EP. 60: NAVIGATING THE IMPLICATIONS OF AI FOR CYBERSECURITY

AUDIO TRANSCRIPT

Rob Boyce [00:00:00] There's a is there a foundation of data and infrastructure and, you know, other things that need to be put in place to be able to leverage AI appropriately and that doesn't change for the security space.

Kris Burkhardt [00:00:18] Hello everyone. Welcome to the AI leaders podcast. My name is Kris Burkhardt. I'm Accenture's global chief information security officer and I have Rob Boyce with me here today. Rob, do you want to introduce yourself, please?

Rob Boyce [00:00:31] Hi, Chris. Great to be here with you, everyone. I'm Rob Boyce. I'm our global cyber resilience lead as part of our, Accenture cybersecurity consulting practice.

Kris Burkhardt [00:00:40] Fantastic. So Rob and I are going to have a conversation today about AI. We're just about at the very tippy top of the AI hype cycle. I think it certainly feels that way. And there's a lot of concerns, and a lot of ideas about AI and its role in cybersecurity, both as a, you know, tool for the attackers and as a tool for the defenders. So Rob and I are going to explore a little bit of that today. Before we dive right in Rob, do you want to, perhaps give us any opening thoughts on AI and you know, some things that you're seeing or hearing?

Rob Boyce [00:01:19] Sure. I guess the interesting part to me is, you know, you're as you're talking about the tip of the hype cycle. I think in at least my career, this is probably the third hype cycle we've gone through for AI, although this is the one that feels the most real to me. So I have, the most optimistic about the possible outcomes. I think there's been truly a lot of focus on the threats associated with AI, but I'm really feeling there's just so much opportunity to use AI on the defensive side and so I think today hopefully we'll get to explore both of those. But for for me, this is really, you know, a great opportunity to drive. You know, what I would say is, you know, more efficiencies, more, you know, true innovation in a space that's been a bit stagnant. You know, I think security operations in general hasn't changed in the 20 years, 25 years I've been doing this, the principles have been the same. The technology may be a bit better, but, you know, there's not been a dramatic innovation in this space. I think something like AI is going to be able to provide that, you know, accelerate that innovation. So I'm super excited about, that possibility.

Kris Burkhardt [00:02:27] Yeah. Well, I sure hope you're right. So let's, as we think about some of the defense, let's start with the attacks, because maybe it's easier to think about some of the defense if we think about some of the attacks that we're dealing with. I want to share one with you, Rob, that you might find interesting, that I saw recently and it has to do with social engineering. I think we've all heard that,AI typically generative AI could be useful in phishing attacks. It sort of seems obvious, you know, tailored phishing attacks where AI gets some prompts from, perhaps your social media it can do instead of the old, you know, phishing campaigns or spam campaigns where everybody gets the same, the same email. Now that can be tailored. AI, of course, a lot of the grammatical challenges that non-native speakers might have and so it can be a powerful tool, but it's really more amplification of the same. But I think it's still important. The attack I want to share is actually a texting attack. So social engineering attack over WhatsApp and so we regularly see our leaders impersonated on WhatsApp and the attacker will reach out to a potential victim pretending to be one of our leaders. Right. Whether it's our CEO or our CFO or what have you, and ask them for either information or, you know, some obscure monetary request. Hey, wire some money to this bank account or strangely, buy me a bunch of eBay gift cards. Right. What have you. They've deckers or. It doesn't always make sense, but this one made a lot of sense and it actually got quite far along. The, attacker claimed to be, again, one of our C-level leaders who was stuck in Hong Kong with some government visa issues, and they needed some money to handle lawyers and all that in this conversation went on for quite some time before the victim, fortunately, ultimately was not a victim before the Would-Be victim realized that it was happening. And I had a chance to actually review the texts, and the text was so good, it really sounded like, it really sounded like our leader's voice. There there were no grammatical errors and ultimately, what gave it away is there were no texting shortcuts, you know, like, you are instead of your right, the letters you are sort of your or, you know, simple things that people do. There weren't like a lot of smiley faces or emoticons or any such thing. It was just pure words and the victim. The Would-Be victim grew suspicious and reached out and called it in. But it was it was quite amazing. This must have gone on. I don't know, for maybe 20 text on each side before our, You know, our employee realized it. So I was just impressed. I mean, I haven't seen that level of sophistication yet and so this was, I think, a sign of things to come.

