Podcast Transcript

The Data Stream: Episode 10 with Craig Hoffman

Date: June 29, 2026

Guest: Craig Hoffman, **Host:** David Sherman, Nichole Sterling

Run Time: 27:07

For questions and comments contact:

David Sherman

Partner
Philadelphia
T: +1.215.564.8380 | dsherman@bakerlaw.com

Nichole L. Sterling

Partner
New York
T: +1.212.589.4282 | nsterling@bakerlaw.com

Craig A. Hoffman

Partner
Cincinnati
T: +1.513.929.3491 | cahoffman@bakerlaw.com

May: Welcome to the Data Stream, the podcast that dives deep into the fast-moving currents of data, technology, and the law. BakerHostetler's Digital Assets and Data Management Practice Group brings you this series, hosted by partners David Sherman and Nichole Sterling, to explore how companies navigate the complex and often intertwined life cycle of data, from privacy and cybersecurity to advertising, AI, and other emerging technologies. I'm Charlie May, and you're listening to BakerHosts.

On today's episode, David and Nichole are joined by Craig Hoffman, partner at BakerHostetler and co-lead of the firm's Digital Risk Advisory and Cybersecurity team. Craig is a digital risk advisor engaged to guide clients on issues related to the use of technology and data. He is well known for using his litigation experience and insights from helping entities address thousands of security incidents to develop prioritized privacy and cybersecurity strategies and effectively respond to security incidents. Visibility to what allows incidents to occur, outcomes of decisions made, and measures taken to improve going forward from security incidents of all sizes enables Craig to immediately forecast what is coming, offer prioritized recommendations, and then operationalize decisions into actions.

Sterling: Welcome, Craig. Thank you for joining us today. I know, or at least I have heard, that you are a diehard Cincinnati Reds fan. I can only imagine that probably builds some pretty high tolerance for things like hope, pain, optimism in the face of evidence to the contrary. What has baseball fandom taught you about counseling clients through bad facts that you just can't realistically spin positively for clients or executive leadership?

Hoffman: I think there is a saying in baseball, you are the back of your baseball card. That means your history of statistics, you might be a little above or below at any moment through the season, but across 162 games in a season, you are probably going to be pretty close to your statistical average. If you look back at security incidents over the last 15 years, we've tracked data and the underlying incidents, root cause, additional factors, have been pretty consistent over that time frame. We're about to see, I think, some change due to technology like AI. But overall, in an area where people are concerned at the start of an incident, that it is a crisis, there are a lot of unknowns, what we do know is that the incident itself will likely follow the same timeline, the same process, the same decisions as we've seen across incidents over the past 15 years.

And being able to go to those stats and show clients for this type of incident, on average, the threat actor's been in the network for 30 days before you learn of it. It will take you five days to contain it. It will take you three weeks to determine what occurred through the investigation. And it may be two to three months before you're in a position to notify individuals. That helps people have context at the start to know, what are we facing? Where are our decision points? And give people benchmarking so they understand, compared to others in our industry and others overall with similar incidents, what should we expect and how do we want to measure up?

In baseball, there is a new stat that people like to talk about, your expected outcome. So those are times where maybe you're a career 250 hitter, but you're hitting 180. I think that lends itself to people thinking maybe we can do a little better with how we handle this incident than others. And sometimes there are better outcomes than others. But in general, you don't have to hit a home run in incident response. Nobody ever gets an award for having the best incident. You only get recognition for mishandling it.

Sterling: On that point, is there some moment when you're dealing with an incident response where you can tell that it's really just become about avoiding mistakes that are going to make things worse?

Hoffman: I would say starting from the back when we have litigation, it is very rare that we're arguing and fighting that there wasn't some underlying issue that occurred that allowed it to happen, where there are facts that would allow someone to say there was a failure to immediate duty of care. So, in the history leading up to it, usually there are less than ideal facts. Most incidents are full with some, if not a lot of less than ideal facts. And our focus is then how do we take where you are at and put you in the best position based on where you start to work through the incident in a way that is aligned with your business priorities and your risk appetite. So, we don't come in with a one-size-fits-all approach and tell you what you have to do. Our goal is to put you in a position to enable you to make choices that are the right ones for your entity, not just choices that you're forced into by facts.

Sherman: So, thinking about that decision path for a minute, Craig, and I'll say that, and I'm speaking for Nichole here also, as Mets fans, and dare I say diehard Mets fans, I and we have absolutely no standing to ask this, but from your perspective, what is the cyber equivalent of a team beating itself? What is the unforced error that takes a manageable incident and turns it into something that just starts spiraling?

