CTL-CISO-Derek-Fisher

Tue, Jun 30, 2026 7:48AM 14:07

SUMMARY KEYWORDS

AI governance, risk management, regulatory compliance, cybersecurity education, incident response, shadow AI, AI governance board, executive leadership, product owners, role-based access control, third-party management, AI authenticity, security servers, cognitive load, burnout reduction.

SPEAKERS

Derek Fisher, Jo Peterson

J **Jo Peterson** 00:07

Hey everyone, thank you so much for joining today's edition of ClearTech Loop. I am Jo Peterson, I'm the CIO of Clarify360 and the chief analyst at ClearTech Research, and I'm here today with Derek Fisher, founder of Securely Built. Hi, Derek.

D **Derek Fisher** 00:23

Hi, how you doing? Thanks for having me on, Jo.

J **Jo Peterson** 00:26

Thanks for taking time to visit. In case you're not following Derek, you should. Derek and his company provide advisory services to mid-market firms for strategic planning, risk management, and regulatory compliance. He designs and implements security programs, including Vol management, secure SDLC, and incident response, in order to help firms strengthen their resilience, and he delivers cybersecurity education and training awareness to advance from, I should say, education and training to advanced technical skills with over 30 years of cybersecurity experience. Derek serves as director of Temple University's Cyber Defense and Information Assurance program. So, welcome, Derek.

D Derek Fisher 01:15

Hey, thank you very much.

J Jo Peterson 01:17

As you guys know, we do three rapid fire questions that are of the moment that security professionals are thinking about, so let's not deviate from the norm. Let's go on with question one. Derek, how do we operationalize AI governance, and who is legally accountable when an AI agent makes an unauthorized decision?

D Derek Fisher 01:42

So that's I mean, that's tricky, because I think one of the biggest challenges is that a lot of organizations don't know where the AI is, where AI is being used, so there's, you know, whether it's being used within sanction platforms, Salesforce, if you're creating workflows within there to automate and create features, and that's great, but then there's also, you know, shadow AI, where people are able to, whether you're going directly to Chat GPT or Claude, or you're creating your own local, you know, agents, and you know, utilizing local LLMs, I mean, this, this is one of those things where it's very difficult to put the toothpaste back in the, in the bottle, right, or the tube, because once these tools are in the hands of everybody, it's very hard to get a grip on where it's being used, how it's being used, and more importantly, how is the company data being used in those systems, you know, one of the podcasts I was listening to recently, I think it was yesterday, they were talking about how, you know, when we had these massive changes in technology, like let's take nuclear technology as an example here, you know, as a collective, as a society, we all decided, like, hey, this is very powerful technology, we really need to get our hands around and make sure that it's not, you know, nobody can build a nuclear reactor in their home, right. And so that's a little different now, because we have a technology that could and could have very far-reaching implications, and yet it's pretty much in the hands of everybody. And so being able to get your hands around that and be able to make sure that we understand how decisions are being made and how we can get, you know, the appropriate controls around it can be difficult, but what that really starts with is having an AI governance board that's in that is cross cutting, so you don't want to just have it in the technology side, but you also needed, you know, members of that governance board that are part of HR, that are part of legal compliance, and have that, like, cross-cutting across the organization visibility into, you know, what is it the organization really, you know, what are the impacts of this technology. Above that, AI governance is really an executive leadership now depends on the organization, who that executive or what that executive leadership looks like, but they should be the ones that are in control of the actual risk appetite at the organization as it relates to AI. So the executive at the top owns the risk appetite, the governance board below that really decides the standards and, and the way that any, like exception use cases, and things like that, that are that need approval, go through that AI governance board, but down at the in the trenches, you really have product owners, the product owners are the ones that are managing the use case outcomes, so for instance, like if you have, you know, if you're utilizing AI in order to, you know, create, I don't know, create better reporting or something like that within a certain business unit, you know, the product owner. Owner of that particular area needs to have the at least in the lower levels the ownership of the outcomes from that AI, but again you have above that the governance board, above that the risk being owned at the executive leadership, so that's sort of the way in a nutshell that that should be divvied up,

J Jo Peterson 05:22

yeah, and that's fair, and that's very thoughtful in terms of the approach. So, I just heard MTP referred to as the confused deputy, which I thought was funny. How do we prevent agents from executing actions that the user should not be allowed to perform,

