CTL-CISO-Louis-Columbus

Mon, Sep 14, 2026 5:59AM 17:57

SUMMARY KEYWORDS

AI agents, cybersecurity, identity management, least privilege, agentic identities, zero trust, governance, risk management, compliance, shadow AI, agent inventory, API keys, credential security, supply chain, automated attacks

SPEAKERS

Louis Columbus, Jo Peterson

J **Jo Peterson** 00:07

Hey everyone! Thanks so much for joining. I'm Jo Peterson. I'm the CIO of Clarify 360 and the chief analyst at Clear Tech Research. And I've got mr. Lewis Columbus here today visiting with us. Hi, Louis.

L **Louis Columbus** 00:20

Hi. Good morning.

J **Jo Peterson** 00:22

Thank you for making time. Oh, you're welcome. In case you guys are not following him, you need to be because he writes great cybersecurity articles. He is a senior cybersecurity contributor, in fact, for VentureBeat. And beyond that, in his free time, he's an adjunct professor at Webster Loyola Marymount University, so appreciate you being here.

L **Louis Columbus** 00:49

Well, thank you for inviting me. I appreciate it. Honored to be here. Thank you.

Jo Peterson 00:53

So, as you guys know, we do three hot take questions that are of the moment about AI security. So let's get rolling with that. First one, Louis. AI agents hold delegated credentials. They make authorization decisions in real time, and they can spawn subagents with their own permission sets. So traditional IAM and PAM tools weren't really built for this. What does a realistic path to least privilege access for AI agents actually look like for a large enterprise today.

Louis Columbus 01:24

Yeah, that's an excellent question and extremely relevant given the situation that happened with Hugging Face and OpenAI, which I wrote about back on july 22. And what is is I think at the crux of that is that privileges were actually violated there, and they were breached through credentials and privileges that should have never been within reach to begin with. And this is the oldest problem in security, but it's the newest one in agentic identities, mainly because agentic identities haven't been designed with that level of security in mind, traditional and I traditional I am and Pam will break mainly because you know it varies from the CSOs I talk about, talk to and talk and write about. A lot of these companies are using IAM and Pam privileges and constructs and frameworks to manage agents, and this even comes down to HR-related elements, and they're trying to squeeze it into a legacy framework. And of course, that doesn't work because an agent is a lot smarter than that. And as you know, as the breakout from Anthropic, they just talked about it what a day ago, where they disclosed that. I mean, agents are incredibly resourceful and will goal seek their way out of a sandbox, given half a chance. So they were, and I think the crux of this is that identity systems in the past were built for one user, one session, and a set of overall requirements or actual privileges that are given, and even in a zero trust relationship where you are limiting those privilege for least privilege access, those technologies and those systems weren't built for an agent. So an agent has a couple things going for it. It's ephemeral, you know, and ephemeral identities are especially challenging for any CSO to deal with. They're also they're delegated. They're over delegated privileges, and then finally they're extremely autonomous. And I did a a panel at Venture Beat Transform earlier this month, and the CSO of of Box was there, and she was saying that that ephemeral nature and that delegated nature of privileges was also brought up by a GM from Rubrik, and together their two comments really I think paint a portrait of of agents in transition, and you know these are two very advanced companies, Box and Rubrik, and they had been trusting the agents, and the CISO of Box put it really well. She said, "We were trusting it up to the point as we continued to move through it, and then it violated the trust, and then we went right back to square one. And I think that that in a microcosm is exactly what's going on with CISOs who are piloting and running with agents right now. That autonomy and that cross-domain capability is really alluring, especially in DevOps environments. But it does come with a price where you have to be able to have a brand new level of of zero trust and least privilege access in an ephemeral environment, which is really challenging to deal with.

J Jo Peterson 04:40

Yeah, totally, and I loved the term that you used, goal seeking, because, and I know this was in your article that you wrote 17,000 tries,

L Louis Columbus 04:52

yeah,

J Jo Peterson 04:52

17,000 tries. So talk about goal seeking,

L Louis Columbus 04:56

yeah,

J Jo Peterson 04:57

like it doesn't stop, and it's not that it's. It's just doing its job, right?

