



Persistent Podcast | Episode 08

Making Multi-Agent Systems Work at Enterprise Scale

Jaideep Dhok and Amit Singh

Jaideep Dhok: Hello everyone and welcome to the new episode of Re(AI)magine Conversations where we explore how AI, technology and the bold thinking are transforming enterprises. In each episode, we sit down with changemakers, industry leaders and disruptors to uncover the real stories behind innovation.

I'm your host, Jaideep Dhok, the Chief Operating Officer at Persistent, and joining me today is Amit Singh, Global Head of Partner Solutions and GenAI partnerships at Databricks. In today's episode, we'll explore agent AI and how the multi-agent systems work in large, complex setups. Amit, welcome to the show.

Would you like to introduce yourself for our audience today?

Amit Singh: Hey, Jaideep. Great to be here and thanks for having me. As you kind of introduced, I lead Global Partner Solutions here at Databricks, and particularly I focus on driving adoption of AI with our joint customers through consulting and SI (System Integration) partners.

And I'm super excited to be here to talk about, AI, as well as multi-agent AI systems. Really looking forward to the conversation today.

Jaideep Dhok: Thanks, Amit. It's great to have you on this episode today. Let's dive straight away into our main topic about agent AI and how to make the multi-agent systems work in large, complex setup.

I mean, Agentic AI is pretty much the top of the town. We all want to adopt Agentic AI into everything in our business processes to begin with. Having said that, we all see that it is in its inferences stage, it is predominantly at the exploratory stage, where our customers are trying to figure out how to have the Agentic AI built-up, forget the complexity of the enterprises.

Where do you start from and how do you really move from a Proof-Of-Concept (PoC) stage to a full-scale production rollout? What's your take on the influencer stage and the scalability for a full-scale production rollout?

Amit Singh: I think everybody is talking about agent AI and we are really in the midst of a new revolution almost.

If you think about it, what started with the Large Language Models (LLMs) with the GPT like two, two and a half, three years ago is now full-fledged getting into the agentic world. According to McKinsey, almost two thirds of enterprises haven't really scaled AI systems across their enterprises yet.

They are thinking about that. In fact, 62% of the respondents from that McKinsey AI adoption study said that they are exploring how to deploy Agentic AI across their enterprise. So, a lot of customers are looking at it. Many of them are starting with pilots and they're thinking about what they can do with LLMs and how best to think about deploying Agentic AI across their enterprises.

I think many of the organizations are starting with taking a specific workflow and figuring out how they can automate parts of it. How they can have or how can they use the power of LLM to really do the planning, do the multi-step reasoning, and then think about executing, not only just kind of in a question-and-answer format, but also how they can take action using this workflow. So, a lot of times they're kind of thinking about the workflow in the sense of how to apply AI to that workflow and how to kind of put in agents into that. In some cases, customers are also looking at completely re-imagining the workflows with AI first approach.

So, I feel like those are a couple of approaches that enterprises are taking, and we are still fairly in the early stages of this AI journey. There are a lot of possibilities, a lot of great opportunity for both the enterprises and the customers, as well as the players in this space.

Jaideep Dhok: That's right. I echo that view completely.

Even at Persistent, we look at this slightly more holistically. So, in order to bring the reality to the hype, what the whole GenAI and Agent AI bring to the table, we have a pillar approach. We have four pillars that support this particular strategy.

The first one is the engineering productivity. So how do you build software applications within the enterprises. How do you know, expedite that particular process?

Second is the productivity in the business operations. You know, be it customer experiences, your back-office operations, and so on and so forth. But underneath both these, you need a strong foundation coming from the overall enterprise readiness to adopt AI, which is where the data becomes very, very crucial.

You know, to have the right data with the right quality, the whole governance, in terms of both the data engineering and the data management.

And then obviously below all of this, you need to ensure that you have a very strong foundation of the responsible AI and ensure that there is whole bunch of governance and guardrails to adopt AI at scale.

