Bridging the GenAI Divide:

Lyzr's Perspective on MIT's State of AI in Business 2025

By Lyzr AI

Table of Contents

MIT's State of AI in Business 2025 reveals a sobering truth: despite $30–40 billion in enterprise GenAI investment, 95% of pilots are failing to produce measurable P&L outcomes.

The report, based on over 50 executive interviews and 300+ project analyses, found that the vast majority of initiatives stall out, caught in proof-of-concept purgatory, unable to transition to production systems that deliver business value.

The cause, MIT says, isn't weak models or regulatory friction. It's **execution.**

MIT's research confirms what many of us have experienced: the majority of GenAI projects don't make it past the pilot stage.

That reality forces a shift in focus, from excitement over models to discipline in execution. Below, we've outlined how we read MIT's findings and how our own work has shaped a perspective on crossing this divide.

1. Execution, Not Model Selection

MIT's finding: The report makes clear that success in GenAI has little to do with model quality. As MIT put it, *"this divide does not seem to be driven by model quality or regulation, but seems to be determined by approach."*

In other words, execution matters more than which foundation model is chosen. Off-the-shelf tools such as ChatGPT or Copilot are already powerful, but most enterprises fail to turn that capability into measurable results. The problem lies in adoption, integration, and governance, not the underlying technology.

Lyzr's perspective: Our experience matches this conclusion. A slightly less advanced model that is embedded in a workflow and owned by a business team will outperform a cutting-edge model that never moves past a demo.

We have seen enterprises fixate on which algorithm to use when the real determinant of value was whether the system could be deployed in a critical process and improved continuously. For us, the focus is not AI sophistication in isolation but business impact. The real bottleneck is the discipline to implement, iterate, and operationalize GenAI at scale.

How Lyzr aligns: Lyzr's approach from day one has been execution-focused. That means working closely with business units to understand their day-to-day operations and success metrics before we write a line of code. Our GenAI platform emphasizes robust deployment pipelines, easy integration with existing systems, and user-centric design.

We measure success in outcomes (e.g. reduced processing time, higher customer satisfaction), not in abstract model metrics. By prioritizing real-world execution, we ensure the AI actually gets into production where it can create value, reinforcing MIT's finding that approach trumps algorithmic glory.

In practice, this execution focus is built into our process. We've partnered with over 1,000 customers using this model, co-creating AI workflows side by side with them. Rather than deliver a black-box model, we collaborate to define the right mix of agents, tools, models, and data needed to solve each specific problem.

Crucially, we layer these AI solutions onto the customer's existing technology stack, augmenting their current CRM, ERP, or other systems instead of forcing a rip-and-replace. And our commitment doesn't end at deployment. Lyzr continues to maintain and support each agent over time, ensuring it remains effective and doesn't drift out of sync with evolving business needs.

Lyzr's Execution-Focused AI Approach

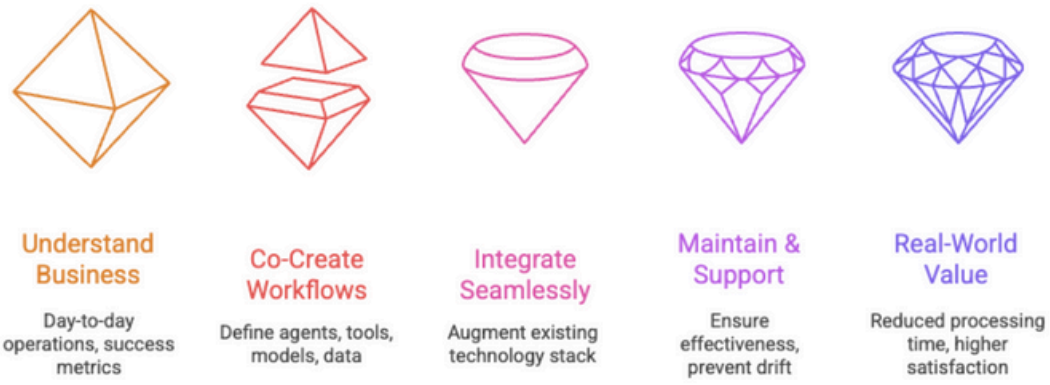

2. Beyond Brittle Pilots

MIT's finding: The report confirms an epidemic of "pilot purgatory." While nearly 40% of organizations have experimented with GenAI tools, only 5% of enterprise AI projects reach full production deployment. Most pilots never scale.

