CTL-CISO-Matt-Sharp

Sat, Apr 04, 2026 3:19PM 18:16

SUMMARY KEYWORDS

AI governance, cybersecurity, shadow AI, workforce productivity, non-human identities, role-based access, agentic experiences, AI defense, SaaS platforms, third-party risk, data collection, collaboration tools, asset management, cloud management, AI adoption.

SPEAKERS

Jo Peterson, Matt Sharp

J **Jo Peterson** 00:07

Thank you so much for joining today. This is clear tech loop. We're on in the know. I'm Jo Peterson. I'm the CIO of Clarify360 and the chief analyst at ClearTech Research. And I've got a treat for you today. I've got Matt sharp, who is the CISO at Xactly Corp. Hi Matt.

M **Matt Sharp** 00:26

Hey, Jo, excited to be here. Thanks for having us.

Jo Peterson 00:29

Thank you for coming. And if you all aren't following Matt right now, you should be, because Matt is a cybersecurity executive who helps software companies grow faster by building enterprise trust, especially as AI becomes core to the product. His focus is security and AI governance that accelerate delivery, strengthen resilience and reduce risk in ways boards and buyers really care about. A little bit about Matt and his background, he spent more than 19 years leading security organizations, and that's across public, private and VCPE backed companies, partnering closely with CEOs, product engineering and revenue teams. He's known for translating complex, complex risk into clear decisions, building scalable operating models and delivering measurable outcomes outside of his operator role. He served as a venture advisor to the Jo y l ventures. He advises founders on go to market strategy, messaging and aligning with buyer expectations. He has also served on advisory boards for Okta f5 Vera code, coal fired, deep watch, trust visor, trust cyber, PRX, NOP, second, others, and he wrote a book, The CISO evolution business knowledge for cybersecurity executives. Yeah, he's got it behind him, so y'all go check it out. And in case you're new to clear tech loop, we're a hot take approach. So that means we ask three questions that are meant to quickly educate our listeners about the AI landscape. And you know that I focus on AI security, so let's get going with the first question to Matt. Matt, give me your thinking around shadow. AI, is it an IT problem, a security problem, both neither. And how are CISOs and CIOs addressing shadow AI in their environments?

M Matt Sharp 02:29

Yeah, great. Two parts. First part, is it an IT problem or a security problem? I think it's both. What we have seen sort of firsthand is interesting. We've seen firsthand that, you know, we have a very strong mandate to deploy AI from a workforce productivity perspective, within the entire VISTA portfolio, but in particular within our company at exactly and what that means is people want to experiment and with some of the new technology, and we want to both make sure that the prioritized investments from a corporate IT perspective are selected and actually leveraged. So we have a series of cohorts of AI tools that are intended to improve workforce productivity. And that suite of tools, of course, has competition, and so one of the things to do to optimize that investment is to make sure that people are using the tools that have been properly, like vetted and sanctioned. The other thing is, if you're going to spend a bunch of money on licenses, whether it's, you know, there's a whole bunch of them, auto bound or Grammarly, or Reggie, AI, or, you know, perplexity, GPT, Gemini, whatever the thing is like, in addition to just making good use of the software investment and spend that you have. You want to make sure that the data you're providing to these different AI providers isn't being trained on and there's a myriad other risks that go along with that. So I would say both, interestingly, so 18 or 24 months ago, we started to sort of tackle this. We found a fantastic partner in Sequoia, backed company that I think is still in stealth, and so our approach has been to leverage them in order to help us. And some of the workflows that we've actually been co designing have included slack oriented, large language model collection, data collection, with individual stakeholders to say, Hey, why are you using this tool? Like just answering that simple question and taking that from like using an agentic approach to take that question off the table, I call that, and then centralize it, and then have some. Classification methodology really accelerates the sort of shortcut to the answer, and it turns out the same question can satisfy both cyber teams and corporate IT teams. So I would say in the context of shadow AI, it's really both. It's a co owned problem for both Mark and myself,

J Jo Peterson 05:23

yeah, and I can see that, I feel like it's a co own problem, but maybe the sunny side of that is, maybe it's getting CIOs and CISOs to have more conversations.

M Matt Sharp 05:35

Yeah, well, yeah, we, I mean, we, you know, we have a lot of conversation already.

J Jo Peterson 05:41

Ya'all do, but not everybody, right?