Now from that perspective, Amit, at Databricks, how do you support your customers in bringing the trust in the overall AI systems?

Amit Singh: Yeah, great question, JD. In fact, trust is paramount when it comes to enterprise adoption of Generative AI or Agentic AI for that matter. At Databricks, as you kind of pointed out, the importance of data is really critical in all of AI.

I mean, if you go back to traditional Machine Learning world, or a kind of early days of AI, a lot of it was pretty much trained on your own data, right? You will build custom or classical ML models and deploy those into production for use cases like prediction, and pretty much all of that was based on the quality of data that you have.

So, if the quality of data is not good, then whether it is Machine Learning or AI or even GenAI for that matter, it's going to give that bad output, right? It's essentially garbage in, garbage out. So, quality of data is absolutely important.

The second approach that we have here is that we strongly believe that it's in the best interest of customers to actually think about building models to where their data is, rather than taking their data to where the models are, right?

So, I think that's where Databricks plays an important role. The way we are, the way we are kind of seeing customers utilize Databricks data intelligence platform is in three ways. Particularly, those are also some of the key differentiators for why customers choose data, choose to build their GenAI, or agent applications on Databricks.

Number one is of course, evaluation capabilities that Databricks offers with the built-in tools and capabilities within the Databricks data intelligence platform. The reason evaluation is super important is that a lot of these models, LLMs and foundation models are non-deterministic in nature.

So that's where customers want that certainty and they want their model output to be close to the best. Right? And they want to do it at a cheaper cost as well. So being able to establish that cost and quality trade-off that's kind of number one.

Second is choice, right? Databricks platform offers choice to customers, whether they're looking to build in any architectural approaches, whether like simple patterns like rag or agent. Or even for example, like fine training or pre-training, a foundation model from scratch. Databricks platform

offers that single case set of capabilities within the same platform so that customers do not have to go out of the platform for any of their use cases and choice.

It is also important that there is no one model. That fits a given customer's use case. So, customers will typically look to rely on multiple models and be able to switch quickly from one model to the next depending on their use case. Databricks offers all the leading models, like we offer models from OpenAI, we offer models from Google. In fact, Gemini models I mean some of the latest model Gemini 3 are a label on Databricks platform. As well as models from Anthropic. Then open-source models from Meta, right? So, if you think about it, we are the only platform out there right now that offers that choice to customers.

And then finally, the third thing is security and governance. You can do all of that security governance model, lineage, data lineage. Governing the models as well as data through Unity catalog within Databricks intelligence platform as well, right? So, these three things combined really differentiate and make it easy for customers to pick and build on top of Databricks data intelligence platform.

Jaideep Dhok: You made two very interesting points. You bring from Databricks perspective models to the data and not the other way around, which is a very comforting thing as far as the risk and compliance groups are concerned, and I know Databricks and Persistent have worked together in multiple engagements.

I remember one such, example where we both worked together for a large Japanese financial services organization. Their whole ask was how do you ensure that they have the data that complies with their internal asks, as well as the regulators also. And not just basically bring AI to the compliance need, but also do that in a very cost-effective manner.

At the same time, Amit, what I also see is the context of the data is very crucial, and that's why I like your comment about bringing models to the data. How do you see Databricks help in terms of ensuring that the context is given to the model, which is where the maximum benefit comes from?

Amit Singh: Yeah, I mean, I think you bring up an interesting point! So, besides the regulated industries like financial services and healthcare and life sciences, context is absolutely important, regardless of the use case. And in fact, we are at a point where these LLMs are pretty much available to everybody, right?

Like anybody anywhere can use these large language models. What really differentiates a customer's use case and how they're going to get value from deploying a use case is by bringing their own proprietary data and feeding that to the model. So, prompt engineering can take you only so far.

It's all about what we call as context engineering now, right? So how do you provide the right context, whether it is two large language models or to the agents or to multiple agents operating together? How do you provide that right context at the right time to when multiple agents are

thinking about planning an activity or thinking about deciding on a specific thing with, particularly in the regulated industries with human in the loop, human being in the middle.