Why? According to MIT, "Most fail due to brittle workflows, lack of contextual learning, and misalignment with day-to-day operations." In plain English, these AI prototypes break easily, don't adapt, and aren't built into the actual business process. They remain science projects. The result: months of effort, exciting demos, but no real business impact, exactly what's driving that 95% failure rate.

Lyzr's perspective: We've all seen the flashy AI pilot that impresses in a slideshow but crashes and burns in the real world. Perhaps it only works on a curated dataset, or it can't handle an edge case so the team loses trust.

Often, pilots are built in isolation by an innovation lab, far from the people and systems that run the business. At Lyzr, we've learned that a pilot not designed with production in mind is a dead end. If a solution can't survive messy real-life conditions, unexpected inputs, system integrations, evolving user needs, it's going to be brittle. The goal should never be just to build a proof-of-concept; it's to build the first version of a product that will be used. That mindset change is critical.

How Lyzr aligns: Our ethos is "no pilot left behind." We co-develop AI agents with an eye toward day-one usefulness on the job. That means early integration with live data sources, enterprise software, and workflow tools, not a siloed sandbox. We also involve actual end-users early, to ensure the AI fits how people actually work (preventing the day-to-day misalignment MIT highlighted).

Additionally, we design pilots to handle the unexpected gracefully. If our AI agent encounters a scenario it's unsure of, it won't just break or make something up, it can defer to a human or fall back to a deterministic rule.

This safeguard prevents catastrophic errors in production and creates a learning opportunity: when a human corrects or handles an edge-case, that feedback is captured and fed back into the system. The next time, the AI is better prepared. By planning for real-world contingencies (and integrating the AI directly into the tools and interfaces employees already use), we avoid the brittle "science project" traps and build solutions that can survive in production from day one.

AI Agents continuous important cycle

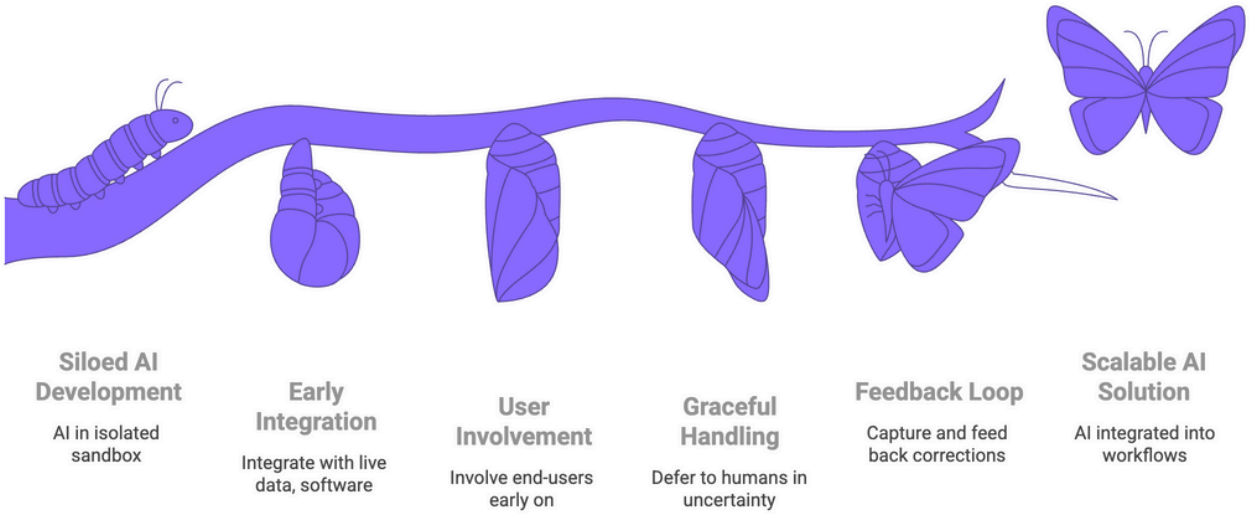

3. Mind the Learning Gap

MIT's finding: Perhaps the most striking insight from MIT's research is that the chief barrier to scaling AI is learning, or rather, the lack of it. They observed that "most GenAI systems do not retain feedback, adapt to context, or improve over time."

