

2026 is the
year *AI moves:*

from **islands of value** to
compounding advantage.

Unlocking business value with agentic AI

Capgemini 





Introduction

For years, enterprises have talked about AI's potential. In 2026, we are finally seeing what it looks like when that potential becomes real: systemic, operational, and truly transformative. The conversation has shifted from *"What can AI do?"* to *"How do we industrialize it responsibly, at pace, and across the entire enterprise?"*

This acceleration is being spurred on by several exponential trends:

Increasing economic viability: model capability continues to double every six months, while at the same time inference costs are collapsing 10x every nine months

- ▶ Increasingly autonomous agentic systems: shifting from co-pilot advisors on the side to digital workers operating deeply in our businesses
- ▶ Mass societal adoption: shifting expectations from consumers on the nature of engagement and transactions

- ▶ AI-robotic flywheels: smarter AI models create more capable robots; more robots generate richer real-world data, powering smarter models for more advanced robots

Fundamentally, AI is now an ecosystem, not a single tool. It is embedded in decisions, workflows, customer interactions, and increasingly in the physical world. And that shift creates a central tension that every enterprise must now resolve:

- ▶ The speed of value creation; enabling fast take up across the business to build competitive advantage

- ▶ The need for uniform control, trust, and safety: to ensure uniform high quality and avoid the proliferation of cost, tooling, and security risk

Winning in 2026 isn't about choosing one over the other. It's about building the architecture to do both simultaneously.

AI isn't another digital wave. It is a fundamental redesign of how work happens, how decisions are made, and how value is created.

AI's industrial moment has arrived

Enterprise AI has evolved rapidly:

2024: Broad and shallow

hundreds of pilots,
limited impact

2025: Narrow and deep

prioritized use cases show
measurable ROI

2026: Broad and deep

AI becomes repeatable, platformed, and
integrated across business systems

Success rates have risen from **5% to 14%**¹, with enterprises seeing **1.7x ROI**¹ on average on first use-case deployments, with compounding ROI on subsequent use cases. This is the year organizations separate experimentation from execution.

But execution in 2026 looks different from what most organizations planned for. Agents are now proliferating across the enterprise in three forms: custom-built agents, co-pilots embedded in existing workflows, and agents native to SaaS platforms. This isn't centralized deployment; it is organic spread across every team, process, and application.

That creates two imperatives: repeatable deployment patterns to scale agents into production consistently, and a control plane to coordinate, observe, and govern agent activity in accordance with enterprise policy, security, compliance, and cost.

¹ Boston Consulting Group (BCG), *From Potential to Profit with GenAI*, January 2024.

From advice to action

For the past three years, most enterprise AI has functioned as a sophisticated advisor: surfacing insights and recommending next steps.

Agentic AI changes fundamentally, moving from productivity tool to enterprise operator. These systems plan, decide, and act: executing across multiple systems, evaluating their own outputs, and looping back to improve. The human role shifts from executor to validator and director.

This is the mechanism that turns pockets of value into compounding advantage: each deployment generates learning that improves the next cycle, increasing competitive differentiation with every iteration.

But just as the opportunity has changed shape, the risk has too. An AI that recommends wrongly is correctable. An AI that acts wrongly propagates errors at speed. This is why repeatability and control are not constraints on agentic capability; they enable sustained value at scale.

This is where productivity gains shift from incremental to exponential.



Why many enterprises still struggle to scale

Despite enormous momentum, scaling AI requires whole-system transformation. The most common barriers are:



1. Rearchitecting your business for AI

Most enterprises are deploying AI into processes designed for humans, automating existing workflows rather than redesigning them. The result is incremental efficiency rather than transformational value. Organizations that automate broken or human-centric processes simply embed their inefficiencies faster and at greater scale.

Rearchitecting for AI means redesigning how workflows between humans and agents, rewriting SOPs to include AI decision points and human validation steps and being deliberate about where human judgement remains essential and where agents can operate with defined autonomy.

New thinking: Leading enterprises are approaching this as AI-native process design: starting from the desired outcome and designing

the human-agent workflow from scratch. They are building agent-ready operating models with clear lines of accountability, structured handoffs, and escalation paths that keep humans meaningfully in control without placing them in the critical path of every decision.

