

Scaling AI Beyond Pilots: Enterprise Approaches to Sustainable AI Transformation

Nitin Anand

Senior Vice President – Client Services,
Hi-tech and Professional Services



WNS
Part of Capgemini

Enterprise Artificial Intelligence (AI) has entered a new phase. After demonstrating value through pilots and localized use cases, organizations are now focused on embedding AI more deeply in day-to-day operations. According to a 2025 survey, 88 percent of organizations are using AI in at least one business function, reflecting one of the fastest adoption curves seen for any enterprise technology.¹ As usage expands, global enterprises and digital-first leaders are increasingly operating AI at scale, ensuring models are continuously supported, data remains reliable and performance is sustained across use cases.

Early adopters are already seeing the impact of this approach. Firms that have invested in integrated workflows and operating discipline around technology report strong returns, with an average 3.7x ROI on AI-related investments across industries and regions.² Sustaining these gains requires the right foundations – from secure data handling and governed workflows to the operating capacity needed to support model training and evolution.



¹<https://www.mckinsey.com/capabilities/quantumblack/our-insights/the-state-of-ai>

²[https://143485449.fsl.hubspotusercontent-eu1.net/hubfs/143485449/2024 Business Opportunity of AI_Generative AI Delivering New Business Value and Increasing ROI.pdf](https://143485449.fsl.hubspotusercontent-eu1.net/hubfs/143485449/2024%20Business%20Opportunity%20of%20AI_Generative%20AI%20Delivering%20New%20Business%20Value%20and%20Increasing%20ROI.pdf)

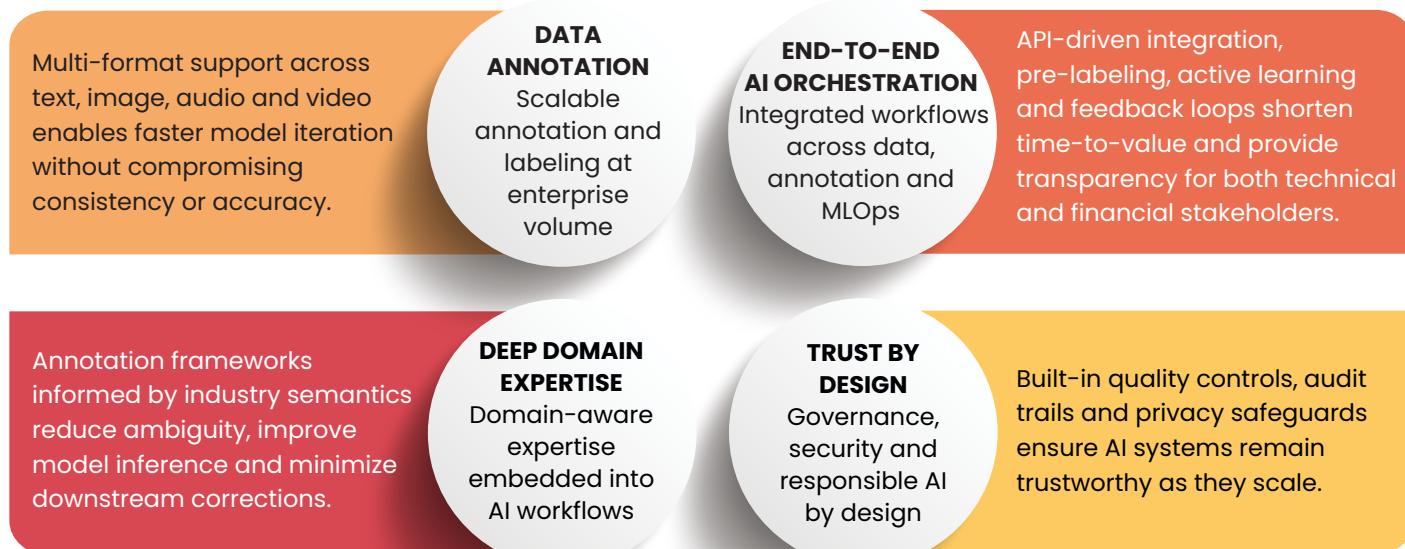


Moving Beyond One-off Builds to Scalable AI Enablement

In the [technology and professional services sector](#), running AI consistently across the enterprise introduces a different set of demands than building individual use cases. What works well in controlled environments must be supported across multiple systems, teams and geographies, often with varying data definitions and review practices. As these differences surface, maintaining consistency in data preparation, [annotation](#) and human oversight becomes central to sustaining performance and enabling reuse.