This "learning gap" isn't about human learning; it's about AI tools and organizations failing to learn and adapt. An AI solution might be static, it does one thing on day one and the same thing a year later, never getting better or smarter from real-world experience. In a dynamic business, that's a recipe for obsolescence. As one summary put it, even a decent resource allocation can't save "an artificial intelligence that can't learn to be more intelligent."

Lyzr's perspective: This finding resonates deeply with us because we've seen the pitfalls of one-and-done AI deployments. A pilot is trained on last quarter's data, set loose, and then... nothing.

No fine-tuning, no iterative improvement, no incorporation of user feedback. Six months later, the model's recommendations are outdated or off-base, and everyone wonders why the AI "failed."

The truth is, **AI projects need to be treated as living products, not fixed deliverables**. They require ongoing learning, both the algorithm learning from new data and the organization learning how to best utilize the algorithm. There's also a human learning

curve: teams need to adapt processes and upskill people to work alongside the AI. MIT essentially calls out that static AI projects will stagnate. Candidly, if your AI isn't getting better with use, it's getting worse.

How Lyzr aligns: At Lyzr, we design every agent with continuous improvement loops. In practice, this means our GenAI agents have memory and feedback mechanisms built in, they're engineered to capture corrections, preferences, and new context from users, then refine their outputs accordingly.

For example, when one of our deployed agents for document processing makes an error and a human corrects it, that feedback doesn't vanish; the agent learns from it and updates its internal model or rules. We also stress contextual learning: connecting the AI to enterprise knowledge bases, past transactions, and relevant context so it isn't operating in a vacuum.

On the organizational side, our deployment teams work closely with client stakeholders to establish a cadence of model reviews and updates. In many cases, we enable the client's own developers or analysts to go "under the hood" of our platform, adjusting an agent's knowledge base or parameters as their business evolves.

This way, AI and the business evolve together. By actively closing the learning loop, we tackle the "learning gap" head on, ensuring the AI doesn't just work on day one, but keeps getting smarter and more aligned with the business on day 100 and beyond.

