Re(AI)magine Conversations

Persistent Podcast | Episode 05

What Will It Take to Reimagine Pharma Workflows?

Abhinav Kohle, Persistent
Bharadwaj Popuri, Takeda

Abhinav Kohle: Hello and welcome to **Re(AI)magine Conversations,** the show about how AI is changing the enterprise, not in theory, not in demos, but in the day-to-day reality of how people work in companies. In every episode, we bring you someone who is living it. I am your host, Abhinav. I am an Offerings Leader for Business Hyper-productivity at Persistent, and today I am looking forward to speaking with Bharadwaj Popuri.

Brad, do you want to introduce yourself?

Bharadwaj Popuri: Abhinav, thank you for having me here. It's a pleasure to be here. As you mentioned, I lead what is called the R&D Hyperautomation team. We build solutions ranging from point solutions to operational AI. The challenge with these solutions is that they have to be GXP compliant, fast enough, and controlled enough for people to trust the output.

To give you some background about myself, I'm an engineer by training. I've spent the last 14 years in Pharma, mostly around clinical trials, operational management, and R&D operations, while finishing my doctorate in machine learning and cancer detection. So most days, I get to play two roles: researcher and operator.

Abhinav Kohle: Thank you, Brad. Let me set the stage, Brad. Pharma is under tremendous pressure right now. You have to move faster, operate in a pretty lean environment, stay ahead of regulators, and give employees and partners the best experience—all simultaneously.

Yet, when you walk into any pharma enterprise, you see that a lot of work is still being done manually—in emails, mailboxes, and similar channels. There are a lot of manual handoffs: approvals sitting in someone's mailbox, data in four different places. These are classic challenges that every industry is going through. At a global scale, all of this creates quite a bit of friction.

Let me take a step back. Forget about all the hype around GenAI and everything that's happening, Brad. From where you sit today, what is the challenge that keeps coming back in the pharma industry?

Bharadwaj Popuri: In most industries, you have systems that don't talk to each other. What you do is you connect them or you create an API to send the data. For pharma, the challenge is twice that

© 2026 Persistent Systems Ltd. All rights reserved.

because not only do we have to bring the data, we also have to bring the audit trail, the decision-making framework behind that.

I'll give you a simple example of a trial master file. One study generates thousands of documents. All of them have to be filed, classified and checked for completeness, and then have to be kept inspection-ready. We moved all these documents into an electronic trial master file years ago. But what's the state of the file? What's expected? What's missing? What's sitting in some subsidiary system that is declared as TMF content, but gets tracked in an Excel spreadsheet or a trial administrator's head or in a Teams message. This is where fragmentation is a challenge within pharma, because it's more than just the data. It's also the audit trail and the decision-making behind it.

Abhinav Kohle: Got it. Okay. Thanks for that elaboration, Brad. To summarize that, it's almost like a speed versus governance problem. The whole tension is around that, in my view, and that's where most of the automation programs have missed the point in the past.

For years—you and I have both been in the hyperautomation industry—we have seen how RPA has been used. We find a repetitive task, try to automate it, eliminate it, count the hours saved, and move forward. That just makes a broken process run faster; it doesn't actually address the real opportunity, right?

If you take a step back, the real opportunity is to look at the whole workflow end to end and see where it stalls today and why a redesign is required. And, in my view, that matters most in pharma. It is a space where you have to consider not only quality, but also compliance and regulatory requirements—and eventually, you are actually touching the patient.

So Brad, take me back, let's say, a specific moment in your career in the hyperautomation space. Was there a deadline that you nearly missed or an audit that you thought, "Oh, I should have... I can't keep doing this in this way"?

Bharadwaj Popuri: The challenge is not one instance or one particular audit. The challenge is the realization that we shouldn't be doing this, like you said, in the sense that we shouldn't be automating a broken process to begin with.

And every time I assembled a portfolio review or a business case for automation, every time I drill down into the details, the second, third level details, one of the questions that would keep popping up is: how can we reimagine this process? Why is it broken and why does it exist the way it exists? Is it because there are too many moving pieces of the puzzle that do not interact well or do not talk to each other? Do we just figure out a new way to do things?

