Re(AI)magine Conversations

Persistent Podcast | Episode 26

Autonomous, Not Unattended: Governing Agentic AI

Manas Arora, Technology Industry GTM Lead - EMEA, ServiceNow
Mukul Pant, Senior Vice President - Europe, Persistent

Mukul Pant: Hello, everyone. Welcome to **Re(AI)magine Conversations**. This is where we have honest, grounded conversations about how technology is actually showing up inside the enterprise, beyond the buzzwords, beyond the hype and into the realities of making transformation work at scale. In each of these episodes, we sit down with leaders who are shaping the transformation journey to explore what is changing, what enterprises are learning and what it really takes to turn emerging technology into meaningful business outcomes.

I'm your host, Mukul Pant. I'm a Senior Vice President at Persistent. Joining me today is Manas Arora, Head of Technology Industry Go-to-Market - EMEA at ServiceNow.

Manas Arora: Thanks, Mukul. Glad to be here.

Mukul Pant: So we'll cover five themes today, Manas.

First, why pilots stall on the way to production. Second, what AI is actually costing and how do we tell the consumption and value apart. Third, what changes once AI can act, not just answer. Fourth is about sovereignty and what it means from a European perspective. And, fifth is what leaders should be actually doing over the coming weeks or next quarter.

So let's start with the big picture. Most of my client conversations this year have begun in the same place. The pilots work, but they aren't scaling. Where are you seeing enterprises in general? Where is the gap?

Manas Arora: I'd start with the fact that enterprises have largely proved that AI can work. There's no question around it. The harder question now is how you make it work across the enterprise and scale it with the right foundations.

What I mean by the right foundations is that once AI moves from pilot into production, it needs the right data, the right business context, the right workflows and clear controls around it. You know, basically the controls of what it can do, how it can decide an action and how it can finally act to deliver an outcome.

© 2026 Persistent Systems Ltd. All rights reserved.

To reiterate, the gap is execution. I'd love to understand from you as well, Mukul, what are you hearing from your customers?

Mukul Pant: This resonates quite well with what we are hearing in our client conversations, Manas. The question has moved from "can it do it?" to "who answers for it when it does?"

So let's just talk about the first theme. It's the move to production. Most pilots look successful. So what changes when you go live?

Manas Arora: I think a pilot creates an unusually clean world. There's usually a small team, controlled data, a bounded task and limited access. But production is a very different beast altogether. Just to put this in context, your typical enterprise has anywhere between 200 to 1,000 disparate systems. And on top of that, there is fragmented data in those systems. Each system comes with its own data model, different security models, and workflows that often terminate at those system boundaries. And that matters because if a person moves inefficiently across disconnected systems, you have a productivity problem, right?

Manas Arora: I think a pilot creates an unusually clean world. There's usually a small team, controlled data, a bounded task and limited access. But production is a very different beast altogether. Just to put this in context, your typical enterprise has anywhere between 200 to 1,000 disparate systems. And on top of that, there is fragmented data in those systems. Each system comes with its own data model, different security models and workflows that often terminate at those system boundaries.

And that matters because if a person moves inefficiently across disconnected systems, you have a productivity problem. But imagine if an autonomous AI agent does the same thing at machine speed without the right context and controls. That creates a real risk for any business.

In fact, we wanted to understand this better - move beyond the theory, and we at, ServiceNow, the Thought Lab actually, did a study for ServiceNow. We call it the Enterprise AI Maturity Index. For the sake of this conversation, I'll only speak about the subset of the technology industry. We found out that AI spending grew 110%, no surprise, companies are investing a lot in AI. But only 16% of those tech companies had built the right foundations around data governance and workflows.

Because of which, and I think that's a big contributing factor to the fact that only 9% of those tech companies were able to move beyond just solving a point problem with AI and achieved real, multi-step autonomous workflows.

If I have to summarize this, the investment is clearly accelerating, but the scale is where enterprises are seeing the biggest challenges.

Mukul Pant: We are also seeing the same thing when we are delivering for our clients. The demo proves that the model can perform a task and production has to prove that the enterprise can trust the outcomes. Yeah. And, you're bringing things together which were never really there.

© 2026 Persistent Systems Ltd. All rights reserved.

So in our conversations, what we are seeing is, clients have a long list of so-called barriers or perceived barriers, such as data security, governance, cost, skills, legacy. Are they really six problems?

