Re(AI)magine Conversations

Persistent Podcast | Episode 01

Re(AI)magine the Future of Data Intelligence with Agentic AI

Gregory Taylor & Sameer Dixit

Sameer Dixit: Hello, everybody. Welcome to Season Two of the Persistent Podcast Conversations at the Core of Change. I am Sameer Dixit, SVP and Head of Data, AI, and Integration at Persistent. In this series, we engage with industry leaders and change makers who are redefining what's possible at the intersection of business and technology.

These are leaders who are sharing real-world insights on transformation, innovation, and impact. So, it's no longer theoretical. It's actually practical. We are seeing a massive shift in how enterprises use data. It's no longer about just storing or moving data from one store or another. It's about platforms which are built to understand, act, and even learn all in real time.

Joining me to talk about this evolution is someone who has been right at the center of all these actions. Welcome, Greg Taylor, who is the VP for APJ at Databricks. I'm glad to have you on the show today, Greg.

Gregory Taylor: It's so good to be here today, Samir. We're pretty excited about the topic and really excited about the partnership we have with Persistent.

Sameer Dixit: So, let's kick off things with what's hot right now, to start with, Can you give us a view in terms of what are the new or emerging technologies in the Data and AI space that you are most excited about it personally?

Gregory Taylor: Yeah, Samir. So, there's never been a more exciting time to be in Data and AI.

If I had to pick just one area that excites me the most, I think it's the convergence of generative AI with structured analytics, particularly the shift that we call it Databricks, from general intelligence to what we think of as data intelligence. The idea that you can move dashboards, static reports to intelligent systems and that understand the context of your business and are able to get you to ask and answer questions in natural language, it's a game changer for us.

We're seeing enterprises wanting to democratize data access, merging technologies like Vector and search RAG; and fine tuning models that are really accelerating that. Another area I'm keeping a close eye on is this rise of how you govern data sharing, particularly with some of the open standards that are coming out.

In places like APJ where we have really complex ecosystems and, you know, multi-party collaboration is super essential. So, the ability to share insights securely and in real time is no longer a nice-to-have; it's really critical now. Finally, I'd like to say that simplification of data architecture through patterns like the Lakehouse adoption, especially for industries like financial services, software and digital natives and healthcare, that's really unlocking new value for businesses.

Because it's not just about the tech anymore, it's about outcomes, faster decision-making, reduce costs, safer AI adoption. So whether it's GenAI, copilots, open governance or these Lake Houses at scale. The theme is clear intelligence needs to be built into every layer. Of your, uh, technology stack.

Sameer Dixit: That's fantastic.

That's a great starting point. See what I'm seeing right now? To summarize what you said. A lot of people come to us at Persistent where they say, is my data ready for AI? And, that's a very abstract concept, right? Is my data ready? You know, we have to be able to convert and translate that into measurables and definite outcomes, so that, enterprises can understand if their data is ready.

And that's where platforms like iAURA — that we have built — come into play.

On that note, you know, I wanted to check with you in terms of what are the new capabilities or tools, that Databricks has launched recently, in this space, and how do you think partners like us Persistent can bring them to life?

Gregory Taylor: Yeah, it's a great question, Samir. And you know, we're pretty excited about all the innovation coming out of Databricks right now, but there are three that I'm really excited about. LakeBase, LakeFlow, and Agent Bricks. These aren't just incremental upgrades. They're foundational shifts in how we're changing Data and AI are delivered to the enterprise.

Let's start with LakeBase. So this is huge. LakeBase is a completely new architecture for databases as it brings everything customers love about Lake House. It's a transactional workload. It's built on open standards like Postgres. It runs serverless and separates storage and compute, which we know is a game-changer.

This means we can scale elastically. It costs next to nothing when idle, and you can spin it up in seconds, which gives flexibility. But what's even more exciting is, it's built for AI so that you can branch databases like code. Let's imagine that giving every AI agent its own sandbox, a high-fidelity data environment to test, learn, and rewind. That's a huge, massive unlock for building next-

generation intelligent apps. Now, Persistent has amazing skills to help developers and customers replace legacy OLTP systems and design entirely new digital experiences around this.

The second is LakeFlow. If you've ever wrestled with brittle ETL pipelines or stitched together four different tools to build a data workflow, LakeFlow is your answer. It gives you ingestion, transformation, and orchestration all in one place. It's fully native, fully governed, and incredibly efficient. This is a game changer for services partners like Persistent. It means now, you can accelerate modernization programs, reduce costs for customers, and deliver much cleaner architecture to your clients.

