Re(AI)magine Conversations

Persistent Podcast | Episode 12

What It Really Takes to Productionize Agentic AI

Jaideep Dhok and Siva Surendira

Jaideep Dhok: Hello everyone. Welcome to Re(AI)magine Conversations. This is where we have honest, grounded conversations about how AI technology is shaping up. We go past demos, the hype, and talk about how AI is really working inside real enterprises. What's working, what breaks, and how do we really have the real outcomes at scale using AI.

I am Jaideep Dhok, the Chief Operating Officer at Persistent. I lead global operations and work very closely to see how we turn AI into reputable enterprise level impacts. Today I'm glad to be joined by Siva Surendira, the CEO of Lyzr. Shiva, welcome to the podcast.

Siva Surendira: Thanks, Jaideep. Glad to be here.

Jaideep Dhok: Great to have you. Before we begin, Shiva, we would like our audience to know quickly about what Lyzr is and more about you.

Siva Surendira: Thanks. Thanks, Jaideep.

Hello everyone. My name is Siva and I'm the founder and CEO of Lyzr. I started my career as a big data engineer with Tesco. Then I built one of the most successful AI-ML startups in APAC, which was acquired by a very large conglomerate.

And then I started Lyzr in 2023. So, I am basically a coder and I run the company during weekdays and in New York, you'll find me sailing over the weekends. These are the only two things that I do apart from taking care of my two boys.

But otherwise, Lyzr is an agent infrastructure platform designed for enterprises with a very key focus towards Productionization. I think that's where we are very different from other agent builder

platforms in the market because our focus is all around not just helping you build agents, but productionize them. We'll speak more about that in the next few minutes, but, that's in a nutshell what Lyzr is.

If we have to explain in three words, Lazar is the "AWS for Agents", highly modular in nature. You'll find all the building blocks in an agent development lifecycle available within Lyzr. And for large enterprises like banks and insurance firms, federal healthcare, etc, you can run the entire agent factory locally within your own cloud or on-premise systems.

So that's, uh, how we've designed laser and uh, kind of helps us to, uh, support enterprises, uh, build and launch agents at scale.

Jaideep Dhok: Great. Uh, thanks Siva. So, today we are going to talk about AI being leveraged and move it from concept to reality at pace. What does it really take to move from mere POCs and experimentations to large scale production grade implementations with tangible ROI?

But before we get into the theme, Siva, I want to start with a very obvious question. Why you talked about "AWS for agents"? Everybody claims that way. What makes your approach so different?

Siva Surendira: Yeah, I think if you look at a customer's journey without a platform like Lyzr, how does it look like? And that's where the answer is.

For example, if you want to build agents, I'll probably look at something like a LangChain and for orchestrations as well, I'll probably even look at combination of a LangChain and an agent. Then for databases, I'll start looking at Pinecone or a VVI eight, etc. for a database. Then after you do a databases, you move on to memory, where I look at platforms like mZERO or ZEP, and from memory, when you move on to say, responsible guardrails, I'll use platforms like Guardrails AI or Glia.

And from guardrails, when I move on to simulations, I have to look at platforms like Braintrust or Arise AI for simulations and evaluations. And from that, when I move on to observability, I'll again go to platforms like LangFuse or LangSmith. And then finally, when I have to give a UI for all my agents, I look at something like Lovable or Rep.

Now there are seven different platforms that are involved for an enterprise to get their agents out to production. In Lyzr, all these are modules. You have all these modules built in one single agent infrastructure platform. And we didn't stop there. We went ahead and also launched Native Voice AI features.

So now for voice, you may have to go to platforms like brand or retail, but we don't have to do that now because on Lyzr you also have voice AI features. So, agents, multi-agent orchestrations, knowledge basis, and knowledge graphs, responsible guardrails, agent simulations, agent eval, agent observability agent, user interface, and voice agents all-in-one single platform.