L Louis Columbus 05:03

Yeah, I mean it's like that age-old story of the scorpion and the turtle, right? I mean, the it is what it is. An agent is what it is, and a adversarial agent is just like that scorpion and the scorpion and turtle, you know, parable. It it just is an adversarial agent. Just is, and that's what it's you know is modus operandi is, and that's the most challenging part about I think handling that, and the fact that agents can be turned into adversarial agents through a variety of different technologies, and even accidentally, I mean, even the Hugging Face CEO, which really kind of shocked all of his adventure. But he turned around and he said, and that was the impetus for the article. Is he said, look, no malice, no no foul. You know, we we realized it was pretty much elevated negligence, but he put it nicer than that. Then he jumped on a plane and went to San Mateo, or looking for or at San Francisco looking for $100 million to be able to underwrite Hucking Face Future Development, which is another word of saying, hey, we'll do damages, and that's another area of this that is swirling around. And there's a fair amount of debate I think going across cybersecurity writers right now in the cybersecurity community is where does the legal ramification or the the the accountability for an agent that goes off the rails like that? Where does that start and end? And there are a lot of different legal precedents. and And one of our contributors, AdventureBeat, has written a really great personal blog on that, and it is the crux of the future of what an AI first company is going to look like. If you extrapolate this out, and you say you have 1000s of agents and you're building an AI first company, which everyone, every CEO who's got a hope of having AI will get on CNBC and say we are an AR first company, you know, with with help of goodness of of shareholder value and stock prices. Yeah, but this is you know this is reality. I mean, these agents will goal seek and they are relentless. So you got to have some really strong beyond guardrails. You got to have really good governance on this.

J Jo Peterson 07:20

Yeah, and that brings me to the next question, and sort of a tongue-in-cheek question, so a little spicy. So let's see where it takes

L Louis Columbus 07:26

us. Sure.

J Jo Peterson 07:27

So 50% of large enterprises have now dedicated AI governance committees. Right? Is that structure actually working, or is it creating the appearance of governance without the substance?

Louis Columbus 07:42

You know, I think that is. I think that's good intentions. And I interviewed one of the senior leaders at Rubrik, and he called he called many aspects of cybersecurity. He called it security theater, and I thought that really struck with me. And it's the rubric article that's out on VentureBeat right now. But I think there's a fair amount of theater going on on boards right now when it comes to this, and it's one of the most fascinating topics to talk to both CISOs about and those that consult with CISOs, I had a great call this earlier this week with a member, a senior member of EEE, and he emphatically said that governance, risk, and compliance is lagging significantly behind AI, and from what he's seen, but there is the implication of having an oversight committee, and a board is typically well schooled in in business business overall acumen and skills. But it is rare to find a CSO, and I've met a few of them who can go into a board meeting and educate a board and tell them, you know, just what's at stake here. And to answer your question, no, I don't think they are. And you know, one of one of the greatest things about about watching this space and writing about it is that juxtaposition of the irony of 50 to 70% of CISOs or CIOs and boards saying that we've got governance, but only 20% are are executing it. That gap, there's a story there, and the answer, the question is why of that. So yeah, you know, it's govern. You know, governance really exists on paper, but it doesn't enforce at runtime. One other really key point I want to bring up is is the CEO of Crosswrite at the last RSAC. Kurt's brought up the point that he was talking to a Fortune 50 CEO, and this agent had started to rewrite security policies on its. To be able to abdicate on its own, to be able to get greater freedom, and only after running an audit of that did they actually unearth that. And it's a true story. And in speaking with other fellow who didn't want to go on the record, but they had actually seen the same thing where agents will start to rewrite policies, start to redefine their ephemeral nature, will define restarting their identities because I know with by modifying their identities they can use the privileged access and the many different controls that they have to be able to move through this. So entering an era of of in the case of Hugging Face and OpenAI, yeah, you could say it's defensible. It was goal seeking to to pass that that one overall benchmark. But still, you know, when you look at goal seeking, goal seeking, the ethics of goal seeking with an agent is is the wild west. I mean, that's unexplored territory, and

Jo Peterson 10:59

It is, you know. It is, and I, you know, I agree with you that there's a story in that gap, and I have a hypothesis, if I may. Sure. I think part of it is that you can have a governance committee, but if you don't delegate roles, right? If you don't delegate responsibility, then it's just paper,

