Scale at Speed™

Whitepaper

Pitfalls to Pillars: Building a Strong AI Foundation for Enterprise Success

Authors

Alex Soejarto
Expert Advisor, Third Eye Advisory

Saurabh Jha
SVP & Global Head - Data & Analytics, Tech Mahindra

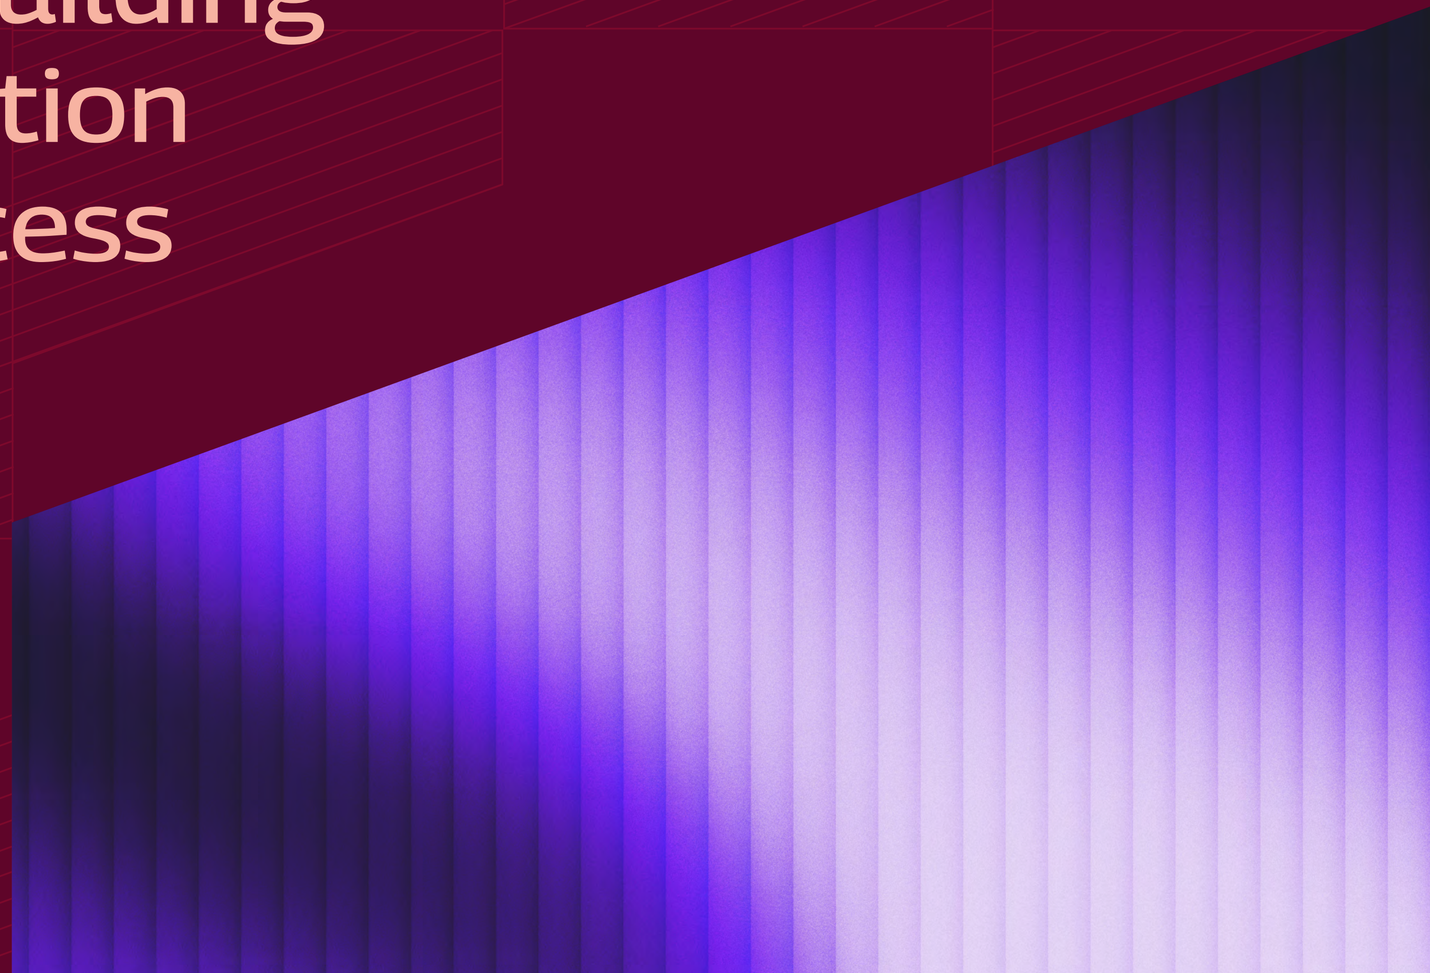

Executive Summary

Rapid adoption of generative artificial intelligence (GenAI), alongside advancements in agentic AI systems, signals a paradigm shift within enterprise contexts. These technologies promise a transformative era of hyper-personalized customer experiences, radical operational efficiency, and accelerated innovation. However, a significant chasm exists between this compelling vision and the reality of execution. Many organizations, in their rush to adopt AI, haven't established a foundation for enterprise scale deployments and are undermined by critical, recurring missteps. This results in a shifting footing as unstable as quicksand. This quicksand approach has three primary pitfalls threatening AI initiatives: the technology-first trap, data and AI function fragmentation, and the neglect of security and governance. Instead, success requires a disciplined shift to a stable foundation.

Creating a stable foundation requires a structured framework. This paper highlights how sustainable AI success is based on adopting an outcome-driven architecture, developing integrated capabilities, and designing security and ethics from day one. Leaders need to move beyond flashy demos and build robust and scalable AI systems that deliver tangible business value and a lasting competitive advantage.

Key Takeaways

- The gap between AI's promise with experimentation and its realized business value stems from foundational architectural and organizational flaws.
- Many enterprises prioritize speed of adoption over sustainable design, resulting in fragmented systems and limited scalability.