So, a lot of that context, if you think about it, comes from the data. Like I said, right? I mean, bringing models to where data is that just makes it so much more powerful. It provides the right context and really increases the level of accuracy and the output that customers in the regulated industries are looking for, and they're able to do that at a cheaper cost.

In fact, we have launched an application called Agent Bricks and. It's a low-code, no-code tool. It makes it really easy for even non-technical users to get started with building agent applications. In fact, multiple agents working together through multi-agent supervisor, right?

So, they can get started. They just point what exactly are the use cases that describe it in a declarative format. And the agents get to work, they provide the right trade-off for them in terms of cost and qualities. So, customers can pick and choose where they want to be in on that period of frontier of cost and quality.

And then they can quickly deploy these agents into production, and they can retrace from there. So that's how kind of Databricks is making it easy for customers to provide the right context and do that even for the regulated industries.

Jaideep Dhok: That's a great point, and I'm very excited about what we can do together on Agent Bricks.

But going back to my original question on agent AI, what I'm increasingly seeing is obviously there is a lot of excitement about Agentic AI and how it can change the world.

But, there are two critical aspects. One is what we talked about so far is the whole data availability or the enterprise readiness to get data to bring that required foundation to do, let AI do its magic. Whatever it claims it can. More importantly, reimagine the processes in the sense that if I'm talking about, to go back to your example, the financial organization that we worked upon together, it is about compliance requirement where we are not only trying to bring agent AI, but we are trying to bring it in the context of reimagining the overall compliance processes internally within the bank.

From your side, how do you see the top-down view of reimagining the processes coming into the way? How is Databricks supporting that particular view for your customers?

Amit Singh: Yeah, that's a great question. If you think about reimagining the processes, let's take this example of this compliance process, right?

What is important in compliance is to be up to date with the latest in regulations, right? And making sure that the bank or the financial services entity stays on top of the compliance. Be it their data, their systems, their people, are they in tune with that. So, one is capturing these regulatory updates on a frequent basis and making sure that you stay current. And second, making sure that whatever regulatory updates are coming, your systems are staying compliant to those new regulations that are coming out. So, the way, if you think about putting this into a process format, there are two ways in which you can do that.

One is you can take a look at the existing business processes that you already have, and those could be any SaaS applications, or it could be any custom-built applications that customers already have, and see how your, compliance agents, as well as your compliance officers are interacting with that application.

Then kind of do that process mapping end-to-end and figure out, and that's where I strongly believe that we've partnered better together with partners like yourself, right? Wherein you bring in that process mapping capabilities, which is really an expertise from a services partner like yourself.

You bring in the, that necessary expertise of getting customers data ready. Because I think getting the data ready, like whether it is in legacy platforms and system migrating that over to Databricks, or if it is like about data quality, or data relevancy, or data accuracy, whatever it is, right?

I think you can do the job really well in terms of getting all of that set up done and that process mapping part done. So that's one way of doing it, which is getting the process mapped out, end-to-end, and then figuring out where AI can play a role, where AI agents can play a role. And automating part of that with reasoning and planning capabilities that latest of AI agents bring to the table.

So that's one way of doing it. The second way is really where you are taking an AI native approach wherein you are kind of completely thinking about reimagining the workflow.

So, let's imagine that the only task at hand is to stay compliant, right? And you do not have any of these systems in place. You're completely thinking about how I can design an agent workflow end-to-end, which is like completely agent native with multiple agents operating at once. You have very specialized agents catering to each of the pieces. And then you have that human oversight that pretty much only comes into the picture for regulatory reasons, for making sure that the decisions are taken accurately and where human oversight is actually needed, right?

So, humans work in tandem with that AI natively imagined process workflow. Which partners like yourself again, can help build, right? I think those are the two ways in which enterprises can think about solving a problem. Like that one is more of a bolt-on approach. I think that's the approach that a lot of organizations and enterprises are taking right now, which makes sense because it's hard to completely rip off everything and get started with an AI native approach.