We just reduce this handshake issues that goes on with all these disparate platforms. Now you have everything in one single platform. So that's the answer. Why Lyzr? Because our focus is to give one unified experience and one central management layer for organizations to build and manage agents at scale.

Jaideep Dhok: Yeah, that's great. It takes a village to build the transformation together, especially enterprises, and which is what Lyzr brings to us so very seamlessly. But let's talk about the adoption. The theme is about how do you scale up from a POC to a production, full grade, full-scale implementation. Where do you see the ball gets stuck when it comes to enterprise adoption to begin with?

Siva Surendira: Yeah. Agent development lifecycle is pretty much like software development lifecycle, and I think Persistent among every other player will know this much better. Because you're practicing it every day. And you would also agree with this point. So, what we realized is like software development lifecycle, things get stalled when they get near to the User Acceptance Testing (UAT) phase, because agents are probabilistic by default. They're not software, they're not a binary code, which means you are not going to get the sign-off from business users unless and until they are extremely satisfied that 1000 test cases were run and they all perform as expected because it's probabilistic.

So, that's where we see a lot of agents get stuck. MAT studies said 95% get stuck. I would actually say practically 99% of the projects get stuck at that situation. So, which is why the real-world simulations can kind of come in handy. Very important.

Take Lyzr for example. We built an agent simulation engine that can run up to 10,000 real world simulations on any agent that you build. You can even build on LangChain and bring it to Lyzr for simulations.

This changes the way agents are tested. Imagine 10,000 simulations for a human to do that. It'll take six months, but we are able to reduce that to 30 minutes max because when you run these simulations, you see the results. And Lyzr as a platform automatically enhances the agent through a reinforcement learning loop continuously, so that all the issues are fixed.

So, we see that the productionization completely gets stuck due to lack of confidence that the users have in getting their agents to, I would say live. We recently saw this in one of the leading customer support AI companies. The CEO had to apologize to a retailer called GAP because the customer support agent that they built was going off track, speaking about Nazi Germany and stuff.

But these are the issues which continue to create doubts in user's mind and business's mind, whether they should take their agents to production. So yeah, that's, according to me, business acceptance. The user acceptance testing is where things get stuck.

Jaideep Dhok: Right. But in terms of the large-scale adoptability, what I have also seen Shiva is, the North Star is probably not defined that very well. And we talked about a deterministic project versus a probabilistic POC, which is what, you know, AI really brings in.

You wouldn't do that for a normal software project. You would put the right business case, you'll put the ROI, uh, you know, you'll figure out how, what business value is this application bringing to me? Or maybe some cost saving initiative, some regulatory compliance standard buckets. You would need to have something very similar around ROI or total cost of ownership and things like that. The real quantifiable business benefits at play, just as you would do for a normal software project, oftentimes getting missed out in the hype and the excitement of adopting AI. What have you seen work best from the Lyzr platform perspective?

Siva Surendira: So, before getting into ROI calculation, I want to highlight another aspect as well; on the productionization part. While we discussed about what happens once you end up building an agent because you get kind of stuck in the acceptance phase.

There is another huge problem that we see in the market, which is designing the architecture itself correctly on the first core because your data is not ready. Your existing systems are not ready. The big enterprise software vendors are not playing ball when it comes to sharing their data with third-party agents.

So, all these are constraints, which means you need to be very innovative in bringing the architecture together, which is another huge issue I see. Where, um, and I would actually say that the next few years, agent architects are going to be roasting the whole job scene because you need process, folks with process reimagination mindset to build this multi-agent systems.

So that is, again, a huge, hindrance when it comes to even getting agents to the testing phase. And moving on to the ROI side most, and I'll be very honest here and candid, the ROI calculation is not done against their existing software systems. The ROI calculation in the agentic AI world is always done against human labor that they end up spending in a particular process.

So, it's brutal, but I think that is the industry standard moving forward. We built a framework which helps enterprises figure out how much human labor they could save, whether in terms of displacing human labor or in most cases it is just increasing the productivity with which the team is working, right?

