CTL-CISO-Anton-Chuvakin

Sat, Nov 15, 2025 6:22AM 16:21

SUMMARY KEYWORDS

Cybersecurity, AI adoption, cloud security, generative AI, shadow AI, AI governance, CISO role, Google Cloud, security innovation, AI tools, risk management, data privacy, AI collaboration, security framework, AI governance framework.

SPEAKERS

Jo Peterson, Anton Chuvakin

J **Jo Peterson** 00:07

Hey everyone, thank you so much for joining. I am Jo Peterson. I'm the vice president of cloud and security for clarify 360 and the chief analyst at Clear tech loop. And if you're with us today, this is clear tech loop, our podcast, and I'm we're here with Dr Anton chuvakan, senior security staff office of the CISO at Google. Hi Anton.

A **Anton Chuvakin** 00:30

Hello there. Thanks for inviting me.

J **Jo Peterson** 00:33

Thanks for coming. Thanks for making time. In case you guys don't know Anton, he has a long history in cyber security, and he is well known and respected in the community. He's been at Google for six years, but before that, until 20 June of 2019 Anton was a research VP and distinguished analyst at Gartner for technical professionals, security and risk management strategies for those teams and a gardener. He covered a broad range of security operations and detection and response topics, and is in credit. And I did not know this about him, but it's kind of cool. He's inventing the term EDR. I did not know that about you. Anton, that's cool.

A Anton Chuvakin 01:16

Well, it's a great achievement. If the date on the calendar were 2013 of course, it was 12 years ago in your life. So I'm not sure I just should be like, super proud of it, because you don't get to invent a major acronym every year, even if you're at Gartner. And definitely it's not after Gartner,

J Jo Peterson 01:33

you should be proud of it. It's relevant still. So yeah, you've got a really interesting job with lots of moving parts. So people are probably sitting here wondering, what you do day to day. Well, tell us, what do you do?

A Anton Chuvakin 01:48

So our team, obviously the CISO at Google Cloud. It's kind of an interesting outlier of sorts, because we do work for a real CISO of Google Cloud, of course, as of today, this role is is vacant. Phil Venables, our CISO of several years have left, and we're looking for a new one, but the team still does what it was created by Phil to do, namely, that we help customers be secure. But notice that some other teams that seem like they're doing field CISO work are kind of more connected sales. We are not. We are really the team that helps Cloud customers to securely Experience cloud we advise, we write, we sometimes advise internal teams, we do a lot of other things that fit into this big miscellaneous bucket. But ultimately, we are here to help Cloud customers, Google Cloud customers, but also Cloud customers in general, to have secure cloud experience. And of course, now that includes AI, which is makes it much, much more fun.

J Jo Peterson 02:45

I have to tell you, hats off to you and your team, because I often read thought leadership pieces, because I think they're provocative and they're thought provoking, right? So, good job, guys. Good job to your team. Well, let's get right into it. In case you all haven't been and visited us with the podcast before. It's a relatively new podcast, and it's a hot take approach, and we're focused on cybersecurity, cloud security and AI security. And so we put out a weekly episode, and we asked our guests three focused questions. And our goal is to sort of educate our listeners, not sort of it is about the security landscape, and it's both from a risk and an opportunity perspective. We do a sound right approach, and it's meant to be sort of a momentary look at what is going on right this minute in the space, because, as you know, it's always changing. So let's dig right in. I'm going to ask the first question of Anton, how can cybersecurity professionals leverage generative AI to break out of traditional tools and tech mindset and drive more innovation and execution in their security programs?

