Re(AI)magine Conversations

Persistent Podcast | Episode 05

Re(AI)magine Data in Motion: Powering the Agentic Enterprise

Bidish Sarkar and Drew Horn

Bidish Sarkar: Hello everyone. Welcome to Re(AI)magine Conversations. Our objective here is to explore how AI technology, and bold thinking are disrupting today's enterprises and helping in the transformation journey. I'm your host Bidish Sarkar. I lead data integration and AI practice at Persistent. In these conversations we sit down with industry leaders and technology disruptors to uncover the real stories behind innovations, whether it's re-imagining how software or solutions are delivered, or how decisions are made, or how we are serving customers — it all takes shape out here. So let's dive in.

Joining me today is Drew Horn, Director of ISV Partnerships at Confluent. Drew, great to have you. At Persistent, we are really energized about this partnership and some of the outcomes we are seeing in the industry.

How we are doing and how we are feeling?

Drew Horn: Oh, great. Well, industry leader, I don't know, that's a tall, order to kick things off with today, but regardless. Great to be here. Excited about the conversation today. Data and AI, great topics!. It's front and center for our customers trying to drive transformation. It's been great working with you and the team, to solve our customer challenges, and we've been keeping busy, so thanks again for having me.

Bidish Sarkar: Awesome. Yeah, and these have been really exciting times and together we are enabling not just agility, but also observability transformation. We are able to make that impact in the enterprises. And that speaks directly into the piece that you are synonymous with — Confluent is synonymous with the power of real-time data streaming. How are you seeing it and more importantly, with enterprise coming into picture, how do you envision all of this?

Drew Horn: Yeah, it's a great place to start. We've been thinking a lot about this and I think the simplest way to put it is that real-time data streaming is operationalizing AI. Traditionally, businesses have struggled to get their AI projects into production with traditional batch-oriented ETL and reverse ETL pipelines, and they're just difficult to maintain at scale, and latency is oftentimes just a non-starter for these types of projects.

If you think about it, this is why event-driven architectures are gaining ground, right? Real-time streaming enables these AI systems to process clean data as it's created and it gives those systems more context, accuracy and agility to act with speed and precision.

Bidish Sarkar: Absolutely. And, you know, we are seeing a tremendous shift in enterprises; the way Agentic AI is now coming in, real-time data is making a tremendous impact.

That brings me to the rise of Agentic AI. How you are seeing Agentic AI being embedded into your ecosystem and what you're doing about it?

Drew Horn: I would say, the interest is massive. Here at Confluent, we already have hundreds of active projects in flight with Databricks right now, and they are generating huge throughput and consumption numbers across both Confluent and Databricks, really to the point where it's driving executive visibility with our partnership.

We know customers are using us for these use cases and many of these customers actually have some use cases in production. Just to call out one example, we are processing over a billion real-time inventory updates per day at a national grocery chain here in the US. And we are delivering all of this data to Databricks in a way that their business users can quickly and efficiently perform forecasting to effectively allocate workforce resources and proactively anticipate product demand, optimizing supply chain operations, and even improving warehouse order to delivery logistics.

What is really enabling this shift is that foundational work — streaming architectures, governed data pipelines, and scalable AI platforms — I mean, think of the different data sets you have to put together in a clean, efficient manner to drive those types of outcomes that I mentioned.

Then as companies like this move to Agentic AI, being able to leverage both quality data for AI with real-time infrastructure is going to be critical. We are super excited with some of the capabilities that are coming out soon to help support that.

Bidish Sarkar: Absolutely. I think you hit the nail there.

I think when we have conversations with our customer, this entire aspect of leveraging data that is available in a real-time fashion for business decision-making is at the top of their mind. But that also brings me to the question of governance, because as you are getting this huge feed of data at a very high speed, and you are leveraging that data, what Agentic AI allows you to do is actually make decisions instantly and execute on those decisions. But then in order for accuracy, the governance

plays a huge role. So, I'd love to hear, how you see governance playing? I know it's a big part in your thinking as well, I'd love to hear your thoughts.

Drew Horn: Yeah, I couldn't agree more. With Agentic AI, I don't think you can get away with it and launch these projects into production without governed data. If you want trusted AI, you need governed data, and that's where capabilities like lineage tracking, access control, metadata management, easy discovery of data products across different teams to help de-couple them to build and launch projects in parallel.