Jaideep Dhok: Yeah.

Siva Surendira: For example, we worked with a very large enterprise who had a 10 member marketing team, and the challenge that they put to us was, "Hey, I don't want to hire more marketing team. I don't want to fire anyone. But I don't want to hire and add more headcounts to the marketing team. I want agents to take up that role"

So, they ended up building 18 marketing agents on our platform. And the same marketing team is now working at 3-4X productivity. So, it's either cutting down labor on certain back-office operations or actually bring agents to really augment the work they do. ROI will have to be calculated either ways, either on the human labor-side or on the productivity.

Jaideep Dhok: So, the ROI quantified quantification is critical, obviously, which is oftentimes isn't done in the same manner as you do it for normal projects. Moreover, Shiva, what I've seen is, more than the ability to bring the change, it is about adopting the change, Lyzr being an agent platform as well. Agentic is not just plain simple AI agents being deployed, it's about how human-to-agent and agent-to-agent interactions also will shape up and should shape up in the to-be world.

What's your take on change management in all this and when you move from idea to production, do you see that is one of the critical success factors for the enterprises to be successful at large scale deployment of AI?

Siva Surendira: Yeah, good question Jaideep. And it's a very sensitive thing to be honest, because we are talking about asking humans, which is us, to change the way we think, change the way we operate.

Jaideep Dhok: Yeah.

Siva Surendira: ChatGPT is a very good example, right? All of us, like I still keep telling my wife to use ChatGPT more than asking me questions for everything. So, and I realized that over a period of time she was frustrated. She was like, I just want to speak to you. I just want to ask you, I don't want to use ChatGPT.

So, that is how humans tend to react. It is the solution. The technology is available, but we cannot just force it on them. They will have to adopt themselves to figure out whether it is useful to them or not. And we see this in the enterprise setting as well. For example, we recently worked with a bank that automated their loan processing for microloans using agents, where now agents are analyzing a loan request from various point of views, and then it sense its recommendations.

But the loan underwriters, they kept rejecting the AI analysis. They kept rejecting planned rejection. I mean, and the head of the business had this issue that, yeah, I think AI is not working because my loan underwriters are saying that it's not doing its job. And we went in and kind of did a workshop with them only to realize that it is just a mental block that they had.

They did, they just didn't want to adopt the technology. So, change management in agentic AI adoption is more about mindset management. But that's where it begins, where we need to sit with them, understand their insecurities, open up technology to them where we bring them as part of building the whole thing. Because your business users, your core team, who will get benefited by the

whole agentic AI automation should not be kept out of the loop. They'll have to be in the loop from day one where they feel that they are also involved in building the system.

So, I'm still heavily pronating towards the human mindset point of view, because according to me, having done all these deployments, that is the major block. I do definitely agree that we need the, I would say, process documents will have to change the way the teams will have to work with software systems, will have to change. That's the regular change management process, which professional firms like Persistent has done very, very well over the years.

Hence, I want to focus more towards the mindset change. Involve business users from day one in building these agents, because then the adoption is going to be very organic.

Jaideep Dhok: Right! Uh, I agree with you. The other side of the coin, which I see is change management in terms of re-imagining the processes.

So, change management in existing staff. Some kind of patchwork is good, but not good enough. Because if you are trying to really think natively in the agentic world, and I share the example all the time, if we were to be born on this planet earth today and leveraging AI as a technology, define something from ground up from scratch without any baggage from the history, we will think about in your example, the customer onboarding for loan for that matter, very differently.

It's just a baggage of the past that kind of, you know, drags us down in terms of our ability to rethink, is where also comes with the way that is also critical component in the agentic AI adoption at scale,

Siva Surendira: What you said was perfect because you need to think like a child all over again. Uh, right.