And finally, definitely the last but not least, what we're super excited about and really innovative is what we're calling Agent Bricks. This is where GenAI gets real. Agent Bricks helps teams build high quality production grade intelligent agents easily. Not demos, not prototypes, but fully managed, secure agents grounded in real enterprise data automatically, evaluated based on LLM benchmarks that are auto optimized. Imagine if Persistent could deploy GenAI agents that automate insurance claims, resolve support tickets, or even analyze contracts, all with speed, accuracy, and trust. All of these solutions are moving us from a world of fragmented data stacks and disconnected AI pilots to a single unified data intelligence platform.

And with the role of persistent, you guys help customers move faster, get incredible value, and deliver at scale. So, the future of enterprise data and AI between the two companies is just getting started.

Sameer Dixit: It's pretty cool. Right. You know, we were at the Databricks Summit, earlier this year in June, and the excitement was palpable right around LakeBase, LakeHouse Federation, LakeFlow, the Unity Catalog, and Agent Bricks.

That brings me to my next question. If I don't ask you this question, this podcast is null and void, and that's agent tech.

From your perspective, where is agent tech already being used in practical, impactful ways? Because everybody likes to talk about it, everybody's talking about it, and then the agent tech is just the, you know, like icing on top. But, your data gonna power it. Right!

Gregory Taylor: Yeah, Samir. GenAI is a buzzword, but at Databricks, it's more than just a buzzword. Agentic AI across our industry, outcome maps, you get real use cases emerging quickly.

So, let's take financial services as an example. Imagine an agent that automatically monitors transactional data flags, anomalies, drafts a suspicious activity report and routes it for approval, all without human intervention unless it's needed. So that's not just automation, it's intelligence in motion.

Things like financial crimes, auditing, financial transactions, all of these being done at scale with complete governance, over the top with human intervention exactly where it's needed, and being able to do it cost effective and really fast.

Another area is healthcare. For example, we're seeing use cases where agents can read unstructured medical notes, read coding classifications and recommend them, and then surface patients who are eligible for a specific care or pathway. All grounded in enterprise data. Fully auditable. Fully governed on your own data. You own it, you secure it. You only share what you want to share. That's the real importance, real data intelligence. It's not just general or what's out in the public. This is your own data. You own it. Don't let anyone take your data away from you, but you need to be able to use it, categorize it, and make actionable insights. And all of this works because we have a strong data and governance foundation. GenAI is a another use case and a number of use cases, but where LakeBase and Agent Bricks comes in without high quality, trusted, well permissioned data.

These agents aren't safe or securable. This is where Persistent has a huge role in helping customers map their data outcomes, industry outcomes to define the right behaviors of these agents, and integrate them into real world workflows, and then the industry outcome maps give us the targets. So, GenAI gives us the means to reach them much faster, but what we're seeing is in a world that's moving where AI won't just support decision-making; it's gonna take action. That's where real value is gonna begin to live, and I think the two of us are in a really unique position to make that real.

Sameer Dixit: That's fantastic. Right! Totally, agree. We are seeing that every other customer who comes to us has to see something around agent tech. I won't say we are struggling, but people have different concepts and views about what agents really means and you know, what you just said right now actually makes a lot of sense. At Persistent, what we are doing is applying a lot of agent tech in real life use cases, you know, to prepare data for AI. So be it code migration, self-filling pipelines, auto documentation, even the conversation analytics, where you do talk -to-my data use cases, is where we are seeing a lot of traction.

Out of curiosity, there's obviously, Agent Bricks, you know, is what you have, what is it that you're doing at Databricks, in the pipeline at Databricks for agent tech?

What are the specific things you're adding into that? How do you see, somebody like us can play a role into that, along with Agent Bricks, there are various other things I'm sure you're doing to make sure Databricks is Agentic AI-ready.

Right. So, If you could dwell on that a little bit.

Gregory Taylor: Yeah, I'd love to. I'd love to dig into that a little bit more, Samir, because you know, Agentic AI is no longer the theoretical as it's becoming embedded in across every layer of the data and AI stack.

Databricks is not just enabling AI agents as a separate feature; we're making them native to the platform through Agent Bricks. This means customers can now build intelligent agents that are grounded, enterprise data context to aware fully production grade. And it's not limited to one workflow. These agents can monitor pipelines, take action, write reports, and respond in natural language, in English and whatever the language that you want to communicate.