It's not just about collecting data, it's about understanding it, securing it, applying the appropriate structure that makes sense for your business, and making sure that it's being used responsibly. And, working together on some of these projects, that's one of the areas where I see Persistent really standing out. Your team has embedded governance into the architecture itself, which I think is mission critical for these companies.

Bidish Sarkar: I think this is where the domain comes alive with technology, right? This is where the collaboration between Persistent, Databricks, and Confluent — all comes together helping customers do this industry-critical or business-critical solutions is so important.

We are talking about industry use cases. You mentioned about the grocery chain, and immediately it reminded me of a couple of use cases. One in banking and financial services, and another in healthcare. On the banking and financial services side, we are working with retail banks processing millions of transactions — credit-card transactions — on a regular/ daily basis. What is happening is the data is now coming in a stream. So, this is where all the technology comes in, right? Streaming data coming in from multiple feeds, feeding into an AI system. Now, what that AI system is enabling them or allowing them to do is make real-time decisions on fraudulent activities; and based on that real-time actions are being taken, which can put a pause to certain of these activities.

The other example is in healthcare. So, we are working with this patient monitoring system in a healthcare network. What this hospital is doing is, it's bringing in data from ICU patient monitoring devices. We are talking of heart rate, oxygen level, ventilator data, infusion data. All of this data is coming in and then the AI system is enabling them to do a real-time alerting to nurses and healthcare practitioners to take actions. This can really save lives, figure out early steps or early detection of cardiac arrest, they can immediately take action.

So, there is an impact on the business side by bringing these technologies together. And the scope — it's just immense of how all of this can work together.

So Drew, how do you see this triangulation between Persistent partnership with Databricks and Confluent and all of this coming together, to make a real difference for business and for customers?

Drew Horn: Yeah, it's a good point. These technology stacks that are solving for these AI use cases, they are complicated. Confluent is just one part of an overall reference architecture, right? You have got tooling, LLMs, your Data and AI platforms like Databricks. So I think it's really important to work with a company like Persistent. I love seeing what you are doing in the field, in terms of packaging up these solutions, delivering best practices, and industrializing the process of building enterprise grade AI. It's a tough balance to strike with both agility and control. If you start an AI project and you have small issues at the beginning, when you amplify the solution, you're going to amplify any problems with it.

So, for organizations that are coming in, they are new to AI, new to Agentic AI, they're trying things out; maybe they know data streaming, but they don't know AI. Or they know AI and they don't know data streaming. It's great to have someone that come in and help support them and deliver those best practices to help them succeed.

Bidish Sarkar: No, absolutely. I think, the complementary skillsets bring to this entire piece — where we are able to stitch the entire domain angle with the power of technology together.

And of course, with Confluence, the entire aspect of real-time data comes in, along with the observability around with the inter-governance associated with it that you were talking about. And this is where we have built an accelerator called Magnify, which incorporates Kafka pipelines on your stack, and enables the entire governance associated with that.

Then, of course, there is leveraging Databricks to bring all these AI models to life and deploying them at a fast pace, right? I think all of that coming together is where the power is. Because business problem needs to be solved, end to end. It's not about solving one piece of it, and that's where all the three companies coming together makes so much difference.

Drew Horn: Agreed. You know, another thing that I just thought about when you were speaking, is the cultural component. A lot of these types of transformations are not just about taking technology and packing it up, and deploying it. It's about changing the way people are thinking and working together, right?

When you break down a data silo, we always say, you're also making structural changes to how teams work together and communicate together. And so again, those can be scary when you're working at an organization. I've been through them myself and when you're partnering with someone who's done this, 50 or a hundred times with other organizations, they have that insight to help coach you through all the squishy stuff like that. It's really helpful.

Bidish Sarkar: No, absolutely, and this is where some of our accelerators come in. As you mentioned earlier, accelerators are, again, tools, nothing else, right? But, the human element there is huge, where our experts can really make a difference. But some of the things that we have done in terms of packaging — that you were talking about and I mentioned about — Magnify, right, which

helps people with tracking latency, throughput anomalies across systems, and all of that is packaged so you don't have to build anything, right? All of that is available. So the entire observability piece, the entire actionable governance piece is done there. And then if you go to iAURA, which is our end-to-end Agentic AI-driven approach for data management for reimagining enterprise monitoring, maintenance, governance, all of that at scale. iAURA sits on top of Databricks fully. That makes a lot of difference.

