

Persistent Podcast | Episode 10

Why Siloed AI Fails — and How to Scale It

Ajay Yadav and Travis Norenberg

Ajay Yadav: Hello everyone. Welcome to Re(AI)magine Conversations, the show where we dive into how AI technology and bold thinking are reshaping enterprises. I'm Ajay Yadav, VP of Sales at Persistent. This podcast explores how enterprises can overcome the challenge of siloed AI deployments by adopting unified context of our platforms, highlighting the journey from prototype to scalable solutions, addressing security, integration concern and showcasing the tangible business outcomes in the future possibility, enabled by innovation like Gemini Enterprise.

Joining me today is Travis Norenberg, GenAI Solution Architect at Google. Travis, would you like to introduce yourself?

Travis Norenberg: Absolutely Ajay, thanks so much for having me today.

Really excited for the conversation. So, like you said, my name's Travis Norenberg. I'm a Solutions Architect at Google Cloud, specializing in Generative AI and in Agentic AI. So, the way that I like to kind of describe it, like in a nutshell, I spend most of my days trying to help organizations really move past kind of that cool demo phase and actually harness, you know, these models and agents to drive tangible business outcomes.

Something that I do love about my job, just kind of on a side note is, you know, it's so fun to see how quickly the technology is moving in the AI space right now. It feels like I go to sleep an expert, only to wake up the next morning as a novice, because all of these new papers get published overnight and you're always trying to stay on top of things.

So, it's a really exciting time to be working with this technology. Like I said, looking forward to the conversation today.

Ajay Yadav: Thanks Travis. With that, I would like to quickly jump into our conversation and would like to address the elephant in the room. As we all know, we are so excited with AI, GenAI, and now the Agentic AI, but most of our time we have seen the pilots are failing and many organizations launching the promising AI pilots. They're not able to scale across the platform. Why do you think the silos problem persist, and can you share any real-world example where siloed AI led to the missed opportunity?

Travis Norenberg: Yeah. Yeah, definitely. I mean, I think that's a really good question, and it is an area I see a lot of organizations kind of struggling with today.

I think, kind of at its core, there's a few different issues actually. And so, when we're looking at these AI initiatives, we frequently see that, these projects are kind of kicked off and launched in isolated silos, just as far as the organization goes. So maybe, there's an application or an agent that's being launched in kind of an isolated team, or maybe it's a specific organization within the larger enterprise. And this makes it really difficult to kind of bring together the holistic view of how Agentic AI, Generative AI is being leveraged.

The consequences that we see here then is that this fragmentation really creates barriers that are making it difficult to share. The insights that these teams are deriving from some of these pilots, POCs and things that they're running, taking it a step further, it's also making it really difficult to automate some of these processes.

I think one of the things that I learned throughout 2025 is 2025 really is the year of Agentic AI. A lot of that power is coming from kind of moving past general semantic, summarization, natural language understanding to actually having AI agents take action on your behalf.

And that's where we get into that automation, you know, of these processes. And I think, as an organization who's starting to kind of dip their toes, so to speak into AI, you really need to take that holistic approach.

The other thing there is, kind of the organizational silo. But I think data silo is another, another huge issue. And this isn't a new issue specific to Generative AI, large language models (LLMs), Agentic AI, this is an issue that has been around for 25-30 years or more. And it's just as important today as it was back then.

And I think, the best way to kind of describe it is like your AI applications, your agents are only going to be as powerful as the data that they have access to. I think it's really critical to take that kind of enterprise-wide approach, really identify what are the key use cases?

What is the data that we have available? Who are the actual human stakeholders who are going to be leveraging these applications and really break down those silos to have that more holistic view.

Ajay Yadav: Absolutely, Travis. These are where most of the organizations, because they are so departmentalized by their functions and their line of operations, they don't come together at times, even for the broader an AI-led initiatives across the organization.

Travis Norenberg: Yeah. You know it's a challenge and I totally understand that. There's a lot of stakeholders, a lot of players that need to be, kind of involved in a lot of this road mapping and use case identification. But I think we're getting to that point.

And the customers that I are really taking that holistic approach are the ones who are being wildly successful with a lot of these pilots that they're running. And actually, just to kind of build on that a little bit, Ajay, what does it really mean for an organization to connect AI to their workflows, to their data, to their teams?

