



# BEYOND THE PILOT: WHAT IT ACTUALLY TAKES TO RUN GENAI AT ENTERPRISE SCALE

## Abstract

Enterprises are investing heavily in Generative AI, yet fewer than 6% report meaningful financial returns. The gap between pilot projects and production-scale impact is not a technology problem — it is an infrastructure, governance and change management problem. This article outlines the four foundational pillars enterprises must build to operationalise GenAI at scale: architecture, data, governance and talent. It offers an achievable roadmap for prioritising high-impact workflows, establishing a shared GenAI platform and embedding continuous feedback loops so AI investment translates into measurable business outcomes.



The boardroom conversation around Generative AI (GenAI) has shifted from [experimentation to execution](#). Today's defining challenge is no longer whether enterprises should adopt it, but how to operationalise it reliably at scale. According to McKinsey, nearly 80% of organisations now use GenAI in at least one business function, yet only around 5.5% report

seeing meaningful financial returns from their investments.

Closing the gap between pilots and business impact requires more than selecting the right large language model (LLM). It demands a coordinated strategy that includes the right architecture, data, governance and workforce readiness. For

enterprises that get this right, the benefits are significant. A McKinsey report estimates GenAI could add \$2.6-\$4.4 trillion annually to the global economy across all the 63 use cases it analysed. Marketing and sales, customer operations, software engineering and R&D are expected to account for 75% of this estimate.

## Why enterprises struggle to move beyond pilots

Most enterprises begin their AI journey with a promising proof of concept in one business unit. The difficulty arrives when those early successes need to [scale across the organisation](#). According to a Writer's report on GenAI adoption, 72% of C-suite executives say their company has faced at least one significant challenge on this

journey. Companies without a formal [AI strategy](#) report only a 37% success rate in implementing AI, while companies with a well-defined strategy report a success rate of 80%.

The underlying issues are rarely about the models themselves. It is most often the absence of sufficient infrastructure,

such as fragmented data environments, inconsistent governance, unclear ownership and a talent base that has not been prepared for AI-augmented workflows. Operationalising a GenAI platform means building something that is designed to scale from the very beginning, not layering AI onto existing constraints.





## The 4 pillars of enterprise-scale GenAI

Scaling GenAI requires investment across four interlocking domains. Weakness in any one limits how far the others can carry an enterprise.

### 1. Enterprise AI architecture

At scale, GenAI becomes an integration challenge. Gartner forecast that global spending on GenAI would reach \$644 billion in 2025, a 76.4% increase over 2024, driven by both vendor infrastructure build-out and accelerating enterprise software adoption. Competitive advantage increasingly comes not from the model itself, but from how effectively AI capabilities are connected to enterprise systems and workflows.

Leading enterprises are designing architectures with Application Programming Interface (API)-first integration, model abstraction layers that reduce dependence on any single provider, reusable prompt frameworks aligned with enterprise knowledge, hybrid or multi-cloud infrastructure and centralised orchestration and monitoring.

A layered architecture typically includes a data layer, model layer, orchestration layer, application layer and governance layer. This modular approach allows enterprises to update individual components without disrupting production systems.

### 2. Data foundation

GenAI systems are only as effective as the data supporting them. According to industry research, 73% of enterprises cite data quality as their biggest AI challenge. Supporting production-scale deployment requires centralised data platforms, pipelines for structured and unstructured data, real-time ingestion capabilities and strong governance around access and privacy.

Retrieval-augmented generation (RAG) has become a standard enterprise pattern because it grounds model outputs in proprietary business knowledge without requiring extensive fine-tuning. But RAG depends heavily on clean, well-indexed and governed data. Without a mature data foundation, even advanced models

produce inconsistent outcomes.

### 3. AI governance and responsible deployment

Governance becomes critical once GenAI operates at enterprise scale. Unlike traditional software, GenAI produces dynamic outputs that introduce risks around accuracy, bias, compliance and security. Effective frameworks include responsible AI policies, audit logging and output traceability, bias monitoring, human-in-the-loop (HITL) review for high-risk decisions and privacy controls for sensitive data.

Organisational governance matters equally. Enterprises need clarity on who owns the AI roadmap, how use cases are prioritised and how AI decisions are explained to regulators and customers. McKinsey highlights that organisations actively managing AI risks are better positioned to scale responsibly — yet, as of 2024, only 18% had established an enterprise-wide governing body with authority over responsible AI decisions.