A Anton Chuvakin 03:54

So I want to start with a maybe admittedly a somewhat negative take, is that I definitely met people who have used the AI revolution as an opportunity to go deeper into the tool and tech trap. So they they now assume that their tools have quote, unquote intelligence, and that they the tool can, in fact, replace the process that's broken. You know, you probably remember the GRC years and the early soar years, where this whole line about, if your process is broken, buying a tool would not fix it. So this was meant to be almost like a axiom, or the given truth, that if your process is broken and you bring a tool to fix the process. The process remains broken. It's just either faster or it's broken in a new way. Now this used to be true, and actually it is still true, but AI sends some people down the path of No, no, no, no, no. The tools we have now with AI are, in fact, intelligent, and they can fix the process. So. I disagree for today. I disagree today in five years, three years, who knows? Ask me again, but I feel like I'm almost like turning your question and saying, no, actually, some people went deeper in the tools and tech, and now they feel like the assumption that tools don't fix broken process is not true anymore. So I'm sort of nervous about this, because I feel like if you bring agentic AI or AI agents to a process that broken, it would just be massively business destroyingly broken. It may be like dramatically more broken than with just classic automation. So that's why I am nervous. But I do see the idea that AI helps people think that test things, test assumptions, gather and structure data. And to me, this is a thinking aid, whether it's a tool or a process. Change that, to me, is kind of an interesting debate to have. But to me, the change is primarily that, well, okay, I have an idea. I have a certain plan. Let me discuss it with Gemini. I was just doing it before this very cold, before this very recording. I wanted to test some idea that I have about using non human identities for securing agents. And I wanted to know what Gemini thinks, and I want to be tested on Gemini, seeded with documents I have. And just like regular Gemini, and I got some interesting ideas, I noted them down, and I would, I would pursue them. So to me, this part is not just I'm going to buy a tool that has AI that will do the work. To me, this is like a collab experience, collaboration experience. You can have it with Gemini, with notebook alum, and, of course, with other leading AIS from other vendors, I'll be, I'll be fair to that. So there is a lot of AI for CISOs that fits into this bucket. And I think, to me, this is becoming transformative, at least from some CISOs I hear about.

J Jo Peterson 07:01

Yeah, that's fair. And you know, we'll have to visit again, because I am fascinated by securing ephemeral agents. There's so much there, so let's move on to the next question. There's so much there's so much shadow AI happening in organizations right now. What is the best way for the CISO to communicate the value of secure AI adoption?

A Anton Chuvakin 07:29

So this is it's interesting because it almost like your question contains an implication that if you do shadow AI, it's insecure, but actually you can set the guardrails so that you can have enough freedom and enough value within the guardrails. So it would not be necessarily shadow AI, as you don't know about it, but it can be the bottom up approach, because it's almost the conundrum that I face sometimes is that we tell people, and everybody tells people, Hey, go and experiment. Go, learn you need to boost up your skills. Ai, skills matter, blah, blah, blah,

this type of line of thinking. So people go and experiment with AI, of course, they would experiment with what they can get for free, which means consumer grade again, Gemini, other tools which you can get for free or for a very nominal cost. So and then later on there the bosses say, oh, yeah, you learned. Now implement, and some of them choose to implement with the same tools. And the question is, whoa, hold on. You're implementing for enterprise. You're not using consumer tools. And sometimes the reaction is, what do you mean, consumer tools? Oh, never mind. It is not the same as you know your CEO wanting an iPhone in 2007 it's a bit different, because there would be enough things where, say, consumer grade AI would learn from prompts, and we want that to be the case, because I want AI to become better and smarter and work better next time. But for many, in many business scenarios, I do not want AI to learn from my prompts, because maybe my prompts has either corporate sensitive data or PII or whatever else. So it starts right there two rows. So when you learn with Shadow AI, and you run shadow AI, even like where your bosses know you do it for this certain use cases, you may be fine and you're not increasing the risk. But then you touch certain other use cases and it becomes either hold on, we have an approved tool. It doesn't have these risks use that. Or in some cases, the answer really is, sorry, no. The answer is, really is no. I mean, we can make all the jokes about how security should not be a party of no, but sometimes really answer is no. You can pick a situation where you would put a very sensitive, personal, regulated type of data in a tool, and the answer is, you just don't do that. There's no real you should not pursue this direction. And I've seen that here. I've seen that elsewhere, where sometimes the answer really is no, but before that, it is almost kind of a maybe catch and release is not very scientific. And not very governance speak. But to me, you do want to observe. You want to know what people do, and you want to have a very rapid process for approving tools, very rapid process for nudging to approve tools. I had a CISO tell me in a hallway conversation that they had data scientists who wanted a certain consumer grade AI, and they were willing to pay for the business version of the same tool, but their company bought Google workspace with Gemini, and he said, Hey, you guys should use Gemini. And their initial reaction was, no, no, we're just used to this tool. But in reality, Gemini can do everything that the other tool does, and more and for their use cases, not in general, for their use cases. And over time, they migrated to an approved tool. It happened to be Google's, but it's actually material to this example. The point is that they might move to an approved tool through a gentle nudging process that started the discovery conversation, and then, sort of like, Hey, this is the road. You can take this road, and what happens after that? There's an exception process. There's this one guy who had this one use case where Gemini was not a fit, so exception, that timed exception, was granted for the use of the other tool. So to me, this is kind of not super magical and not very AI driven, but it is the process that has worked compared to governance of software as a service 10 years ago, the only big difference, no, not the only one of the big difference is that you have to run faster. You have to run at much higher speed. We did an RSA panel, and that RSA panel, this came up from all the panel members. Hey, governance process. You have good ideas, but if you run this on quarterly meetings, sorry, buddy. No, go. It's the same as saying no if you say yes, but in a quarter, it's essentially a no today, and you will get bypassed, and much risky options would be in de facto use.