2. Talent and culture

AI maturity is no longer about hiring specialized teams; it's about creating enterprise-wide fluency. Employees across functions must understand how to work with AI agents, orchestrate workflows, validate outputs, and escalate exceptions. This requires a cultural shift from technology-led change to behavior-led transformation.

New thinking: Leading enterprises now treat “AI fluency” as a core competency, similar to digital literacy a decade ago. They are deploying AI companions for every employee: not just as tools, but as coaches. These systems teach workers in context, provide real-time guidance, and help them adopt new workflows. AI becomes not only a productivity driver but a cultural accelerator.





3. Data readiness

Most enterprises underestimate how deeply data fragmentation slows AI maturity. When data lives in silos, different systems, inconsistent formats, and varied governance standards, it creates fragmented intelligence that limits every downstream AI capability. Fewer than 20% of organizations report high data maturity, and even fewer treat data as a continuously evolving product rather than a static asset.

New thinking: The most forward-thinking enterprises are pursuing three data shifts simultaneously. Rather than waiting for perfect data,

they are using AI to make data ready in real time through active metadata ecosystems that continuously classify, clean, and improve pipelines. They are also rebuilding data as AI-ready products: structured, documented, and governed specifically for machine consumption, not just human use. And they are building a semantic layer above it all, enabling agents to search, discover, and access the right data in context, with role-based permissions enforced automatically. The result is a data estate that doesn't just support AI; it actively enables it to scale.

4. Trust, safety, and governance

Only 27%² of organizations fully trust autonomous AI decisions, a signal that trust is not a by-product but an engineered outcome. Traditional governance models, built for static software, break down when systems learn, evolve, and interact with other agents. Explainability, traceability, safety by design, and ongoing human oversight are not optional; they are foundational to responsible scaling.

New thinking: The emerging pattern is “governance as code”: controls that are automated, embedded at the point of action, and enforced consistently across every agent regardless of its origin or form. Rather than humans retroactively reviewing AI behavior, governance becomes a continuous, real-time layer built into the same pipeline agents operate in. Trust is created not through manual oversight, but through verifiable, auditable action at scale.

² Capgemini, *Rise of agentic AI: How trust is the key to human-AI collaboration*, February 2026. https://www.capgemini.com/wp-content/uploads/2026/02/AI-Agents_web_160226-1.pdf





5. Regulation and sovereignty

Increasing regulation and geopolitics now demand that organizations operate predictable, transparent, and governed AI within their control. As AI systems become agentic and autonomous, regulators expect enterprises to demonstrate not only compliance at launch but ongoing assurance throughout the model lifecycle. Organizations also need to design and build for sovereignty, ensuring they retain control and continuity in the AI they operate.

New thinking: Leading organizations are treating both compliance and sovereignty as architectural principles

rather than constraints. They are building compliance-ready pipelines that automatically track lineage, log decisions, classify risks, and produce audit evidence as a by-product of normal operation. On sovereignty, the emerging response is designing AI systems to be portable, privately deployable, and vendor-independent from the outset, rather than retrofitting controls later. Together, these shifts turn regulatory and sovereignty pressures into a foundation for more resilient, trusted AI at scale.

6. AI security and adversarial risk

AI introduces a security challenge that is both new in kind and accelerating in scale. As a target, AI systems face threats traditional security models were never designed to address: prompt injection, data poisoning, and agentic attack surfaces where an agent with system access becomes a high-value vulnerability. Third-party models and SaaS-native agents introduce supply chain risk that is difficult to assess and monitor. But AI is also making every existing threat more dangerous. Adversaries are using increasingly powerful models to discover and exploit vulnerabilities faster than human teams can

respond, generate social engineering at scale, and automate exploit development that previously required significant expertise.

New thinking: Security needs to be embedded at the model and agent level, not just the perimeter. Adversarial testing and AI-specific red teaming should be standard before any agent reaches production. Equally, enterprises should be deploying AI defensively: using models to accelerate threat detection and automate response against adversaries who are already using the same tools.