Hoffman: You can look back at incidents that received criticism for how they handled it or mistakes made during the process. And they fall into a few categories. There are some technical things. In days past, you'd have somebody on the security team uploading malware to VirusTotal and doing it from an account connected to the company. So, it becomes known early on that the company had an incident and what malware was involved before the company has said anything. Communications more often are areas where there is some kind of unforced error. And it goes back more than a decade ago to people making videos of the CEO. I think people learned that that is probably not the thing to do, and you don't see companies having their CEO get in front of it and put out a video and talk to their customers.

I think if you go back to the stats, what is really helpful and it helps when we talk to our clients, both before an incident and during, if you know you're one or two days in, and on average it is five days to contain a network intrusion and three weeks to complete an investigation, you can then explain if you're trying to reassure people or rule out certain things and talk about an incident early on before containment or before the investigation is complete, there is a pretty good chance you are going to learn something during the investigation that means what you said early on was wrong. That is why early on you can talk about method. The method you are using will be accurate because you know what you're doing. I think you have a window where you can talk about method and then allow yourself time to finish the investigation. Then you can talk about the underlying facts.

There are some simple things. There was a recent event where there were two phases of the attack and during the second phase of the attack, the threat actor defaced the technology provider's site with their extortion note. Shortly after that, the company took down its service and put up a message that said it was down for scheduled maintenance. Those are little things that probably don't change the course of how an entity is viewed in a significant way, but they're little things that allow people to lose confidence or lose trust. So, I think the focus is on, in a way, handling the communication like product recall. You have certain things you have to share. You have certain legal obligations. Do it in a straightforward, normal way that looks like all the other normal communications about incidents, and don't do things to attempt to stand out or overly reassure people, and you have a better chance of managing through it effectively.

Sherman: Well, it sounds like we can almost look back to using expected outcome, as you mentioned, as a metric to standardize things that go well, things that don't go so well. To that end, I don't know if you agree, but I feel like many times some of the critical unforced errors really happen in those first 24 to 48 hours. And during that window, leadership is dealing with a crisis and they want answers but as you were suggesting, what they actually need is a decision path. You started alluding to this, but how do you build that decision path when the facts are changing? And how do you keep clients on a path of making good decisions based on incomplete information to lead to that, to your point, expected outcome?

Hoffman: I think it is critically important to show the method you are using to manage the incident. If people understand that there is a reliable, trusted method being used to determine what is occurring, mitigate impact, and investigate and restore if there is business continuity impact, they will trust the process and work in the process. If they don't hear a method, if they hear chaos, if they hear things like, we're in a really fluid, rapidly changing environment, that suggests to people that they should be each individually taking on their own ideas and thoughts on how to manage it instead of working on the method.

In order to get people to trust the method, they have to know about the method usually ahead of time through the incident response plan process and tabletop exercises. If they know there is a method the company is going to use and is executing it, Executives can settle in, support the process, be prepared for the critical decisions they will need to make, and you have a much better outcome usually when you follow that process. A good example of that is discussing the kill switch approach.

So, in a scenario where a threat actor is in your network, you haven't contained it yet, you're uncertain of your ability to contain it quickly and prevent further or additional business impact, one of the discussions you can have and illustrate through a tabletop exercise is, do you delegate authority to the CISO to hit the big red button and contain by shutting down the network? Does it have to be a discussion with the executive team before you do that? Or is there some other way you want that to occur? When you've discussed things like that ahead of time, people can then follow that method, follow the protocol, and you end up with better outcomes.

Sterling: That just makes me want to ask if you have control of the kill switch for the firm. But I'm not going to. I'm going to ask after an incident. So, moving past everything that we've just been talking about, you usually end up having some sort of long fix everything list that the client is going to see. How do you get through that? How do you prioritize or help people pick the changes that they should really focus on? Because there is a tendency really to become overwhelmed when you have too much of a list and just do nothing.

Hoffman: An effective strategy has been to have three categories of those changes or fixes. One is immediate containment items, things you are doing to stop the attack. The second is short term, the third is longer term. And while you're in the middle of an incident, it is really hard to take any time apart from the containment, the investigation, the analysis process to look forward. But if you don't start looking during the incident at what your short-term and longer-term enhancements are, you miss the opportunity to be better prepared to address questions from stakeholders internally and externally.