M Matt Sharp 05:43

Yeah, but that's right. And I do think tendency is, oftentimes for cyber teams to want to own their technology

J Jo Peterson 05:53

Yes

M Matt Sharp 05:54

stack and not to sort of share and so I think one of the things that we've done at exactly that I think has been really quite appreciated, is like we have a an asset management tool, cyber asset management tool called lucidum, that can provide penetrating insights across our environment, just from an asset perspective, and we've shared that with our corporate IT team, and we've got wiz on the cloud management plane, and We've actually, when we came in, part of our success criteria was to take wiz foundations training, and we have 15 or 20 people. Now, my team is not 15 or 20 strong, right? But so we have stakeholders across the business that have actually taken it and can leverage the insights from these cyber tools and so, and that pattern kind of repeats itself, you know, DLP, email security and on down the list. So we've been very collaborative, I think, in that regard, with both our CTO and our CIO.

J Jo Peterson 06:50

I like that. That's a cool approach. So I am just so intrigued. I don't know if you are, but I'm so intrigued with non human identities and caring for them from a cyber perspective, as well as from just a basic inventory perspective. As you talk to other CISOs and other CIOs, what are some of the ways you're seeing those executives enable NHIS?

M Matt Sharp 07:21

Well, I mean, I think this is a long standing problem, right, where you have customer identity, you have the back office, like, all of the cloud identities and API keys, and you have the SSH and PKI challenge and so, like, this is just another iteration on it. And the, I think, I think, from my perspective, what has happened in the broader market has been, we have, for the last 20 years, failed to get Role Based Access Control to work at scale. And we've tried, and we've gotten better, for sure, and we've layered in SSO and MFA, and we've gotten the authentication piece by and large, I think correct, but the authorization piece is very challenging, and there's a whole bunch of reasons why that's challenging, and now the assertions that we make when we give agents access to do things basically say, well, you're going to inherit or adopt the broken role based access paradigm that existed before and and so I think that's been a real challenge. And I think kind of, as we've seen the big push towards agentic, what we've seen is this forced acknowledgement of the reality that role based access was never actually implemented correctly in the first place, yeah, and so I think, like, if you look at the evolution, it was, hey, large language models are here. We embed them into our product. Now we're going to have automated workflows. We're going to marry the two, and we're going to call that agentic, and these businesses are going to start taking actions. I think a lot of people are sort of slowly dipping their toe into how aggressively can I allow these things to take actions? On one hand, the business wants the upside of like cost compression and efficiency, and on the other hand, the business is afraid of those agents doing stupid things, because there's a lot of evidence that they will do stupid things. And so my take is, my take is that from a non human identity perspective, there's probably, like five or six different technical approaches, and it really depends on what you're authorizing and approving from how you're allowing AI to be embedded into workforce tools and also into the value delivery mechanisms in your business. So what I mean by that is we're a SaaS provider, but not everyone writes software for a living, and part of the value that we deliver to customers is agentic experiences for sales, performance management, um. Um. And so we need to be very careful about what the agents in our software can actually do and can't do. And I think the real point where this non human identity stuff becomes dangerous is about to happen. It's on. We're on the inflection point. And I think people can see it's not about being able to cleanly cascade, permissions from an initiating human's perspective, it's when these start to be initiated, instantiated and action autonomously without a per like a designated human's identity behind it. Once you have that sort of agent to agent without the ability to cascade a human identity permission schema behind it. The question is, well, so now we've sort of unhinged ourselves from the from the grid, or from the I think about pain right at some point, without the plane, you can't tell where you're at, and I think we're about to remove the Cartesian plane of identity, and we're just going to be floating in the ether of identity. And I think that's the thing that scares people.

J Jo Peterson 11:08

Yeah, it scares me too. And I think, and I think about it, but I thought about it. I've thought about a different way, I guess I've thought, and I like the way you framed it. That was great. I guess I think about agents creating agents right, and that ephemeral agent, as it were, out there, and if we're giving permissions to agent a and Agent A is giving permissions to agent B, how do we know when those permissions are turned on or turned off, or are they turned off? Are they just floating right? Also, all those things sort of make me go, you know, there was, there was a phrase in the 90s, things that make you go, Hmm, well, I think, I think it's back.

M Matt Sharp 11:51

Yeah, I think so, yeah, where's the band, the band? I feel like that's Millie Vanilli. But maybe that's not right.

J Jo Peterson 12:01

I don't know. I don't know who's saying that. That's great. So when you hear the term AI defense, what comes to mind for you?