To address this, many organizations are shifting from project-centric delivery to an operating model that leverages AI-as-a-Service (AlaaS). Rather than re-building data pipelines, annotation processes and validation mechanisms for each initiative, AlaaS standardizes how AI is developed, trained and supported. This approach allows AI capabilities to scale more predictably, while giving organizations a structured way to manage cost, quality and governance as intelligent technologies become embedded across business workflows.

Figure 1: Core Elements of the AlaaS Operating Model



Data Annotation as the Learning Backbone of AlaaS

In an AlaaS model, data annotation becomes the mechanism that allows AI systems to remain reliable once they move beyond controlled pilots into enterprise production. As AI systems are consumed across multiple use cases and exposed to evolving data, scalable annotation ensures models can be refined, corrected and adapted without disrupting operations. This makes annotation central to how AlaaS maintains accuracy and relevance as enterprises scale.

STRONGER DIFFERENTIATION

in AI-led services, as re-usable annotation standards and learning loops accelerate deployment of new AI use cases across clients.



IMPROVED DECISION RELIABILITY as annotated data anchors AI outputs to validated ground truth, reducing ambiguity in knowledge-intensive and client-facing processes.



For firms operating in complex, information-rich environments, this creates tangible impact in

HIGHER OPERATIONAL PRODUCTIVITY

as automation is able to scale without requiring proportional increases in manual review or re-work.



MORE CONSISTENT CLIENT EXPERIENCES

as models trained on domain-aware annotations adapt better to context, language and process variation across industries.