7. Total cost of ownership

Enterprises routinely underestimate the true cost of AI at scale. Inference costs grow rapidly as agent usage expands, and the economics of a pilot often break at enterprise deployment. The full cost picture includes fine-tuning, evaluation, monitoring, human oversight, and integration – easy to overlook when a use case looks compelling in isolation. Treating model selection as a capability decision rather than an economic one compounds the problem; deploying frontier models uniformly is a fast route to unsustainable AI economics.

New thinking: Enterprises managing AI cost most effectively apply FinOps discipline from the start: estimating token costs at design time, building model-tiering strategies that match capability to task, and treating cost as a first-class architectural constraint. Internal cost attribution for AI usage is emerging as a best practice, making consumption visible and driving accountability at the team and use-case level.



The new enterprise playbook: firm vision, focused progress, fast, flexible scale

Enterprises achieving real, repeatable AI value are not succeeding by chance. They follow a disciplined, pattern-based approach that blends ambition with pragmatism. The fastest accelerators anchor transformation in bold vision, validate it through targeted execution, and institutionalize it through platforms and reusable capabilities.

1. Firm vision

Winning enterprises begin by reimagining the business end to end – from R&D and manufacturing to supply chain, sales, service, and corporate functions. Instead of seeking isolated efficiencies, they ask: what would our business look like if AI were embedded in every decision, workflow, and customer interaction?

New thinking: The most successful organizations treat AI not as an efficiency lever but as a strategic redesign of operating models. They are shifting from process-centric to agent-centric value chains, where autonomous systems orchestrate work across human, digital, and physical environments.






2. Focused progress

Rather than modernizing everything at once, leading organizations start with targeted, high-impact use cases that demonstrate immediate value. These are not science experiments; they are strategically selected force multipliers that prove feasibility, demonstrate ROI, and build confidence.

New thinking: The emerging pattern is microtransformations: small, modular workflows where AI delivers disproportionate leverage. These wins provide templates that scale across functions, markets, and geographies.



3. Fast, flexible scale

Once early value is proven, focus shifts to scaling with speed and discipline. Organizations build minimal viable data foundations, shared platforms, and reusable components that allow teams to stand up new AI capabilities rapidly and consistently. Each success accelerates the next, reducing marginal cost and increasing marginal impact.

New thinking: Enterprises are moving toward platformized scaling, where agents, APIs, governance controls, and integration patterns become reusable building blocks. The most resilient go further, designing for modularity and interoperability from the outset. In a world where model capabilities double every six months, architectural flexibility is not optional.



4. The flywheel effect

Done well, this creates a self-reinforcing flywheel: early wins build confidence; confidence unlocks investment; investment accelerates capability; capability drives adoption. Over time, AI becomes a core organizational muscle – instinctive, repeatable, and responsible.

New thinking: The flywheel strengthens when AI is embedded into change management itself – used to teach, onboard, and operationalize AI. This is where organizations shift from “doing AI” to becoming AI-native.

The path forward: *system-wide change*

Scaling AI is not a technology initiative; it is organizational redesign across leadership, teams, engineering, operations, and learning.

This is not about doing AI projects. It is about building an AI-enabled enterprise. And the winners are acting now.



Final POV: *the advantage belongs to the early industrializers*

The question is no longer, “Can AI transform the enterprise?”

It is, “Do we have the architecture, culture, and governance to scale it responsibly – faster than our competitors?”

Because in 2026, the advantage does not belong to early adopters; it belongs to early industrializers: those who can harness AI as a governed, trusted, organization-wide ecosystem with agents at the core.

Author



Craig Suckling

Global Head of Agentic Scale

Let's continue the conversation, reach out to me at raise@capgemini.com

Sources

¹ Boston Consulting Group (BCG), *From Potential to Profit with GenAI*, January 2024.

^{2 & 3} Capgemini, *Rise of agentic AI: How trust is the key to human-AI collaboration*, February 2026.

About Capgemini

Capgemini is an AI-powered global business and technology transformation partner, delivering tangible business value. We imagine the future of organizations and make it real with AI, technology and people. With our strong heritage of nearly 60 years, we are a responsible and diverse group of over 420,000 team members in more than 50 countries. We deliver end-to-end services and solutions with our deep industry expertise and strong partner ecosystem, leveraging our capabilities across strategy, technology, design, engineering and business operations. The Group reported 2025 global revenues of €22.5 billion.

Make it real | www.capgemini.com