We're also designing all new capabilities with LakeFlow and LakeBase to support generic use cases from the ground up. Whether it's an agent managing a transactional system or optimizing data orchestration in real time, this intelligence is built in natively. With Persistent, we're taking it even further. We're embedding these agents into real customer environments, be it agentic managed services to proactively resolve issues, or insights agents that give business users confident, conversational access to payments data. These live production-ready deployments, not just pilots. We can take some of the agency in the connective layer between data decisions and action.

And really, this is just the beginning. We're building into the fabric of everything that we do. And so the real benefit for, anyone that's thinking about partnering with the two organizations is what's the customer value they get out of this, they're gonna get speed, governance and world-class innovation, which gonna put them in front of their competition and make their customers wanna do business with them.

Sameer Dixit: That's fantastic. This is where innovation comes from, the platform itself. The power of Databricks speed is something that we could bring and, you know, together we could, make this happen. The other thing I was thinking, you know, agentic, and this is like, a problem in itself. On one side, you have agents which are supposed to be autonomous. On the other side, you're talking about governance, which means less autonomy, more often known as, ethical or technical or all those things, right?. So as we scale up agent tech, first, there are humans who are worried to take decisions based on the data that they have.

And here you're talking about agents who are gonna take decisions **autonomously**, which is a deep word, and, you know, people use it, just like that.

What guardrails, ethical or technical, do you think of? Do you believe they are essential as we scale up Agentic AI? I know how Databricks is scoping up to that, because that's one thing that I hear a lot. Okay, let's do this. When there's getting, getting these platforms into enterprises and use cases is really a challenge because then they are to deal with the security and a lot of other things, which, you know, gets into this. So, what is your view here?

Gregory Taylor: It's a critical question. Honestly, it's one we have to lead with so we can't bolt it on in the end. When we talk about Agentic AI, we talk about systems that have the capability for potentially autonomous action, and that brings incredible opportunity, but incredible responsibility.

We think about a number of principles, the first principles, transparency. People need to know when they're interacting with an AI agent, what it can and what it can't do. They need to understand how decisions are being made. This builds trust in a regulated industries. It is non-negotiable. The second is control and oversight. Even the smartest agents need guardrails. Human in the loop design with clear escalation paths and continuous monitoring are essential not just for compliance, but to ensure these systems are working as intended.

The third, which has been critical to Databricks and everything that we do, is governance around the data.

We call it data governance. It's with a product that we call 'Unity Catalog'. It is at the core of everything we do. These agents are only as good as the data they're grounded in. If the data is biased, incomplete, or improperly governed, the outcomes will reflect that. This is where the Databricks platform plays a really big role because we give customers the tools to manage data quality, lineage, and access at scale with Unity Catalog.

So if we think about Persistent, you guys brings real value here by helping customers operationalize their governance, not just talk about it. Whether it's enforcing role-based access to finance, or ensuring healthcare agents, or aligned to kind of privacy regulations. The ethical design has to be embedded from day one. So ultimately, we're not just building intelligent systems, we're building hopefully responsible ones. And that starts with asking the right questions early and keeping humans in the loop throughout the whole process.

Sameer Dixit: That's fantastic. I like your perspective — you know, on one side you are basically talking about governing the data itself using Unity Catalog.

And there, you know, is something which is design time, which people are always good at. But runtime is where it all matters because enterprises are moving and what you do at design will obviously change at runtime. And that's where we could play a role. And then the other is the whole agent layer where you want to bring more transparency, more trust, and all those things which are these two separate areas.

I love how you articulated that in terms of how Databricks is suited for looking at that. Right!

That brings me to my next question. Looking ahead. So we are building a lot of solutions around Databricks and IBS is one area that we have and I'm sure there are more. There's a Databricks marketplace and there are so many other ways that we could, you know, work together and filling those, I won't say white spaces, but those areas. Obviously, on the platform which you want to take to the production.

Firstly, a company like Persistent could play a huge value because of the leadership that we have in the data space, right? So how do you see Persistent and Databricks working together to create more

solutions like iAURA and, you know, you leverage the marketplace and leverage the other programs that Databricks has in this space.

So, you know, with that, let me just take a step back a little bit in terms of, a little bit crystal gazing of what is the future. I know, I know it is hard because what you say now may not be true, four weeks or six weeks down the line. But still, let's take a shot at it, over the next 12 to 18 months.

