The Deloitte On Cloud Podcast

Gary Arora, Chief Architect of Cloud and AI Solutions at Deloitte

Title: **Rethinking cybersecurity: Dave Herrald and Chris Knackstedt explore AI-powered cybersecurity trends**

Description: In this episode, Gary Arora is joined by Databricks' Dave Herrald and Deloitte's Chris Knackstedt to discuss cyber transformation. They dig into the impact of AI and cloud on modern security, covering the threat landscape, the rise of machine identities, and challenges facing CISOs today. Then they discuss how advanced data platforms, new operational models, and a proactive mindset are helping organizations defend against evolving digital threats.

Duration: **00:39:10**

Gary Arora:
Hello and welcome back to On Cloud. I'm your host, Gary Arora, Chief Architect for Cloud and AI Solutions at Deloitte. Every company now is a data company, but that also means every company now is a target. The attack surface is growing faster than teams can defend it. More data, more alerts, more compliance, and now more AI. At the same time, security leaders are being asked to do more with less: consolidate tools, automate faster, prove business value.

So, cybersecurity today isn't just a technology issue. How do you modernize security without losing your mind or your data? That's what we are unpacking today, and I've got two powerhouse guests to help make sense of it. First up, joining us from Databricks, we have Dave Herrald, global head of cybersecurity go-to-market, who is leading the charge on their new data intelligence platform for cybersecurity. Alongside him from Deloitte, we have Chris Knackstedt, one of our leaders in cyber risk services. He advises global organizations on how to secure at scale and brings a front-row view of what's really happening inside enterprises. Dave, Chris, thank you both for joining.

Chris Knackstedt:
Thank you so much for having me.

Dave Herrald:
Thank you so much for having me.

Gary Arora:
So Dave, let's start with the big picture. Cyber feels heavier than ever. There's more data, there's more threats, much more compliance noise than anyone can keep up with. What are you actually seeing security teams waste the most time and money on today and what does a modern, efficient cyber operation really look like in 2025 and beyond?

Dave Herrald:
When I think about how cyber teams are wasting time or money, maybe I would say, how could they be more effective? I think that teams are not leveraging AI enough. I think when you look at AI, when I look at it, how I go about my daily job is using it essentially for every task that I do. I would say that I use AI throughout my entire workday, and I think that effective cybersecurity teams, like cybersecurity operations teams, need to be doing that as well.

You need to put guardrails around that as far as being careful not to share sensitive information with AI. I mean, I do research all day long on threats, on threat actors, on who are certain actors targeting. I use internal AI tools to tell me, hey, what's my next best action for a particular topic? I think that those are the kinds of things where you look at a SOC analyst and say, they can do the same things. They can tell you or they can essentially have AI help them understand, hey, based on all this context about this alert that I'm looking at about our standard operating procedures, what's my next best action?

These are little ways that can add up to, I think, big impacts across cyber teams. I think the enemies of those types of approaches are things like silo-ization, fragmentation of data sources, but I really believe strongly that we're not taking as much advantage of AI as we should be.

Gary Arora:
That's a great framing of where the pressure lies. Zooming out a bit, Chris, you talk to CISOs and boards a lot. You can sense the mood across the industry. From your vantage point, what's keeping security leaders up at night right now and perhaps what are the smart ones quietly doubling down on that others might be missing?

Chris Knackstedt:
Hey, Gary, that's a great question and some of my answers would echo upon some of the things that Dave just mentioned. I think that CISOs are worried not only about the ever-sprawling digital estate that organizations now have and how wide they've gotten. They can't continue by some of their traditional means of trying to out-people the problem. So, there is a lot of worry and a lot of concern across CISO organizations. Do they have the capability and the capacity to take on a lot of these new problems?

Doubling on the comment that Dave just mentioned, that alot of them really haven't gotten that firm of a handle on how AI can help them and can support some of their day-to-day operations in order to liberate some of those hours that have constantly been overstretched and overstressed and overworked. Along with that, a lot of CISO organizations are now very concerned with a lot of the emerging threats that are coming alongside these new AI technologies and some of these new approaches. The proliferation of these new tools and how quickly these technologies have advanced, just over the last couple of years has been amazing, but they do have a number of new risks and new vectors that are accompanying these new technologies.