So, that's the realization, the workaround of moving documents, moving data, those are all point solutions that will get you from A to B. But, like you said, pharma is under a lot of pressure. We no longer exist with zero interest capital. There's more pressure on pharma companies to get their drugs to the market faster. So, this no longer becomes like an efficiency conversation, it becomes more of a strategic conversation where it forces leaders to reimagine the processes, instead of

© 2026 Persistent Systems Ltd. All rights reserved.

thinking about it from the lens of "Should we automate this or can we automate this" to "I think we should look at the process and say, is the process how it stands today optimized for automation or AI?"

Abhinav Kohle: And that's what we see in our discussions with other customers at Persistent as well, Brad.

The challenge that we are seeing is more of the orchestration, the reimagination part of it. That's what is something that we need to relook at, in my view. These are core orchestration problems which now AI can solve in a very different way. The hard part isn't doing that task. It is almost like bringing the right data, the right people, the right controls, checks and balances in place.

That's what we want to do, and that's what brings me to the next thing, right? Why reimagine and not digitize and not automate? Brad, can you give me an example of a process in pharma that genuinely is ready to be reimagined? Not just automated, but completely reimagined the way it has to be done in the future.

Bharadwaj Popuri: Yeah, there's a bunch of them, but my favorite example is study startup. So pharma's been optimizing this process for the last 20 years, I would say, and there are different pieces to that process - there's a sponsor, there's a CRO, there's a site, there's legal, there's ethics, there's IRB, and all of these different institutions that make it part of that workflow and work chain. But, every leg has a different owner and there's a different handoff at each interaction point.

Rather than optimize, I think we need to reimagine this process, and finally, I think we are ready as it is. Each site should be live and rather than tied to a system or someone's inbox, we are now able to have systems of records that have APIs or MCPs that can talk to each other. The models like LLMs can read all the unstructured documents, and then we can have a human in the loop to make the decision that is based on experience and nuance.

The state of that system is ready for reimagination. I'd say, five years ago, two of those three were missing, but now I think we're ready for that process to be reimagined ground up.

Abhinav Kohle: I think the other thing is that most of the organizations buy the vision. Where they get stuck, Brad, is the first move, right? You look at hundreds of processes to be automated. Mining tools will tell you, "Oh, we have 500 automation opportunities out there." Is that really true? Is that really something possible and doable in a span of a year? Because budgets usually run on a yearly basis, right?

A thumb rule that we have been using is, to look at highly manual effort processes. Are there several approval points in the process? Somebody is constantly checking? Where is this stuck? And then, of course, what's the biggest business outcome, right?

All of this put together is what your right candidate is for automation in the new era. With tokens and all the consumption that's coming up, ROI becomes central to what use case you want to automate.

© 2026 Persistent Systems Ltd. All rights reserved.

So Brad, let me ask you something very specific. What is the most boring win you would have been actually proud of? Something that never shows up on the limelight, never comes on the slide, but genuinely changed someone's Tuesday.

Bharadwaj Popuri: A good example of that would be with the trial master file, right? So essentially, the manual process - as it existed back in the day - required someone to grab emails, documents, notes, and then put them into a portal. It was the job of that coordinator to make sure that portal stayed up to date with all the documents.

And if you think about it from a process perspective, it's actually very boring and it can get to you. It can wear a person down doing this, repeatedly. So the bot that we built basically did all that. It grabbed all the documents from the different intake portals and put them into one central portal.

I tend to look at it like it just did not change someone's Tuesday, it changed the way they look at work, right? Rather than, "Oh, now this is what I have to do for the next eight hours," it frees them up to do something more that will tap with their experience, give them a more enjoyable work environment, rather than just doing a boring, repeatable process that could have been automated to begin with.

Abhinav Kohle: Each of these programs, Brad, also assumes that data is always at the right place. It's connected, it's ready for AI, and that's not reality on the ground, as we all know, right? It is scattered across multiple systems of record and that's where the challenge lies today. How do you make agents context-aware?