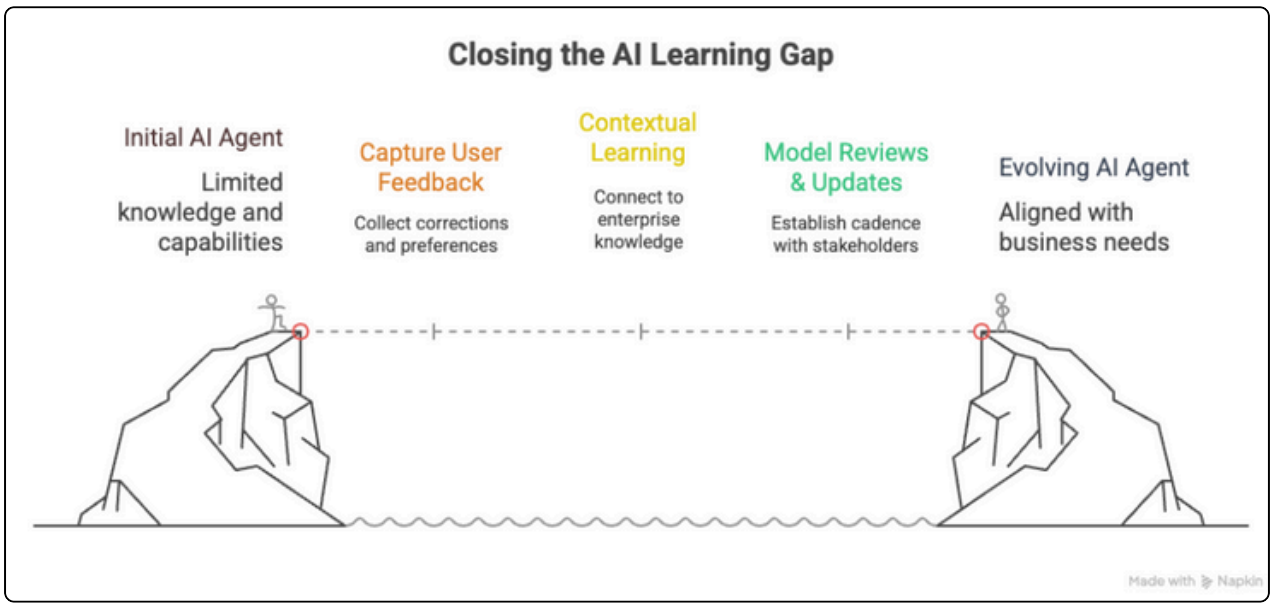

4. ROI Is Hiding in the Back Office

MIT's finding: The report highlights an "investment bias": more than half of enterprise GenAI budgets flow into customer-facing pilots, while the clearest ROI has often come from operational areas like finance, compliance, and IT .

MIT points to back-office automation as delivering some of the most measurable returns, citing examples such as eliminating business process outsourcing and reducing external agency costs. The conclusion is not that front-office projects have no value, but that organizations may be underestimating the scale of impact in less visible parts of the business.

Lyzr's perspective: We see both sides of this dynamic. Sales and marketing pilots can be powerful when designed with the right KPIs in mind, improving lead quality, accelerating content workflows, or deepening personalization at scale. These projects often create momentum and executive attention that open doors for broader AI adoption.

At the same time, MIT is right to note that back-office functions are where AI can quietly unlock millions in savings and efficiency. Invoices, claims, compliance reports, and vendor contracts rarely make headlines, but they drive tangible cost avoidance and speed. For us, the key isn't to elevate one over the other, but to recognize that ROI looks different depending on where you point GenAI, visible growth impact in the front office, and often more immediate cost or cycle-time benefits in the back.

How Lyzr aligns: Our platform is designed to support both. We've built and deployed marketing and sales agents with key clients, from lead generation to content production, but we also deliver pre-built templates for functions like HR onboarding, compliance monitoring, and finance reconciliation. One client automated vendor contract review, cutting cycle times and reducing legal outsourcing; another used Lyzr agents in marketing operations to accelerate campaign execution without adding headcount. By treating front-office and back-office use cases as complementary, not competing, we help organizations pursue both the visible wins and the hidden efficiencies. MIT's data is a reminder to broaden the lens: don't overlook the operational goldmine, but don't abandon the innovation momentum that sales and marketing AI can create.