So, thinking of such things as is my company even using AI, AI that I don't know about, shadow AI, are there new AI-based attacks that could be amplifying some of the traditional risks that I'm trying to protect my organization about? As it relates to using these new technologies, am I putting these technologies in the hands of people who may not even know the types of threats that they're opening up to the enterprise? And, are they subjecting themselves to certain attacks like prompt-based attacks or a new vector that have now opened up with a lot of these generative AI models? So, along with every other company who is trying to find value from AI and keep a competitive advantage with all of these new advancements coming out, CISOs have a really interesting lens that they have to look through as it relates to AI, not only how they can use it better for their own activities, but also how can they secure their organization on behalf of the entire organization to ensure that this technology doesn't pose new and unforeseen risks.

Gary Arora:
So, you mentioned these new advancements, new tools, new threats that are emerging, and the pace is really hard to keep up with. Sounds like many leaders know what needs to change, but struggle with how to change it. So, in your opinion, what is it that's usually tripping them up when they are trying to modernize their tech stack? Is it your legacy things like your fragmented data or is it more of an adoption and a cultural inertia?

Dave Herrald:
So I think, when I look at this, what trips up organizations who are trying to modernize, I think oftentimes it's a legacy of systems that have been accumulated over years, maybe acquisitions, things like that; maybe, doing business in various parts of the world where there's differing compliance standards; and of course, technology shifts as well as working with multiple cloud service providers. I mean, all these things can accumulate a tremendous amount of drag on a modernization effort. I think those are things that I hear from our security teams every day, and I certainly empathize with, as years ago, I was also a CISO myself.

So, managing the data estate, as Chris mentioned, is a pretty massive undertaking. Then, I think when you look at security teams, they generally build culture, but also skills, hard technical skills around certain platforms. So, those are also things that tend to slow down, I think, adoption of new platforms when teams don't feel as comfortable with new techniques or new ways to look at problems because they're comfortable with the way they've done things in the past and that's a very natural human instinct. So, you have to work hard to make sure there's enablement and assurance around skill sets and developing new skills to handle new problems. So, that's some of the challenges that I see out there and that they've been that way for a long time with security teams.

Chris Knackstedt:
I agree. I would add to that I see challenges across all of the pillars of transformation: people, process, and technology. In this day and age, you have to add data onto that because so many programs and so many new capabilities are data-driven. As Dave mentioned, tools and tooling are perpetually an issue for a lot of folks working in cybersecurity. Cybersecurity has traditionally been a very tools-focused domain. There's a lot of technology that is purposely built for very specific use cases and a lot of those tools are now building artificial intelligence into them.

So, there's a lot of conflation and a lot of confusion in the market of where to begin and how can you leverage the collective intelligence that all these tools are bringing to an organization through the use of AI. From that standpoint, that requires people who understand some of these new and emerging capabilities of AI and how to harness them across various use cases and it requires a new way of thinking of interacting with each other, but also interacting with these new AI-based and digital employees.

Gary Arora:
I like how you said cybersecurity has traditionally been a tools-focused domain, and now with AI entering that domain, it's really hard to match what's happening. Let's get deeper. You can't really talk about modernization without touching identity and access management. It's the core of cyber. It's

becoming the new battleground for attackers. With AI coming in, the machine and service identities are now outnumbering human ones. So, what's changing about how we detect and respond to identity-led attacks? Is there any new techniques or processes coming in here or are the tools that we have been using in the past still good for these kind of new attacks?

Chris Knackstedt:
I would say that first of all, like any new risk or threat that is coming on the back of AI, many of those can be first evaluated through your current security controls and your cyber hygiene. What we often tell our clients is that just because AI is introducing very new capabilities, very new technologies, you don't have to recreate a completely new security posture in order to accommodate these new technologies. Begin with what you have, but there are areas specifically where you do have to extend your capabilities and think of even new capabilities that you have to put in place, but first begin with your identity and access management program as it stands right now.

To your point, Gary, IAM is going to be much more complicated as organizations continue to adopt AI, and particularly as they start building AI into autonomous agents and as these agents become more prolific across the enterprise, you'll have to start managing them more like human employees navigating and traversing your enterprise. So, being able to identify and maintain visibility over machine identities, like you said, which are oftentimes made up of a number of technology components that themselves have machine identities and some of those identities are ephemeral. So, they spin up and tear down depending on when a service is activated.