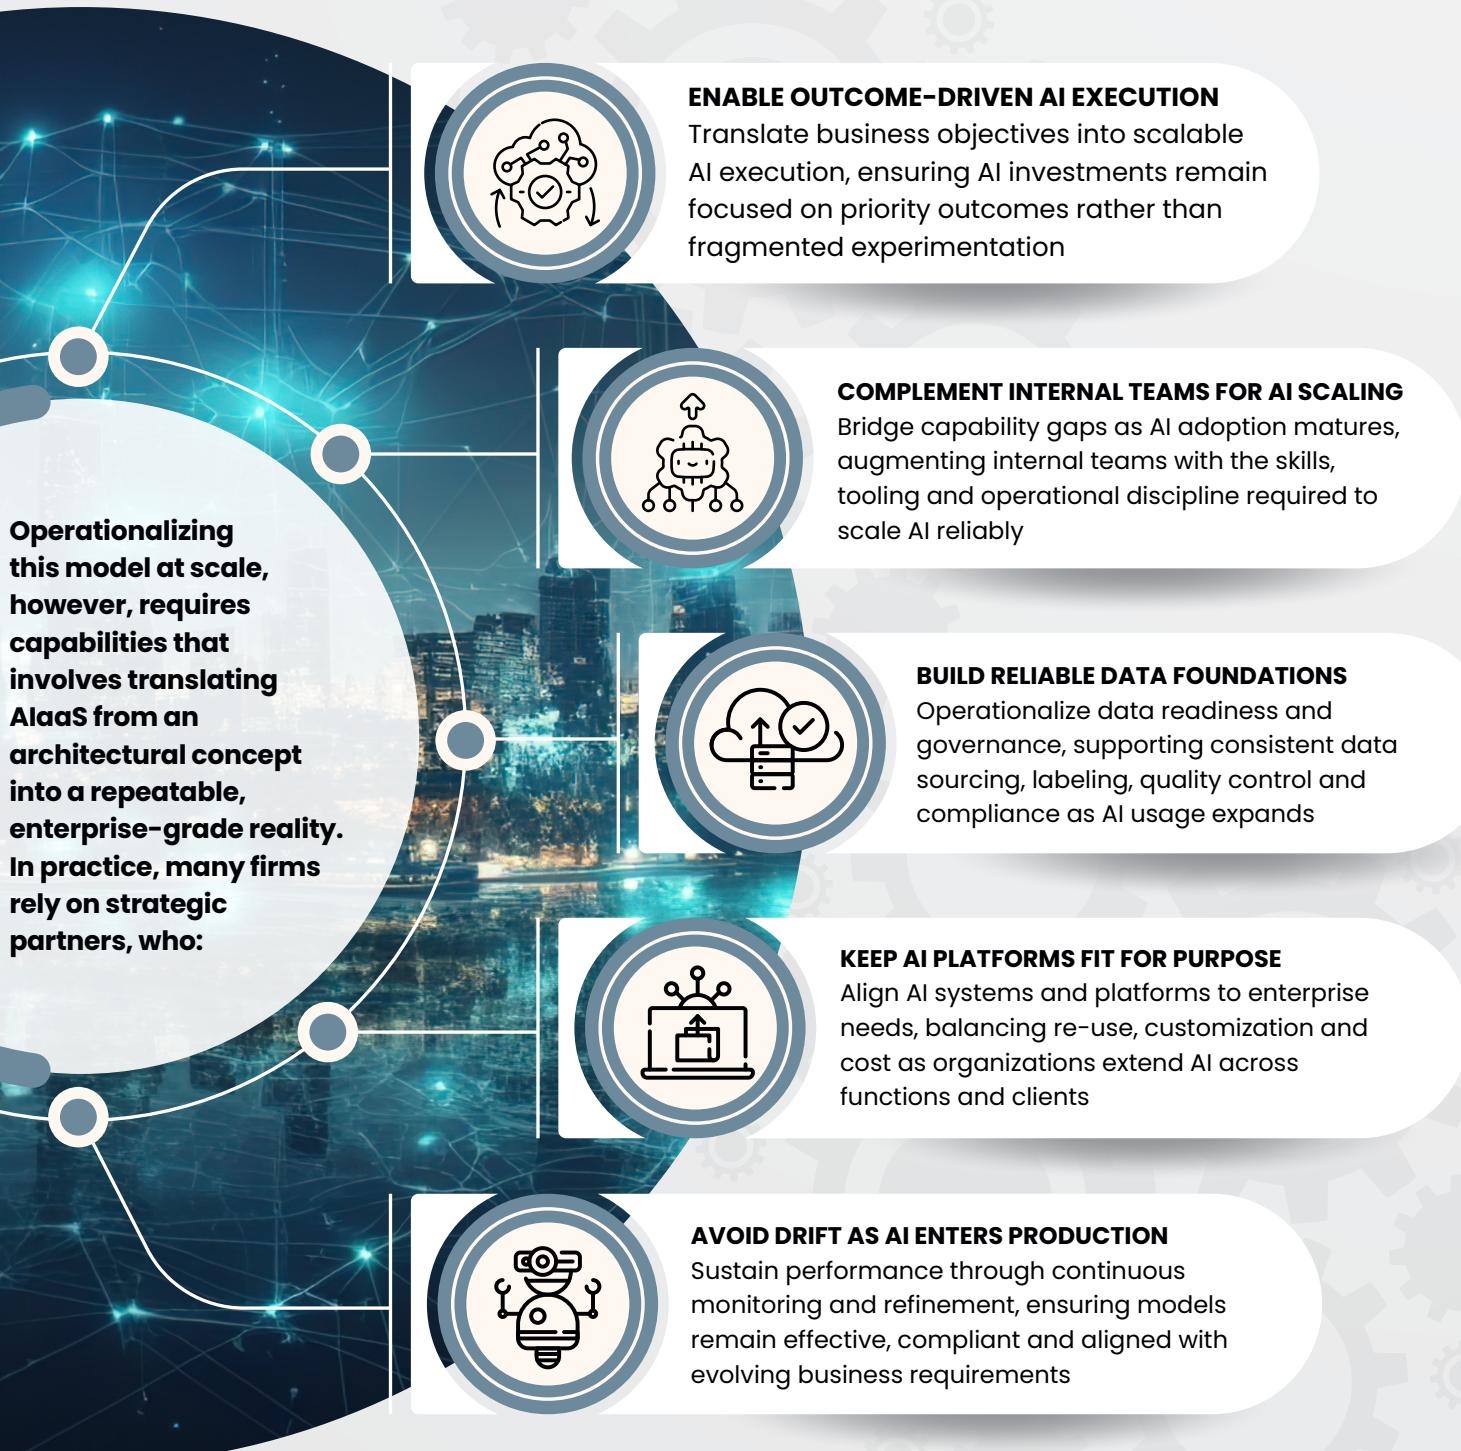
Human-in-the-Loop (HITL): A Designed Control and Improvement Layer

AI and machine learning can automate and accelerate many decision-driven processes. However, for AlaaS to operate reliably at enterprise scale, human expertise must be embedded as a structured operating layer. The governance, judgment, output validation and feedback that the HITL approach provides becomes a direct input for model improvement, enabling reinforcement learning and continuous optimization.