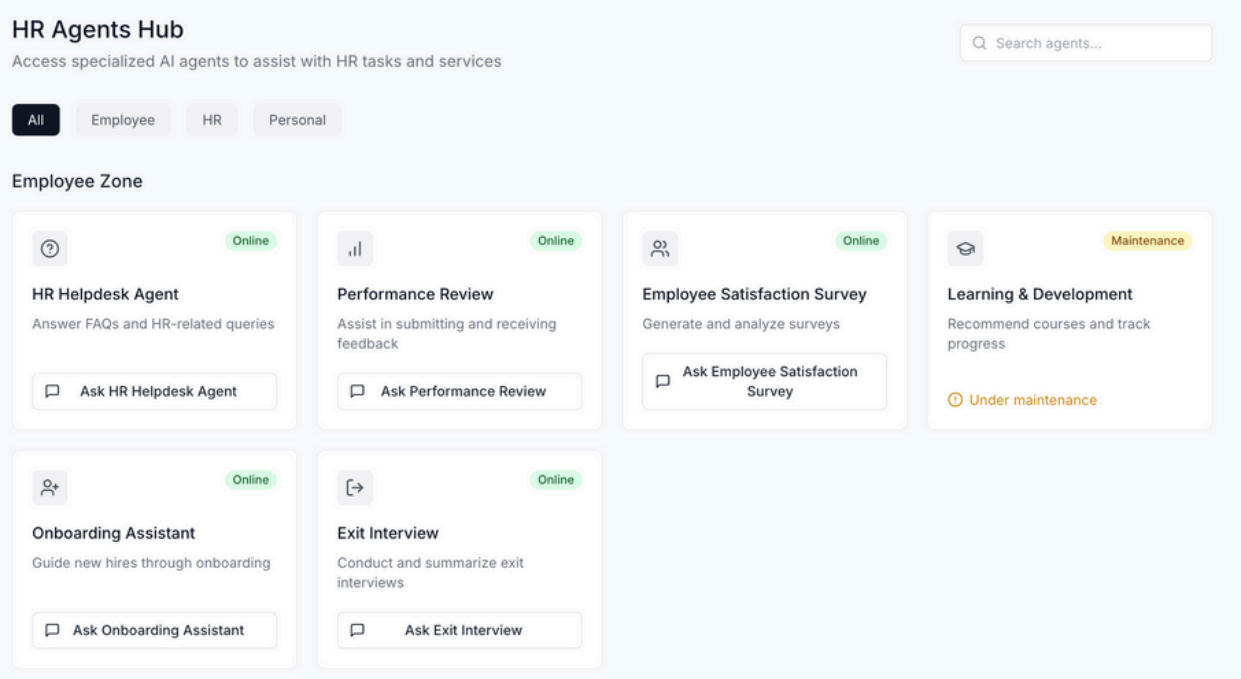

5. Collaboration Beats DIY (The Partnership Edge)

MIT's finding: Finally, the report highlights an "implementation advantage" enjoyed by companies that partner with external AI experts versus going it alone.
According to MIT's research, "External partnerships see twice the success rate of internal builds." In fact, they found that vendor-partnered AI initiatives reached deployment roughly 67% of the time, vs. only ~33% for in-house efforts. That's a huge gap. Why the difference? Successful companies treat AI vendors not just as software providers but as true partners – co-developing solutions and holding them accountable to business outcomes. This approach brings in specialized expertise and shared responsibility.

By contrast, purely internal projects often struggle with limited talent, slower learning, and no outside pressure to deliver results. As one MIT author noted, the standout AI successes often "execute well, and partner smartly" as part of their strategy.

Lyzr's perspective: This point hits home for us, for obvious reasons, Lyzr often plays the role of that external AI partner. We've seen the difference a collaborative approach makes. Enterprises that try to build everything from scratch internally often underestimate the complexity (integration, maintenance, evolving the model, etc.), and they reinvent a lot of wheels.

Meanwhile, those that bring in a focused partner can leapfrog years of trial and error. But there's a caveat: not all vendor relationships lead to success.

The key is shared ownership and accountability. MIT's finding that partnerships outperform internal builds holds true when both sides commit to clear business KPIs and honest feedback loops. From the founder's perspective, I also recognize that external partners must earn that trust, not just through transparency and adapting to the client's processes, but by putting real skin in the game. And for us, that isn't a throwaway phrase. It means we're betting that the agents we build will actually go live in production, not sit in pilot limbo. It means we're committing the resources, time, and expertise required to get there, even when it's harder than expected.

How Lyzr aligns: Lyzr was founded on the principle of co-building with clients rather than simply delivering software. One concrete way we do this is by embedding what we call Forward Deployment Engineers into a project, think of them as an extension of the client's team, working hand-in-hand to adapt our GenAI platform to the client's unique needs.

We don't just drop a tool and leave; we jointly define success metrics (e.g. "reduce average handling time by 30%" or "cut outsourcing costs in half") and we stick around to make sure those goals are met.