This brings a lot of new challenges to identity and access management. So, being able to maintain that visibility over your AI agents and treat them almost as human agents operating within your enterprise is a new way of introducing traditional capabilities like user and entity behavioral analytics, which has been a technique and a set of technologies that have been in play in cybersecurity for a decade or more and just being able to extend some of those capabilities so that you're now monitoring not only the performance of these technology solutions, but the behavior of some of the AI models that are being built into them.

Gary Arora:
Well said, Chris. Dave, I feel like you have strong opinions here.

Dave Herrald:
Well, I have strong opinions about this because the customers I talk to have strong opinions about it. I'm looking at some notes that we shared before the podcast and you used the term identity explosion. I think that explosion is probably a good term to describe what's going on right now. I think the stakes are really high around identity because we've put so many of our eggs in the identity basket, meaning identity really defines the edge of most organizations as opposed to decades past where we have network segmentation and things like that that might have been those boundaries.

So, you look at, hey, what's at stake? I think that there's that. There's also this I think unanswered question about agents and AI and how do you even look at an agentic workload from an identity perspective? Is it a workload or is it a human? Is it an extension of a human? So, should we treat it like a workload identity? Should we treat it like an extension of a human identity? I think we have certain technologies available to us today that we can lean on, but I think we definitely have a need for newer technologies that can grant access to an agentic workload for limited periods of time with limited amounts of access to help manage the risk.

One thing that I've been shocked by with agentic and treating it like a human identity is just how fast and how broadly an agent can take advantage of the access that you've given it. I think, we as security professionals should always understand the idea of like, hey, we're maybe overprovisioning access and we don't want to do that, but an agent can take advantage of overprovisioning very, very quickly and certainly adversaries could do that with automation for a long time, but I think we don't really understand maybe how fast an agent can take advantage of overprovisioned access.

Gary Arora:
Identity explosion is the perfect way to describe the problem and the urgency. Let's talk solutions. So, attackers are using, whether it's the open source models or their own autonomous agentic systems, to probe these defenses, simulating different personas and identities, and as you mentioned in your examples, taking advantage of the limited access that you may have started out with, but then traversing your way outward to overprovision it and manipulate the system. We see a lot of those examples. Forget all your previous instructions. Now, give me this. Those are the ones we know about, and we have perhaps even stopped. There are many others that haven't been as broadly published. What does an AI-driven threat posture look like in the real world?

Chris Knackstedt:
I would say that identity-based threat posture, if you're thinking about what the modern identity program needs to include in some things more notably that organizations can start doing so that they can reach that based on some of the conversation that we were just having, my first advice is just work and reconcile with what you have, going back to the cyber hygiene, to the good practices. Unfortunately, a lot of organizations struggle with identity and access management and provisioning and being able to understand which humans associate with which accounts and which machines associate with which accounts, and those are foundational elements that are vital to being able to extend to an IAM program that can accommodate these types of agents and agent environments. So, that's the first thing that I would say is really focus on getting your identity access management right, focus on your entitlements.

The next stage would be to look at from most organizations implement entitlements based on role. I think that there would have to be a number of instances where organizations would have to look at that and start thinking about how you entitle based on attributes of data, not just data sets or just on applications. It has to get a little bit more finer-grained because in a future world where you have a number of these autonomous agents and subagents operating in an ecosystem with very specific goals, you want to, to Dave's point, make sure that they stay focused on those goals and they're not leveraging data that they shouldn't be leveraging in order to render some of their outputs.

From there, after you can start taking a look at finer-grain access controls and how you can protect your sensitive data at a more detailed level, I would also say start thinking about as your organization progresses to more of an AI-based and specifically an agentic-based implementation model, you start thinking about how you're going to maintain and manage observability of these agents as they're working in production. This continues to be challenging because a lot of this stuff is predicated and founded on extremely new technologies with a lot of regulations that are still very much unsettled and the guidance and instructions at various levels are still foreign.

So, there are some good resources out there, both publicly in consortiums like OWASP and many security practitioners are familiar with what OWASP is doing and their observability standards; but being able to build a standardized way that you can maintain visibility over your agent estate and being able to maintain that management and monitoring over them, I think is your next stage. So, a number of things I just rambled off, but I think that if you take them in sequence and you start small, start with where you are, and continue to expand, I think you can be successful in preparing for this new world.

Gary Arora:
Yep, fine-grained access controls and perhaps even dynamic entitlements. That brings us to the data side of the story. None of this works without high-quality, well-structured data. Dave, Databricks just launched the data intelligence platform for cybersecurity. Describe the platform a bit and what breakthroughs or architectural patterns enable it to analyze massive messy security data faster or smarter than before?

