CTL-CISO-Joanna-Wiggum

Tue, Sep 08, 2026 9:02AM 12:07

SUMMARY KEYWORDS

AI security, zero trust, cyber attacks, social engineering, cyber crime investigations, AI governance, least privilege access, red teams, multi-factor authentication, incident response, outsourcing, cyber policy enforcement, AI agents, cyber threats, cyber intelligence.

SPEAKERS

Joanna Wiggum, Jo Peterson

J Jo Peterson 00:07

Hey everyone! Thank you so much for joining this week's episode of Clear Tech Loop. I'm Jo Peterson. I'm the CIO of Clarify 360, and I'm the chief analyst at Clear Tech Research. And I've got an amazing guest today. I've got Joanna-Wiggum, who is the founder and agency principal of Countervail Intelligence. Hi, Joanna. Hi. Thank you for making time.

J Joanna Wiggum 00:29

Hi.

J Jo Peterson 00:29

Thank you for making time.

J Joanna Wiggum 00:30

Thank you for having me.

Jo Peterson 00:32

And thank you for your service, guys. Before we started recording, I found out that Joanna is an Air Force vet just like me. So there we go.

Joanna Wiggum 00:41

Thank you for your service.

Jo Peterson 00:42

Oh, thanks. So, Joanna, before we get launched into our questions about AI security, you found it Countervail. Tell me why you founded it and what business problem you're trying to solve, as well as the URL where folks can find you if they want to visit.

Joanna Wiggum 01:01

Thank you very much. My tagline for Countervail is enterprise-grade cybersecurity made personal, and the idea was actually the fact that most cyber attacks are psychological. Two thirds, approximately two thirds of cyber attacks are all social engineering based, and I thought if I could solve that problem either through proactive review of programs, policies, and things like tabletops, exercises, or through training, workshops, and general advice, then that problem could be reduced. And specifically, the way I'm tackling it is teaching people how to think like the adversary. So to understand things like your social media posts are informing the people who are using that information to target you, you know, posting something like working from home today sounds like it's no big deal, but if you've already been on the radar because you happen to hold a position of privileged access, even if you're a help desk worker, that can potentially put you and your business or your whole company, whatever that is, at risk, and especially given the the landscape in cyber right now, with multiple wars that are you know ongoing or not ongoing, and we have advanced persistent threats that are fighting in cyberspace. With that escalating, it felt like the most important thing to be doing right now. And the other part of my business is doing cyber crime investigations, where people who get defrauded, scammed, doxxed, whatever it is, harassed online often feel like there's nothing they can do because when they report it to law enforcement, which is really most what everybody knows to do, most of those departments don't have a cyber team. Most of them can't investigate. So I offer a service where I do the investigation. I package it and I give it to law enforcement or refer to legal counsel, depending on the situation.

Jo Peterson 02:45

That's super interesting. That is, you're right. What do folks do? So it's scary, right? And that's it's much needed for sure.

Joanna Wiggum 02:57

Yeah. Oh, and they can find me at countervailinigence.com.

Jo Peterson 03:00

That's what I was going to ask. So thank you. Perfect. All right. So let's launch into our question. As you know, everyone knows it watches three questions, and and we're out. So first question: AI agents hold delegated credentials. They make authorization decisions in real time, and they can spawn sub agents with their own permission sets. Scary traditional IM and PAN tools weren't built for this. So, what does a realistic path to least privilege access for AI agents actually look like for a large enterprise today?

Joanna Wiggum 03:35

Zero trust, actual zero trust. So that means that agents cannot delegate their own authority, so you need to have an outside agent that has only this one task that's able to delegate, while the other agents cannot create their own credentials. They have to get permission from exterior, and that's in general how AI should be built of like independent, isolated machines or agents that operate on a certain set of tasks, so I say that, and then I'm gonna I'm gonna actually give the honest truth about it, though, which is even if you give zero trust, it is not a guarantee that you're not gonna have a massive situation on your hands that ends up being a major problem, and the reason is because in most security programs, and I'm leveraging my experience of of running red teams for big tech for enterprises as well, and understanding that most of the time people haven't done the basics, and I mean basics like they don't store passwords in a secure way. Often they'll put them in share drives, or they'll reuse passwords, or they won't enable multi-factor authentication everywhere. And so when you have credentialed problems like that that are just prevalent in every organization, and you have agents that are independently thinking and trying to solve problems. You're going to find that they're going to leverage those low-hanging fruit basics haven't been covered to achieve the objective that you've given them, and they don't keep guardrails on themselves like a person might because of moral issues or thinking because of you know. Community problems, or they want to keep their job, that that doesn't occur to an AI agent. So there is a brewing problem as as folks still haven't laid the foundations of good security, and then they put AI on top of that, and you're now just creating more problems than you're solving.

Jo Peterson 05:16

You make such a good point because my mind went to the recent anthropic sandbox issue, and I'd read somewhere. Don't know if I 100% got it right, but I'd read that there were 17,000 tries over that weekend. So the AI agent is not bad. It's just trying to get its job done. It's just doing whatever it can do to get its job done, right? 17,000.