D Derek Fisher 05:42

so you know, I think sometimes about AI agents specifically, as, as I know, this is like an oversimplification, but as a privileged account, right? It's an account, it's a non-carbon user that has the ability to, you know, make changes in the environment. I mean, we're not that it's not that far removed from like a service account, with the exception that you know these are far more powerful, have far more access, can do more, you know, with with the service account, you sort of have a, you know, single direction or bi-directional type of interaction that you can really kind of shape that with that traffic looks like in that interaction looks like with you know with with MCP in particular, where you have clients and you have, you know, servers or clients and and services, you know that the MCP is is you know making that connection to it's a little bit different because tools can be added on the fly, so let's say you know, for if we're, if you take Claude as an example, that just happens to be my, my tool choice, you know, I can ask it to generate a Word document out of a text, you know, that that it output, it's utilizing MCP in the background to, you know, connect to a service to generate a Word document, you know, that's a lot different than what we would typically see in, like, a service account type of activity, but it really comes down to, like, shaping that access and putting the guardrails around it, so you know, having the appropriate RBAC control, you know, role rule-based access control, or role-based access control, having also a registry of allowable services, like, here are the services that this, you know, that this client is allowed to be able to access, in terms of, like, you know, you can't just create, you know, connect to a service on the fly, like it has to be in an allowable set of services that that MCP is allowed to connect to, you know, putting those those guardrails around in terms of how it interacts with those services, and also having auditing in place, so you know that, like, okay, here's what it did, you know, here's what it connected to, here was the way it interacted, and so that you're able to lease peel back the onion at some point, should something go wrong, that you're able to at least audit that interaction, and then really just having that transparency around how the MCP, you know, how it's making those connections between the services and the client, and knowing what types of, you know, again, that sort of ties into that auditing, but like, what is that? What is it? How is it thinking about making those connections, how is it thinking about, you know, meeting the goals of what the client is asking, so it really comes down to again shaping that connection, having a solid auditing around, you know, what it's doing, and having that visibility into it, and then really just having those guardrails in place, shaping the type of traffic that it's allowed to do.

J Jo Peterson 09:08

That's a great answer. Thank you. One of the things I think about, and I wonder how we solve for, is there's more and more MCP servers coming into everyone's environment somehow, some way. How do organizations verify the authenticity and security of these third-party servers coming in,

D Derek Fisher 09:31

and this comes back down, and I don't mean to, like, sort of like tie back down into like what's sort of been security classic or whatever, but I think you know it comes down to third party management, like we've, like we've been doing third party risk management, you know, the first thing is, you know what, what business need is this particular, you know what business need is this, is really solving, like what problem are we solving with this, yeah. Because you know, and again, we would ask that question in a third party risk assessment, is like, do we need this, like, do we not already have this capability, is it something that we really need to go external for, or use, you know, another service before onboarding them, so really, you know, what is the, what is the value, and what is the problem that this, that this is that this particular case is solving, and then it comes down to, you know, understanding. Do we have an internal team that's actually capable of being able to manage this, because again, looking at like a third-party risk assessment, like, okay, we can get this tool and bring it in, but do we have people that have number one, the knowledge to be able to to manage this tool or utilize it appropriately? Do we have the ability to provide proper oversight for it, and so forth, and then also, you know, compliance and regulatory, more importantly, like the regulatory aspects of utilizing that service, you know, is there concerns in that, in that aspect. Do we have concerns about how our data is being used? Do we have concerns about, you know, the way they're, you know, the way they're processing the data, the information we're getting back? Are we, you know, doing critical or making critical decisions based on on the interaction that we have there, so you know those are sort of, again, not the, not the sort of put, put in the same bucket as like 3/3-party risk assessment, but it sort of is that right, we should be treating it in the same way, AI is not AI, is a novel way for us to sort of tackle technical problems, but at the same time it's a, it's a tool, it's a technology tool that we, we should be leveraging as much of our kind of learned experience and learned knowledge to be able to apply, in this case, it, yes, it's more powerful. Yes, can do more things, but in essence, we should be utilizing what we already know from a security standpoint on on how to better get our hands around it

J Jo Peterson 12:11

right. And I asked that, and you went exactly where I think you needed to with that one. Is I want to draw out the idea that I don't think people are asking the question right. It's just for whatever reason there's some sort of disconnect. And sure, large companies have lots of talented folks, but the smaller you get in terms of companies, the less bench you normally have, and not that those people aren't smart, but they're overworked, and they're, they just may not have enough resources, and I don't think sometimes the questions getting asked, What do you think?

D Derek Fisher 12:51

Yeah, agreed. I mean, you know, I sort of think about, and I was thinking about this the other day, that we've been complaining for so long, you know, in in cyber security and technology, and in general, that like we have people that are just burned out, you're working many hours, you're never getting ahead, you're never catching up, you know, this is our moment to utilize these tools, whether it's, you know, AI, particularly AI, to be able to hopefully reduce that cognitive load and hopefully reduce that amount of stress and burnout in the people that we, that we employ, and you know, hopefully AI is going to be a solution to that, and I hope that that is the direction we're heading. Sorry, I hope that sort of answers your,

J Jo Peterson 13:42

it was perfect. It was, it was a perfect answer, and you and I are thinking about it the same way. So, that's that was great. Well, it's been great to have you visit. Thank you so much for taking time to visit with me. And, guys, we'll catch you later.

D Derek Fisher 13:56

Thank you.