I know this is a very large period, but you can change it if you want to. What changes do you see happening in the data and AI space? Do you know how Databricks is planning for that future? A little bit, you know, understand to put your limb out there and, you know, take a shot at this, knowing perfectly well that you'll not be held for this.

Gregory Taylor: Yeah, it's a great question because the pace at which everything's moving 12 to 18 months seems like a lifetime away. And by no means am I answering on the behalf of Databricks, but you know, what do I see and what is my opinion? I've been wrong many times. So, you know, there's a possibility that my LLM may not be, you know, predicting all the data points in the future. But I see three areas that I think are important to the data landscape as it change. You know, first I think AI is become embedded in every workflow.

I think we're gonna move from pilots and dashboards to really agentic systems which help make decisions, take action, and constantly learn safely and at scale. And we believe this, probably gonna blur the lines between operational systems analytics entirely. And that's why, you know, we're starting to invest very deeply in things like AI.

Second, openness and interoperability is gonna win. Customers want choice. They want flexibility. They want control. I think what you're seeing now is Databricks is anchored on open standards. Things like Spark, Delta, LakeFlow, or Lakehouse ML, Flow Postgres. Giving Enterprises like a future-proof foundation, we believe that Open is gonna win over time, versus closed systems. So I think that's the direction that a lot of companies are going that you know, you can start to see that. And then third, I think customers are gonna hold us much more accountable for outcomes. And that's gonna take center, center stage. It's not gonna be about tech stacks, it's gonna be about business value through reducing risk and improving customer outcome or experience driving growth in their business.

And that's why we're focused on kind of these solution accelerators, industry outcome maps, co-innovation with partners like Persistent. As for how we're preparing, we're continuing to build a platform that's intelligent by design, governed by default, and open in every layer. And then we're working with partners like Persistent who can take that innovation and deliver it with precision, speed, and scale.

I feel like the next 12 to 18 months will belong to those who can move fast, but with trust, and we feel like that's the future we're building toward.

Sameer Dixit: You know, everything that you said about, and all the other aspects of what the future is ought to be. I think this is outcome-based is something that we hear a lot, significantly service-as-a-software is what this whole agentic is about. We are on a journey which is gonna be hugely transforming and not just for customers, but for everybody, right? Including Databricks, you know, as to watch the platform power, people like us in terms of what should you take, and openness is something that we hear.

Very, very much. Right. When we talk with customers, though, they would go with the stack. They really want to have the choice that if they want to go to a different thing tomorrow, they could. So, this whole loosely coupled architecture where you have got these different Lego blocks making a solution and where you can rip and replace it, rip is not the right word, but replace something with there is a better way to do it.

I think that always wins hands on and I always built on the same stack. I know Databricks is, that's why we love Databricks because it is open. There is no lock-in. That's what people want. They want to be able to decide that they're not married into something and they will have the choice to get out if things don't work out or whatever. The reason could be anything, but you must know you can do that.

So with that, you know, how do you see Databricks and Persistent partnership evolving from here? In terms of all that you said, solutions, outcome-based things, as a platform, you will give everything that is required to build an outcome, but the outcome is to delivered by somebody who builds the service. How do you see our partnership evolving from here?

Gregory Taylor: Yeah, great question. First, we need to anchor on what we're doing together and some of the things that are, you know, what we call able to be consumed by customers and, and I feel like iAURA is a strong example of how Databricks and Persistent can co-innovate to deliver like tangible business outcomes. It's built on the Databricks platform, optimize for scale, governance performance, and it's already helping enterprises turn data insights into action. What excites me is the potential to use to, to scale the iAURA. As a repeatable accelerator across many different industries like financial services, software sector, you know, potentially healthcare, where speed to insights is critical. It fits naturally with our Brick Builder solutions program. So I'd say we need to see you formally, kind of get that into the Brick Builder and, co-sell it and develop it globally. I think that's really important because I don't see that as just a solution. It's a blueprint for how we build and scale the next generation of AI outcomes together.

And I think, let's take what we have there and extend it out. My encouragement to the team over at Persistent is push us to make sure it becomes a Brick Builder and then push us to help you take it out to market.

Sameer Dixit: Sure, Cool! So that's nice. So, you know, Let's have some fun, a rapid-fire round! You know, nothing controversial here, so don't worry.

I'll have five questions, and just be frank, on the buzzer there's nobody that you're competing with, So, let's go. The first thing that comes to your mind when I say Agentic AI.