AI native approach is probably the longer-term approach. Some organizations are starting to think about that. That's really the future. I think a lot of these workflows, over a period of time, will get reimagined completely with an AI native approach.

Jaideep Dhok: I agree with you completely. In fact, our platform, what we call as iAURA, does precisely that, and when we go to the customer, we just don't try to take the bare technology as a mirror solution to them. We try to translate that into a business process, understanding and then helping them to reimagine the processes, obviously powered by Databricks platforms underneath. So, I completely agree with you on that. One of the key observations I'm also seeing is when you have to scale up these Agentic AI solutions further, the ability of the teams, both on the customer side as well as from our side, to think about the overall ROI is becoming very crucial. Why am I doing this reimagination of the process? Why am I trying to bring Agentic AI to create an end user impact or create a better employee experience internally within the organization and so on.

How critical do you think the whole ROI is to the whole agent AI based transformation?

Amit Singh: Oh, ROI is absolutely critical! A lot of initial adoption of AI tended to focus on reducing the cost and increasing the employee productivity. I think the conversation started from that.

And I think now it's moving towards how you can use AI to potentially also have an impact on the top line, right? So not just thinking about the bottom-line cost savings, as well as like making your employees more productive, but how you can deploy. The AI capabilities or these agents to generate additional revenue for you.

That could take different shapes and forms depending on what industry a customer is in, right? Whether they're a B2C, B2B, it can mean different things to different people. But that's where the conversation has evolved. So, I think ROI plays an important role. Absolutely.

I think if you think about ROI, it has two pieces. One is the investment side. A lot of times customers initially focused on that cost side of things. The second is the return part of it. The return can come from reducing the cost, but it can also come from increasing the revenue. So, if you can deploy these multi-agentic systems and have an impact on the revenue, and at the same time reduce costs, you can address both the numerator and the denominator, and that can really have a meaningful impact on the ROIs.

So, ROI is absolutely important, but I think the way customers are thinking about ROIs is changing.

Jaideep Dhok: That's great. Now moving on, what I'm also seeing essentially is the need for us to bring organizational change. So while the top-down mandate is great, the technology is great, the ability for me as an individual, as a participant in the overall agentic transformation. The way I would change my way of working potentially, if that is what is required, is very crucial. You talked about

how the input data is important that comes from multiple systems across enterprises, but then every system and the people and the processes behind them need to change as well.

How do you see that being taken up as a part of the wider rollout of Agentic AI implementation?

Amit Singh: That's a great question. In fact, what I've seen since the advent of LLMs and like all of this, what do you call the adoption around Generative AI and agenda is that a lot of times customers want to invest in it and they will invest in it.

And I mean, I think we saw that POC to production blocker, right? But even when customers make it all the way to production, the employee adoption of that AI is still not up to the mark. That's where a lot of customers struggle. So, one is that POC to production leap. The second leap is scaling that AI across the enterprise and making sure that their employees are using it.

So I feel like that's where, uh, partners like yourself, uh, with services capabilities, and, Particularly capabilities around organizational and change management can play a critical role. And I've seen, uh, and I mean I've seen your capabilities, in fact, uh, you make it easier. And that could be about, uh, like doing organizational change management consulting, as well as making sure that you are enabling and training the employees on the customer side on a frequent and regular basis.

So, I think that's super important, right? But your job is still not done. Once you have deployed AI into production, I think the actual job starts from there, to make sure that our employees are taking and reaping the benefits of that deployed AI. The good part is that AI has now made its way all the way to like a regular consumer, right?

AI and applications like ChatGPT, Gemini, Anthropic, right? Cloud. All of these are available to everybody. These organizations have made it easy for day-to-day consumers in any corner of the world to use AI. So, the good thing is that consumers are becoming familiar with AI and they're using it for their personal use cases. So that helps. I think that helps narrow down that learning gap for enterprises when let's say the same user is using that in their enterprise or in their work-setting. I think that's the good part, so that learning and change management effort, particularly with a technology like this is shortened.