That's when creativity comes. We recently had a very large bank; I cannot name them. I was super surprised when they came with this thesis. They said, today, we try to onboard users on our app. They get about 40,000 downloads a month in their app. But apparently only 2% of them convert into a customer.

The rest of them drop off, but the bank actually came back and said, can we have agents onboard customers wherever they are? Which means can the entire onboarding process happen via WhatsApp or SMS or email instead of asking them to onboard on our app? It was very refreshing, to be honest, because for a very large bank to think completely differently, where they believe that with agents we can have the customers get onboarded wherever they are.

They don't have to necessarily visit this rigid system that we have. So, again, a good example of what you write, what you said, you need to reimagine the whole thing from scratch. Which needs a bit of boldness as well. And I think leaders are getting there with agentic AI.

Jaideep Dhok: So, be more like a child, don't get constrained.

I mean, you're not afraid of the world because you don't know how bad the world is. Perhaps you don't relate your history to kind of, you know, the education as it, the people say, kind of ruins your thought process. It's a very critical point. But talking about example, talk to us about one or two real life example where this thought process employed effectively, Shiva, and you had some real-life production grade scaled up operations and internalization of AI.

Siva Surendira: Two examples I would like to bring. One is in the specialty insurance space. There is this insurance firm who kind of write / underwrite insurance for cruise ships, for example, right? For an offshore oil rig, because each of these are in individual. I would say elements on their own. There is no comparison to the other things that they have.

Every cruise ship is very different today. The issue that they have is like 20% of the loan underwriters that, sorry, insurance underwriters that they have are really good in their job. And the rest, 80% of them are like average because when it comes to specialty insurance, you need to think of various scenarios.

Every single time it is like a consultant, McKinsey consultant doing their job analyzing what is a cruise ship? What is it made of? What all aspects should I look at? And we worked with them and built a dynamic multi-agent system, which based on the asset that is being underwritten, automatically spins up to 20 different expert agents which actually analyzes a cruise ship or any object from various point of views.

Think of a cruise ship, for example. We think about a ship engineer. A maintenance engineer as another point of view, a cruise ship. Uh, mechanic, etc. But we are also thinking about, uh, weatherman. Where is the cruise ship going to work? Be between New York and Miami? So, what's gonna be the weather there? So, weatherman and a point of view analysis.

These are all various expert point of view analysis that are done by agents today, which they then auto review; their own work. And finally bring a very detailed overview of what are the risks involved in underwriting insurance for a cruise ship? And more importantly, it also suggests what is the premium that they could charge.

So, we were able to do this and show them, and I asked them, okay, what do you think? How do you think it's going to help you? Do you think your insurance underwriters will be able to now complete things faster?

What, what the CIO said was very different. The CIO said this is going to improve the bar of all my underwriters, my top underwriters will be even better. But more importantly, my average underwriters will now work at the bar, which was already set by these top underwriters that they had. So, it's a very different outlook.

We were thinking about, okay, now the speed of underwriting is going to be good. The quality of underwriting is going to be good. But they said they're looking at it very differently. I'm not going to fire or let go of any of my underwriters. In fact, the overall underwriting process is going to be extremely strong compared to our competitors. So that's a good example of how this helped, in the insurance space.

But what I'm really proud about is another use case. This is one of the largest hospital chains in Netherlands. We had a call come in from the chief. The heart surgeon who works in the cardiac IC unit in that hospital. We got on a call, and we spoke, and we realized what she was trying to address. She's been using Chad GPT and that's how she got the idea. She said, Siva, I spend 12 hours a day actually in surgeries, but I have a huge set of patients in my cardiac IC unit. And there are the resident doctors and nurses who are tending to them continuously. If they have to decide on any medication or any change in medication, they have to wait for 12 hours until I come back from surgery.

And between surgery, I do get breaks and I want to see if there are agents that could do a multi-point of view analysis. And then help the resident doctors and nurses with immediate, I would say actions that they could take but keeping me in the loop. So, she wants to be the human-in-the-loop where it can go to her handheld device.

