look at what additional risks do I have and what might be an appropriate investment to further reduce or manage those risks. If you're in certain segments of healthcare - e.g. pharmaceuticals, medical technology - they tend to be more common targets for the advanced persistent threat. Those are the nation state types of actors who are usually engaging in some sort of espionage, where they're more interested in stealing information than they are in stealing patient data or shutting down systems and extorting you. That's a whole different problem.

Assuming a cyber incident does happen, are there backend strategies you can put in place proactively to mitigate the cost of these attacks? The burden of downtime, is there a strategy for a redundancy plan, switching to paper or something like that?

I think what you're focused on is our ability to respond to the attack itself, and then our ability to recover from it. And to the extent that we have put into place appropriate procedures, technology, and that we have secure backups that can facilitate recovery, it can reduce downtime and reduce a lot of the cost associated with ransomware. And going back to the tabletop exercise, having gone through those tabletop exercises, the leadership team is then aware and understands how prepared the organization is to respond to circumstances like this, having controls in place that can limit the spread of the ransomware, like endpoint detection and response capabilities or tools that can limit the impact of a ransomware attack. There's a number of things that an organization can do to reduce the risk. And there's different controls like network segmentation, endpoint detection response, and others that I mentioned that you can deploy to, to limit the spread. But then how do we recover, and do we have appropriate backups in place? And this too is becoming more and more complex because the organization started recovering from their backups, and then the ransomware actors started loitering

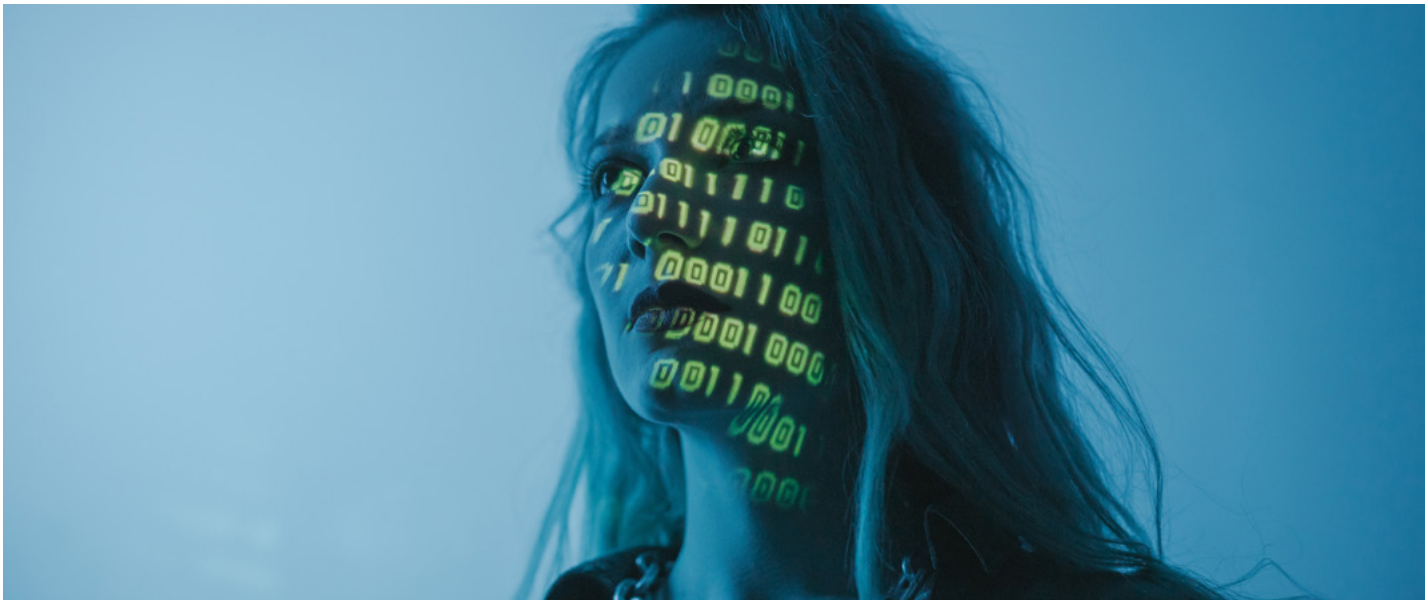

so that the ransomware would get caught in the backups. Then you wouldn't have secure backups. Your backups will be corrupted as well. More frequently, we'll see them get inside the network and they'll first steal electronic protected health information - for example steal images from a surgery - then threaten to post the images if they didn't receive the ransom, or steal family contact information and threaten family members for payment.

Do you see any changes from a regulatory perspective that will make it easier to comply with regulations? BAAs get signed all the time. Some of these documents feel like they're quite difficult to follow to the letter of the law. It feels like you're vulnerable to cyber attacks, and then you're at risk of not being able to be in compliance with the true letter of the law such that you'll be in a bad position if something happens.

As a former trial attorney, I was shocked that it took this long to see a rise in class action lawsuits for damages related to cyber incidents. There were some reasons for the delay - in particular, in many jurisdictions you need to show actual damages in order to prevail, which is somewhat difficult in some of these cases. But, nevertheless, what we're seeing now is it's inevitable.

There are going to be one or multiple class action lawsuits following a breach. What you'll see as soon as there's a publicly announced breach, there's also basically advertisements going up from law firms saying that we're investigating this breach, and if you were impacted call us today. That is just rampant right now. And, some of these suits can be quite dramatic. Now, how do you protect yourselves against that? Well, given the current regulatory world it's difficult to stop anyone from filing suit against you. However, it is my belief that they're going to have to show negligence. Is the breach itself evidence of negligence? Well, maybe, maybe not, right? If so, to the extent that I can demonstrate that I did have a reasonable and appropriate cybersecurity program in place, that I was doing the things that I should have been done, and despite that this thing happened, maybe I can limit my culpability or liability as part of that. It's really difficult to write legislation with any specificity that some folks would like. That's part of the challenge with cybersecurity legislation. The risk of third parties is one of the greatest risks out there for healthcare organizations. If you look back over the last few years, the largest breaches in healthcare have been third party service providers. Usually technology providers. And then it has a kind of cascading effect across the industry because their clients or customers are impacted as a result.