Gregory Taylor: I honestly think about how do I actually open my bank account a lot faster and I want someone to come and come and tell me, all the bills that I need to pay tomorrow. And I want someone to basically do that in a governed way, so I trust it.

Sameer Dixit: Alright, so your agentic AI is very personal. Alright, what's harder: is it harder to align teams or is it harder to align data?

Gregory Taylor: It's absolutely harder to align data in my opinion. Too many silos, too many things out there.

Sameer Dixit: What is the one technological trend that you wish would just disappear?

Gregory Taylor: The concept of big data. I wish it was gone because it, it's the biggest inhibitor I think to folks, you know, moving forward with LLMs and Agentic AI is the cloud that was around big data for many years.

Sameer Dixit: Okay, that makes sense. The Hadoop and, you know, the, the whole, I think the 2006-07. If you launched an AI startup today, what space would you go after?

Gregory Taylor: I think healthcare is a space I'd be after the ability to make an impact on humanity. With some of the capabilities and the data available now, I genuinely believe we're gonna solve and cure many of the diseases in record time. We're gonna create a healthier population, and we're gonna detect things a lot earlier.

I genuinely believe that's what I'd be looking at.

Sameer Dixit: : That's a fantastic, I was just talking about this with somebody else, you know, a few days back and think most people, this is an area to bet in. Not just a longer life, but a good-quality life. Both are important.

A last one here right now. Tell me about Databricks in one sentence.

Gregory Taylor: Databricks in one sentence. Databricks wants everyone on the planet to democratize Data and AI. We want to have a lasting impact on the world by giving people access to Data and AI, which delivers value.

Sameer Dixit: Fantastic. That should go on your website, Greg. I hope somebody from your marketing team listens to this.

This was well summarized, so thank you very much, Greg. Thank you for playing around, uh, for that rapid fire, especially. More importantly, thank you for your thoughtful insights, both personal and from a Databricks perspective.

Fantastic conversation. Everything from platform innovation to agent tech to responsible AI to what's next for the partnership.

Gregory Taylor: Thank you, Samir, for allowing Databricks to be part of this podcast. You're driving thought leadership and doing so many amazing things in the market. I hope we can come back in six to twelve months and all those things that we talked about are true, and true for our joint customers.

It's just a real pleasure being here.

Sameer Dixit: Great. Thank you Greg. For everybody else, thank you for listening to this Persistent podcast to hear perspectives about shaping the future of business and technology. Please subscribe and share this series with your network.

And if you have a story to tell or if you would like to hear from a guest which is of interest to you, please do drop us a note at **podcast@persistent.com**. We are excited about the road ahead and look forward to building what's next together with you.

Thank you very much.

Re(AI)magining™ the World

About Persistent

Persistent Systems (BSE: 533179 and NSE: PERSISTENT) is a global services and solutions company delivering AI-led, platform-driven Digital Engineering and Enterprise Modernization to businesses across industries. With over 25,000 employees located in 18 countries, the Company is committed to innovation and client success. Persistent offers a comprehensive suite of services, including software engineering, product development, data and analytics, CX transformation, cloud computing, and intelligent automation. The Company is part of the MSCI India Index and is included in key indices of the National Stock Exchange of India, including the Nifty Midcap 50, Nifty IT, and Nifty MidCap Liquid 15, as well as several on the BSE such as the S&P BSE 100 and S&P BSE SENSEX Next 50. Persistent is also a constituent of the Dow Jones Sustainability World Index. The Company has achieved carbon neutrality, reinforcing its commitment to sustainability and responsible business practices. Persistent has also been named one of America's Greatest Workplaces for Inclusion & Diversity 2025 by Newsweek and Plant A Insights Group. As a participant of the United Nations Global Compact, the Company is committed to aligning strategies and operations with universal principles on human rights, labor, environment, and anti-corruption, as well as take actions that advance societal goals. With 468% growth in brand value since 2020, Persistent is the fastest-growing IT services brand in 'Brand Finance India 100' 2025 Report.

USA
Persistent Systems, Inc.
2055 Laurelwood Road, Suite 210
Santa Clara, CA 95054
Tel: +1 (408) 216 7010
Fax: +1 (408) 451 9177
Email: info@persistent.com

India
Persistent Systems Limited
Bhageerath, 402
Senapati Bapat Road
Pune 411016
Tel: +91 (20) 6703 0000
Fax: +91 (20) 6703 0008

Persistent
Re(AI)magining the World

www.persistent.com