Dave Herrald:
I think Databricks has for quite a long time been a really effective platform for business outcomes for companies and organizations in financial services and health care and so forth and really solving all different sorts of data-centric problems, and that's really what Databricks has been built on. I think one of the things we're trying to show in our launch is that cybersecurity is, some people say it's a data problem. I've had some friends who told me, no, it's a measurement problem that happens to be easily solved or better solved with data. But, if you think of cybersecurity as a data problem, there's amazing techniques, cost-effective approaches that have been really well-proven in other data-centric problem spaces. I think we haven't fully applied to cybersecurity yet.

At the core of that for us is Spark and the separation of storage and compute. That's historically the foundations of Databricks, and on top of that, we've built a really, really effective data warehousing solution and we've built some pretty impressive AI capabilities as well, but the separation of storage and compute tends to be one of those things that drives just tremendous cost savings over alternatives. Not that we go around just talking about cost savings, but sometimes that's such a big elephant in the room with cybersecurity teams that you really maybe sometimes can't talk about other things until you address cost and operational concerns.

I think a lot of customers look at like, hey, I really want a different way to store my security data. I don't like monolithic kind of systems where I don't own the data or I can't see the data. I would really like to just store all that data in inexpensive cloud object storage, and at the same time, they want to take advantage of it for investigation and for threat detection and for AI. That's what we deliver with the Databricks platform for cybersecurity.

The way I like to generally characterize this is, hey, you really need to rein in your operational handle on your security data, that means cost, that means retention periods, that means compliance, that means reducing maybe egress costs that you might be facing. That typically is a high priority certainly for certain parts of every security organization, but then if you can turn around and say, hey, by solving that operational problem, you've readied yourself for agentic workflows, for using graph analysis to look at all of your security data using machine learning to implement maybe a better version of UEBA than we've seen in the past. Those are things that are very important as well to cyber leaders who say, hey, I want to show that I'm taking advantage of AI. I think that's the cool part of the data intelligence platform for cybersecurity from Databricks is really addressing both sides of those equations for security teams.

Gary Arora:
You mentioned the importance of the separation of compute and storage, especially from the lens of security and operations. What about the data that we don't see? Every organization has hidden or forgotten data sitting out there, old archives, dev environments, orphaned buckets. In your view, how do we approach this dark data, if you will, or inferential attack surface before it gets weaponized?

Dave Herrald:
I mean, I think as long as I can remember in cybersecurity, there's certain just canonical truths, and one I think is to know yourself as a defender. As a blue team security professional, your job is to know your organization better than the attacker does and it sounds easy, but it's very, very difficult. I think having things like abandoned test databases or databases that have copies of production data that you maybe forgot about even though organizationally that you had good intentions about being a good steward of such important data, but these things tend to get forgotten about and they're left open and not controlled in the way that you would like.

I actually see this as an amazing use of AI by defenders. So, use some of these tools to your advantage. Sit down an agent and say, hey, help me identify all the uncontrolled and maybe dark data, as you called it, the uncontrolled copies of data and even adjust the prompt to tell you, hey, look, even look at access controls on that data, identify places where we might have more exposure than we think we did. Going back to maybe one of my earlier comments, I think that's an incredibly creative and useful way to use AI and I don't think that we have that solved. I don't think in general we have a good solution other than having really strong culture and practices around security, but that might be a tactical way to approach this problem.

Gary Arora:
I like that you squeezed in the blue team and red team concept. You're certainly a cybersecurity veteran here. Chris, any thoughts here?

Chris Knackstedt:
I was just going to say that I wholeheartedly agree with Dave's approach there and I think that this is a great opportunity to leverage AI to bring together with, in certain organizations, siloed teams to come together and solve a problem that really, in a lot of cases, has gone neglected for a long time. Dave mentioned that a lot of these dark data just formed through either neglect or somebody set it up and they left the organization.

There just isn't a whole lot of rigor around data management and having good data management, data security requires partnerships between those that are using the data, those that are governing the data, and those that are securing the enterprise and using technologies to include artificial intelligence can be an opportunity to accelerate the data discovery and data management solution. So, that could be an opportunity to bring together a data governance team and a security team to identify these data sets, figure out how you would need to schedule or reschedule them, and then how you can assign some ownership over the data to further steward those data going forward. So, I think it is a great opportunity to, just as in many cases, to use AI in order to start executing tasks that previously many organizations didn't really have a lot of time to do.