M Matt Sharp 12:11

Oh, my God, it's so deep. So first, so this is the way that I think about the world. Somebody says, AI. I think to myself, are we talking about machine learning AI, or are we talking about Gen AI, right? And then the next thing that I do is I say, well, in what context? What's the surrounding context? And so for me, there's really only two contexts which we've touched on. Both it's value creation or it's workforce productivity, those are the things, right? And then once I start to sort of dissect from there, my mentality says, like, let's go down the workforce productivity, because that's where we started. There's open source, there's plugins, there's IDEs and plugins and IDEs, like cursor and the like. Then you have a handful of things, like, a handful of users have deployed a llama on their local workstations and are hosting the large and small language models locally. And then we have all of the things embedded in our SaaS platforms, so all of the agentic experiences that we're creating so many of our competitors and obviously the broader landscape are producing. And then that leans into more of a third party risk perspective. And then you start to think about MCP as sort of taking center stage. It's a lot. And then, not only that, but like you think about the back office or corporate IT, we're starting internally. And I think this pattern sort of replicates across many CIO orgs. You have backhauled A bunch of, we'll say, core system information, CRM, ERP, whatever that data gets centrally hosted in a data lake, presumably. And then you overlay large language model to interact with the structured and unstructured data. And then you connect it to a meaningful collaboration interface, like teams or slack or something like that, right? Like that. Pattern is being repeated pervasively, I think across all people, because I think cioc, there are a handful of agentic investments that are going to be useful, but if they can build and centralize maybe 70 or 80% of it, then why pay for every vendor to reinvent the wheel, right? Um, yeah, so I think, like, rapidly, depending on which are we, are we pulling on the browser extension, the IDE, the open source piece, the SAS piece, the like, there's a whole, there's a whole myriad number. And guess what my budget like, when they when, when, when the team said, Hey, rapidly adopt AI. They didn't say, and, oh, by the way, here are, here's budget for nine new vendors, each of which wants a, you know, whatever the number is, right?

J Jo Peterson 14:51

And you, you drove the bus exactly where I was hoping you would. Because this is definitely one aspect of it. I feel like we're going to enter into. To another platform phase, right? And people like you and I that have limited budgets are going to go, Hmm, where can I get the most bang for my buck?

M Matt Sharp 15:12

Yeah, I think, like, naturally, we're going to have

J Jo Peterson 15:14

to right. We're going to have to right. We're going to go, whoa. Well, this guy's got this and this guy's got this, but who's got, like, the most of it, whatever it is, right? So it'll be interesting kind of see the race for dollars sort of happen. And I think we're going to see it over the next 12 months. What do you think?

M Matt Sharp 15:33

I think it takes longer? I think there's a if you look at the amount of capital that's deployed in this AI space, it's a ton. And if you look at the number of innovative, disruptive vendors, there's a lot, because vibe code is another thing we haven't even talked about. Vibe code created the ability to rapidly get to working software. And I think the long term profitability of those vibe coded entities providing immediate value to customers becomes very difficult. I think it's a hard slog once you've got the the code to actually maintain that in the long term. And so I think I don't know if all investors have that clear. And so I think there's a huge amount of capital that's been deployed. So we're going to see, we're going to see a handful of folks gain a number of customers and then get sold as features to broader platforms like Paulo and Okta and whatever. And then, I think separately, there's going to be a handful of folks who secure early revenue and then implode because they can't. They just don't have the maturity to outstrip their, their their runway, their their cash burn. And then I think, I think there's probably one more category in here. And I lost my train of thought, but I think that 18 months is too fast to consult

J Jo Peterson 17:01

maybe, maybe, I mean, I'm just probably, you're right. It took a longer, you know, I started, I started architecting cloud workloads in 2009 and I'm amazed at even today, right? So all these years later, how many people are still on prem right? I mean, it's just lifted and shifted to the cloud. Yeah, right. So all that, so you're right, the window is probably longer than I'm thinking. I'm just probably being too excited about what's coming maybe. So anyhow, thank you for taking time to visit with us today and share your thoughts. I'm sure it's going to resonate with a lot of the audience that we've we've had join us, so I hope to have you back on the program at some point.

M Matt Sharp 17:51

Yeah, that would be great. It was a fun chat. I appreciate the, you know, stimulating conversation. And yeah, folks, I guess, find me on LinkedIn. Happy to connect and continue the conversation there as well. Awesome. Thank you. Matthew, bye. You.