If she's in between surgeries, she doesn't have to come all the way from the surgery building to the cardiac ICU building. Very, very practical thinking that she had. So, we borrowed this, Stanford STORM algorithm that Stanford, uh, released, wherein again, any particular health record could be analyzed from various point of views, like a cardiac surgeon, a cardiology consultant, an MD and etc., So that the resident doctors and nurses get a pretty detailed analysis of what they could do for that particular patient in real time. They can then feed in their input and this goes to the doctor via hand-held device so that the surgeon can look at what is the diagnosis. First of all, she could validate that. Then she can see what the doctors, junior doctors are thinking, and then she can give her input.

And because of this, we are able to cut down this whole dependency system. But more importantly, every feedback that the surgeon provides gets recorded in a state memory, and that continues to enrich the overall decision-making aspect of the system because the system is learning continuously on how the surgeon is thinking, and this reduces the corrections required for future diagnosis as well.

This project is so successful now we are speaking to other departments within the hospital, like pulmonology and others. So, this is something that I'm very proud of on what we could do with the AI and these are just two among 1000+ examples we have. I want to bring these two to the audience today.

Jaideep Dhok: That's very interesting. Uh, direct lifesaving, in fact, if I may call it that way. And I can relate to something that we did, with Lyzer for one of the major banks on the credit memo generation. This is the commercial banking credit memo and credit decisioning is an extremely complicated task.

One, it is full of risks. So, when you underwrite for hundreds of millions of dollars of loan, that's a very different ballgame. And I remember us doing this on the platform, to understand what it took to actually document the process, ensure that you reimagine some of those aspects, basically, grounds up at times.

And it was a tough exercise to be frank with you. And then getting to the level where you make cure standard operating procedures, amended in the midst of the regulations to ensure that you have what is done using Lyzr as a platform completely, uh, auditable. Compliant with the regulations, internal as well as external, and then derive those business benefits at the end of the day.

And what stuck with me, at the end of that particular engagement, Shiva, if you recall, was the journey from assisting the underwriter to automating certain processes to going to the autonomous way of execution. So, from assisted where humans are still the drivers and the agents are, or maybe one agent, if not multi-agent, are supporting them to automating the same workflow, a slightly more AI-driven manner to the autonomous execution.

That journey was fabulous. What it led to was, interestingly, the impact it had, the example which you gave about this hospital change. You went from the heart or cardiac ICU to the other groups. That organizational impact basically is the most exciting aspect of this. And now that takes me to the next topic.

One of the most interesting discussion that you and I had in the past was about this organizational wisdom. People call it AGI, Artificial General Intelligence. It can come whenever it comes, hopefully in our lifetime. Siva, yours and mine! But you called it as OGI, Organizational General Intelligence, which is quite exciting. Talk to us about that.

Siva Surendira: Yeah. Uh, I coined the term OGI, Organizational General Intelligence, and probably that emanated from me being a Big Data engineer. I've always architected Big Data systems for Tesco, for Google, for capital, and all these large enterprises in the past, and I've seen what kind of problems that we ran into when we had petabytes of data in front-offices because making sense from all this data with weak data structures is a huge, huge task. Tesco was world's second largest retailer when I was working with them, and the amount of data that we had to deal with.

So, with agents, this problem is going to come back all over again. We are talking about petabytes of data being produced every year from these agents. An organization like Persistent might actually run a million agents in the near future, right? It's not going to be tens or hundreds actually. And that's when we realized, okay, what if all these agents that are working in different departments, can be brought into one central knowledge craft, as forming a central brain of the organization. It is like, say, Sandeep Kalra wants to understand what is happening across the industry. What does he do, which within the organization he brings all the leaders into a room and brainstorm and to understand how do we, because as a curious brain, he will be wanting to see the patterns he would want to understand and correlate data from multiple leaders.