Manas Arora: I hear the same list, Mukul, but I don't think these are six separate problems. To me, they all point to the bigger challenge, which is about creating the right environment for AI to operate across the enterprise. If I have to put it down, it starts with business context. A general AI model may understand what a server, a customer or a payment is, but it doesn't automatically understand what those things mean to your business. To simply put it, AI needs to understand your business. For example, it needs to know which servers support a critical service, which customers have special contractual commitments and what payments need additional approval. Without that business context, even the most advanced AI model can struggle to make the right decision.

Secondly, I do feel there is data. I think we've moved beyond the idea that simply giving AI more data will make it better. What really matters is giving it the right data. We've heard about terms like data lineage, data provenance becoming more and more important. Is the agent accessing the right data with the right permissions? Where is the data coming from? Is it the latest version? All of these things actually matter a lot while making those decisions.

Mukul Pant: Can I just pick up on identity? Because this is one where we are seeing clients are least prepared.

Manas Arora: Yeah.

Mukul Pant: What I mean by identity is that once an agent can act, it effectively assumes an identity, right? What can it access? What can it change? And critically, is it acting on its own authority or on behalf of a person? Most of the firms haven't yet figured this out. It almost feels like shadow IT days again, except this time nobody needs to buy anything, right?

Manas Arora: No, exactly. I think that's where governance has become really important, because once agents start taking actions, it's no longer enough to simply control what they can access. You also need to know what they're doing, what decisions an agent made, what information did it use, who authorized the action. All those questions need to come, and in fact, if something goes wrong, can you step in, stop it or even reverse what the agent has done? These are the controls enterprises need to think about as AI moves from just answering questions to actually taking action.

I think that brings us back to the bigger point. The challenge isn't putting AI into every system. Most of the platforms today have started to offer AI out of the box. It's about creating a consistent operating environment where business context, data, identity, governance - all those things come together to make sure the AI is working effectively within the organization.

You see this from a service provider perspective and I'm seeing this from a product perspective. Where does the real complexity lie in your view, Mukul?

Mukul Pant: See, the model is the most visible part of the solution. The hard work is always connecting it to the existing systems, securing the data, defining access, handling exceptions and making it supportable in production. I think that's where we see challenges or issues. And technically

© 2026 Persistent Systems Ltd. All rights reserved.

speaking, successful pilots do not even have a proper operating model. The business has the use case, technology owns the platform, security would come in to have access. You would have risk, and then finance would see the kind of the bill that you would get. Agents make all these things harder because system isn't just generating an answer. It can even trigger a next viable or tangible action.

Manas Arora: Exactly. I think agentic AI is not just a technology shift. So many companies have evolved their go-to markets around AI. They're offering new products and services in a differentiated way. It typically changes how enterprises operate and how AI is being managed. When AI moves beyond just giving answers and taking actions, there's a need for clear ownership and accountability, and it becomes more and more important as you scale with AI.

These are no longer theoretical concerns. We are already seeing discussions happen around a kill switch - if the agent goes rogue or something goes wrong, stop the work immediately? I think, it all rounds back to the fact that how an enterprise can have more control over AI, including discovering shadow AI, which is again turning into an increasingly bigger problem for enterprises.

This is why I think the operating model is just as important as the technology. Before enterprises can confidently scale agentic AI, they need to have clear ownership, defined responsibilities and the ability to step in when something goes wrong.

Mukul Pant: Very true. Moving on to our second theme. Enterprises want people to adopt AI. But as we are all seeing, everyone is concerned about unpredictable bills. How should the leaders think about it, Manas?

Manas Arora: I think adoption is important, but consumption is not the same as value. We're increasingly hearing stories of organizations exhausting their AI budgets in a single quarter, not because they're creating more business value, but because they're consuming more AI.

I think, the big part of this comes down to how fundamentally the economics of software has changed. Looking at traditional software, the CFO or the business always had predictability in terms of how many licenses, users, infrastructure you were paying for. But with agentic AI, that fundamentally has changed. You're not just paying for just one interaction. An agent might make several model calls. It might connect to different systems or even involve other agents, using an orchestrator or a multi-agentic workflow to complete a single task. So what looks like a simple request can end up costing much more than expected to a business.

The enterprises need to be smarter in the way they use AI. Not every task needs the most powerful and expensive model. This is a conversation that has been coming up, that how do we get the ability to use different models for different tasks. And then there are those simpler tasks where you can use more lightweight, cost-efficient models. The real challenge is finding the balance between performance and cost.