Rob Boyce [00:05:46] That's super interesting. It's funny that we usually able to detect phishing or social engineering attempts because of the mistakes in the text and this time you recognized it because it was too good. It was no mistakes. I mean, it's so fascinating. So we have to a medium of being able to use proper text line and things like that. But yeah, I honestly because I think, you know, the cases that we're working for, clients are very similar. We're really starting to see much higher fidelity, and efficacy of social engineering attempts leveraging AI and I think when I think about all of the ways that we've seen it in the wild so far, that is the number one way. We haven't really seen, you know, a lot of you know, attacks being generated by AI outside of social engineering or threat actors using AI to accelerate their attacks in the way that we, you know, thought we may by this point. I think on the offensive side, on the threat side, other than social engineering,
you know, there's a couple of things that worry me a little and a couple of those are, you know, we're seeing at least one threat actor group now move to, Python based, ransomware and, you know, launching, you know, creating malware in Python, which is not new on its own. But, you know, we've done a lot of research on our side to try and create malware and leveraging ChatGPT and gen AI in general and it's not very high. You know, the efficacy of it, creating it with C or C plus plus is very low. We're not it's not very successful, but when you do it with Python, it accelerates the success immensely. So it's allowing less sophisticated threat actors to use things like Gen AI to create Python based malware, at a much higher success rate. And so I think that's going to, you know, almost lower the bar of entry for threat actor sophistication and if they can start doing leveraging gen AI to create that. I think that's been super fascinating. The second avenue that I'm more worried about than what we've seen, although we've seen some evidence of this is, you know, there's been so many attacks, data leaks, attacks over the last two years as threat actors have moved from, you know, ransomware, pure ransomware to ransomware and data, extortion to now a lot of just data theft with extortion. There's so much data leakage sites use with a lot, a lot of data around organizations and we're starting to see threat actors index that data, and use that data, weaponize that data for secondary and tertiary attacks. And so if you take if you think about all the data that's out there that just talks about business processes, billing schedules, things

that organizations may find benign because it's not regulated data, there's no fines associated with it. But when you think about the implications for threat actors to leverage that data in the way that you described through social engineering, it creates such a compelling, story, for higher fidelity, social engineering attack, leveraging that data. So I think this is going to be a really interesting way for threat actors to basically generate their own lens for organizational data, to use that to be able to, you know, to launch really sophisticated, secondary tertiary attacks against organizations.

Kris Burkhardt [00:09:13] I think you're absolutely right, Rob. I mean, there's so much ability now with AI to learn from data from different voices, from different styles out there, that it can really be tuned to attack an organization.I think it's critical. I think it's really critical that we start to rely on other tools, possibly AI, which I think is where we need to go next, to defend against. That's right. You can't rely as an employee or an individual. You can't just rely strictly on your own interpretation, be it a voice or a video. Even in some cases, I've seen some really good deepfake videos, or certainly text, in this world. So maybe we can transition now a little bit to, the offensive side of things. So, Rob, to that end, maybe you can share some insights, where we are now and you know, what the future looks like from a defensive side. For AI. There's a lot of promises out there. What do you think?

Rob Boyce [00:10:22] There's a lot of promises out there, sure Kris.I truly as I said earlier, I am very optimistic about the opportunity here. I do think, you know, that we all thought as a community we would be , we would have moved a little faster than I think we have right now. I again, I just go back to how tangible ChatGPT made AI feel for everyone that we thought that it was going to move much faster and what we're realizing truly is that there is a foundation of data and infrastructure and, you know, other things that need to be put in place to be able to leverage AI appropriately and that doesn't change for the security space. So I think we're seeing a lot of large companies, Kris, that are making good strides. I don't know if you were at RSA with me or not. I kind of remember our black hat. Every security company is now an AI company, every single one. So I do feel the level of investment is there the level of interest, the level of commitment. But truthfully, still, I'm waiting to see these super tangible outcomes that are have been promised. And I think, again, it's very early days, like we're talking about AI, like it's been around for years. This is what, six months into this journey really in a meaningful way. So we shouldn't be too surprised that we're not maybe seeing super tangible outcomes, but, you know, but the promise is there. I do think the biggest area for me that I'm most excited about is in the security operations space. Because as I was saying previously, this is a space that just has not had a lot of investment in and a long time to change fundamentally how we do this. It's always been about aggregating data, looking for anomalies and responding and I think it's been a very manual process so far.

Kris Burkhardt [00:12:12] Yeah, I think you're right. I think it's got some of the fundamental characteristics that make it, you know, good for you know, AI can see it's got structured data right and in many cases, we've already seen, there's been a little bit of progress there with some of the automation. Right, some of the sort tools. Right. So we've seen some of that. But I agree with you, Rob. I think it's ripe.