Louis Columbus 11:24

right

J Jo Peterson 11:25

It just kind of is. It's just it's just paper, right? And so, I think that's maybe a piece to that puzzle. I don't know that it's a whole puzzle, but it's a piece

L Louis Columbus 11:35

to it. I think you pretty well summarize the current state of it. You know, it it is on paper. I think there's good intentions, but it is moving so quickly, and it will be rare to find an organization who will double down on that and get identities and I am for agentic AI in general under control. There's you know there's every once in a while you run into a situation, or you meet a CISO that that is making progress there, and an interview with David Levin, who's at Amex GBT, and he anchored governance in NIST, and he has a cross-functional committee that reviews AI at every phase. And David Levin's interview is a little bit older on VentureBeat, but I encourage you to read it. He's very, very pragmatic. And then Jerry Geisler was kind enough to allow me to interview him. He's the CISO of Walmart, and he pairs velocity with governance instead of trading one for the other. And Sam Evans, also of Clearwater Analytics, he's a CISO there, and he talked about shadow AI In detail, and the way he sold shadow AI as a solution to governance with to his own board, and he was a young young in the position at Clearwater Analytics when he did that, and that's a little bit of an older story. But these each of these CISOs are obviously trading on their own unique strengths. Jared Giesler is is you know multi decade a multi decade Walmart employee and really a fascinating discussion about identity and the future of identity as it relates to Walmart and its entire business model and across the way that it is managing everything from physical assets, supply chains, and of course their online presence.

J Jo Peterson 13:26

I need to go back and read that because that sounds like such a good read. Okay, last question: If a CISO in this audience that's listening to us talk today is starting from zero on AI agent governance, so that person has no inventory, no agent-specific controls, no monitoring. Maybe they just walked in. They just walked in the door, and they're like, "Oh, wait a second, we're a little bit behind here. What is the single most important thing they should do in the next 30 days?

L Louis Columbus 13:56

You really have to go hard on an agent inventory. I mean, you really got to go with that with everything you possibly have because agents spin up, and I take briefings from really interesting startups, and I also take briefings, of course, from the larger players as well. One startup who formerly comes out of Israeli intelligence had found a technology that would scan for agentic identities, and they came from the hospitality industry. and their Their value proposition is to go in and say, okay, how many agents do you think you have running to the IT team and maybe even to the board? And they'll say, oh, you know, 700 to 800. They'll come back, and there's like about 10,000 agents running. I mean, there's like a 10x multiple, and so being serious about getting agents and copilot pretty much under control. And there was a study from CrowdStrike recently where they detected 1800 agentic applications. Across 160 million endpoint instances, and even our recent VentureBeat Pulse data showed 50% 4% of enterprises have already had an incident or near miss, mainly due to inventory lapses. Some of the most fascinating technology coming out of startups right now, admittedly out of defense, is is finding these blind spots, and like I said, that gap gaps are very fascinating in cybersecurity in general, but in this area specifically because there's this massive, you know, vulnerability that companies aren't seeing until they do that. So I think discovery is got to be a high priority, followed by ownership and triage. Ownership being don't vacillate on who owns what. Don't allow whatever constraints there are to to not place ownership on what happens with those agents, and then triage what's going on with them, and at the very heart of all this is starting with API keys. You have to look at overshared API keys or orphaned agent credentials that could be hijacked. Really tighten down and do an audit of API keys and agentic credentials. And it goes without saying, share shared resources. You know, hear all the time about breaches starting because of a shared Dropbox opening that happened to have passwords on the shared drive or on the shared virtual drive that was just been open for decades, you know. Somebody would, so I mean, attackers are getting very, very savvy about how to find these things, and it's beyond the overall scanning for different ports. And there's some really exciting research coming next week from a major player that they had a briefing yesterday and under bargo for I think journalists from all over the world were on it and the response time of these attacks I mean they're like in seconds now and and that report was it's really good and it talks about how supply chain attacks are executed within milliseconds or within seconds all this is automated, and never before has been a need to have governance with teeth in it that can can actually enforce least privilege access and and slow this down.

J Jo Peterson 17:34

Great points. Thank you for all the knowledge you shared, and thank you for your time. And everyone, thank you for yours, and I hope you'll visit with us again.

L Louis Columbus 17:45

Well, thanks for having me on. I appreciate it.