CTL-CISO-Josh-Copeland

Sat, Sep 27, 2025 12:00PM 15:36

SUMMARY KEYWORDS

Cybersecurity, AI adoption, security controls, hypothesis testing, knowledge synthesis, decision making, creativity catalyst, shadow AI, privacy impact, data classification, synthetic data, AI governance, incident response, cross-functional councils.

SPEAKERS

Josh Copeland, Jo Peterson

J Jo Peterson 00:00

Jo, thank you so much for joining. We're here today with clear tech loop. We're on the move and in the know, and I've got a wonderful guest today. Josh Copeland, Hi, Josh.

J Josh Copeland 00:17

Hi Joe. How are you today?

J Jo Peterson 00:19

I'm doing great. Thank you for making time to visit.

J Josh Copeland 00:22

Yeah, absolutely,

J Jo Peterson 00:23

yeah. Appreciate it. In case you guys don't know Josh, he is the director of cybersecurity at crescendo. Josh is an Air Force veteran like me. Thank you for your service, Josh.

J Josh Copeland 00:35

Thank you for yours.

J

Jo Peterson 00:38

Josh has 20 years of cybersecurity experience y'all, and he holds a Master's in cyber security management from Tulane, and he's an adjunct professor there. So how is it? How is it teaching young cyber minds? Josh,

J

Josh Copeland 00:53

it definitely keeps me engaged and keeps me looking at things from a very different set of lenses. It's always interesting to see how folks who truly grew up in full internet, you know, high speed broadband, native folks, view problems compared to, you know, the folks like me who have, you know, a lot of gray in their beard and, you know, were around when it was dial up, and how we look at problems differently, and how those different perspectives lead to different outcomes.

J

Jo Peterson 01:23

You know, it's so funny that you said dial up, because can you ever get that sound out of your mind?

J

Josh Copeland 01:29

No, not at all.

J

Jo Peterson 01:29

You can't unhear it, right?

J

Josh Copeland 01:31

It lives rent free in my head for the rest of my life.

J

Jo Peterson 01:34

How about it? And then, how many nights did you come home from work and there was an AOL? Dial up, um, CD on your doorstep.

J

Josh Copeland 01:45

I'm old enough to remember getting the floppy disc in the mail.

J Jo Peterson 01:48

Okay, that's, that's great. Yeah, they were, they were working hard. There was a period of time over that one year or so that maybe, like, three nights a week I'd come home and find one was, was great so, but we're revisiting a different time, because people are like, Why AOL, dial up. What's that? So one, just get right into the questions here for you, because I'd love to hear your take on these things. First one, how can cybersecurity professionals leverage gendered AI to break out of that sort of traditional tools and tech mindset and drive more innovative thinking and execution in their security programs?

J Josh Copeland 02:33

I kind of look at this where there's four real ways that you can really leverage AI from a cybersecurity perspective, to really accelerate your cybersecurity program. One is accelerating hypothesis testing using AI to create and model attack scenarios, create synthetic malware variants, simulate adverse, adversarial TTPs, and really allow you to get that attacker mindset that you're not going to really know what's happening next, because it's not your experience. It is tied from all the experience everywhere in the entire world. A second kind of silo from that is using knowledge synthesis and siloed expertise. You know, unfortunately, in cybersecurity, we still have a lot of silos. There's still a lot of I have an engineer who does this thing. I have an engineer who does this thing. They don't really talk to each other, and using AI just to take documents, logs, threat intelligence reports, and build that together and produce kind of a holistic view from a enterprise level, is a great way to get some kind of contextualized information. Then you take that, move it into contextualized decision making support. You pair Gen AI with your risk data, your security leadership teams can receive kind of a narrative justification for their security decisions. You know, why did I do this? What's the real root cause? What's the justification? You know, you're shifting from what does the tool say, to, why is this the best course of action, and what data supports that? And then kind of the final one is just having a creativity catalyst, you know, being able to generate playbooks, tabletop exercises, draft policies, you know, get kind of those best practices that you spend hours and hours scouring the internet for. Of you know, I need to do a tabletop exercise for something, and I want a scenario I haven't used before, right? Really great to go on to a Gemini or a chat GPT, and kind of give it, give me a scenario for a tabletop exercise that will last about this long with these pieces, my tool stack, and really kind of treat Gen AI as a co strategist rather than just another tool in your stack.