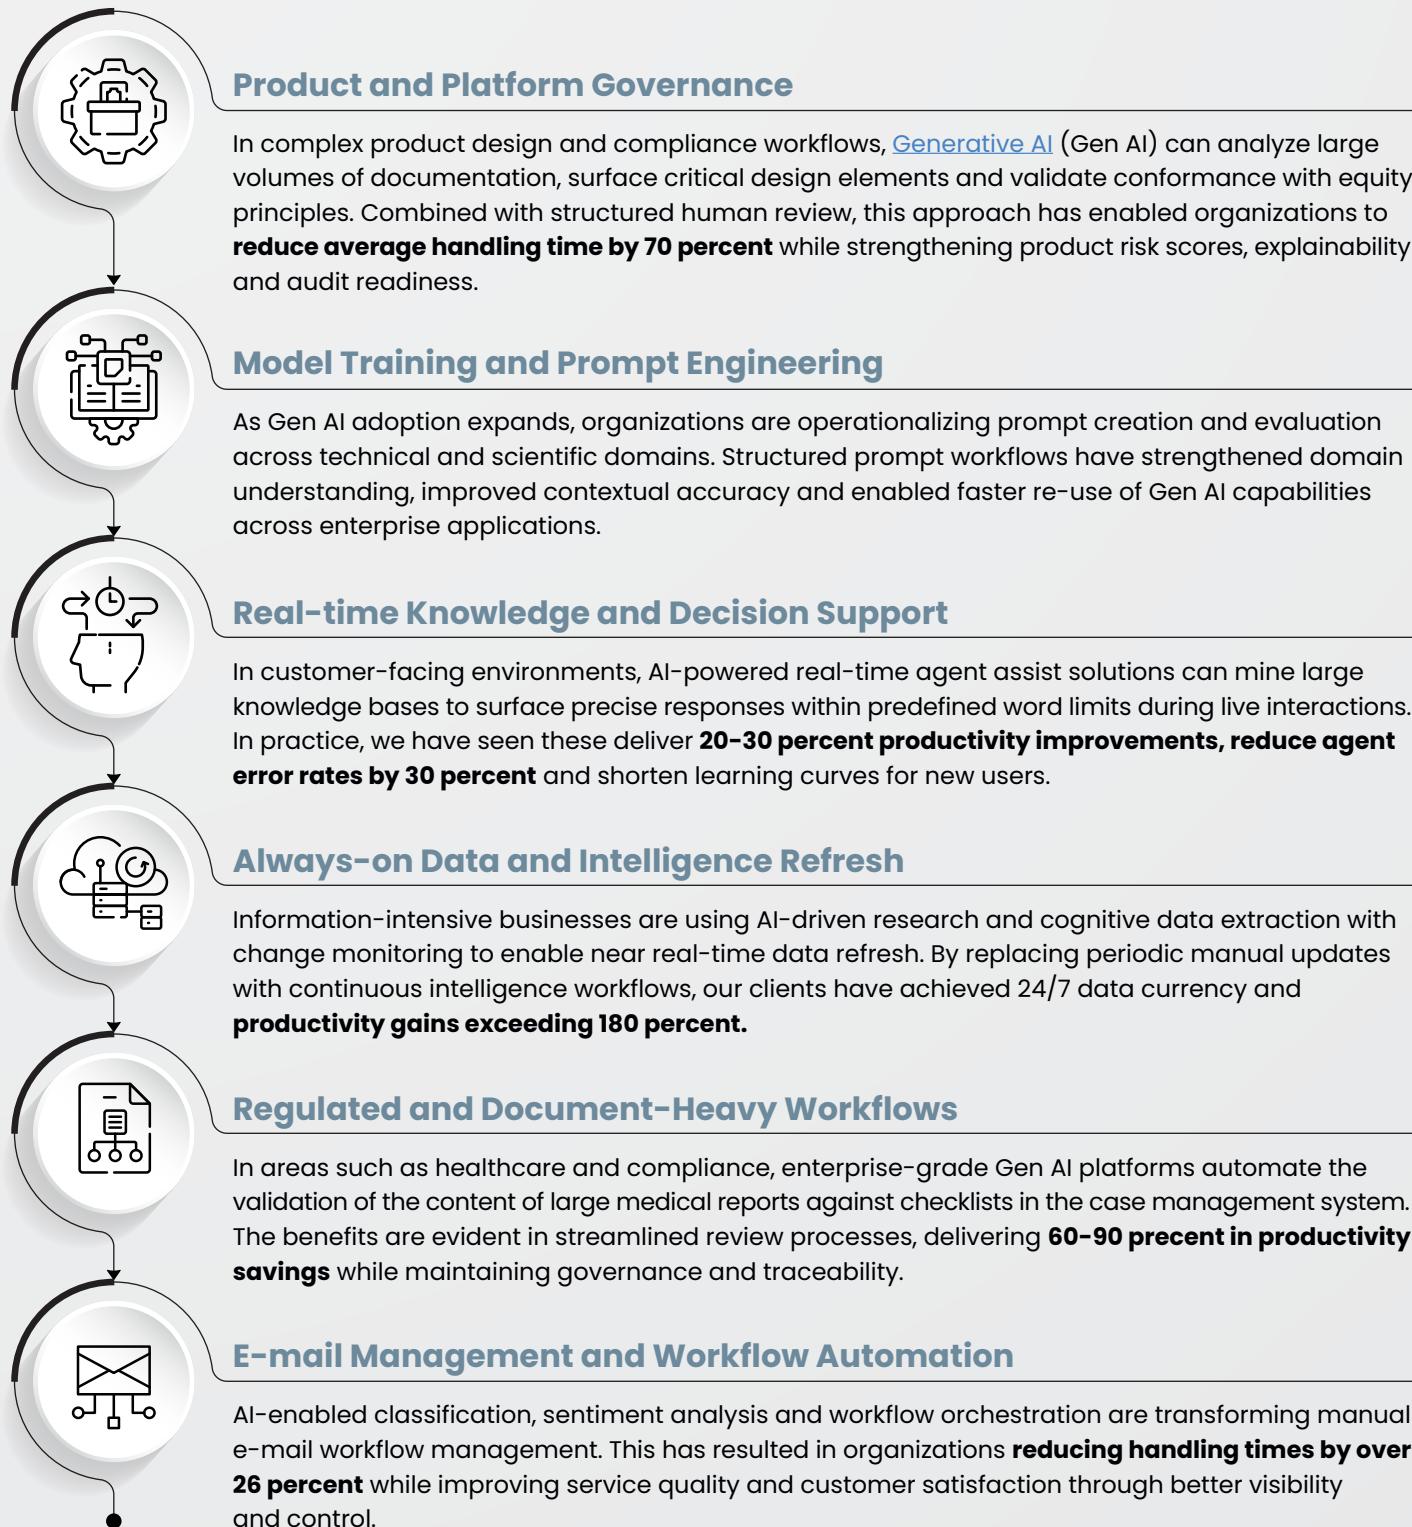
Closing the Loop Between Operations, Learning and Governance