- Lasting AI success requires a disciplined, outcome-driven approach anchored in integration, governance, and ethics.
- The future of AI leadership will be rooted in executing systems of trust and accountability.

Content

Introduction

Autonomous systems are reshaping industries, and GenAI is creating opportunities for innovation at a pace never seen before. The vision is clear and compelling: intelligent automation driving new revenue streams, flawless data engineering workflows, and complex business processes managed with "hands-off," human-free efficiency.

Yet, beneath this surface of excitement, a dangerous reality is taking hold. For every AI success story, countless organizations struggle with fundamental execution gaps. They are investing heavily in AI initiatives that fail to move from pilot to production or, worse, fail to move meaningful business metrics. According to Forbes, 62% of employees consider AI overhyped in 2025.[1]

This disconnect between vision and reality stems from a fundamental flaw: organizations are building on a quicksand of unstable footings. In the rush to capture AI's potential, many are repeating the mistakes of past technology waves, prioritizing speed over a sustainable foundation. Tour extensive work with Fortune 500 enterprises and innovative startups, unpacks the recurring missteps that undermine digital transformation and presents a disciplined framework for building an AI foundation that lasts.

Five Questions to Ask Before Scaling AI

1. Are the AI initiatives explicitly tied to measurable business outcomes?
2. Do the data and AI teams share accountability for achieving those outcomes?
3. Are governance, ethics, and security embedded from the start—not added as a layer later?
4. How are fairness, transparency, and compliance ensured in the models and data pipelines?
5. Do the organization's technology platforms support scalability, collaboration, and continuous improvement?