I know you guys are in the field working with customers on this every day. So just kind of curious, are you seeing a big difference between those organizations using a collection of AI toolkits or smaller POCs around the organization versus those who really are adopting kind of that true enterprise platform?

Ajay Yadav: Great question, Travis. And that's definitely experience we have on the field. The AI tool can give you a good amount of power and a leverage to build what you want to build, but unless we are trying to stretch the entire workflow together, they will still be solving 1%, 2%, or maybe 10-20% of the problem.

They might not be solving for a hundred percent of the organization. Toolkits can solve specific problems, but they rarely scale. And the scale is because we have not solved for the underlying platform, whereas a unified platform where we have connected data, we have defined workflows, and we have people who are coming together for a broader cost will always be a challenge.

So, in a unified world, you will see when people coming together, they try to solve for the bigger problems and have better ROIs and all that needs definitely an integration. And integration is not from a technology perspective, but integration is also from the human point of view where all thoughts, wider goals need to come together and have to stretch a wider story.

Travis Norenberg: Yeah, absolutely. I couldn't agree more. And you know, it's interesting with a lot of the customers that I interact with. The ones who are successful with a lot of these projects, kind of these early-stage pilots, POCs, it's been very interesting to see how quickly that enterprise adoption does take place.

And, you know, kind of thinking about it, it's like, okay, you go from an organization who's got maybe 1, 5, 10, or 50 agents maybe running within their organization. Once they start

seeing that value, the adoption just accelerates rapidly. And I think having that scalable enterprise platform is really critical. If you're in a position where maybe you're managing a thousand agents that are out in your fleet, across the organization, and that gets to be a little bit more challenging.

So, yeah, I definitely can agree. I'm seeing the same thing. It's having that enterprise level platform. I hate the term, but that single pane of glass, uh, to really understand here are the agents that we have out in production, here's what we're testing, uh, here's kind of the metrics that we're measuring against to kind of validate, like is the performance where it needs to be? It is really critical for success with a lot of these projects.

Ajay Yadav: Thanks Travis. And let me switch gears here. As we just spoke about, it has to be a unified pane of glass. But most of our time when we are trying to remove the silos, we need to think about the broader security, broader scalability, and the wider integrations. One of the biggest hurdles as we see in moving prototypes to a production is specifically comes when we start dealing with.

The organizational level security policies, we need to take care of the data privacy as well as the integration across the different systems and networks within the organizations. How are you leading organization addressing these challenge service?

Travis Norenberg: Yeah. I think that security is absolutely critical.

And I do kind of pile that into the scale kind of bucket of challenges, you know, and when I'm looking at a lot of these projects, I think there's kind of three main pillars of scale that we're seeing. The first that we're seeing is really need. Robust infrastructure to host and manage all of these applications and agents that you're running.

The second piece there is security. I mean, you have to have strong governance, you have to have auditability. You need to make sure that you've got those compliance controls in place for these applications.

And then the third pillar, we talked about this a little bit already, but is really that integration with your organization's data or enterprise apps.

So, really kind of focusing on the security aspect, I think when we're thinking about security and privacy with AI applications and agents, it really can't be an afterthought. I think we need to start these projects with security top of mind. It really needs to be built in from the start. And that includes, you know, obviously the security piece, but governance, auditability, compliance, these are all really critical in that enterprise environment.

And something that very candidly is a little difficult with large language models. It's that they're very non-deterministic. So, it's a little bit more difficult to be able to point back and say, okay, here was my original prompt, here was the output, how did we actually get to that output?

And having those controls and guardrails in place is, kind of, table stakes at this point. Again, especially in that enterprise kind of segment. So, you know, I think a lot of these orgs that I'm working with, are investing in these centralized platforms to make sure that they have all of the tools available to them for auditability and compliance, governance and obviously security.

And I think when we're looking at integration into the organization's data, which again is really critical for the success of these projects, the large language model is only going to be as applicable as it is to the data that it has available to it. I think from a security perspective, that also opens up a lot of watchpoints for us.

So, you know, things like understanding access control list in various third-party systems that you might have integrated into an agentic system is really, really important. Ensuring that as a user, I only have access to the data I'm supposed to have access to, just doing things like, ummm, just basic authentication of agents. You know, really giving like an agent an identity, again from a compliance governance perspective is really critical. So, I don't know, that's a kind of a long-winded way of answering your question, but I do look at security as it needs to be part of the conversation from day one. It really is critical that we build these kind of frameworks around these applications to ensure that we have that ability to ensure that things are governed in a proper way that we're compliant, etc.