When AlaaS, data annotation and HITL are orchestrated into a single operating system, organizations establish a closed-loop model for AI execution. Data flows into annotation, human validation informs feedback, models improve iteratively and governance is enforced through traceable workflows and quality controls. This integrated approach allows AI to move beyond experimentation and deliver sustained business value with speed, reliability and confidence.



Real-world Execution: AI Operating Models Across High-impact Workflows

When the operating layer is in place, AI becomes an engine for measurable productivity. The following use cases from WNS illustrate how treating AI as a continuously managed service, rather than a one-off deployment, changes the ROI equation.



AI at Scale Demands Ownership

As [AI adoption](#) accelerates, leaders across the AI, technology, finance and operations domains face a defining choice. Some will consume intelligence as a service, others will shape it for their domains and a few will take responsibility for building and sustaining it as a core capability. In each case, long-term success depends less on deploying individual solutions and more on establishing operating models that treat intelligence as a managed asset.

Organizations that invest early in scalable data and annotation capabilities will be better positioned to move faster, re-use AI across use cases and adapt as requirements evolve—without sacrificing governance or trust. Those that do not will continue to scale experimentation, not outcomes. The future of AI leadership will belong to enterprises that are prepared not just to deploy intelligence, but to own its behavior, performance and impact at scale.

[Talk to our experts](#) to explore how your organization can translate AI ambition into operational efficiency and long-term business value.



About WNS

WNS, part of Capgemini, is an Agentic AI-powered intelligent operations and transformation company. We combine deep domain expertise with talent, technology, and AI to co-create innovative solutions for over 700 clients across various industries. WNS delivers an entire spectrum of solutions, including industry-specific offerings, customer experience services, finance and accounting, human resources, procurement, and research and analytics to re-imagine the digital future of businesses. WNS has 66,085 professionals across 65 delivery centers worldwide, including facilities in Canada, China, Costa Rica, India, Malaysia, the Philippines, Poland, Romania, South Africa, Sri Lanka, Turkey, the United Kingdom, and the United States.

To know more, write to us at marketing@wns.com or visit us at www.wns.com

Copyright © 2026 WNS. All rights reserved.