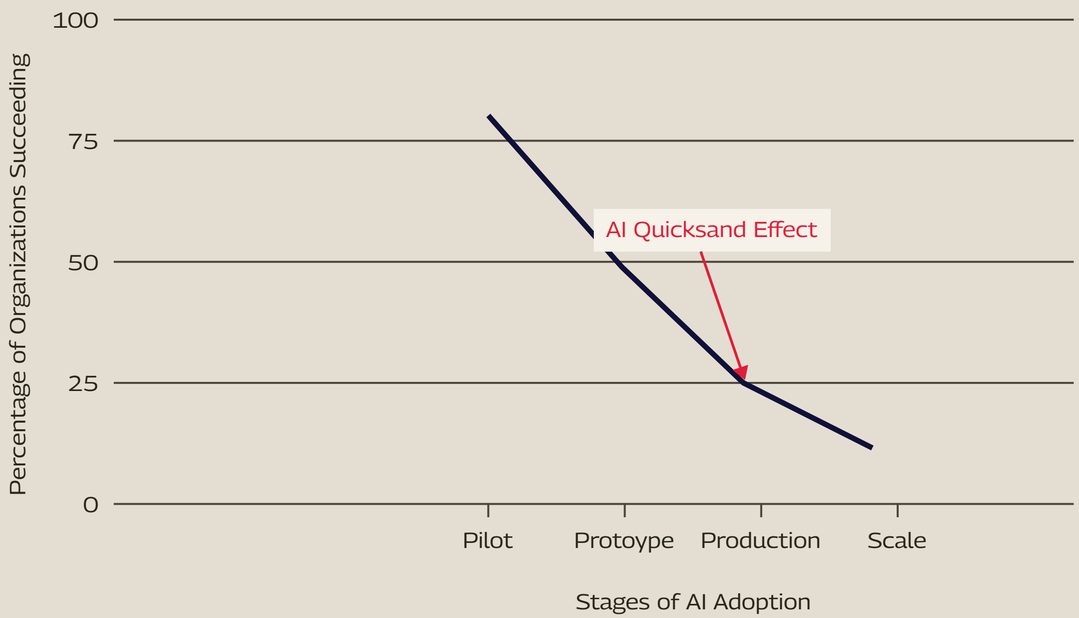

Core Challenge

The gap between AI's promise and its real-world value stems from three pervasive challenges. These pitfalls represent a fundamental misalignment between technology, people, and process.

Pitfall 1: The Technology-First Trap

The most common challenge is the "solution looking for a problem" syndrome. Organizations become fascinated by technology, whether it's the latest cloud data platform, a new visualization tool, or a novel AI model. This leads to:

- **Elaborate Data Swamps:** Expensive data lakes are built without clear use cases, quickly becoming costly and unusable.
- **Unused Analytics:** Complex dashboards are deployed, but nobody opens them because they don't answer the right business questions.
- **Models Without Mission:** Advanced machine learning models are created as impressive technical achievements but fail to drive or inform a single business decision.

This approach is fundamentally flawed. Technology must be an enabler of business strategy, not its driver.

Pitfall 2: The Data-AI Divide

The second critical error is organizational fragmentation. Many enterprises treat their data infrastructure and AI development as separate, siloed functions. This creates what can be called "broken data supply chains". According to BCG, "While 80% of CEOs and managers say they use GenAI a few times a week, only 51% of frontline workers do."[2]

In this siloed environment:

- Data engineering teams and AI/ML teams work in isolation.
- There is a lack of shared accountability for business outcomes.
- Solutions developed in a pilot environment cannot be scaled to a production-grade system.

AI is not a separate function; it is the pinnacle of an integrated capability stack. You cannot build a powerful AI system on a fragmented, broken data foundation.

Pitfall 3: The Security and Governance Crisis

Perhaps the most alarming trend is the relegation of critical functions to an "afterthought." In the race to innovate, organizations are treating security, responsible AI practices, data privacy, model integrity, and ethical use as items to be "bolted on" later.

This approach is not just risky; it is unsustainable

- Repeating Past Mistakes: It mirrors the early, chaotic wave of cloud adoption, where prioritizing speed over security led to massive technical debt and vulnerability.
- Eroding Trust: A lack of robust governance, bias detection, and fairness enforcement exposes the organization to significant regulatory, reputational, and financial risk.
- Building for Collapse: A "bolt-on" security architecture cannot scale with trust and transparency. Organizations that master this early will not only mitigate risk but also gain a significant competitive edge as regulatory frameworks inevitably evolve.

Organizations that fail to build systems that are secure and ethical by design are building a foundation destined to fail.

Framework for Sustainable AI Success

To close the vision-reality gap, organizations do not need more flashy demos. They need organizational discipline and rigorous architecture. Capturing AI's transformative potential requires three critical strategic shifts.

Pillar 1: Adopt an Outcome-Driven Architecture

Success in AI, as in all data and analytics, begins with business, not technology. An outcome-driven architecture reverses the "technology-first" trap. This methodology follows a disciplined path:

- Identify High-Impact Use Cases: Start by defining specific, high-value business challenges or opportunities.
- Establish Clear ROI: Define measurable success criteria and clear key performance indicators (KPIs) from the outset.
- Engineer the Solution: Build the technology stack and data pipelines specifically to meet those precise business objectives.

This approach applies to every level of the data stack, from basic reporting to advanced generative AI. It is the fundamental principle that separates an AI-driven cost center from a true competitive advantage.

Pillar 2: Develop Integrated Capabilities

Organizations that are winning with AI have broken down the silos between data, AI, and the business. They recognize that data engineering, machine learning, and business strategy are interconnected disciplines, not parallel tracks. Achieving this requires a profound structural and cultural change:

- Cross-Functional Teams: Form integrated units where data engineers, AI researchers, and business stakeholders share accountability for business outcomes.
- Integrated Platforms: Implement platforms that enable seamless collaboration across the entire AI lifecycle.
- Shared Goals: Align all teams around the same business metrics and objectives.