Ajay Yadav: Absolutely. And to add on this, uh, as Persistent, we deal with a lot of BFSI customer, healthcare life sciences customer, where there is a lot more regulatory and compliance requirement on how to access the data, how to process the data, and who can access the data.

Putting these guardrails in the place, building that framework, which we call as the responsible AI framework, even before we start building an agent or building the entire platform for the agentic orchestration, we need to really address these concerns first before we can talk about the actual implementations.

Definitely, it's like the highest priority in every organization. They can't really deal with any agent without having a security in place.

Travis Norenberg: Absolutely. And especially in those highly regulated industries like you mentioned and having the responsible AI framework, and I think that is such a good practice for any organization out there is really document how are we using AI, what are the goals?

What are our specific guardrails, compliance standards that we need to meet governance, you know, things that need to be addressed.

But kind of starting these projects with that more high level understanding of like, okay, as an organization, how do we want to use AI? What are our internal kind of guidelines around using AI, I think is really critical.

So, uh, love to hear that your team at Persistent is kind of leading with that workflow of just identifying like, okay, high level, before we even start building the cool stuff, let's just understand what is our goal with AI, and how can we, kind of document that to keep that as our north star, as we continue to build new applications and new agents.

Just to switch things up, one question that I had is like, you know, again, your team is so closely intertwined with customers. You're working with them, building some awesome things. So, from your perspective, really working with these clients, like what tangible outcomes are organizations that you're working with, realizing as they move beyond these silos that we've talked about, whether it's data or organizational and what do you see as the future holding for some of these organizations that are kind of embracing this new approach?

Ajay Yadav: Very good question, because while we can build the great things, but unless it makes sense from a business ROI, there is no real interest or a return from a customer perspective while building and scaling these agents. Can we build a workflow leveraging the agent that help us to reduce or optimize our overall operational cost, or it enables the faster decision making process? For example, today, if I have to wait for a customer or for our engineer to give me certain data inputs that I can consume to make my business decision versus can I take that decision while talking to my data right away without having an engineer engineering the data in between?

And I can quote with a very recent example. We are working with a credit company, and the credit company processes loans for different cars. Their bigger concern is, today, we have a high relevance on a particular team whose job is to crunch all these data, how many loans were filed, how many approved, at what rate?

But it is always that team. If the team is limited by capacity or overload by different work, we are not able to scale how we can build a system where people can talk to data at the same time, should have the access limited to their roles. I don't want to expose my data to every single person. So, the measurable ROI comes in when we start unlocking the productivity, when we start expanding the use cases for the wider organization.

Adaptability. At times it unlocks new product because it gives you a new thinking dimension. At times, it opens up new services because you might not be looking at your

data that might help you to figure out where a customer is mostly coming back, and if there is an avenue to generate a new service model or even the entire business model.

We have seen how the brick and mortar or the traditional workplaces today getting challenged by the digital natives and now the AI-enabled organizations. So, I believe enterprises who are embracing a unified AI, who are embracing the scalability, who are embracing the overall integration and a wider net for a wider goal of business that lead their industries in innovation and growth, that's where we have seen most of the customers measuring the ROI from the customer experience to operational efficiency to the wider innovations in the business models or in the customer service model itself.

Travis Norenberg: Yeah, I love what you said there around, you know, just bringing these use cases to actually showing the productivity gains. And I think that's where 2025, for me at least, has been really exciting, is moving past kind of this standard like prompt and response sort of model. But actually having agents take action on your behalf. I think that's kind of where this thing is going.

And it is really exciting to see the enterprise adoption and just some of the really exciting use cases that are being created. I always laugh because I like to describe, you know, right now just in GenAI in general, it's kind of the wild west right now. I mean, there's new use cases being built everyday that it's the first time somebody's building this. And obviously, like new technologies and stuff are coming out at such a quick rate, it's really exciting to see some of these use cases that these organizations are starting to adopt internally.

Ajay Yadav: Great Travis. And with this, as we are wrapping up the podcast, I would like to pick your brain and would like to understand, how do you see Gemini Enterprises being used across industries like banking, insurance, healthcare, retail? Your perspective with some of the real-life use cases that you might have solved for the problems for the businesses, and where do you think Persistent can help the client leverage Gemini Enterprise to drive transformation in their organization?