Joanna Wiggum 05:44

Yep, I think

Jo Peterson 05:46

the kind way to put it

Joanna Wiggum 05:47

is like it's a like a dog with a bone, right? It's just going to do the same thing. I think the the more nuanced way is to say it has a personality and it's a narcissist. It's only going to think about itself and what it wants to achieve.

Jo Peterson 06:00

Oh, I hadn't thought about it with a personality before. That's that's gold right there. Okay, love that. And you said something else funny before we went on camera. What was it? Share it.

Joanna Wiggum 06:12

Oh gosh. Well, not all cybercriminals are super intelligence. Sometimes they eat the crayons instead of coloring with them, and it's really fun to come across those.

Jo Peterson 06:22

Yeah, that's great. I love that. Thank you for the color there. So, number question number two: 50% of large enterprises now have dedicated AI governance committees. Is that structure actually working, or is it creating the appearance of governance without the substance?

Joanna Wiggum 06:43

I again, this is experiential, and my experience doesn't necessarily reflect the the ecosystem as a whole. But in my experience, it's more of a appearance thing than it is actually doing in solving the problems. Just like policies in general, if you write a policy, it doesn't mean that the policy is being enforced or that it covers everything that it needs to cover. Just because you've dedicated a headcount to something doesn't mean that you're actually resolving the issue that you dedicated the headcount to. I think that every industry that's that's going in this direction, every company organization that is deciding to embrace AI needs to really think clearly about what the problem is that they're trying to solve, and then analyze if AI is going to solve that problem at all. And then if they decide that yes, we have a problem that you know, like let's say the problem is scale. They're building something quickly. They want to have AI assist them in building that because they cannot resource it with humans. Okay, that's fine. So for example, incident response would be a really good example of that. It is really hard to staff a soc. It just is, and when you do staff it, it doesn't mean that like you have people who get tired. You have to do this 24/7 cycle. People miss things, and if you can scale better, especially when you're combating AI with AI. So, okay, you have this problem. You need to staff your SOC. You want to staff it in part with AI, keeping a human in the loop, which is very, very necessary. But okay, you want to solve this problem. Well, now you have the issue of implementation, and as we've talked about, because that low hanging fruit has already not been established, building AI policies on top of already poor policies is not going to help you. Just because you dedicate half of your team to governance doesn't mean that they're going to effectively execute on that. So I think that we need to really get down to bare bones, hit the basics, make sure that you have something that's already scalable. Like how are you already writing your policies, and how are you checking that those policies are being enforced? And then if you've already solved that problem, now dedicating to building upon that is great. But if you haven't solved that base layer, you're you know just dedicating people to something doesn't mean that you're going to have the right answer at the end.

Jo Peterson 08:49

Yeah, that's fair. Third question: Yeah, CISO in this audience is starting from zero on AI agent governance. That means no inventory, no agent-specific controls, no monitoring. What is the single most thing that you would tell him or her to do in the next 30 days?

Joanna Wiggum 09:09

Outsource. I wouldn't try to teach yourself. Here's a.

Jo Peterson 09:19

I'm sorry, that was funny. You didn't even hesitate. You're like your goose is cooked outsource.

Joanna Wiggum 09:25

You missed it. Outsource. There are certain things that every individual is going to be an expert at, and if you're already a CISO, you have established yourself in information security. And if you miss the bus to understand, or or your organization has missed the bus to to have AI implemented, I would strongly recommend that you outsource to a company that is already showing that they're doing this well. And I would pick one that that is well established. You know, I I wouldn't necessarily discount a startup, but I would say that look to the folks who have already proven that they're doing this well, and they put good guardrails. In place, and then leverage their experience to build it.

Jo Peterson 10:05

Yeah, and even if you're just stepping into the role, you just got there, right? And you're looking around, and you're like, "Oh, this is a problem, right? Up online. To your point, you're already behind the eight ball, so it doesn't even say anything about you not having a skill set, not knowing what to. You just don't have time. Time is your problem.

Joanna Wiggum 10:26

Yeah.

Jo Peterson 10:27

Right. So that's really, I think that that's fair advice in certain situations. And you find yourself just dropped in somewhere, you know, and you could already be behind.

Joanna Wiggum 10:38

Yeah. So so the alternative is, you know, because typically the resource demands are time and money, and you have to decide if your time is worth it. I'm-I'll use a personal example. When I launched my company, I decided I was going to teach myself to build my own site. It took me 150 hours to get it to where I thought it was like doable. So, and at the end of that, I looked at okay, what what is my time worth, and what was the cost and sweat equity that generated that? When I could have just paid somebody who's already an expert and spent all that time in generating an actual product that would have served my customers better, rather than creating a website, which really was so. I just I think that in general that applies to most people in most roles is you have to decide what you're going to be an expert at, and then outsource where you're not already an expert.

Jo Peterson 11:25

Yeah, that's that's super fair, and and doesn't say anything about your lack of skill or it's just a time situation. Yeah, so so lovely talking to you. So my perspective, it was fresh and clever, too. So yeah. So definitely have you back another time if you're open to it.

Joanna Wiggum 11:48

Oh, I would love to come back any time to talk to you. Any excuse.

Jo Peterson 11:52

All right. Great. Well, guys, thank you so much for joining, and we'll catch you next time.