J Jo Peterson 11:52

Yeah, and I'm really you know, you're right. It was sort of a little bit of a trick question, but you got that right away, and I'm glad that you sort of bifurcated, for people, the reality of what's happening today in the workplace, because this is exactly what is going on today, right? We may we may think it's something different, but this is what's happening. So last question, share your thinking on establishing a strong AI governance framework. You already sort of talked about the fact that you have to go faster. But besides going faster, what are a few other ways that organizations should be thinking about this, and what tactical measures can they put in place to get it moving?

A Anton Chuvakin 12:36

So I think that the you do need the discovery process. This is kind of one of the fun ones, because you end up with knowing what people use today, and I mean, it's, it does start from shadow, AI indeed, but at the same time, you do need this for some other use cases, because you may have teams that purchased with the approval of their manager. And you have, I think, in one of the blogs, we called it shadow AI type two, where a certain part of the business approved it. So it is business approved, but it is not top down approved. Yeah. So, so if you, if you have that, you still need discovery, because you may be sitting in the security office, in CISO office, in detection response, or in a sock, and you discover a certain tool in use, and you confront the team, and they say, Nah, it's an approved tool. My manager approved it. So, so discovery applies not just to shadow AI, but also to this type of business, but not really. The other bit is that, of course, the speed that I mentioned, I think the framework is, don't, don't go on a quest for the best framework. Have a framework to me, the the advice is not go and request it on top of the mountain for 30 years assuming that you will get a framework. No, just pick a framework and run with that, because you would adapt. You that You would be agile, and you would have the well, this is allowed. This is not allowed. This allowed in this particular way. If you do this here, make a slight change, and it becomes allowed all that is kind of at the center of it, and then you do center on the use cases quite a bit, because you may say, Can I use AI for x? And you I'd be like, well, what are you doing with x? Are you doing more this or more that? If you're doing more this, actually, even consumer grade is fine. Do whatever you want. Don't care if you do this. Be careful. Here's the tool. And if you try to do this with AI, by God no stop. I mean, here's my lawyer friend. Hello, lawyer. Make them stop. There would be, yeah, we just did a podcast on lawyers take on an AI agent, and it's really hilarious. We did not, but it was so. And then escalation, I think that you would have these certain decisions that people may want to bypass, and they would go up. And it is possible that big, the big, big boss, would say, No, your exceptions. Fine, go do that. And then, of course, his, he owns the risk, which is true. For pretty much every other it exception down all the way back to my CEO wants an iPhone, but it applies here as well. You will have an escalation point. And then after that you monitor. I mean, the you would obviously not lose the detection response angle. You would do a lot of, you know, classic monitoring, from traffic sniffing and the URL web access, watching, and a whole bunch of other things. So that's always the final step that you would you would observe. If you don't observe a lot of other stuff you've done, you may not have actually done it. You may have you may think you've done it, but you haven't really. That's somewhat short and humorous take. And of course, I've missed a few things. I've missed people sitting in committees for hours deciding some of the things. But as long as it happens rapidly and the decisions are made, you may be fine.

J Jo Peterson 15:53

Yeah, that's good advice. It's moving fast. We're all learning. So I appreciate you taking time to visit today, insightful as always, and I'll have to have you back to talk ephemeral agents, because that'll be really cool. But thank you for the time today. Sure, thank you. You.