Rob Boyce [00:12:38] Yeah and I think even with SOAR, Kris, now that you brought it up, I think the thing that's funny to me is I actually compare SOAR an AI and Gen AI frequently because of the fact that SOAR had the same promise coming into the security operation space and had not the impact that we were hoping and because I truly feel that we underappreciated the level of foundational elements that have to be put in place to be able to do automation in a meaningful way, the process have to be defined. The data has to be in place. You know, being able to even have the interconnectivity to orchestrate all the tools together, I think was a lot of a much bigger challenge in most organizations thought, and this is my same worry with gen AI and AI in general, is that we are under appreciating the level of work that needs to be put in place to be able to make these models, work effectively.
Kris Burkhardt [00:13:33] Yeah, I think you're right, but while the reward is there, if we can do it

right, there's so much demand for, you know, technology. Sorry. Security technology, folks right out there. There's, you know, there's we always hear there's 2 million open roles worldwide. You know how true that is or not, but certainly, we know we're missing people, right? We need more and if AI can help fill some of that, that'd be just of enormous benefit. So I think, you know, you've got the demand, you've got a good space here, so let's hope that some of the big guys can go fill it or maybe there's a little guys are going to go for it, become the next big guy. Right, we'll see. I think it could be really good.

Rob Boyce [00:14:24] Yeah. I mean, just based on the research that we've been doing, you know, we've been able to find that just being able to use gen AI to automate the generation of incident reports. I mean, just that simple task saves so much time from from SOC analysts and that's something that Gen AI is perfect for. Right, and then you can also start using it to look through trends of those different incident responses, like how did we respond and mix the after action a little bit faster. So it's not all about just using AI to find these unknown unknowns and these, you know, one in a million different attacks, but it's truly just to fundamentally automate the routine activities that have been difficult for us to do from a SOC analyst point of view and I think that allows the SOC analysts to also become a little bit less technical. Right now, to be a SOC analyst, you need to known 5 to 15 different technologies to effectively do your job, whether it's a SIM, an endpoint, an orchestration tool, whatever. But if we can use AI to do a lot of that orchestration of technology for us, then the SOC investigators or the SOC analysts become more like traditional intelligence analysts, more like traditional investigators, like almost like a detective being able to ask questions of the tools, get responses back, and issue commands where now, instead of having to know five different programing languages specific to five different technologies, you need to know the English language and the English language can then be able to prompt and control all of that through, you know, just natural language processing.

Kris Burkhardt [00:15:58] Yeah. Everybody can have their very own pocket sherlock. Right? I think it'd be great. You know, one example that we're already doing, when we see an event that's interesting, that we want to go investigate. Right. We have some sort tools that automatically pull all of the information from different sources that would be useful to investigate that event right. Feed a few things. But what those tools don't do is they don't really do any pre analysis on it. They don't do any connecting of the dots. They just kind of here's your raw information dump. Right. And I think I would love to see AI kind of, it's almost the other side of it somewhere you were talking about. It's the early investigation summary, like, hey, here's everything I learned. Here's how it fits together so far. I think there'd be an obvious case for it. I don't know if the tools today can do that, but, I'll just raise my hand and ask for it to whoever might be listening out there. It would be lovely to see and it's just one small area of security.One thing that I also get asked a lot is more on the, I'll say, the prevent side of things like, you know, how is AI can help us prevent attacks and see you know, interesting behaviors on networks before they become a problem, etc.. Are you seeing Rob, are you seeing any tooling out there like that or any, I'll say promise tooling because we know it's not real yet.

Rob Boyce [00:17:34] Yeah, for sure. I think this is going to be, you know, the super interesting area as you can start to almost predict causation of different threats to be able to then, orchestrate protections in an automated way. So I think we are we are hearing that this should happen. But as you can imagine, it's again, much more complicated than probably most people think. To understand the aspect of causation, you know, this threat, this new even if we use, say, the SolarWinds attack that we saw maybe two years ago now, even though I still feel like I live that nightmare every day, being able to understand in an environment, you know, where does this apply and how do we remediate it and how do we orchestrate those remediation in a faster way? I think the opportunity is there. We're hearing, you know, again, some big players in this space talk through that. We're seeing some emerging players come up who are talking about that. Again, I think it's just the fact that, you know, so many organizations have between like 30 and 75 different security tools. Again, it makes it a little bit more complex because of that integrated nature that has to the orchestrated nature that has to take place, but which we're starting to to see some of that, emerge. So I'm again, very excited because I think that's where the true power of this is going to make a huge difference for us, is it's helping with the prevention aspect.