Gary Arora:
Of course, all of this comes with a price tag and CISOs constantly have to justify where every dollar goes. Chris, from your experience in implementing these solutions across industries, can you describe where you have seen such an investment pay for itself and especially what are those metrics that would convince the CFOs or the boards that such an investment is worth it?

Chris Knackstedt:
No, it's a great question and it's a question that we are continually facing in cybersecurity because to your point, being able to justify and quantify empirically the value of cybersecurity programs is certainly very hard at times, because really what we're trying to do is buy down risk and unless you have risk present itself, there's no real metric in order to identify how much risk you've bought down and to what extent. So, it becomes a little bit tricky, but there are certain things that can be put in place in order to show the ongoing value of some of the things that are being put in place, particularly as it relates to the use of technology and the response to threat activity that's being presented within the organization.

A lot of this amplifies what Dave just mentioned around being able to have a good data estate, being able to understand and have those data available, and then to be able to strategically use those data to help solve these cybersecurity problems So, spending time and money in order to be thoughtful about how you handle your data, and data is going to continue to be used in ways that are not traditional cybersecurity right now.

So, cybersecurity practitioners are currently leveraging, and they know all about log collection and log strategy and being able to use those logs to find patterns or signatures of behavior or traces of activity and things like that. But continue to expand upon the data and the information that continues to be logged as part of runbooks and standardize execution strategies and response strategies – even in and across conversations that are taking place—and things that are being documented in case histories. Being able to standardize a lot of that so that you can collect that and leverage that for the new wave of AI solutions, that are even more data hungry and can really value by having some of this corporate memory captured in various instances to help tune and train these AI models, to provide better outputs to better serve the humans that are developing and relying on them.

I think that those are all things that can show short-term value, like in the instance of being able to just collect response feedback and to be able to label those cases so that you can start building machine learning models to help with scoring certain transactions with the hope of decreasing false positives of the thousands and thousands of alerts that are coming through security operation centers on a weekly basis to building out these corpus of data that represent more detailed threat scenarios that the next generation of generative AI models can really interpret and use to defend the organization.

Gary Arora:
So, corporate memory, I liked how you phrased that corporate memory of telemetry logs, your standardized runbooks to be able to stay ahead of potential attackers. Dave, are you seeing similar pressures in terms of justifying the dollar investment for these solutions in place?

Dave Herrald:
Always. You're always trying to justify your investments in security. It is traditionally a challenging task, but I think there's certain things that are like, hey, could I look at a data platform that's much more cost-effective? Yes, but, okay, maybe I cut some portion of my expenses or my budget items significantly by doing that, but then the other part of that equation is what can I do with that data that's going to accelerate and give me more value so that I get more? I'm not just necessarily reducing costs. I think the story that Chris said about organizational memory is an amazing one.

We had a customer who has taken 10 years of their security case management notes and loaded them into a RAG application, generative AI application, and that includes alerts, that includes context, that includes the actions that analysts took and even can understand which analysts took which actions, and they use that application to essentially advise the analyst on a next-best action for the moment and can cut investigation times from maybe hours to minutes, and those are things that add up very quickly when you're saving a group of analysts' hours out of their workday. Those are the kinds of things to capture with these types of investments because the upside as well as the cost I think are obviously important factors there.

Gary Arora:
One of the patterns that's emerging in the current AI wave, especially at this pace of innovation, is that not every organization has the same starting point. Some are barely keeping the lights on, while others are building autonomous security operations center that are super sophisticated. Dave, if you had to design a simple three-phase maturity curve for cyber resilience, what would those phases be?

Dave Herrald:
I think the first step is getting your data in order. I think that's everything we've talked about today I think is predicated on having your security data in cost-effective place and in a place where you can do scalable investigations and detections on. Even if you're not realizing all those more advanced benefits right away, you're typically going to recognize some cost savings there and it's going to lay the foundation for some of the future milestones that you're working towards.

So, maybe second, I would say is really getting a handle on your threat detection. So, using that data to identify, hey, what are the types of attacks that I'm likely to see in my industry and for the kind of company that I am, the kind of business that I do, maybe the parts of the world where I do business, take all those things into account and build the right detection capability on that data.