Now, what do you do in your role? You actually bring all your teams together and you do the same. And it cascades all the way down to the, to probably the interns and junior employees. But in case of agents, you don't have to bring them into a room. What if they write all of this in real time into a knowledge graph, a living, breathing system that is finding out patterns in real time?

And we actually put this to work. What we did, we took the customer support agents and the product management agents and for one of the firms that we work, and we plugged them together. And it was fascinating to see what happened. The customer support teams were, obviously, they support the tickets, the agents were supporting the customer issues, and the OGA was able to analyze and find out the top 10 tickets that are issues that it always addresses.

On the other hand, the product management agent is something that sits through all the product calls internally. It then assigns work or sends the action items to the end of the call. So, the product management agent knew what are the active features that the product teams were working on. So, OGA was able to come up with an insight that these are the top 10 issues that your customers continue to complain about, but your product team is actually working on number two, number five, and number seven. They're not even thinking about fixing number one, number three, number four. So, it was able to bring that insight, which was very revealing of the gap that the organization had in terms of the organization across organizational knowledge.

Then they realized that the customer success leader is actually not speaking to the product leader as much as they should ideally do. And this is a simple but very powerful example of what OGA could do just with two departments. I imagine what OGA could do if there are, if all the departments are being fed into this central graph and you're analyzing and entire time. So that is the bigger vision that you're moving towards.

Jaideep Dhok: That's supremely exciting. I mean, when I think about it, Siva, in the context of enterprises, I mean, if you take a bank. A bank is a bank is a bank. Their regular entity, each one of them cannot do anything more than what the regulatory ask for.

Same product, same services, same interest rate, same reporting. What varies is how do they treat the customers, which comes from all things you talked about, you know, in terms of the operating models. How the departments talk to each other or don't talk to each other for that matter, and so on. Which is where OGI brings that clear differentiation because that's the one that really sets one enterprise different than the other in the same space.

You talked about the agents, and well, I'm looking forward to when we develop millions of agents basically. But the flip side of it is so many agents and the pace at which they get developed and deployed at scale is tremendous at times. So, the biggest aspect is what you talked about right in the beginning of this podcast, the governance, the responsible AI, the whole guardrails.

I'll probably summarize this as the overall safety of AI or safety of agents in the larger scheme of things from the larger platform perspective. And in general, what do you see as the agent safety being put to use, as a stepping stone for wide-scale enterprise level adoption?

Siva Surendira: Yeah, very good question and very important one.

Responsibility that Persistent holds in bringing out and recommending the best solution to the customers as well. So, there is this evolving concept, which is nothing new, but I think there's a new term to it called agent Harness that's happening in the market. So, what does it, what the concept is, what cons, is pretty simple.

Agents are probabilistic, as we all know, because they're covered by large language models. They are like the CPU. Agent Harness is nothing but all these guardrails and instructions and binary rules that are outside of these agents that restrict agents. It ensures that they do that exact work, not go off-track and into other divergent systems.

So that's what Agent Harness is all about. It all started with simple responsibly guardrails like, say toxicity controller, prompt injection manager, PAA, data manager, etc. Then it evolved into hallucination management modules, and now we are talking about step functions as rules that can sit outside of the core agentic AI system.

So, long story short, we should not move away from experimenting and launching agents at scale because, that is how you will know what works and what doesn't work. But a large enterprise should think of systems. Like Lyzr with the help of expert professional services, support from Persistent when it comes to implementing these harnesses around the whole agent, uh, deployment.

That's the only way the organization can still move at speed and compete with their other competitors or leave them behind and move forward. But at the same time have this comfort of safety built-in from day one.

Jaideep Dhok: Talking about Lyzr and Persistent, I know we have done multiple engagements together. We are one of the happy adopters of Lyzr.

Within Persistent as well, we run some of our operations on your platform. Siva, from larger standpoint, what do our clients have to see as a partnership between the two organizations? What do they get out of this?

