

Death by Data? p. 16



40+
YEARS
SINCE 1982



How GNX Redefined BlackBerry's Future: The Acquisition That Transformed Research in Motion p. 44

NetSuite's Global Vision: Building the Intelligent Enterprise p. 19



V. Rajaraman: The Teacher Who Built India's Computing Mind, No More p. 72

www.dqindia.com

VOL XLI No 12 | DECEMBER 2025 | ₹100

DATAQUEST

CyberMedia

THE BUSINESS OF INFOTECH

THE PRIVACY RESET INDIA CAN'T IGNORE

DPDP reshapes how businesses handle data, consent, breaches and governance



BEYOND THE CLOUD, INTO INTELLIGENCE

AWS powers silicon, agents and sovereign AI for enterprise growth

MATT GARMAN
CEO, Amazon Web Services



Pilot or Paradox: Where are you parking your AI today?

Fragmented data, model pluralism, lack of a fabric, not enough skills, model economics, model volatility and the blank page syndrome- everything matters when it comes to making sure that an AI pilot does not end up as a paradox. And whether you are in that '5 pc' club?

By Pratima H

It's surprising. But also not so much when we see how a recent MIT study found that 95% of AI pilots are never able to scale beyond proof-of-concept. There are so many struggles around fragmented data, lack of integration, unclear ROI, and gaps on explainability and compliance that the AI Pilot Paradox is as real and full of irony as many up-and-new technology dead-ends. In this interview with Deepak Dastrala, Partner & Chief Technology Officer at Intellect Design Arena Ltd, we get to drill deeper into the many pieces that complicate this puzzle. It's not 'hell' yet but it's quite 'black' for sure. Let's face this jigsaw.

Why does the AI Pilot paradox even exist? Is it a region-specific or vertical-specific or size-specific pattern?

This 'pilot paradox' is something I see everywhere; it's not limited to one region or vertical. I see the same pattern in banks in the Middle East, insurers in the US, and wealth firms in Europe.

“ The '5 pc' ones treat the pilot as 'version 1 of a product'. There's a product owner, a clear roadmap, and a dedicated budget for iteration. The key question isn't 'Did the POC work?' but 'What did we learn, and what do we roll into v2, v3, and v4?' **”**



DEEPAK DASTRALA
Partner & Chief Technology Officer,
Intellect Design Arena Ltd

“ Key elements like integration, security, data readiness, and change management are postponed to ‘after the pilot’- which means they often never happen.

It stems from a fundamental mistake: organisations treat AI as a temporary experiment, not as a core product or system. Pilots get scoped as ‘tech demos’ rather than as real changes to how work actually gets done.

Key elements like integration, security, data readiness, and change management are postponed to ‘after the pilot’- which means they often never happen. You end up with misaligned teams, business, IT, ops, and risk, and no single end-to-end owner.

While the friction might be higher in sectors with heavy regulation and legacy systems, like BFSI, the underlying behavior is the same. AI is kept at the edge of the organisation, never truly wired into the core.

If (As found by a MIT Media lab) report, only five per cent AI pilots make it to production, what are these five per cent doing right?

From what I see, the successful per cent share a common playbook. First- They start from a business journey, not a model. They pick a specific process (like complaints handling, KYC refresh, or submission intake), define two or three hard KPIs, and design the AI around that journey. Second- They treat the pilot as ‘version 1 of a product’. There’s a product owner, a clear roadmap, and a dedicated budget for iteration. The key question isn’t ‘Did the POC work?’ but ‘What did we learn, and what do we roll into v2, v3, and v4?’. Third- They build on a platform, not a one-off stack. They use an ‘AI fabric’ like our Purple Fabric, for data, knowledge, agents, and governance. This makes the second and third use cases significantly faster and cheaper to deploy.

What else do they do right?

They plan for integration and governance from day zero. They design exactly how the AI agent will sit inside the existing workflow, how it will authenticate, how it will be monitored, and what controls the risk and compliance teams have. They also invest in change management. They explicitly design how their teams (RMs, underwriters, ops) will actually use the agent, how performance is

measured, and how feedback loops will improve the system over time.

What kind of business functions generate better/more tangible/speedier AI outcomes - and are these more on the side of cost-savings or productivity?