So Brad, where do you sit at, how do you actually play out in the real world? I mean, do you actually try to fix the data problem in the use cases that you have seen in your company? Do you try to address the data problem first? Do you try to address the process-first problem? How do you tackle that in the context of pharma?

Bharadwaj Popuri: The best framework that we use that works for us is something that we leverage called people, process and technology. So, folks tend to look at this as a data problem or a process problem or a technology problem. They forget that people are a big component of what drives the uncertainty and the variability.

I think specifically to your question, whether fixing data first or fixing the process first, I think the challenge of looking at individual problems is, if you fix the data, nothing gets shipped, right? Because you need to first clean the data, you need to make that accessible, you need to make that contextualized and ready for agents to consume.

While you're doing all that foundational work, there is no outcome. But if you fix a process, but don't fix the data, you're still going to speed up your decision-making but not address the foundational data problems. So in theory, people tend to say, fix the data or fix the process, but in practice, I would say neither is correct.

© 2026 Persistent Systems Ltd. All rights reserved.

For each field, what you need to understand is that whole workflow of things, right? Take the unstructured part of pharma, which is all the documents, treat them as data, keep them ready in a consumable format. Then reimagine your process to say, "Is this process ready for AI or automation?"

Because end of the day, if you don't ship anything within a year, you're not going to have any backers for the solutions that you build, whether it is addressing just the data problem or the process problem. What I urge people to do is to think about it from that integrated lens, which is leverage the people, use their experience and nuance, address the data problems, and then build on that quickly, so that way you're able to see a quick ROI.

And do it iteratively. You don't have to fix the whole thing. Build an MVP, leverage the small gains, and then build on top of that.

Abhinav Kohle: Thanks, Brad.

So, let me move on to the next section. What we have seen is that in a manual process, the next steps essentially live in someone's head, and they have to remember that. In an intelligent workflow, the next steps are basically visible, guided as part of the system, and that's what AI will change in the future, right? Old automation used to follow the rules that were written by somebody, and it was not smart enough because it was written and mapped to the T, right? Whereas in the new workflows, you will have the process define the next steps.

In pharma, it matters quite a bit, in terms of trust, oversight, explainability, all those basic features that are already required by everybody. So, given all of this, and if you want to really make the agents context-aware, what workflows do you feel, Brad, are ready for the next steps and where do you think the context layer will bring added value to these workflows? Which of those workflows will add more value when a context layer is also provided to the agents?

Bharadwaj Popuri: So, specifically for the context layer, I think it's the data problem wearing an AI costume and pretending to be something else.

You can get the model right, but models are now replaceable and interchangeable based on your use case. You need to still address the data problem, and you need to address it in a structured way. Do you know what your entities are? Do you know what your sites are? Do you know what your standard operating procedures are? Do you know what state they reside in? And do you know the history? Because for pharma, audit trail and history are very important. Then you also have to think about it from a different perspective, which is, what is the persona and identity of the person asking the questions, and what are they allowed to see?

In summary, for context to work, you need to look at entity, state, history and the identity of the person, and that's how you're able to build successfully in a GxP...

Abhinav Kohle: ...environment.

© 2026 Persistent Systems Ltd. All rights reserved.

Moving on to the next thing, right? Where do you think agentic AI should be applied and where do you think you should not even look at in the current context of things?

Bharadwaj Popuri: Agentic AI is great, in the sense that it helps automate a lot of the workflows in ways that were not possible before with just RPA or workflow automation. But there's also a sense of overselling the capabilities. People tend to think, "Oh, it'll read, draft and then it'll decide," and that's where I urge caution.

AI is phenomenal at reading and summarizing. It's phenomenal at drafting something. But without that human nuance, it's a dangerous decision-maker. So the value is to keep AI in reading, summarizing and then drafting, but then let the decision be made by a human in the loop.

Abhinav Kohle: That's a very nice way to put it, Brad. You said AI is a phenomenal reader and drafter, but a very dangerous decider. I'm going to steal that. It's a really powerful way to frame it.