As the adoption grows, businesses need to move beyond measuring how much AI they are consuming. They need to start measuring what they're actually getting out of it in clear business

© 2026 Persistent Systems Ltd. All rights reserved.

value, in clear KPIs. What is the outcome of AI? Was the change really worth it - is the question that the businesses need to answer.

Mukul Pant: I think we are also nudging our clients who are seeing a trend that it is just not about token, cost optimizations. It is about cost per verified outcomes, not cost per user, not cost per token. If you can't really count the verified outcome, you can't really fund the next phase of your transformation journey.

That brings us to our third theme. What changes when AI can act? I mean, governance conversations have been so far about hallucination, bias, wrong results.

Do agents change that?

Manas Arora: I think, quite significantly. If you compare it with generative AI, the biggest concern was often that the system might say something wrong, which had its own complications and legal implications for businesses. But then with agentic AI, the risk is much greater.

The system can now take the wrong action and potentially cross multiple connected systems and at a speed where just watching the agent work manually will not help. We need smarter systems to do that. And that changes how we need to think about governance. It is no longer enough to have policies that simply describe what AI should or shouldn't do. Most importantly, and you need to build the controls in the system itself, the controls that define what an agent can access, what systems can it work across, what decisions it can make, and most importantly, it should define when a human needs to step in.

And then there's the other aspect. When something goes wrong, simply knowing an incident has occurred is not enough. There are great telemetry and observability tools in the market, but I think they're just a part of the story. Enterprises today need to understand what it means in terms of their business. Which business process got affected? What is, the potential impact of the breach or the incident? How to contain the issue before it spreads? If I have to simply put it, finding the risk or the vulnerability is important, I'm not denying that, but the ability to act on it is even more important because you have to control the blast radius before it becomes too big.

I would also want to go back and understand, Mukul, where do you see that exposure in client environments?

Mukul Pant: That usually happens to be places where nobody's in fact looking. The agent with the most reach, many a times, isn't the one that businesses have built themselves. They've arrived switched on in a tool or in a SaaS product that they already own.

The question then becomes, do we have to then have a person involved in every action?

Manas Arora: No, I think there are different kinds of terms that the industry uses - human in the loop, human on the loop, human out of the loop. I think there is an important difference between keeping humans in control or as against making them a bottleneck. For example, if an agent is doing

© 2026 Persistent Systems Ltd. All rights reserved.

a complex or a high-impact task, human judgment always needs to be involved before that task is completed or that decision is taken.

But if, on the other hand, somebody has to approve every single action, especially when AI operates at the speed that it does, you lose much of the benefit of the autonomy.

The better approach is to let people set the rules of what an agent can do, what it is trying to achieve and where it needs to get human approval, and then let the agent operate within those boundaries by making sure you're able to see what it's doing and you're also able to control what it's doing without a bigger impact.

I call it VAR - Visibility, Audit and Remediation. These are three terms which are extremely important to keep in mind when agents are acting. It starts with every action that an agent takes should be visible. That's the V. There should be a clear record of what happened with clear ownership. That's the audit part. And where possible, there should be a way to reverse it. That's the remediation piece. Because ultimately, autonomy should never be about the human losing control of what the AI is doing.

Mukul Pant: It matches very well with what we see in clients' scenario. Autonomy is something that a workflow earns, not something that comes switched on. People have to approve at defined stage gates. They have to then govern the loop or step in by exception. And then the selected workflows act on their own with a full audit trail.

Manas Arora: Yeah.

Mukul Pant: Autonomous should never mean unattended and it should never mean unaccountable. Autonomy, we agree, needs to be earned.

On our fourth theme, and a very European one, I'm sure you're having those conversations, Manas. Is sovereign AI mainly about keeping data within national borders or how are you seeing clients adopt?

Manas Arora: I think, data residency is certainly an important part of sovereignty, but I think sovereign AI goes beyond just where the data is being stored. It is about having control over the AI value chain. When you look the data, your AI models, infrastructure that those models run on, you have to have control over the value chain.

Data is just a part of it, as I said. This is where we are seeing enterprises increasingly take a hybrid approach to sovereignty. For example, in sectors like government, banking, healthcare or critical infrastructure, there are workloads which are sensitive. Data needs to be protected and AI actions need to be tightly controlled. That's where sovereign-first approach makes sense. On the other hand, not every workload needs the same level of control. For less sensitive or lower-risk tasks, enterprises can still benefit from using global AI models and cloud platforms.