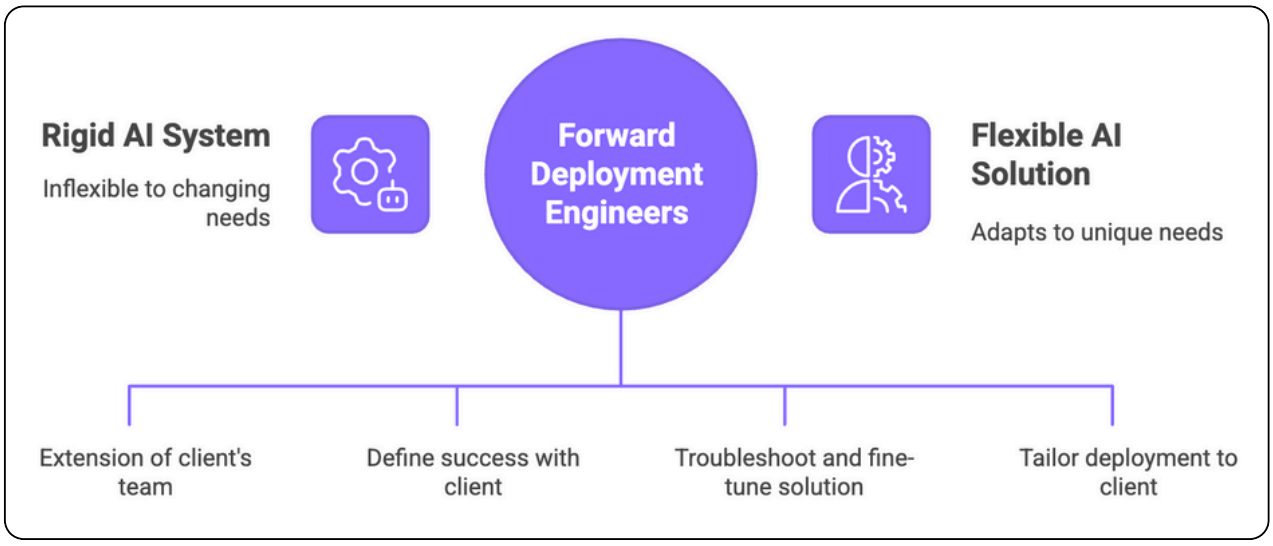

Lyzr's co-building approach

Our Forward Deployment Engineers act as always-available AI advisors along the way , essentially AI consultants to troubleshoot, fine-tune, and continuously improve the solution. They help solve the "rigidity" problem by ensuring the AI system stays flexible as requirements change or new edge cases appear. Many clients tell us it feels like our team becomes part of their company, and that's exactly our aim. When a client engages Lyzr, they're not just buying software; they're gaining a committed partner.

Importantly, we built Lyzr to be a flexible infrastructure (not a one-size-fits-all app) precisely so we can tailor each deployment to the client's processes. Our platform plays nicely with existing toolsets, we'll often deploy an agent right into your current workflow or IT environment rather than ask users to log into a new interface. By meeting clients where they are, we speed up adoption and minimize disruption. The result is a shared victory: when the AI solution works and delivers value, both the client and Lyzr win. And if it doesn't deliver, we take that just as seriously as the client does, if not more.

6. Safe and Responsible AI (Guardrails and Data Security)

MIT's finding: Although the report emphasizes execution, it also implies that trust and safety are prerequisites for that execution to matter. Most users will embrace automation only if their data remains secure and proper guardrails are in place.

In follow-up interviews, MIT noted significant skepticism toward new AI vendors in high-trust or regulated workflows; many leaders said they'd prefer to wait for known, established partners to offer AI solutions rather than risk an unknown vendor who might not understand their compliance needs. In short, if GenAI is going to scale in the enterprise, it must be enterprise-safe and reliable.

Lyzr's perspective: We completely agree, no enterprise will fully deploy an AI solution if it risks leaking data, violating compliance, or generating unpredictable outputs. In fields like finance, healthcare, or law, a single AI mistake can mean regulatory trouble or lost customer trust.