Siva Surendira: I think we, to start with, we are super excited with the energy that we see across everyone at Persistent.

Today, Lyzr is at a spot where we work with Big Four. We work with management consulting firms. We work with literally all the top players in the market, and Persistent is one of the top players in the

market. And the energy is coming across from, say, a banking team or healthcare team, the business leaders, is encouraging for us and what we also started understanding is, I would say the thought process that Persistent is putting in recommending the right solution to the customer.

Because as an enterprise, you can get a wrong advice, and you may end up losing two years of your important time. This, we are at a very important place in history because enterprises are going to be evolving much faster. It is not like adopting an SAP system, wherein you could decide either today or five years from now when it comes to adopting agent AI.

It has to be today! It has to be now! And that's where you need a partner like Persistent who is ready to take this, these risks, along with the customer. So, we are definitely super excited, and we are, I think, just getting started. To be honest, 2026 is going to be a very big year. We are also leaning towards doing a few joint research projects with Persistent this year. Because it's all about setting standards and templates for enterprises to adopt.

So, it's going to be huge unlock for enterprises, starting this year onwards. Definitely looking forward to the partnership.

Jaideep Dhok: Oh, same here. Absolutely. Shiva, as we come to the close of this podcast, we talked about 2026. What's in store for you in the rest of 2026?

Siva Surendira: 2025 was toying with agents. Kind of in prototypes to figure out what is happening and Jaideep would also agree that a lot of enterprises, CIOs now have mixed opinion about agentic AI because they burned their fingers, money, time in 2025. But 2026 is when the real agents moving to production will happen and this will change the sentiment all over again.

We are already seeing success like in just in couple of weeks where we see systems going to production. So, 2026, the meter will move towards productionization. There will be even more pressure from the board, from the executive leadership team onto the IT teams and business teams to actually show ROI.

There is no hiding unfortunately, for large enterprises, for CIOs and for business leaders. You will have to show ROI because your peers are showing ROI. Companies who are executing well are able to show that they're able to save, they're able to move faster, they're able to provide better customer experience to their customers, and that's all that matters.

Eventually, your customers will have to feel the difference. The end customers I'm talking about. So, we are going to be hovering towards Productionization this year. And the secret sauce is, I would say two things. One is choosing a platform like Iyzr, which is a full-stack agent infrastructure platform, and choosing a right partner like Persistent.

And the other one is getting the agent architecture right from day one. Do not underestimate the power of what a good agent architecture can do for that particular workflow automation. Because you have to start reimagining instead of duct taping AI on top of existing processes. So, these are the two big themes for 2026, and we are super excited about it.

Jaideep Dhok: Oh, I agree. I agree. As we wrap up this podcast, one thing stands out is that 2026 is going to be more towards the adoption of AI at scale. We are going to see more and more successful large-scale implementation, because it has got all the ingredients it takes now for us through platform like Lyzr and the partnership from someone like us who brings the domain, the context of the problem that we are trying to solve.

So, if you have focus on the coverage on the topics you talked about, the agent architecture, the ROI, the guardrails. More importantly, the change management adoption basically, and coverage around that. That's where we are going to see a broad-based adoption of AI more successfully. We are seeing between both adoption quite a few very interesting at-scale production implementations.

So, it is to both the organizations and the partnership to bring a lot more client success in 2026. Siva, thank you so much for being part of this podcast and sharing some really interesting thoughts and examples of how to scale from POC to production grade implementation of AI.

Thank you very much for that. I appreciate your time on this podcast today and look forward to the partnership going forward.

Siva Surendira: Thanks for hosting us. Thanks Jaideep.

Jaideep Dhok: Thank you Siva. And for the audience, if today's conversation's spark in new thinking, follow the Re(AI)magine Conversations podcast with me and share this episode with your network.

And if there is any other topic you would like us to explore, please write to us at **podcast@persistent.com**.

Until next time, stay curious. 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 26,500 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