CTL-CISO-Marcus-Cylar

Tue, Jul 28, 2026 5:39AM 16:59

SUMMARY KEYWORDS

AI security, least privilege access, IAM tools, GRC space, security awareness, shadow AI, system inventory, confidentiality, integrity, availability, cyber fundamentals, workforce training, compliance, culture of awareness, human in the loop.

SPEAKERS

Marcus Cylar, Jo Peterson

J **Jo Peterson** 00:07

Hey y'all! Thank you so much for joining this session of Clear Tech Loop. I am Jo Peterson. I'm the CIO of Clarify360 and the Chief Analyst at ClearTech Research, and I've got a wonderful guest today, mr. Mark, mr. Marcus Cylar. Hi, Marcus. How are you?

M **Marcus Cylar** 00:26

Hi, Joe. I'm well. How are you?

J **Jo Peterson** 00:28

I'm good. Thank you for taking the time to join.

M **Marcus Cylar** 00:31

And thank you for thank you for having me. Oh, you're

J Jo Peterson 00:33

welcome. If y'all aren't following Marcus, you should be because he's a consultant and thought leader in the GRC space. So, as everyone knows, we're a hot take on AI security, and I'm going to be asking Marcus three questions, which is our normal format. So, let's get going. Marcus, first question. You ready?

M Marcus Cylar 00:56

I'm ready. All right.

J Jo Peterson 00:59

AI agents hold delegated credentials, make authorization decisions in real time and can spawn subagents with their own permission sets. Traditional IAM and PAM tools weren't built for this. What does a realistic path to least privilege access for AI agents actually look like for a large enterprise today?

M Marcus Cylar 01:20

So before before I answer that question, I want to offer some immediate pushback. Not not to say that you're wrong about the traditional IAM and Pam tools not being built for this, but the the question I want to ask is: Is it so much that those tools are are not sufficient for for for for where we're going with AI agents, or is it just the the the fact that we're not keeping things in a proper perspective as far as what AI agents are capable of doing and what kinds of controls we need to be able to have on them that are even independent from the tools that we have in place right now. So that that's the first question that I ask, and and and the reason that I think about that is you remember Watson, right? I'm sure you know Watson from IBM. So I think about when I think about Watson. Basically, Watson crawled so that all these other AI agents that we have now, you know, it's been what 20 years or so since then. So, so Watson crawls so everyone else can run. That we all the tools and all the agents that we have today, right? And so, I remember when Watson was on Jeopardy, right? So when so when Watson was on the Jep on Jeopardy, did the rules of the game change for Watson. Watson was playing against human and human contestants. the The name of the game didn't change. It was still Jeopardy. The objectives didn't change. The competition didn't change. Right, and so Watson wasn't playing a modified version of Jeopardy. wasn't playing a watered down version. He was playing the the the machine was playing Jeopardy, the same Jeopardy that we've all known and loved for for years. That's the game that Watson was playing, and and and so and and so that's what I think about when I think about artificial intelligence, and and I think about you know when I when I think about answering your question about least privilege assets, right? I I think about the fact that if we think about why the tech is intimidating in the first place, because we're looking at the newness of AI, and and we're we're scrambling to try to whip up these new rules and these new standards, and and we're and and we need to emphasize more than ever, I think, the fundamentals of security, and and and I might be you might you might find me saying that a lot, and but we really need to lean back into the fundamentals and making sure that we're stewarding those fundamentals well. I mean, we're we're overall security has not changed. We're we're still supposed to be protecting the confidentiality, integrity, and availability of systems, right? Separation of duties that's still important. Role-based access control that's still important. Making sure that we're we're being conservative with our permissions and and and and having a good handle on how we handle approval workflows and things of that nature, and and and how we handle monitoring and and and logging and making sure that we're we're doing those things properly. Those are the fundamentals that are still important. And then my favorite, which I which which is my passion and my focus, is security awareness. Security awareness training is is still important and is more important than ever, right? And not only is it important, but because the stakes are higher because of the potential for damage that that AI can can can can wield because it's so much more. I think that puts. More of an importance on the fundamentals, and more of an importance on awareness training, so that we can build a more knowledgeable and prepared workforce, right? And so, finally, to to to hopefully to answer your question, I can't really answer in detail what what what a what a realistic path to least privilege access is, but what I can say is I believe that that that path starts with a passionate return to the funnels of cyber fundamentals of cybersecurity.

J Jo Peterson 05:31

Well, you bring up such a good point because I've been thinking about AI awareness training as well. Remember, 1015 years ago, we started giving employees phishing training for their emails, right? And I'm I'm waiting for the company that comes out and says, "Hey, here's some AI awareness training for y'all, because everybody, you AI has become ubiquitous, and so many folks, to your point, in the workforce, just like email, are using AI on the daily, and yet, are we doing a good enough job as enterprise leaders teaching our folks about how to use AI properly? I wonder,