But again, coming back to what you were saying, this tool helps customers. But at the end of the day what is important is hand-holding, helping with change management. That adoption piece is about decisions being made and decisions being executed, and that entire piece, end-to-end, is where all the benefits come in.

So, that gets me to this entire thing around vision and innovation. How do you see our partnership evolving to create industry impact going forward?

Drew Horn: I think a big piece of it is some of the patterns that we're seeing across industries, in the business outcomes these organizations are looking to achieve with all of these AI capabilities that are becoming available. I think with Databricks, and all of these projects that I mentioned earlier, we're starting to see a bunch of patterns.

Some of the key ones that I've noticed, hyper-personalization is of course a really common one across multiple verticals; fraud detection in financial services; we're seeing predictive maintenance in manufacturing and energy; supply chain optimization, I mentioned that a bit earlier, but that also applies beyond retail, right? This is a horizontal use case.

So in automotive, think about building cars; and in pharmaceutical, think about building medicines; and then efficiency gains in logistics and delivery. All of these use cases require the ability to source data from several disparate systems, clean that data at the source to create reusable AI-ready data products, and then deliver those data products in real time to platforms like Databricks for AI and Machine Learning, advanced analytics to again power those business outcomes that we're starting to see, that are a huge opportunity for organizations to generate real customer value.

Bidish Sarkar: Absolutely. Some of these features that you are talking about that both Confluent and Databricks are releasing, and we are seeing the stumbling blocks have been lineage quality, policy enablement, access management, all of that. Those are getting solved automatically by these platforms.

And then, we have always seen freshness, monitoring, drift detection, data quality metrics, model drift detection, all of that, right? The platforms kind of enable us to make those investments on actual use cases; like you are talking about in pharma, in healthcare, in financial services. I think, the fraud detection crime that we were talking about, hooking up all this, connecting with all these devices, all

these end systems, and then leveraging that data for real-time analytics. It's amazing opportunity.

At Persistent, the way we are looking at is creating those use cases for business impact and would love to collaborate, continue this collaboration with you and Databricks together. Any closing thoughts?

Drew Horn: I couldn't agree more.

You mentioned at the start of the call, what's really exciting in the data and AI space is that there's a shift coming, and we are all working to react to it and prepare for it. And that's moving from insights to autonomy.

The future isn't just about dashboards and taking data from a source, cleaning it, applying AI, machine learning, or whatever, and deriving those business insights for leaders or whoever to then react to. It's about creating a closed feedback loop. And I think, AI agents are going to be a big piece of that closed-loop feedback.

I think we are going to move from people starting to adopt and use data-streaming platforms like Open Source, Kafka, and Confluent, and see it more of as a must. It's a two-way street of taking data from its source, generating insights and then connecting that back to applications in the operational state.

It's really an omnidirectional data flow, if you think about it. Another analogy I'll use is when we moved from monolithic applications to microservices. We used to have a big tightly coupled system. You'd had a bunch of modules that were part of a single codebase, and then, we realized, shoot, we need to take these different modules. We need to break them up into their own microservices that don't just talk to one other microservice, but talk to many. And that's really where Kafka and Confluent got their start is. Organizations looked at embedding us, and allowing this technology to underpin an event-driven architecture for these microservices. And it's looking to me like Agentic AI.

These are going to be smart microservices. As opposed to having one or two Agentic AI applications that are helping do real-time orchestration, there will be dozens, hundreds of them. That's where I think it's going to be really critical to have a real-time data streaming platform to connect all of these systems together.

And with Persistent's engineering depth and accelerators like Magnify and iAURA, I think we are setting the stage for a new kind of enterprise, one that's truly real time with governed data and intelligent by design. That's my take.

Bidish Sarkar: That is great. And, I couldn't have agreed more.

You know, it's very funny. We talk about prediction models. Prediction models have been there for

a while. What is so interesting today is these models continuously get data in real time, learn from it, adapt to it and then actually, execute based on changes that are happening. And that's going to make a difference in terms of business decision-making, in terms of how new solutions would be much more effective.

Drew, thanks so much for taking the time to have this conversation with us. To our audience, thanks for tuning in to Re(AI)magine Conversations. In today's episode if we sparked new thinking. Please follow us and share it with your network.

If you have a story to tell us or a guest that you think we should speak to, please drop us a note at **podcast@persistent.com**.

Until next time, stay curious.

Drew Horn: Thanks, Bidish. Appreciate it. Looking forward to what we build next together. Have a great one.

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