SPONSORED CONTENT | WHITE PAPER

Breaking out of pilot purgatory and scaling agentic AI across the enterprise

How Tech Mahindra & Google Cloud operationalizes enterprise-grade agentic AI systems

CIO

SPONSORED BY

TECH mahindra

Google Cloud

Contents

3 **Executive summary**

3 **Introduction**

4 **Defining autonomous delegation**

5 **The scalability wall: Structural barriers to autonomous operations**

6 **Industrializing autonomy: Tech Mahindra and Google Cloud framework**

8 **Cross-industry production deployments**

9 **Moving from pilots to production**

9 **Beyond the pilot: Achieving industrial-scale autonomy**

10 **Endnotes**

Executive summary

Although 2025 promised agentic artificial intelligence (AI) breakthroughs, global enterprises faced "pilot purgatory": Only 15% had scaled experiments to enterprise-wide deployment, according to Foundry's "AI Priorities Study 2025."[1] The rest struggle with fragmented pilots, stalled by integration, governance, and orchestration gaps. What's holding them back isn't algorithmic capability but a lack of industrial-scale orchestration.

This white paper advocates for a shift in AI strategy from human-assisted "chat" interfaces to autonomous agentic delegation. By integrating Google Cloud's high-fidelity intelligence layer with Tech Mahindra's proven implementation frameworks, organizations can resolve the governance and integration gaps that currently impede scale. The path to market leadership depends on operationalizing these agents as core capabilities. Enterprises that master this will secure long-term leadership and build the disciplined agentic foundations required for true autonomous transformation.

Introduction

In 2025 agentic AI was expected to go mainstream. The reality, however, has fallen well short of enterprise ambitions as deployment has stalled at scale. Most organizations are trapped in pilot purgatory, where appetite is high but outcomes remain limited to fragmented experiments.

Foundry's "AI Priorities Study 2025" reports that only 15% of organizations have achieved enterprise-scale AI (up from 8% in 2023); the remaining 85% are still running isolated pilots.[1] They most often cite integration, governance, security, and in-house expertise as their primary barriers. Governance frameworks that cannot scale and unclear returns further discourage deeper investment.

This stagnation reflects execution failure, not technology limits. IDC and Lenovo research indicates that roughly 88% of AI pilots never reach production, with most stalling on data, process, and infrastructure readiness.[2] An MIT study on enterprise generative AI (genAI) found that about 95% of pilots fail to deliver

measurable profit-and-loss impact, due to integration, data quality, and governance gaps.[3] Together, these findings make it clear that the core problem is rarely model capability; it's execution and orchestration at scale.

Legacy stability-oriented architectures clash with the demands of autonomous learning, leaving most initiatives stuck in proof-of-concept limbo. To bridge this gap, enterprises must resolve orchestration complexity and rethink how work is assigned to AI systems. Capitalizing on the potential of agentic AI requires a fundamental shift in perspective: moving from assistance to delegation.

The cost of inaction is a widening "intelligence gap" that will soon define the market leaders of the next decade.

Defining autonomous delegation

Traditional AI systems are inherently reactive; they analyze data and provide static recommendations, but at the critical final step, they still require human intervention.

Agentic AI departs from this "assistance" model. It introduces a system that reasons about enterprise data, plans multistep workflows across applications, acts autonomously within guardrails, and learns through continuous feedback. When successfully deployed, it shifts the interactions with AI from assistance to delegation.

The impact of this shift is already visible across high-stakes industrial applications:

- **Financial services:** The deployment of agentic architectures accelerates document processing for loan applications, replacing manual validation with autonomous high-precision workflows.
- **Healthcare:** Clinical documentation agents update patient records in real time during consultations.
- **Manufacturing:** Industrial agents run predictive maintenance that schedules repairs and procures parts before equipment fails.

The scalability wall: Structural barriers to autonomous operations

Every major cloud provider now offers AI platforms with agentic capabilities. But the challenge is operationalizing these systems within complex enterprise environments.

To move beyond this experimental stagnation, leaders must resolve four systemic friction points:

- **Opaque ROI frameworks:** Although pilots generate impressive technical demos, they often struggle to demonstrate concrete financial impact. Without validated return on investment (ROI) frameworks, initiatives stall at proof-of-concept, unable to justify the investment required for full-scale production rollout.

- **Integration complexity:** Enterprise technology stacks involve dozens of interconnected legacy systems. Agentic AI needs to operate seamlessly across this entire architecture, which often leads to long procurement cycles and exhaustive security reviews.

- **Governance gaps:** Trusting agents with autonomous decision-making requires robust control frameworks. Most organizations currently lack the technical policies and audit trails to govern AI systems at scale.

- **High initial investment:** Building production-grade agentic AI requires significant up-front investment in infrastructure, specialized talent, and implementation. In a climate of extreme spending scrutiny, initial costs cannot be overlooked.

At this point, the central question for leadership shifts: What is the architectural blueprint for directing autonomous agents across a fragmented enterprise?

Industrializing autonomy: Tech Mahindra and Google Cloud framework

Along with Google Cloud, Tech Mahindra overcomes the scalability wall by synchronizing enterprise-grade intelligence with industrial-scale implementation capabilities.

This partnership provides a unified operating model that moves agentic AI from isolated experiments to core business functions.