Pillar 3: Design for Security and Ethics First

In the AI-driven economy, trust is not a feature; it is the fabric. Security, governance, and ethics cannot be afterthoughts—they must be embedded into the foundational architecture from day one. A "security-first" and "ethics-by-design" approach involves:

- Treating Governance as an Enabler: View security and ethics as enablers of innovation, not constraints on it.
- Robust Frameworks: Implement comprehensive data governance, bias detection, and model integrity capabilities.
- Enforcing Accountability: Create organizational processes that ensure fairness and accountability across the entire AI lifecycle.

Organizations that master this early will not only mitigate risk but also gain a significant competitive edge as regulatory frameworks inevitably evolve.

AI Maturity Assessment Framework

Maturity Level	Characteristics
AI Beginner	Technology-first mindset, siloed functions, limited business alignment, and governance as an afterthought.
Integrated Practitioner	Partial data-AI integration, emerging governance structure, pilot-scale deployments, and inconsistent metrics.
Sustainable Leader	Outcome-driven strategy, unified data and AI value chain, ethics and security embedded by design with measurable business impact.

The Path Forward

The AI revolution is not a future event; it is the competitive landscape of today. Capturing its undeniable potential, however, requires more than just enthusiasm and investment. According to McKinsey, AI leaders unify data, AI, and security under a single business outcome charter.[3]

The winners in this new economy will be those who do the hard, foundational work of building AI systems that are:

- **Outcome-Driven:** Focused relentlessly on business value.
- **Robust:** Built on a supply chain that integrates data and AI.
- **Secure and Ethical:** Grounded in trust, transparency, and accountability.

Ultimately, the question is not if AI will reshape your industry. The question is whether your organization will be positioned to lead that transformation or be left struggling to catch up.

References

1] Drenik, G. (2025, October 14). AI's promise vs. reality and why 62% say it is overhyped. Forbes. https://www.forbes.com/sites/garydrenik/2025/10/14/ais-promise-vs-reality-and-why-62-say-it-is-overhyped/https://report.telekom.com/annual-report-2024/management-report/group-strategy/data-ai.html

2] Boston Consulting Group. (2025, June 26). AI at work: Momentum builds but gaps remain. https://www.bcg.com/publications/2025/ai-at-work-momentum-builds-but-gaps-remain

3] McKinsey & Company. (2025, November 5). The state of AI. https://www.mckinsey.com/capabilities/quantumblack/our-insights/the-state-of-ai

About the Authors

Alex Soejarto

Expert Advisor, Third Eye Advisory

Alex is recognized as a thought leader in business technology strategy. He has a unique view of the market as an enterprise strategist, a former service provider leader, and a recognized top influencer analyst. His business technology expertise spans Organizational Change, Enterprise Applications, Cloud, and Analytics emanating from his 16 years at Gartner as a VP of Research. He has been awarded Analyst of the Year by IIAR. Alex has a talent to synthesize market forces that will impact a client's competitive advantage in how it operates. He is the founder of Out Athlete Fund, avid triathlete, a mentor with StartOut, and a Board member of Cornell Pride

Saurabh Jha

SVP & Global Head - Data & Analytics, Tech Mahindra

With over 24 years of global experience, Saurabh has worked across India, Europe, the UK, and the US. He leads Tech Mahindra's Data and Analytics (D&A) practice, which helps enterprises strategize, design, implement, and deliver data and analytics, cloud-based data, and AI-related transformation initiatives. He has wide experience ranging from setting up new teams and practices, planning and executing go-to-market strategies, leading global alliances, and advising customers on effective alignment between their business goals and the latest digital technologies. Previously, he held strategic roles at Oracle, KPMG, and Mphasis, where he advised clients across industries and spearheaded regional expansions.

About Tech Mahindra

Tech Mahindra (NSE: TECHM) offers technology consulting and digital solutions to global enterprises across industries, enabling transformative scale at unparalleled speed. With 149,000+ professionals across 90+ countries helping 1100+ clients, Tech Mahindra provides a full spectrum of services including consulting, information technology, enterprise applications, business process services, engineering services, network services, customer experience & design, AI & analytics, and cloud & infrastructure services. It is the first Indian company in the world to have been awarded the Sustainable Markets Initiative's Terra Carta Seal, which recognizes global companies that are actively leading the charge to create a climate and nature-positive future. Tech Mahindra is part of the Mahindra Group, founded in 1945, one of the largest and most admired multinational federation of companies. For more information on how TechM can partner with you to meet your Scale at Speed™ imperatives, please visit https://www.techmahindra.com/.

TECH mahindra

www.techmahindra.com
www.twitter.com/tech_mahindra
www.linkedin.com/company/tech-mahindra