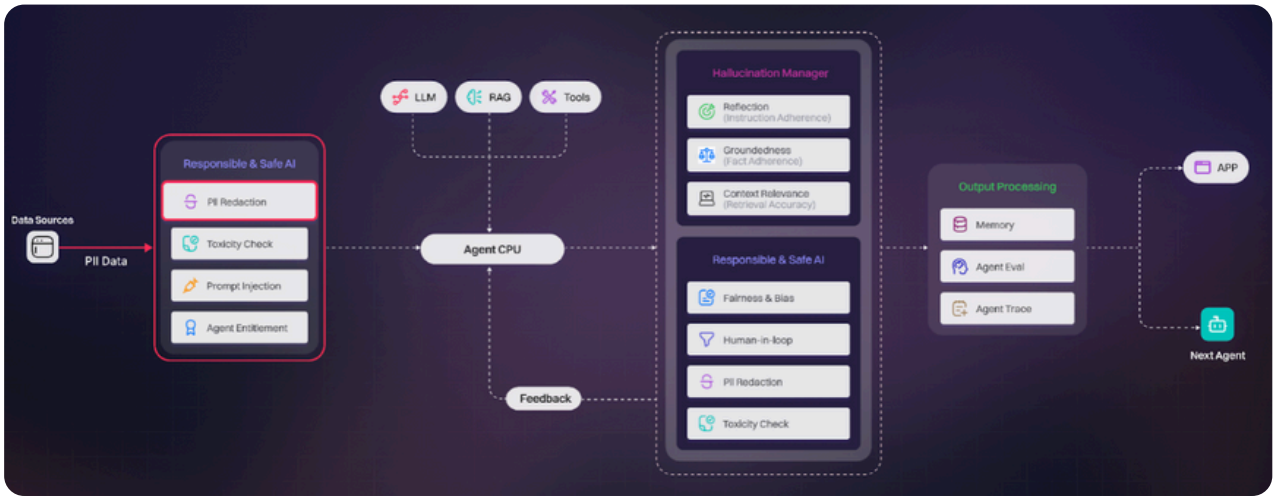

Fig: Lyzr's Responsible AI by Design

This addresses the data leakage problem directly: unlike teams pasting proprietary data into public ChatGPT (where it might inadvertently be used to train the model or be exposed), Lyzr ensures your data stays protected and private. We also implement automatic guardrails at the agent level: personal identifiers or other PII can be detected and redacted, and we filter outputs for toxicity or policy violations to uphold brand and compliance standards.

Beyond data security, we tackle the issue of AI “hallucinations” and reliability. Lyzr’s platform includes a Hallucination Manager, essentially a set of tools and workflows that double-check the AI’s responses against trusted sources or logical rules.
For critical use cases, our agents can cross-verify facts (e.g. comparing an answer to a knowledge base or database) and will gracefully defer or ask for clarification if they’re not confident, rather than spew a random guess.

We’ve modularized these safety components so that they’re always running in the background of an agent’s reasoning process. The goal is deterministic, controllable outcomes: the business should feel that the AI will do no harm even as it moves fast.

Finally, we know that winning enterprise trust isn’t just about technology, it’s about transparency and credibility. That’s why we’ve pursued industry certifications and partnerships to show that we meet high standards. (For example, Lyzr is an AWS Jumpstart partner, just as a hypothetical illustration, and has been vetted by global SIS

like Accenture.) We also provide full audit logs and explainability for our AI decisions, so stakeholders can always understand why the AI did something.

By addressing data privacy, safety, and accountability upfront, we remove a major barrier to GenAI adoption. In effect, we make “safe scaling” possible: executives and IT leaders can sign off on deploying the AI, because they see that it’s not a loose cannon, but a well-governed tool. When you don’t have to worry about leaks or wild outputs, you can focus on execution, which, as MIT showed, is where the real divide is bridged.

Conclusion: Accountability in AI, Lyzr’s MIT Guarantee

The MIT study makes one thing clear: closing the GenAI Divide isn’t about moonshot research or unchecked hype; it’s about rolling up our sleeves and taking a disciplined, accountable approach to AI in business.

The companies that succeed with AI focus on execution, insist on continuous learning, invest where it counts, and partner for the long haul. In short, they treat AI not as a magic wand, but as a strategic capability that must prove itself with real results. This ethos is exactly what we stand for at Lyzr. We’ve seen first-hand what separates the failed pilots from the scaled successes, and it comes down to those principles above.