The intelligence layer: Google Cloud

Google Cloud establishes the high-fidelity reasoning foundation required for autonomous operations. Integrating intelligence directly into enterprise workflows ensures that agents remain grounded in governed data through three core capabilities:

- **Contextual reasoning:** Utilizes Gemini Enterprise and Vertex AI to enable agents to reason contextually across complex data sets while remaining anchored in BigQuery and trusted data pipelines.
- **Systemic connectivity:** Integrates agents with core enterprise systems, including enterprise resource planning (ERP), customer relationship management (CRM), and devops, through secure application programming interfaces (APIs) and connectors, to generate insights and execute actions in a controlled, auditable manner.
- **Inherent security:** Leverages built-in governance and adoption frameworks to ensure that autonomous actions comply with organizational security protocols from inception.

The outcome layer: Tech Mahindra

Tech Mahindra translates this technical architecture into measurable business impact.

By applying specialized engineering and domain expertise, the partnership ensures that autonomous systems deliver sustained value:

- **Industrial-scale engineering:** Utilizing domain-specific accelerators and AI engineering to move agents from proof-of-concept to high-volume production environments.
- **Comprehensive governance:** Implementing robust safety guardrails – including identity and access controls, cost governance, explainability, and regulatory compliance – to manage the risks of autonomous delegation.
- **Strategic translation:** Converting technical capabilities into specific business outcomes by aligning agent behavior with localized operational goals.

Here are two particularly compelling examples of this in action:

- **Industrializing manufacturing automation**
  A multisite U.S. conglomerate successfully transitioned from localized predictive maintenance to an autonomous enterprise-wide ecosystem. Earlier, it had to operate across heterogeneous operational technology data, site-by-site variations, and strict change windows.

  The agentic system addressed this friction by monitoring equipment signals in real time; flagging failure patterns; and triggering governed workflows, including ticket creation and part procurement. Tech Mahindra replaced a single-plant pilot with standard deployment patterns, monitoring, and guardrails scaled across facilities.

- **Optimizing telecommunications customer support**
  A global telecommunications leader utilized Gemini Enterprise–enabled agents to transform frontline customer service into a context-aware intelligence environment. These agents pulled data from disparate billing, CRM, and network status systems to provide real-time guidance without compromising data sensitivity or producing impractical answers.

  While Tech Mahindra handled the underlying system integration and access controls, the resulting architecture maintained accuracy and compliance even under massive production volumes.

Cross-industry production deployments

The shift toward agentic architectures is currently reshaping the global industrial landscape. Financial services organizations deploy multiagent systems for document processing and fraud detection, automating loan queries, expense validation, and anomaly detection. Insurance companies use agentic AI to personalize customer engagement, proactively identify upsell opportunities, and streamline claims workflows.

Life sciences firms run pharmacovigilance agents that autonomously classify and prioritize adverse events, shortening response times while strengthening regulatory compliance. Manufacturing operations implement industrial AI agents to analyze equipment data, predict failures, and schedule interventions before breakdowns occur.

Each deployment shares common characteristics: clear ROI metrics established up front, integration with existing enterprise systems, governance frameworks that ensure compliance and auditability, and measurable business outcomes within weeks.

Moving from pilots to production

The pathway from localized experiments to an enterprise-wide agentic ecosystem necessitates four core capabilities that provide the architectural discipline required to secure a sustainable return on intelligence:

- **Layered risk mitigation:** Tech Mahindra implements responsible AI frameworks with guardrails, policy enforcement, and monitoring to manage autonomous behavior without compromising innovation.
- **Cross-functional governance:** The formalization of unified policy frameworks that scale across business units and geographies ensures consistent agent behavior irrespective of deployment location.
- **Embedded explainability and auditability:** The integration of transparent reasoning chains into every agentic action. Systems log decisions, track data sources, and provide clear reasoning chains, enabling regulatory compliance while building trust.
- **Decreased time-to-value:** Rather than enduring lengthy custom development cycles, organizations deploy prebuilt accelerators through proven frameworks, tailored to industry-specific challenges, reducing implementation timelines from months to weeks.

Beyond the pilot: Achieving industrial-scale autonomy

Agentic AI represents a fundamental shift in how enterprises operate. The divide between market leaders and legacy firms will soon be shaped by the ability to transition from simple assistance to autonomous delegation. Those that remain stuck in pilot mode risk a widening intelligence gap that threatens enterprises' long-term competitive edge.

The hurdle has shifted from the availability of technology to the discipline of execution. Success is about moving beyond demos to production, isolated experiments to enterprise-wide deployment, and theoretical ROI to measured outcomes.

Tech Mahindra's partnership with Google Cloud provides the operating model necessary for this transition. Gemini Enterprise embeds enterprise-grade intelligence into real workflows, and Tech Mahindra provides rigorous implementation expertise and governance frameworks for global stability.

For organizations ready to move beyond pilots, the road map is clear: Prioritize high-impact use cases with measurable ROI, deploy proven governance frameworks, scale systematically across business units, and measure outcomes continuously.

Mastering this orchestration requires a pivot toward a future in which institutional intelligence is the only barrier between market leadership and structural obsolescence.

Learn how Tech Mahindra, together with Google Cloud, can accelerate your organization's agentic AI deployment at techmahindra.com.

Endnotes

1 From hype to reality: AI adoption gains traction in 2025. Foundry, 2025. https://resources-cs.cio.com/wp-content/uploads/sites/6/2025/03/R-ES_AI-Priorities_Foundry-brand_2025.pdf

2 88% of AI pilots fail to reach production – but that's not all on IT. CIO, March 24, 2025. https://www.cio.com/article/3850763/88-of-ai-pilots-fail-to-reach-production-but-thats-not-all-on-it.html

3 MIT report: 95% of generative AI pilots at companies are failing. Fortune, August 18, 2025. https://fortune.com/2025/08/18/mit-report-95-percent-generative-ai-pilots-at-companies-failing-cfo/

CIO SPONSORED BY TECH mahindra Google Cloud