And coming out of an incident when you're talking to people, and this is more acute depending on what industry you are in, if your discussion is only, we identified it, we took steps to contain and eradicate and investigate, and here is what happened, that meets a basic requirement. But if you then are able to say, based on what occurred, here are the things we were doing over the next three to six months to strengthen our security measures, and they tie back to what allowed the incident to occur, that is your opportunity to rebuild trust and confidence and to change the narrative from being about the underlying incident to what you're doing going forward. And especially if you're, let's say, a technology company with business customers who are affected by the incident, if you don't anticipate that those questions will come immediately, you end up unprepared for your first engagement with your business customers, and you make the rest of the process of keeping them as your client harder because you didn't anticipate what their basic questions would be at the start.

Sherman: If it is all right, Craig, I have a quick follow-up on that. Let's say we meet a client for the first time in a proactive capacity, meaning before an incident occurs, and they come to us and they have a robust or what they believe to be is a robust incident response plan, et cetera. In that scenario, where do we see proactive planning really fall short when it is not grounded in an incident or in direct experience in the incident response field?

Hoffman: It may be over-planning in terms of a really long incident response plan with a lot of content that isn't focused and practical. I've seen 50, 75-page, 100-page incident response plans. And day in and day out during an incident, companies aren't pulling out the plan and using it as a step-by-step guide. The value of the incident response plan is the effort that goes into putting it together, identifying the internal team, identifying the key third parties you will engage, knowing how to bring that internal team and third parties together, and knowing how to bridge the gap from the core team to the executive team, and for public companies to the audit committee or board. Being able to anticipate the teams, how they get together, how they communicate, and making sure it is built on a reliable method

like a incidence and response method. You can have a short but very effective functional plan that can be used.

A good example of something that is often missing would be classification and use of the classifications. You can be very simple, low, medium, high, critical. That classification can drive decision-making on when you escalate to the executive team or audit committee or board. It can support your cybersecurity materiality analysis program, and it can support how you decide when to engage with third parties. There are simple things you can do. You don't need a really long, detailed, prescriptive incident response plan. You need a reliable method. You need to know your involved parties, how they work together and have a core functional approach for how you manage the incident embedded.

Sterling: So, you mentioned AI earlier, and I'd like to ask you a little bit more about that. It feels like lately we've been seeing a never-ending AI hype cycle, but one that has increasingly focused, I think, on a bit of the incident response world. It can sound doomsday level scenarios, things like autonomous AI finding and exploiting vulnerabilities within systems or networks on its own, or highly convincing deep fake calls, videos impersonating executives, all of this type of stuff. What are the risks that you are actually seeing right now and what does feel inflated still within the AI hype?

Hoffman: If you asked me this question six months ago, I would have been more in the camp of it is hype. When we were doing our 2026 BakerHostetler Data Security Incident Response Report, as we were collecting our 2025-year incident data in December of 2025, very few of our incidents had AI as a significant component of the incident. But as we went into data collection and report writing, January, February, March, we started seeing AI show up more often. I think we've moved out of the pure hype into definitely still some hype, but starting to see the reality of AI being involved in security incidents more often.

Historically, we would see AI used to write a better phishing e-mail, or for some smaller impacts or smaller methods of access, some type of vibe coding. Then we started seeing stories of the Agentic AI scenario reported by a couple of companies. Then with the Anthropic Mythos announcement, that led to people worried about Bugmageddon. I think some component of Bugmageddon will occur. And I think it is a scenario companies are just going to have to gear up and get ready to withstand. But the Bugmageddon will have an initial phase and then should die down.

I listened to a well-known security firm go through their 2026 incident response report, and they talked about incident classification, and they suggested that in their coming year report, they will likely categorize Agentic AI as a separate type of incident. I think we're moving out of the mostly hype into the early days of AI being a significant part of security incidents. I think one of the things we may see, in 2025, our incidents very rarely involved malware as a part of the incident. A lot of it involved identity and access credentials. If tools like Mythos are going to be used to find multiple vulnerable conditions and then connect them to the skills and tools to exploit those conditions, that will likely mean involvement of more

malware in an attack which makes the incident harder to investigate. It makes the investigation longer. It makes the investigation more expensive. But I think in ‘26 and ‘27, some of the trends we saw of investigations completing faster and being less expensive may reverse course. If vulnerable conditions are exploited by AI tools and that malware is used more often, I think we’ll see the first responder industry struggle a bit with the volume of some of these incidents.

Sterling: So, do you think then that AI changes the core risk picture, or is it just really amplifying some of the issues that were already there in the incident response world?

