accenture

INNOVATION DAY: LET’S REINVENT SECURITY WITH GEN AI

VIDEO TRANSCRIPT

Daniel Kendzior: Thank you so much for joining us for our Accenture Security Innovation Day. My name is Daniel Kendzior and I'm the Global Data and AI Security Lead for Accenture.

I'm joined by some close partners and colleagues who would love to share a bit of our experiences in the exciting times of AI – what we've done from an internal perspective, what we see from an industry trend, and then what we're doing with the esteemed partners that are in the room, as well as a number of our clients.

So just wanted to start with a little bit of context and we'll also kind of frame it in – you know a number of us were here last year – so, kind of what have we seen in the world of AI year over year and what do we project about as we look at the next 12 months. When we kind of walked the conference floor last year at this time, I think what we saw from a lot of cyber practitioners was really kind of a learning motion, trying to understand what AI was evolving into, how that was going to impact cybersecurity, what were some of the new emerging threats, risks, etc.

We've pulled together this research, and this is actually pretty fresh. It's come from the first quarter of 2025, but kind of really trying to understand what are our clients seeing in the past 12 months, and what we see is, unfortunately, we still see a bit of a split, where there's a lot of non-technical executives that kind of understand that there's risks related to AI. There's a lot of cybersecurity practitioners that know that there are new things that they need to do, but there's still a bit gap there in terms of actually covering off on what are the new capabilities, controls etc., to make sure that as the business is looking to put new AI workloads into the organization, you've actually done some diligence up front to make they're secure.

What we've seen work well with a number of our clients is really up level the existing tool chests that they have. And so, we think about how do you take the existing risk assessment frameworks and you up level those for the world of AI? How are you properly evaluating the projects, the partners, the ecosystem providers, your service providers, from a risk perspective? What do those new policies and standards need to look like? And Kris can talk a little bit more around how we do that at scale with Accenture.

And then once that part is in place, then if you kind of look on the right side of the screen, these are a lot of the more technical guardrails that start to go in. So, things like AI firewalls, being able to do testing of AI workloads, which is very unique compared to traditional app dev. And then being able to monitor and detect risks on an ongoing basis. I think the monitoring component is extremely unique and one where there's a lot of monitoring that traditionally goes in place for an IT system. But when you think about AI, you have to think about it both from a functional perspective. Are we seeing things like functional drift or model drift? Are we see security vulnerabilities emerge? How are we leveraging security posture management capabilities to understand across that full stack of an AI workflow? What threats and vulnerabilities remain, and then being able to test that, obviously, on an ongoing basis, knowing that the AI system that you have today, likely within six months, is going to be heavily changed, right? Likely from an architecture perspective or just new models that are coming in.

Damon McDougald: I just want to set the context of what an agent is. Agent can mean many things to many people. It is something that has a goal, an objective. It's autonomous. So those are the things when I think of it. You give an agent a goal or an objective, it's autonomous, and as Daniel mentioned, too, you come back when it's done and interact with it. Instead of like today, if we're using ChatGPT or Co-Pilot, you're asking it a question or prompting it or, or telling it to do something, it comes back and does something like, then you got to say some more stuff, and it's conversational. So, the opportunities that it presents are boundless, and at Accenture, the opportunity that we see is tremendous. From the challenges that we've had in the security industry for many years, we can't find the right head count, the right skills. Every organization has 90 to 100 plus tools to manage security operations that they have to operate and maintain. All those tools are siloed. Different teams, different skills. There's no end-to-end use cases there. And what we envision and what we are building right now and what have built some early use cases around is not replacing security software at all but living between the spaces of security software and fusing those end-end, use cases that can drive material outcomes from a delivery standpoint.

So, what you're looking at behind here is a very high-level architecture of what has been built and what we're continuing to build, and it's built on top of Accenture's AI refinery, which is built on the top of NVIDIA and has all the infrastructure and the bells and whistles. So, it has the ability to declare agents, it has ability to use whatever model, it has ability customize model, and then it has our security brains. As Daniel alluded to, the 20 plus years of Accenture experiences shoved in there, and our industry experience all in that brain right there. So, the agents are smart from the beginning. And two things that are critical, and if these acronyms are new to you all, don't be surprised because some of them are just a week and a half, two weeks old. A2A, which is a new standard for agent interoperability that we work with Google in partnership with and then we just announced something earlier today called Trust Huddle, which is Accenture's AI Refinery interoperability. So, for us security practitioners, it allows you to publish an agent in a market, like a card – it's actually called an agent card in the protocol – and it just declares what the agent can do. It has an O-Auth endpoint to it that you can use to authenticate and get a bearer token for, and then you can encrypt the communication between, so it's built with security in mind for this. And so that's critical for agent operability, which is the composer layer of our solution that creates the workflows.

And so, what we envision is that we will build some agents and our partners are going to build their own agents. And we will fuse those together in our orchestration layer securely to be able to drive those end-to-end workflows. And then another critical thing, another acronym, MCP, Model Context Protocol. It's been hot for, I don't know, three or four months now. And that gives the ability for agents to talk to dependencies, talk to tools, in an English-like, natural language fashion, where you don't have to know the APIs. You don't have to know the database calls. You don't have to know if there's a file system sitting back there. It gives us the information right away so we can do some very powerful things. That would be the integration layer. And then the last layer is the brain layer for this thing.