As a final thought, I believe accountability is the thread that runs through all these findings. If you’re deploying AI, you should hold it (and your team and vendors) accountable to deliver tangible value, otherwise, why do it at all? As the CEO of a GenAI infrastructure startup, I feel this accountability personally.

That’s why we introduced the MIT Guarantee program at Lyzr, not as a sales gimmick, but as a commitment to our belief that production AI should deliver results. The premise is simple: if a company partners with us to build a generative AI agent and doesn’t get that agent to production, we pay them $50,000. Why $50,000? Because that’s how confident we are. It reflects the real time, resources, and deployment costs we’re willing to take on ourselves, and it shows we’re putting our reputation on the line. We succeed only when our customers do.

In closing, the era of experimentation-for-experimentation’s-sake in AI is ending. Business leaders are rightly asking, “How do we become the 5% that get it right?” The MIT report gives a data-backed roadmap, and Lyzr’s experience echoes it: nail the

execution, build learning into the system, target real ROI (even if it's not flashy), and don't go it alone.

Most importantly, demand accountability, from your AI, from your partners, and from yourself. If we all do that, we won't need to talk about 95% failure rates much longer. Instead, we'll be sharing the stories of AI projects that drive sustainable, scalable success across the enterprise. And at Lyzr, we'll be proud to have played a part in those stories.

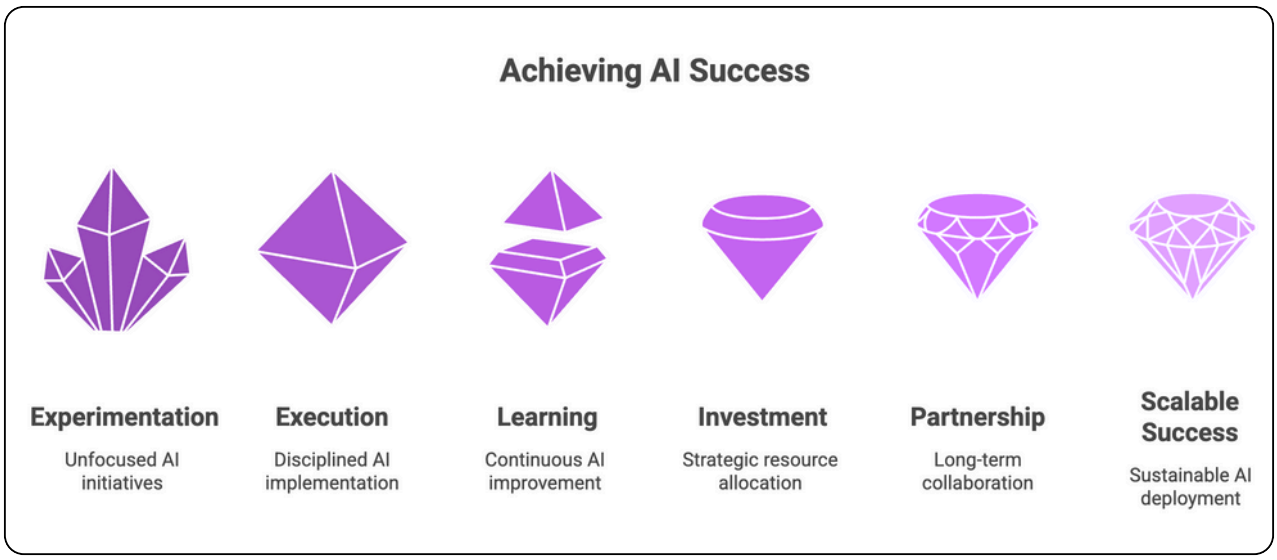

Sources: *MIT N/A NANDA, The GenAI Divide: State of AI in Business 2025[19][20]; Tech.co – Adam Rowe, MIT Finds 95% of Enterprise AI Pilots Fail to Boost Revenues[21][13]; Unframe AI Blog – "95% of AI Pilots Fail. Get on the Side of the 5% That Scale."[22][23]; Fortune (via Tech.co) – quote from Aditya Challapally[18].*

Thank You!

Book a Demo