Kris Burkhardt [00:19:02] Well, I'm certainly excited to see that. You know I'll just share one, one final thought with the group, and I said it at the beginning, I'll say it at the end. I do think we're at the top of the hype cycle, but I'm not pessimistic. I remain optimistic about where AI is going. I think we all need to set our own expectations and be a little patient. But if we want it to get where we want it to go, I think we all need to, invest in some way in AI, whether it's by trying it out, whether it's by buying a little product or whether it's by developing a product. If we want to see this go where it really needs to, there's a lot of interest and a lot of momentum and I think it's important that we don't lose that and as we've seen with other promising new technologies, give up on a little bit too fast, right? Because I think we'll get there, but we're not there yet. Rob, how about you? Any closing thoughts for our listeners?

Rob Boyce [00:20:03] So many, you know you just brought up a really good point, Kris, about making sure that people are experimenting, and I think that's super important. You know, we had just finished doing a CEO survey recently and it was like 95 plus percent of CEO's see, you know, gen AI is playing a role in their future strategy, but only, you know, 30% feel like they have the foundation to be able to make it effective or to make it safe and so this is just showing me that the intent of using it is there, but there's a bit of maybe fear, uncertainty of how and I think that's a super important concern to have, because we've already seen incidents where organizations have moved a little too fast and they've lost source code, right. Like they've lost some other IP and I think it's really important that, you know, organizations think about how to build secure, responsible AI and even if you're not building security use cases, security has to be at the heart of everything you're thinking about doing in this data and AI strategy, because, you know, it's so new and another, yet another emerging technology that people are wanting to go so fast, for good purpose. But maybe, they need to also, I mean, they must also consider the security implications of that and make sure they're creating the trustworthy, secure, AI as well. So I think that's going to be super important for companies to understand.

Kris Burkhardt [00:21:28] Yeah, I couldn't agree more like any new tool, just to kind of restate a little bit what you said, like any new tool, we have to be careful with it and we have to treat it securely and the data has to remain secure or that, you know, there's no point in it at all. I will just make our lives worse, not better.

Rob Boyce [00:21:47] Correct. Yeah. I mean, AI is bringing us amazing opportunity, but also introducing new risks that we need to be aware of and I think that's the same with every emerging technology. But I agree with you, I'm super optimistic. I think, you know, one more item that I'll mention and so I think one more interesting, item that we have that talk to you about, Kris, is workforce and workforce. You know, there's a lot of talk about AI is going to take jobs and I think in the security space, I don't well, I mean, I don't think that's true in general, but in the security space, I think what it's going to do instead is it's going to allow people it's going to reduce the barrier to entry in security. I think there's this large misnomer that to be in security, you must be technical. And I think, you know, with technologies like AI, people will be able to reduce that technical aspect and be able to enter into the security workforce, much faster, because I think security to me is, you know, about being logical. It's about, you know, how to solve problems. It's not about all the time, you know, hands on keyboard and I think if you can remove some of that, requirement by using technologies like AI it's going to get a whole different perspective into the security workforce that we have not yet seen and I think it's really going to make a huge, impact in our ability to defend.

Kris Burkhardt [00:23:08] I agree with you. That's a great point. If it can enable people who don't have specific skills but do have pretty good brainpower, you know, take on a role like this,I think that's excellent and I think you're right. What you've made me think of another, another thought as well. I think companies I think sort of the converse, the converse of that is companies that organize their data well and have good clean, hygienic data sets and and pretty good, I'll say, business functional rules in mind I think, will be really well placed to take advantage of AI. Right. That's you know, we've heard it said and I know we Accenture have studied at all levels. You know, you have to have a good data foundation to make this all work the way you want it to and we've kind of spoken about security and security often has a good data foundation set in and of itself. You know, companies around good security have that. But I think, you know, the implication is even broader. So I think it's yet another reason, to get your data organized, if you haven't already, good data governance and all that is just going to become critical to success in an AI world. Well, Rob, thank you for the great conversation. I appreciate it and all of your excellent insights on AI and security. I'm looking forward to our next conversation about it when we're declaring the AI victory and, you know, five years from now, right?

Rob Boyce [00:24:51] That's right. When we have robots fighting robots and our job will be a little different, for sure. Thanks, Kris.

Kris Burkhardt [00:24:59] All right. Thanks.