In my experience, three areas show the quickest, most tangible results. First. Operations and Shared Services: This includes functions like complaints, customer service, KYC/AML refresh, and submission intake. The outcomes here are very clear: 30-70 per cent reductions in manual effort, faster cycle times, and fewer errors. It’s primarily a cost and efficiency play with very clear before-and-after metrics. Then comes the area of Knowledge-Heavy Decision Support. Think of underwriting, credit reviews, portfolio analysis, or sales support for RMs. Here, the value is productivity and quality. You get better, more consistent decisions and faster onboarding for new staff. This often translates to new revenue through better cross-selling or more accurate risk selection.

Finally, there is Risk and Compliance. This covers areas like policy interpretation, regulatory mapping, and surveillance. The main outcomes are lower compliance risk, faster responses to regulatory change, and better traceability. While the early waves are often framed as cost and speed, the sustainable value, especially in BFSI, is the combination of productivity, enhanced risk quality, and new product possibilities.

The same MIT study also noted that external partnerships reach deployment about twice as often (around 67 per cent) as internally built efforts (roughly 33 per cent). Can we explain that?

I see a few clear reasons why a good partnership has a higher ‘deployment hit rate’. Pattern Reuse is where external partners come with ready-made patterns and accelerators. You aren’t starting from a blank page each time. Then there is platform maturity (a mature platform has already solved for multi-tenant security, observability, agent lifecycle, cost control, and rollback). This gives it a far higher chance of passing internal scrutiny. Also most internal teams are stretched thin with business-as-usual changes, cloud

“ The problem is that fragmented data, tools, and ownership are why organisations end up with 10 pilots and zero platforms.

migrations, and other data initiatives. A good partner brings a blended team of applied AI, MLOps, domain, and product skills in one package. Plus, an external partner can often articulate risk, compliance, and controls more clearly, simply because they have gone through multiple regulatory reviews elsewhere.

That said, the best outcomes I've seen are hybrid.

How crucial are adequacy of data for training, existing automation etc. to the success of AI projects?

They aren't 'nice to have', they are the pre-conditions for moving beyond pilots. You don't always need perfect data, but you need sufficient, accessible, and well-understood data. In many of our Purple Fabric use cases, 60-70 per cent of the effort is just clarifying where data lives, who owns it, and how it can be used safely. Also, Process Clarity matters. If the underlying process is undocumented and inconsistent, AI will just amplify that chaos. Even simple workflow automation or a clearly mapped process makes AI deployment much cleaner and safer.

What about integration readiness and internal ecosystem?

An API-first, event-driven architecture and clear patterns for connecting to core systems dramatically reduce the time from POC to production. If architecture, security, risk, data, and business all operate in silos, every AI initiative will get stuck in approvals. You need a basic AI operating model with agreed-upon standards to reduce that friction.

Can solving fragmentation issues and skill-set availability help in a big way?

Absolutely. That's one of the biggest levers you can pull.

On the fragmentation side, the problem is that fragmented data, tools, and ownership are why organisations end up with 10 pilots and zero platforms. You have to counter this with a unified AI fabric where data, knowledge, and governance come together, alongside clear ownership of journeys, not just individual applications. On the skill-set side, the gaps are very real. You don't just need data scientists or prompt engineers. You need a blend of domain, AI,

engineering, and product skills. You need people who think in terms of use cases, metrics, and adoption, not just models. When you address fragmentation with a platform and address the skill gap with cross-functional AI pods and upskilling, the success rate of AI initiatives improves dramatically.


Are formats like AlaaS an answer in some way to this problem?

AlaaS is a great starting point. It lowers the entry barrier and gives you access to powerful models and scalable infrastructure. But on its own, AlaaS does not solve for your specific domain knowledge and policies, complex integration and security architecture and internal operating model and change management. AlaaS is part of the answer, but only if it's anchored in your enterprise context, not just exposed as another external API.

Anything else that CIOs should think of and plan better here?

Some reflections can help here. Like - Define your AI operating model early. Who owns the platforms? Who owns the journeys? How do you prioritise use cases, what are your guardrails, and how do you measure value?

Also- Treat governance and trust as first-class concerns. This means implementing runtime controls, like ISO 42001-style policies, roles, and risk assessments, not just writing documents.

It will also help to plan for model volatility and economics. Models and prices will constantly change. You must architect for 'model pluralism' and cost observability from day one. That said- Don't try to jump straight to full autonomy. Start by assisting your people, then augmenting them, and only automate the portions where confidence and controls are strong. And most importantly, treat AI as a long-term capability, not a series of projects. You are building an AI-native foundation, a platform, skills, governance, and patterns, that needs to support dozens of use cases over the next three to five years. 

pratimah@cybermedia.co.in