J Jo Peterson 04:46

Yeah, and you that's, I love the last one the best, because then it feels like you're getting the full benefit of AI. So next question, I think, is totally. Partner for you, given that you work for an AI forward company, I'm really eager to hear your commentary on this. How can organizations embed security and privacy controls into AI model development without slowing down innovation?

J Josh Copeland 05:18

So allow that, and you're going to hear some things that are seem like no brainers, but they're truly critical pieces of you know, shift left security and privacy. You know, build it with that in mind. You know, embedded into your threat modeling, your data classification, your privacy impact assessments, and use not only the AI to help you get that in early, but also as validation as you go through. You know, it's easy to fix things early than it is. Try to bolt security on after the fact. Build your guardrails. Know what your requirements are and have the appropriate things around it. You know, a lot of organizations now have folks that are, well, say, is going rogue and using, you know, Gemini or Chachi BT or Claude, and they're not using a sanitized version of that, right? So if you want to be a for looking company, buy the tool, get the appropriate enterprise license, get the zata zero data retention agreement with them. You know, this way you can actually do some DLP related to it. You know, if you're a Google stack company and you have Gemini purchased, have your folks use that, you can do things like, Do you know DLP data leak or data loss prevention, identify when they're putting things that they shouldn't be putting in there. Build appropriate education around that. One of the great things that you can do with AI is get synthetic data and use differential data for privacy. You know, let your engineers use data sets that are nearly identical to what the real data set is, but it's synthetic data that's been created by your AI, so you get the look, the feel, and all of the advantages of real data, but none of the worry that I'm now putting legitimate data in a pre production system and then really looking at just your model problems and audibility. You know, know what models you're using, why you're using them, and track, you know, who's using what model for what approach. And that allows you to kind of rapidly expand what you're doing. Because I can look and say, I'm using this model for this capability, and now I have somebody else in another silo in my organization that's doing something similar. So I can say, based off of this, I can do this. So it's kind of that mind shift, shift of mindset where, you know, cybersecurity isn't the gate at the end. It's kind of an invisible exoskeleton supporting the rapid expel experimentation around all the things that you do.

J

Jo Peterson 07:58

I like that answer. That's a great answer. So kind of a curveball question on a scale of one to 10, one being not a problem at all, 10 being it's, it's like, you know, top of my list problem shadow AI.

J

Josh Copeland 08:16

So shadow AI, to me, would probably be a solid eight for most organizations, okay, I can't control what I don't know I have. And the fact of the matter is, everybody's using AI one way or another. If you would informally pull your users in your company, they're all going to say that they've used chat GPT, that they're going to use Gemini, that they use Claude or dolly or big journey, insert whichever one you want to use here to do something, and then ask them if you've done anything related to work. And they're going to go, Well, yeah, I used it to help clean up a little bit of my email or something like that. So now you have risks that are associated with that, and it's far better to get in front of it. Give them the tools that you want them to have. Give them training on how to use AI appropriately. Give them guardrails to say, here are the type of things that you're absolutely allowed to use AI for. You want to, you know, work on your grammar. You want to have it generate ideas for you. You want it to go ahead and do quality assurance capabilities, yes, oh, you want to do personnel reviews, or you want it to use it in hiring. No, that's absolutely not going to happen. You want to put phi or PII or credit card information in No, that's a red flag offense. We're going to fire you on the spot if we catch you doing that, but giving them those guardrails up front and having the tooling around it that they're not going around you to use the things like anything else in security, the more friction you put into using something, the likelihood of someone going around that. Is going to happen. I'll show my age and go back when Microsoft Exchange first put in the ability to do file blocking, and you would block zip files. It took all about 24 hours before someone figured that they could go I can change it to a dot txt file and just tell the person on the other end to change it back to a zip, and now I've lost that security capability, right? People are smart, and then, on the other hand, we're also extremely lazy, so we want to do things the quickest, fastest way possible. So finding that happy medium between the level of security that we need, but also having the least amount of friction for the user, is really that kind of sweet spot you want to get to when it comes to building approaches around AI.