Hoffman: I believe there is a chance it could change the core risk, but not yet. What you will see AI do is find things that have been methods and root causes for the past 15 years, but they will do them faster, better, and at a broader scale. The challenge there is threat actors mostly don’t need AI to be able to do that. You can have a discussion with someone about how should you prepare and guard against deep fakes. You don’t really need a deep fake to effectively social engineer someone. You just need to call them up and pretend to be the IT help desk. And if you’re skilled at all in social engineering, you don’t need a voice that sounds like the IT help desk voice.

Right now, you don’t need the power of AI to compromise an entity. It can certainly help you. So, we’re in a bit of a timeframe where the underlying methods, the underlying causes often have simple day in and day out root cause behind them that tie back to things companies have focused on for years and just struggle to effectively, continuously address. What AI will do in the short term is make bad guys better at exploiting those underlying issues.

Sherman: So, to that end, as we address AI, say AI risks going forward, as you just mentioned, companies or entities are certainly aware of the risks, or maybe they’re not aware, but are going to continue to see actors using AI against them. or Agentic AI against them as you outlined. But there is a second category of AI risk that perhaps is going a bit less, I’ll say it is a bit less prevalent, which is an entity’s own use of AI creating exposure. Of those two buckets, again, actors using AI against entities, and then the second being an entity using own use of AI creating exposure, which of those two categories causes you more concern?

Hoffman: When I talk with companies about risk related to AI, there are really three categories we talk about. One is AI being used to attack you. And as we just discussed earlier, a lot of the underlying things that AI can be used to exploit your security program should already be designed to address. The second category would be your enterprise tools that use AI and the third category would be employees using unapproved forms of AI. I’m more concerned about category two, in terms of when you’re comparing category two and category three, enterprise tools or employees misusing unapproved tools, I’m more concerned about compromise of enterprise tools.

If you take an average employee and let’s say they wanted to do something for a presentation and they decide to use ChatGPT and they put some sensitive

business data in ChatGPT, that often is something you can address. If you go back to classifying incidents as low, medium, high, or critical, often those are low. There are usually ways to work through things like that. And an employee who puts a spreadsheet in an unapproved enterprise tool that has sensitive employee data, there are ways to work through those. But if you have an enterprise AI tool that you are connecting to critical internal systems, the AI tools have not been around for in some cases, months, let alone years. They may not have gone through the same series of secure code development life cycles and updates.

You may be in the arms race in the competition to use AI, installing a new, maybe less than effectively tested program. If that AI tool has a vulnerable condition that gets exploited, you're not connecting it to your most significant assets. That gives them the ability to either take a significant amount of data or shut down something that is important. To me, that is a more risky scenario than an employee putting 10 names and social security numbers in a non-approved tool.

Sherman: Got it. So last question, setting aside AI, since I know we spent quite a bit of time there. What should organizations be paying more attention to right now? What is that one thing you'd tell every in-house counsel to pay attention to today?

Hoffman: It is not new or different or unique. It is risk assessment and governance. If you talk to security experts, you might hear them say, right now is really dynamic. You have to be fluid. Things are changing rapidly. When you hear things like that, doesn't reassure people and that doesn't allow people to have an approach or a method to address the risk. And I think taking that, it is fluid, it is dynamic, it either leads to people throwing their hands up and saying, well, why try? Because I'm just going to be overrun by AI.

If you look at the tools you're using and your program and your data and use some data about incidents occurring, you can come up with what you think your most significant risks are. You can look at your existing controls and decide how effective do we think our controls are at mitigating that risk and that can help you identify areas where it makes sense to spend more time and effort improving your controls and your approaches to further mitigate your more significant risk. And you can continue to update that quarterly, every six months or so as technology and use of technology develops. That helps you have a consistent approach to it.

And it is also something you can use to explain if something does go wrong. Here is what we tried to do to avoid this. Maybe we weren't perfect, maybe we didn't get it completely right, but it wasn't because we didn't try and do it the right way. That kind of explanation goes a long way with third parties like regulators. It may not help you as much in litigation, but it at least arms your, if you need to engage an expert to talk about your standard of care, it gives them something to use to say, this company had the right program. They were trying to do it the right way. Maybe it wasn't perfect, but that is not the standard.

Sherman: Fantastic. Well, look, I don't know. I find that somewhat comforting, which is to say, if we keep the governance and risk-oriented fundamentals in place and stick to process, odds are you'll have a better outcome. As I said, I find that somewhat comforting at least. But Craig, thanks so much. This was a great conversation. Really appreciate your time.

Hoffman: I enjoyed it. Thanks for having me.

Sterling: Thanks for joining us.

May: If you have any questions for David or Nichole, their contact information is in the show notes. As always, thanks for listening to BakerHosts.