Travis Norenberg: You know, for most of this conversation, I think we've, we've been kind of talking high level agents, large language models, you know, Generative AI. But bringing up gen, uh, Gemini Enterprise, I think is really important for the sake of this conversation, just as you know, kind of going back to that enterprise platform for agentic and large language model adoption.

We're seeing a lot of very cool end-user use cases that are powered by Gemini, Gemini Enterprise, as kind of that single hub that an end user within a company can go into. You have access to obviously the Gemini models for search summarization, content generation, but in addition to that, we're seeing a lot of these kind of like process automation types of use cases, which is really exciting. And again, being able to do that in one platform

where you go in, you're like, here's my whole list of agents that I have available to me. I'm connected into all of my enterprise data stores. We're seeing some really interesting things there.

So, just a couple of examples here. Um, in banking, this is actually a area of huge adoption for us and doing things like agentic fraud detection, uh, is a use case that we're seeing a lot of right now. And even like going a step beyond that and providing personalized financial advice.

So if you're a financial advisor working with a client and, you are trying to analyze their portfolio and identify, um, you know, diversification strategies, for example, we're seeing some really cool use cases on that front. Really powered by kind of that breaking down of silos, bringing in the relevant data within the organization to really have that kind of personalized experience is extremely powerful.

Another area that we're seeing a lot of adoption is insurance. So, doing things like streamlining, like claims processing, and again, like customer service answering, answering questions about policies, how can we service these policies for our customers? Identify gaps where, maybe, there's some additional coverage that might make sense for that particular customer, is another really great use case. So, I think, kind of looking at those two in particular, it's really bringing that personalized experience to your customer all powered on these agentic workflows.

Ajay Yadav: Absolutely. And with certain customers, our experience, and I'll just give you a very recent deployment we made for. One of the food and beverages company, they wanted to empower their field sales agent with the real-time intelligence. And they use one of the third-party data warehouse where they have lot of data about their billing systems, about their distributors, about their products and whatnot. But it includes a heavy data in changing to provide that information to the field sales individuals and that too in real time. So, what we did, we deployed an agent which directly talks to the data on that third-party data warehouse, and from that third party data warehouse now these salespersons can directly look at what kind of products or SKUs this particular distributor has consumed, or what does a history from a transaction plan look like? They're clean from the payment, they're clean from the purchase, and what kind of potential this market holds, because they bring the third-party data as well, while running different digital ads or different physical promotions and whatnot.

So, they combine that power and then they are making their field sales rep intelligent in real time, that consume the information to have more meaningful conversation with your distributor, eventually increasing your sales or the top line. While selling or cross-selling or upselling based on the real-time data insights based on how the customer behavior is trending today, what does the predictive model look like and whatnot.

This is really cool, but at the same time, this was like, okay, if I have to be on a field, I need to raise a request at least a week ago. To figure out. And in the week, I'm sure a lot of things might have changed. So, this is the power that AI is enabling to the customer today.

And definitely, as Persistent, we bring that deep domain expertise while working with these industries. Similar use cases for the healthcare BFSI that we have been solving, for our customers, ultimately helping them to partner with an organization who understand the business domain, who has great expertise in the integrations like we spoke about, how critical it is, who have great expertise in the data.

As we spoke about the underlying platform remains the whole, or the partner who understand why it is most relevant to have, scalable secure foundation models to expand on your engine to the AI journey. So, in short, I'd see a lot more industries and how Persistent can be a real partner of choice for some of these customers.

Travis Norenberg: Yeah, totally. That's a really good use case too. I mean, I think getting into like the real-time data processing as well is such a great use case and it's something that historically has been pretty difficult to kind of manage and now we're really empowering everybody to have access to data, to have access to these tools that are going to help make decision making easier. And so, I love hearing about use cases like that.

Ajay Yadav: With that, we come to conclusion and I would like to take this opportunity to thank you Travis, for sharing your insight.

Travis Norenberg: Yeah, absolutely. Thanks for having me today, Ajay. I really enjoyed the conversation.

Ajay Yadav: Thank you. At Persistent, we are committed to helping our clients reimagine what's possible with AI. Connecting technology with people and process for lasting impact.

And thank you for tuning into Reimagined Conversations. If today's episode sparked new thinking, follow the show and share it with your network.

Have a story to tell or a guest you would love to hear from. Drop us a note at podcast@persistent.com.

Until next time, stay curious, stay inspired. Thank you.

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 26,500 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)magingTM the World