I think going forward, when once the applications are deployed in production, the rollout and adoption will continue to increase as well with, of course, support from the board, from the leadership, all the way to the top-down approach, right? Going all the way to every single employee across the organization.

Jaideep Dhok: Yeah. And thankfully the technology is also there to help us. Like I remember a common engagement that Databricks and Persistent did, together for a leading scientific instrumentation leader in the U.S. We solved the problem of an enterprise-wide data fragmentation, multiple lines of business, multiple systems of records, and then data coming from so many sources.

So that just took, you know, so much enormous time to basically bring it to a stage where it could make meaningful insights. And then something that the downstream systems and the teams can leverage. Using Databricks platform, we use the AI to process data better. It was a metadata-based intelligence that Databricks platform helped us to gather.

And that basically helped us alleviate some of those concerns and friction points as far as change management is concerned. And that was a great success for all of us, together.

Amit Singh: Yeah, no, that's a great point. And in fact, what I've seen with, I think you mentioned iAURA, right? Uh, that IP of yours, that solution Persistent has really delivered value to a lot of our joint customers. I think the best part of that is you have really decoupled the headcount. I mean, a lot of times customers think about like, when they think about services organization, they think that they will bring an army of people to support and the deployment and the implementation.

I think we have decoupled that with that IP. You have used AI as part of iAURA IP itself. And that gets the job done faster. That gets the job done at a much faster pace with fewer resources. And in fact, like I was reading about one of the studies that HFS research has published, Persistent has done a great job of decoupling from that headcount based growth model that a lot of these services organizations are known for.

I think that is a proof of how you're doing that with tools like iAURA and leading with the IP rather than with the headcount in many of these cases. So, I think that helps organizations get to their data readiness faster, and then ultimately get to the AI implementation and adoption at a faster pace as well.

Jaideep Dhok: Absolutely! We are very deeply invested into that and very proud of what we are doing along with Databricks. What I'm really excited about, personally, is the pace at which these things are moving, what seem to be a slightly tedious transformation to adopt is getting increasingly more efficient and effective, you know, every single day, every single week.

So, the pace of change, the maturity, that it is gaining across the board is what excites me the most. And obviously, Databricks, as a platform, which is becoming more and more AI native, is the right tool for the enterprises to adopt AI technology better. From next 12-24 months perspective, what excites you more as far as the overall AI journey is concerned, Amit?

Amit Singh: Yeah, it's hard to say. I think like there are quite a few things that are happening and as you said the pace of innovation in AI overall is just so rapid that it's hard to keep up for a lot of organizations and enterprises and that sometimes can lead to mean analysis paralysis or customers not making the decisions or waiting to commit to something. I'm very excited about Agentic AI and the evolution that is happening there. I think, like I said, we are in very early stages of what is possible with Agent AI.

I think there is a lot of talk about displacement of workforce. I feel like that's a little bit overblown. We'll see what actually happens, but, I think we'll see the complete reimagination of the business processes and workflows.

Jaideep Dhok: Yeah, no, absolutely. One thing I can say for sure, that Agentic AI is here to stay. Every passing day makes it more realizable and thanks to the strong platform features that Databricks brings in, it basically just makes the overall AI intake and AI transfer journey a lot smoother, which is what I'm sure all of our customers are looking forward to.

So, thanks Amit. And with that, we bring a wrap on today's episode. Thanks for tuning into Re(AI)magine Conversations. If today's episode sparks new thinking within you, follow the show and share more with your network. If you have any story to tell or a guest you would love to hear from, please do let us know at podcast@persistent.com.

Until next time, stay curious and stay inspired.

Amit Singh: Thanks, Jaideep. Thanks for having me, and I really enjoyed the conversation. I look forward to driving a real change and adoption of AI with our joint customers.

Re(AI)magingTM 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 26,000 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)magingTM the World