#### 4. Talent and change management

Technology alone does not drive enterprise transformation. McKinsey data shows talent and skill gaps account for 46% of the obstacles executives cite when explaining why AI initiatives fail to scale — ahead of resourcing constraints and technical complexity. Building AI fluency across the

organisation, not only within a data science team, is a structural requirement. Successful enterprises are investing in AI literacy programmes for technical and non-technical staff, embedded AI specialists within business functions, clear usage policies and reskilling initiatives tied to workflow redesign. As

Writer's research found, employees who use AI platforms at work are twice as likely to champion broader AI initiatives internally. Treating AI adoption as an organisational transformation — not a technical deployment — is what separates enterprises that scale from those that stall.

### A practical roadmap for operationalisation

With foundations in place, execution follows a structured path.

#### 1. Define success before building anything.

Every use case should connect to a measurable business outcome such as improved customer satisfaction, reduced operational costs, faster releases or higher productivity. Enterprises generating the strongest AI returns focus on business value rather than experimentation volume.

**2. Prioritise high-impact workflows.** The highest-value enterprise applications cluster around customer operations, software engineering, knowledge

management and compliance automation. Rather than spreading efforts across dozens of disconnected pilots, enterprises should identify two or three domains where AI can fundamentally reshape a workflow and build deeply there first.

**3. Build a shared GenAI platform.** A platform-based approach creates the consistency and governance needed to scale. Shared infrastructure enables centralised monitoring and access controls, reusable AI services across business units, standardised compliance processes and faster deployment of new use cases. The

GenAI platform becomes the operating layer through which the enterprise continuously governs, deploys and improves AI capabilities.

#### 4. Establish continuous feedback loops.

GenAI systems drift — models update, data distributions shift, user needs evolve. Production deployments require ongoing monitoring for output quality, model drift, security risks, bias and business impact. Enterprises that embed evaluation into AI operations from Day One avoid the costly retrofits that come from treating AI as a set-and-forget system.







## Where enterprises are capturing the most value

Organisations generating measurable value are embedding AI deeply into operational workflows rather than treating it as a standalone productivity tool. McKinsey's analysis of AI high performers found they are nearly thrice as likely to have fundamentally redesigned workflows as part of their AI efforts, rather than simply layering AI on top of existing processes.

In software engineering, AI-assisted coding tools are accelerating development cycles, automating testing and improving developer productivity. In customer operations, AI systems are reducing response times and improving service efficiency. In knowledge management, natural language interfaces are making enterprise information more accessible

across departments.

Across all of these, success depends less on model sophistication and more on the quality of integration. Clean data, effective prompts, workflow redesign and human oversight remain the primary drivers of enterprise outcomes.

## The next phase of enterprise AI

Gartner projects that by 2028, 33% of enterprise software will include Agentic AI capabilities capable of handling complex, multi-step decisions. For enterprises, this means operationalisation is becoming a long-term strategic capability, not a short-term innovation project. Organisations building scalable AI foundations today are simultaneously preparing for more autonomous AI systems in the future.

The competitive advantage will not come from adopting AI first. It will come from operationalising it effectively, responsibly and at scale. That requires simultaneous investment in architecture, governance, data maturity and workforce readiness.

The window to build that advantage is real but not permanent. Enterprise AI spending is scaling at a pace with no precedent in modern software history. The organisations

that will lead in 2027 and beyond are those that treat operationalisation as the primary design constraint today. That means investing in foundations, not just features. It means governing from the start, not retrofitting guardrails after incidents occur. And it means measuring what matters most: business outcomes, not model benchmarks.

## How Infosys BPM can help

The Infosys [Generative AI in Business Operations](#) platform comprises a suite of customised, responsible design frameworks and ready-to-use BPM-focused solutions.

We have a robust blend of business and technology resources. Our domain expertise and multidisciplinary workforce combined with our cutting-edge technology assets drive AI transformation

quickly and at scale. Our extensive repository of pre-trained models and AI assets can accelerate the time to market and ensure superior outcomes.

For more information, contact [infosysbpm@infosys.com](mailto:infosysbpm@infosys.com)



© 2026 Infosys Limited, Bengaluru, India. All Rights Reserved. Infosys believes the information in this document is accurate as of its publication date; such information is subject to change without notice. Infosys acknowledges the proprietary rights of other companies to the trademarks, product names and such other intellectual property rights mentioned in this document. Except as expressly permitted, neither this documentation nor any part of it may be reproduced, stored in a retrieval system, or transmitted in any form or by any means, electronic, mechanical, printing, photocopying, recording or otherwise, without the prior permission of Infosys Limited and/ or any named intellectual property rights holders under this document.