Most people, today, picture AI as a chatbot on the side using RAG solutions, and that's the limit of what the industry is solving for today. With agentic AI, things are changing, but the real value is somewhere else. It is where AI sits inside a workflow. The workflow is what keeps it safe. The controls, the approvals, the audit trail, all of that is part of the entire workflow. AI does the reading, as you said, AI does the drafting, but when it comes to decisions, the process is governing it, and it decides whether the agent should take that call versus if a human should be involved in that decision.

So, it's a kind of a dial, in our view. And how you calibrate that dial is something that we need to learn through iterations in the UAT cycles and then figure out how to manage the dial when it goes into production.

So moving forward, Brad. Looking ahead, and this is a little bit of a look forward type of a discussion, right? Assume that we are two years down the line, agents have started to reimagine the process, the whole agentic automation has taken leap forward. The processes have started to get reimagined in the new way of things.

In that state of mind - let's say we are there, Brad - what are the two top things which you think that your teams will still do by hand, even two years down the line? What is one thing that you think will still be manual, irrespective of what advances happen in this industry?

Bharadwaj Popuri: I think in three years, manual document handling would go away, and that would be something that we would be able to fully leverage agentic AI and orchestration capabilities for. The other thing that doesn't get talked about, but is a persistent problem within pharma is status chasing, which is someone has to ask for what is the status of this document or initiative or program or project. I think that will go away, too. I think we will finally be able to either deploy agentic OSs or agentic orchestration to the point where it becomes proactive.

So, let's say you are an owner for a process. We will have agents notifying you when it is done, when it is broken, when it needs your attention. That way the human in the loop still continues to make that judgment call.

© 2026 Persistent Systems Ltd. All rights reserved.

I think one thing will continue: at the end of the day, we humans are accountable for the decisions we make, and that judgment call will still reside with humans. I think that's a healthy way to leverage machine learning, AI, or agentic AI for process and routing, while keeping humans accountable for decision-making.

Abhinav Kohle: This has been a genuinely useful conversation. Let me pull the threads back again, right? To summarize what we've discussed today - in general, what we have seen is that modernization has stopped being an efficiency project anymore, and it has become a strategic one for all companies, including pharma.

It is not a technology-first problem, as people keep talking about. I think all of us need to look back at the problem first, reimagine the process first and then take the next steps on it.

The line I will take away from this conversation is that, in a manual process, the next step actually lives in someone's head. The work for the next few years is getting it out of there and into the system visibly, traceably and without losing the control that pharma runs on.

Brad, thank you so much. Genuinely appreciate you being on the call. If there is something that you want to summarize, please go ahead.

Bharadwaj Popuri: No, thank you, Abhinav. I really enjoyed the conversation too. If there's one thing I would like your listeners to take away, it is this: make sure you leverage AI for all the reading and drafting, but keep the human in the loop and accountable for decision-making. Make sure you know where the dial for AI exists, what it is set at, and who is allowed to control or turn it. Doing so will allow us to continue building the trust we have with patients, regulators, and all the other stakeholders. It will also help us make progress where it counts.

So, thank you for your time this morning.

Abhinav Kohle: Thank you, everyone for listening to **Re(AI)magine** Conversations.

If today's episode sparked new thinking, follow us on our network. If you have a great story to tell or a guest you would like to hear from, drop us a note at podcast@persistent.com.

Until next time, stay curious and stay inspired.

© 2026 Persistent Systems Ltd. All rights reserved.

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 27,500 employees located in 21 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 Best-in-Class 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.
1731 Technology Drive
Suite 700, San Jose
CA 95110
Tel: +1 (650) 481 9180
Email: info@persistent.com

India

Persistent Systems Limited
Bhageerath, 402
Senapati Bapat Road
Pune 411016
Tel: +91 (20) 6703 0000
Fax: +91 (20) 6703 6003

Persistent
Re(AI)magining™ the World

© 2026 Persistent Systems Ltd. All rights reserved.

www.persistent.com