M Marcus Cylar 06:21

and and I think that is a legitimate question and concern, obviously. But I I would say that the best way right now, because the the tech is still evolving. So my background was electrical engineering. I did that in undergrad, and then I kind of stepped away from the tech space for a long time. I was in Christian ministry for for over 25 years, and I just kind of got back into the tech space about four or five years ago. And so I remember when I was coming up in electrical engineering, we always heard about Moore's law and and and how that cycle is, you know, every 18 months. But we're we're so far past Moore's law right now. Like things are changing. It seems like every 18 hours, right? And so, I think that the best way right now, as things are still changing, as we're still in the the wild wild west kind of portion of the the the the evolution of AI. While we're still kind of in that space, I think the best way to do AI training right now is to shore up regular awareness training, and and and and the reason why I say that is because if if you have a and and I know that you've you've talked on on on other shows, I've I've had a chance to I've had a chance to to watch a few, and I know we've talked about you know shadow AI and and the the impact of that on the business. If you have an an informed and if you have an aware workforce, then you're going to have people who are less likely to to try to implement shadow AI on their own, but then also you're going to have leaders who don't allow for the conditions under which, and hopefully this makes sense, under which shadow AI persists in the workplace anyway. So what I mean is, not only do you have a workforce who is informed about the the dangers of having shadow AI in the workplace and and what that could do, but then you also have a workforce who who who doesn't turn into the Department of No, but they they're actually they're they're they're people who work with their employees and and and create the type of environment where where where new things can be tried, new things can be tested, and and they they they they present that environment, they curate that environment for that to take place in a safe way where where not only your employees not encouraged to to go to shadow AI, but there there is there there there's more. What's the word I'm looking for? There there's more.

J Jo Peterson 09:34

They're more open to it. Yeah.

M Marcus Cylar 09:36

Transparency. There's more transparency around the the the tools that people are using, and I just think that that yes, we do need to have specific AI training, but that training is going to change so much that you know by the time you get that training, it could be outdated, and then you know, and so it it needs to be on. Going, but before you even get there, we need to we need to shore up our our just regular awareness training, and we need to make sure that we're not doing it as a means to an end, as as a means to just fulfill your your compliance requirements. But it's something that you're doing because you've established a culture of awareness at your company. Awareness is not just something that you're doing to check a box, but awareness is something that you're doing because you've made a decision as a company that that is something that is of paramount importance. And and I don't think enough businesses, at least from where I'm sitting, I don't see enough businesses doing that, so we need to have first things first before we really get into the nuts and bolts of of AI. I think though the the companies who are going to really thrive in this space moving forward, I believe are the ones who are going to have the greatest handle on the fundamentals and are going to be the companies that have the best culture of awareness and the best culture around foundational principles.

J Jo Peterson 11:13

Fair enough. All right. Well, so I'm going to sort of condense the second and third question because they think they're they're close. If a CISO in this audience is starting from zero on AI agent governance, no agent-specific controls, no monitoring, what is the single most important thing they should do in the next 30 days?

M Marcus Cylar 11:35

So again, I'm I'm going to go back to to fundamentals, and and this is going to be a fundamental answer, but really getting a handle on your system inventory.

J Jo Peterson 11:46

Right. So

Marcus Cylar 11:46

before you could even before you could even discern your AI inventory, you have to have a good system inventory, and and you need to make sure and figure out what your overall system inventory is from soup to nuts. Everything in your environment, everything in your organization, you have to figure out what you have and make sure it's accounted for. Then you go through that that inventory and see if there are agents that are in place that within those tool sets that, and you need to evaluate the permissions that are in place. See if there are agents attached to those, and see what permissions they've been given. See if any of those permissions need to be revoked, and and and and also make a note of who owns those controls, who who's responsible for them, either as an individual or a team. You have to figure that out as well, and then and then finally, you have to sit down with everybody in the company, and and this is something that is going to require. And I'm going to go back again to having a good security culture. So if you have the right culture in your workplace, then what I'm about to say is something that's gonna. I think it's something that's more likely to happen because it's gonna take a level of trust. You're gonna have to sit down with everybody in the company, and depending on how big your company is, you you might need to scale this conversation with a survey, or there there might be you might have to do it like by department or something like that, and it might take some time, but but you need to you need to ask everybody in the company somehow, and and and really ask them to be honest and tell you, and and we're going to the shadow AI thing again. What AI platforms are you using right now? What agents are you using right now that haven't been approved for company use, and and and again, I'm I'm speaking to culture right now because you need to have a culture in place where where people can be honest, and where you have convinced the people that you're talking to that there's not going to be any retribution for your honesty right now. There's not going to be there's not going to be any fallout from you admitting yes I've been using some some shadow AI yes I've been using tools that we have not yet approved as a as a workforce and and and when you're asking for this honesty and you and you get people to to be honest with you you just need to make sure that you have made sufficient promises to them and let them know that okay, there's nothing that's going to happen. You're not going to lose your job. You're not going to lose any kind of bonuses that may be on the table for you being honest. We just need you to be honest right now because we need to have a handle on everything that's going on so that we can make decisions moving forward. And not only do we need you to be honest as employees, but but but but as the leaders of the company, we need to listen to you as you're telling us these things. We need to take them as signals as to okay, these are the things that our teams are thinking about because any. Any use of shadow AI is just indicative of these are these are the tools that the workforce wants to use, and and that just means that these are needs that have been identified. These are needs that are on the table. How do we meet those needs in a way that's going to embrace their concerns, and how do we present ourselves as not being the Department of No, right? But also making sure that we are being receptive to what our workforce believes needs to happen. How can we work with them to implement these things in a safe way that is going to again secure the secure the the confidentiality, integrity, and availability of the systems in our place.

M Marcus Cylar 15:48

So again, it's something that that speaks to culture and it speaks to the honesty is going to have to be in place in order for you to get the data that you need in order to To have what you need to be able to make the changes that that need to happen, and so I guess that may not be the single thing. It might have been, I think, that might have been three things that I said there, but those are the things I think need to happen within the next 30 days. That's what I would do. That's the direction that I would go in if I were a seesaw. So

J Jo Peterson 16:22

there's nothing there's nothing wrong with empathy and human in the loop, right? So absolutely new new take on human in the loop. But thank you for taking time to visit today. Thank you, audience, for joining. Marcus is making us think maybe a little bit differently, so that's always a good thing. I hope you have a lovely rest of your week, and see you next time.

M Marcus Cylar 16:47

Thank you, Jo. I appreciate

16:48

it.