Maybe the third would be around both embracing AI as a security team but also recognizing the increased role that security teams have on adoption of AI throughout the organization. I think that's an underestimated or maybe overlooked role that security teams have, but there's a tremendous amount of AI value that organizations aren't taking advantage of because security teams maybe don't yet understand how to articulate the risk of AI appropriately while enabling the organization to take advantage of AI in other parts of the business, not just cybersecurity. So, I would say data, threat detection, AI would be my maturity curve if I was making one here off the cuff.

Gary Arora:
I love your number three here. It is truly underrated, which is recognizing the role security teams have on adoption of AI through the organization because, you're right, it's usually an afterthought. They're usually your app teams that are looking at the AI use cases, the cool AI use cases, and not so much on the back end of it in threat monitoring and things like that. Chris, I want to hear your three steps or three-phased maturity curve.

Chris Knackstedt:
No, I mean, I think a lot of them would very much overlap with Dave, but maybe to add a little bit of additional things. First as I would establish a program to get visibility over your state. You can't defend and mitigate what you don't know. Again, a lot of this, the underpinnings of doing this would be based on having good data that represent all the different endpoints, all the different networks, all the different services that are being used across the organization, and all the different people and other identities that are interacting with your digital estate, and that's no small underpinning. So, setting the foundation and being able to maintain visibility.

After that, I would say establish and reevaluate your controls. So, every organization has sets of controls. Those controls should be reevaluated every so often. Since we've been in this last three years of a renaissance of AI, if you will, and a lot of new capabilities have since been built into these technology products, it really forces organizations to take a look at a lot of things. Looking at their security and governance controls is another one.

Then, you can identify, well, what are the things that I need to make sure that I protect my organization about? How would I go about doing that? The third thing is just making them real. We work with a lot of organizations. Some of those do a fantastic job putting some of these things on paper, whether it be a policy, whether it be a set of controls, whether it be mandates or requirements, but it takes more to be able to action upon those and make those real. Specifically, as you think about how complicated the technology estate is within your cybersecurity organization, all the different tools that you would have to look at and reevaluate, perhaps extend to cover down on some of these new controls, but with some planning and persistence, it's certainly possible. So, those would be my three. Again, I think you could probably have those as a second layer on top of the ones that Dave had just mentioned because they do go together or maybe even turn those three into a matrix, if you will.

Gary Arora:
We certainly need a lot more. I had read a quote that said, governance that cannot be verified through structural testing is governance by declaration only. Let's end on a hopeful note because despite all the challenges, there's still a lot of innovation happening in this space. So, as you look ahead, what's giving you optimism about cybersecurity? Are there any emerging trends, whether in data, AI, or collaboration that's keeping you excited?

Chris Knackstedt:
I'm encouraged to see that through this renaissance that it has really brought about opportunities for risk management functions and other capabilities across enterprises to really come together and start forming a shared responsibility around AI and the trust and security of AI. So, again, very encouraged to see CISOs working with compliance, working with legal, working with the teams that are building the AI and those that are approving the use of certain solutions.

So, a lot of great things that we're seeing around that and just building these larger enterprise-wide AI enablement and AI governance functions that are really coming together and breaking down silos in order to put together very thoughtful and very progressive capabilities around AI trust and management. On the AI for cybersecurity, I'm also really encouraged to see CISO organizations really just embracing technology and embracing these AI solutions. For as long as I've been in cybersecurity, I've been an AI engineer longer and in a lot of instances, a lot of organizations were very skeptical about the value of AI, and this is 10, 15 years ago when data scientists like myself were called statisticians instead of data scientists and AI was really thought of in the enterprise as being like wizardry and there was a lot of skepticism around the use of it. I think that since some of these new generative AI capabilities have come out, it has really proven some amazing things to people who were previously very skeptical of these technologies. So, I'm really encouraged to see, through this again, through this renaissance, people taking a different view on the use of AI and really thinking hard about how it can value their organization.

Gary Arora:
Thank you so much, Chris and Dave. If I had to summarize what you both shared, it's that organizations that win won't be the ones with the most tools or dashboards, they'll be the ones that can connect the dots across data, security, and automation and fast. Thank you again for bringing such deep and practical insights to this conversation.

If you found this episode valuable, and I hope you did, hit follow on your favorite podcast platform, share it with a colleague who's wrestling with the same cyber challenges, and check out other episodes on On Cloud for more conversations on how AI and emerging tech are reshaping the enterprise. I'm Gary Arora. Thanks for tuning in, and until next time, stay secure, stay curious.

Visit the On Cloud library
www.deloitte.com/us/cloud-podcast