I don't think it is about the choice between going sovereign or going global. It is about creating the right balance in your operating model and taking advantage of the broader AI ecosystem which is

© 2026 Persistent Systems Ltd. All rights reserved.

available. I increasingly see this is happening, barring a few industries which completely want to go sovereign, given everything that they do is sensitive and has PII data.

Would love to understand, Mukul, what are you hearing from European clients?

Mukul Pant: Residency does come in, Manas, but lock-in comes in more often. If your context, your workflow logic and your audit trail is live in one model or in one cloud, then you have moved the dependency. You've not removed it. I mean, one thing that we are also hearing clients talk about is AI factory, sovereign AI. How does AI factory fit into all this? I would love to understand how terms like AI factory and sovereign AI are being used interchangeably? Is that the case?

Manas Arora: I think they're closely related but different. To explain, an AI factory is about building the capability to develop and run AI at scale. It brings together computing power, infrastructure, data, models needed to move AI from experimentation to production.

On the other hand, sovereignty or sovereign AI is about control. It's about where that AI runs, how the data is protected and who ultimately has authority over the technology and how it is being used. I think, and I think that distinction is important, simply building a data center and installing GPUs automatically doesn't give you a sovereign AI. It'll give you a sovereign infrastructure, but you still need the right governance and operating model to have real control over it.

If I have to just summarize this into three parts, the AI factory gives you the capability to build and run AI, sovereignty determines who controls it and workflows and governance around it turn that capability into trusted business outcomes.

Mukul Pant: What should the enterprise leaders be doing about it? If we have to bring all this together, what foundations does an enterprise need to scale AI with confidence?

Manas Arora: It comes down to four things, and I'll keep it short. Sense, decide, act and secure. AI needs to sense the right data and business context, decide what to do within the guardrails set by the business, act using trusted workflows, and underpinning all of that is secure - making sure you have visibility, control and the ability to act over what AI is doing.

That's how I put it.

Mukul Pant: To finish off practically, if the CEO or CIO were to do only three things differently over the coming weeks or quarter, what would they be? Give me the first one. I'll try and throw in a third one.

Manas Arora: Sure. For enterprises that are trying to get on the curve of agentic, I think they need to stop measuring AI activity and start measuring business outcomes.

For the enterprises that are actually trying to scale AI, we've seen a lot of success with customers that pick up one workflow that really matters to the business and look how AI can improve it end-to-end, keeping the business KPIs and business impact in mind. Whether it is saving time, reducing cost,

© 2026 Persistent Systems Ltd. All rights reserved.

including customer experience or whatever KPI the business is trying to impact. That's what will tell you whether AI is creating the real value.

Second, I think it is absolutely impossible for businesses to scale AI without having the right foundations. I've touched it towards the beginning of this conversation, and I'm going to go into this as we are summarizing it. Make sure your AI has access to the right data and understands the business context.

It also needs clear boundaries around what it can and it cannot do. I think building agents is becoming easier. It's not the problem that the industry is facing anymore. The real challenge is making them work reliably within your business and deliver meaningful outcomes. Especially for the businesses that are trying to create a moat or a differentiation in their go-to-market using AI.

I think the real competitive advantage will be there when you really have your foundations set so that you can scale at super speed to match the pace of the market. That's what will turn AI capability into lasting business value.

Mukul Pant: Very true. Let me just add in a third one. I think we have to start putting cost, risk and value in a similar dashboard or a view. Leaders need to be able to see which agents exist, what they can access, what they cost and what outcomes they produce. And you have to measure cost per verified outcome because that's how you would be able to prove and move forward in your transformation journey.

If I summarize our conversation today, the five themes, we talked about production failing on operating models, not the models themselves. Value needing a unit that you can measure. Autonomy has to be earned. Sovereignty is about control, not just location. And the place to start, making this all safe, secure is one workflow that matters, one workflow at a time.

Thank you, Manas, for your time.

Manas Arora: Thank you, Mukul. It's been a pleasure.

Mukul Pant: And thank you for being part of Re(AI)magine Conversations.

© 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 29,500 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 agentic business 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 has been named one of America's Greatest Workplaces in Tech 2026 by Newsweek and Plant A Insights Group. The Company is also a constituent of the Dow Jones Best-in-Class World Index. Persistent has achieved carbon neutrality, reinforcing its commitment to sustainability and responsible business practices. 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 a 22% YoY growth in brand value, Persistent has been recognized as the Fastest Growing IT Services brand globally in the 2026 Brand Finance IT Services 25 report, among the world's Top 25 IT Services brands and ranks as the 12th strongest brand.

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

www.persistent.com