J

Jo Peterson 10:47

That's a really solid, fair answer, and it just makes me think I need to have you back, because I want to, you know, dig into AI security frameworks versus traditional cybersecurity frameworks, because there's differences, right? And I think that's a whole nother conversation, so I hope you'll agree to come back and talk about that, because I think it's super relevant to what folks are experiencing right now.

J

Josh Copeland 11:16

Yeah, absolutely. Anytime you want me at your disposal.

J

Jo Peterson 11:20

Well, lovely, because I think there's a lot to dig into and that folks would benefit from your thinking. Let me ask you the last question, and I feel like this is a balanced question. So you tell me, how should the CISO be thinking about AI adoption in terms of its use to secure emerging threats, but also from an organizational governance perspective.

J

Josh Copeland 11:44

So that's a great and fair question, so I'll kind of break it down in a couple different ways. One, you know, look at it from a strategic dual lenses. You have to position AI both as an asset to protect and something that you can wield within your organization. So you have to look at it. What risk do you have? What type of integrity that you need to have for it? Is there a chance for data poisoning? Is there a chance for injection prompt? And you need to enter, kind of integrate that into your Enterprise Risk Register on both sides of the house. And you kind of look at that and look at and go, you know, I can absolutely use AI driven security operations, you know, do anomaly detections, automated triage, incident summarization, you know, to free your humans. You know, those, those big, smart, wet brains that we have to do the really high order stuff that we really want them focused on, you know, having your Sim fire off, you know, 100 false alerts is great because you wanted to fire off enough alerts that you catch the real thing, but there's always going to be those false alerts. Using AI to help reduce that noise. Great capability, and then pair those deployments with human in the loop to avoid, you know, automation biases. You know, ultimately, the big thing here is you always want to put the human in the loop somewhere. Because I like to tell folks, you know, AI will not testify for you in court, nor will AI serve your jail sentence. So you got to put a human somewhere in there. Because ultimately, if you try to give it all the control you're still ultimately responsible for, that which kind of leads in that governments and accountability establish cross functional AI councils include legal ethics, security and, most importantly, your business unit leadership, because they're the ones that are going to want to do all these cool things. Says cybersecurity has a unfortunate reputation as being the Department of No, we're going to say, No, you can't do that. It's not secure. And we have to change that thinking, because if we are the Department of No, they're just going to go do the things without us, and then we're going to find out after the bad thing happens. So go with them. Define what your explainability threshold is. You know what? What do you want to do? Why do you want to do it? And What advantages does that provide to us as a business? And how can I help you secure that in a way that not only do you get benefit from your line of business, but we have some reasonable assurances that it's not going to do anything crazy and then kind of do that total resilience thing, you know, if you assume the compromise of an AI system and build contingency plans, you know, what are our fallback models? Do we have kill switches? Do we have retraining protocols? Do you have an AI instant response book? You know? Have you figured out? What do you do if you integrate AI into your key business processes? What happens if you have to turn the switch off? So AI, adoption is not about plugging a model in. It's about governing an emerging capability that touches potentially every part of your enterprise and your enterprise's risk surface.

J

Jo Peterson 14:56

Man, y'all heard it here first. What a light bulb. Thing to say, AI Incident Response book, like, like, if y'all aren't thinking about that yet, you should be. That's a great piece of advice. Josh, thank you for that, and thank you for taking time to visit with us today. And let's do have you come back and let's just talk frameworks, because so important, so thanks. Thanks for the time.

15:23

Yeah, thank you so much. Jill.

15:25

You.