Kris Burkhardt: We have a lot of things, again, you could see here, policy, guardrails, leadership support, integration, all of those things are important to make this work. So don't think that the way to do this is, hey, I'm just going to try a few things, and it looks good, I want to put it in production. You need to think about the responsible portion, you need to think about how it's going to use your data, you think about who's going read the output. All those types of things. Those are important considerations. I say that not because anybody here doesn't know it, but I'm telling you that because that's how we live it. That's how do it day in, day out.

Security controls, I think these are pretty standard, frankly, I think we're just on the cusp of better AI security controls. They're not a lot different than security controls for other things that you have in your environment. Don't overthink it. Make sure you're doing the basics. You've got the hygiene right. You're doing things like infrastructure security. You're doing things like access control, typical stuff. Get that right. The rest of the security will come along as we move forward. So don't sweat it, but again, do the basics, dip your toes in, make sure you're coming along. I don't think this is super useful for this group, but we did spend some time figuring out what good AI use cases are for us. So, we looked, you can see, we looked across some of the different work areas that my team does, and we chose different areas to focus on. I'm going to zoom in on three today that we did that I think are going be of interest to this group that are actually in use today.

First one risky public site classification. At Accenture, we have a lot of developers that do a lot development. They create sites, they may create sites as client assets, they may create sites for sales, they may just create lab sites –a lot of those for various reasons do need to be internet-facing and they find themselves on the public internet. We struggle to keep up with all of that and make sure that it's all secure. One of the things that we started doing that we can do with AI that we couldn't do before is we now have a way to go through and classify all those sites. We look at those sites. We've trained AI based on a number of rules as well as sites themselves to go look at, to tell us which sites are actually high-risk which we don't need to worry about. So then we can have a human go through and look at the high-risk sites, to determine what security controls really need to be put on those, and also which sites we can safely ignore. So, this was something that would take us 15-20 minutes per site when a human was doing it. Now we can do it in less than one minute per site with AI and we have a much better hit ratio and we're able to just really just cover a ton more ground. I know not everybody has this problem, a lot of people sit behind VPNs. We don't have that luxury. So, this really helps keep us safe.

Damon McDougald: Right now, organizations have a handful of agents. Five, 10, 15, 20. But pretty soon, we are going to say five work, why not 10? 10 work? 20. 20 work? 100. 100? 1,000. 1000 work? Million. We're limited by compute and imagination here. So, we're going to see the scale that is going to be amazing. And from a security standpoint, scary, but from a practitioner standpoint who want to see 'doing stuff', as Kris mentioned too, meaning a lot of work that we don't like to do as humans, and then we make mistakes doing that work a lot over and over again. Agents are perfect to do that stuff. That's where I get very excited. BT, what do you want to add to that?

Brandon Traffanstedt: So I love the last component because humans are, well, we're all human here, I think. Sometimes I feel human, sometimes I don't. But humans are amazing at making unstructured decisions on structured data. At least that's what an engineer would call creativity. So, depending on where your mind sits, it's the same. To Damon's point, the goal should be, when we look at practical, not only agentic approaches, but practical AI approaches, to take a bunch of unstructured data, structure it nicely, draw some inferences so that a human can now come and be creative. So, we want to make sure we're not slowing this down, while also giving our humans more time to focus on what makes us just generally awesome. The other thing that I'll underline, because you got me excited, you said the thing, you said zero standing privileges.

If we look at the average human in an organization, typically they interact with around 55 different applications, whether it's web stuff, whether it is infrastructure, whatever it may be: 55. If we look at the average agent, number one, we need to replicate the amount of services a human interacts with, so 55, magnify that to Damon's point. We have a decision point in that workflow. Decision point A might be that every single one of these workloads, every single one of these pieces of infrastructure, every single of these services has a credential associated to it and we have to deal with that. 55 multiplied exponentially potentially - I was not trying to rhyme there but it kind of worked out like that - means that now we have the challenge of credential bloat. By moving towards concepts like zero-standing privileges, ephemeral access, whether these agents work more like humans or more like machines. Things like spiffy S-vids as an example, if we get to ephemerality, we have less credentials to bring under management, which sounds kind of crazy coming from me who works for an organization that originated in credential management, but we have a decision point to modernize there based on the inference of the technology on the agentic side. So, we can still go down the path of traditional cred management, but at some point, the level of scale and the amount of identities that need to be created might get out of hand. So, it's an opportunity to rethink how you might look at something like ephemeral access, not just for agents, but for people and more traditional machines too.

Daniel Kendzior: Perfect. Thank you so much to my fellow panelists. Thank you to the clients and the partners. At this point, we'll open it back up and we'll go to the demos and feel free to have further conversations. But thank you so much.

Copyright © 2025 Accenture
All rights reserved.

Accenture and its logo
are registered trademarks
of Accenture.