

## AI AGENTS

# How Agents Are Accelerating the Next Wave of AI Value Creation

By [Vinciane Beauchene](#), [Nicolas de Bellefonds](#), [Djon Kleine](#), [Rich Lesser](#), [Amanda Luther](#), [Tom Martin](#), and [Daniel Sack](#)

**ARTICLE** DECEMBER 15, 2025 8 MIN READ

AI is transforming the way work gets done, but it hasn't fundamentally altered the way most companies operate—yet.


That's about to change.

We are entering the era of AI agents. While previous waves of AI focused either on logic and optimization or creativity and synthesis—the “left brain” and “right brain” of predictive AI and generative AI, respectively—the next wave will involve agentic systems that take on entire workflows, applying judgment shaped by a company’s own institutional knowledge. (See the sidebar “The Frontal Cortex of AI.”)

## — The Frontal Cortex of AI

If predictive AI is the left brain for logic and optimization, and generative AI the right brain for creativity and synthesis, then agentic AI serves as the executive function that turns creative probability into business impact.

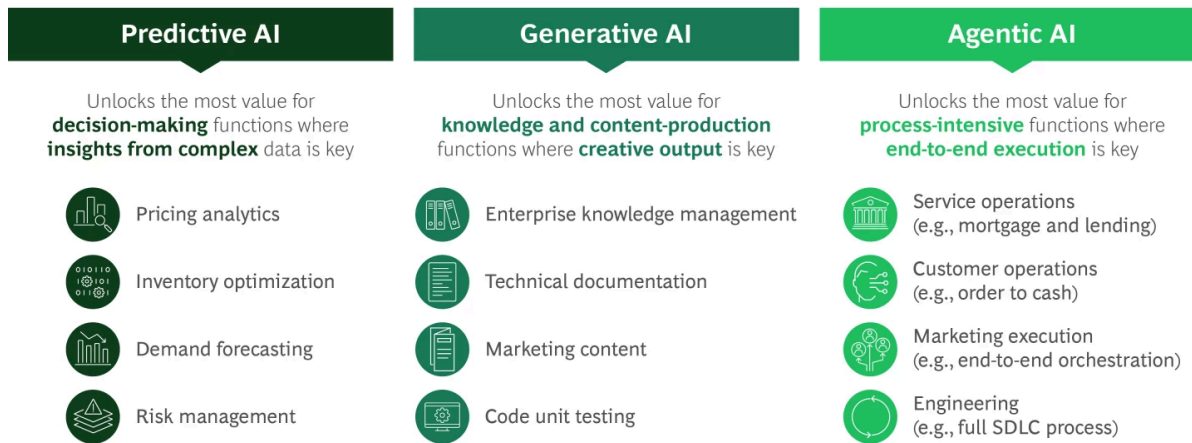
Agentic AI is the Executive Function that Connects Predictive and Generative AI



Source: BCG analysis.

Predictive AI opened value in decision-making functions. Generative AI opened value in knowledge and content production functions. Agentic AI is now opening value in process-heavy functions where execution defines performance.

# Agentic AI Unlocks Value in Process-Heavy Functions



Source: BCG analysis.

Note: SDLC = software development life cycle.

These systems do more than analyze and create. They observe, plan, and act on their own. They learn continuously. And when embedded in the business, they take on the steps that slow organizations down, whether coordinating internal workflows or engaging directly with customers and partners.

Early adopters are already showing what this looks like in practice:

- A shipbuilder cut design effort for a key component by 45% and lead time by 80% by using orchestrator and worker agents to run a multistep design process across more than 14 systems.
- A lender reduced mortgage approval time by 20% and freed half of its mortgage-assessor capacity with an agent system that handles flagged applications end to end.
- A payroll provider resolved anomalies automatically through a supervisor agent supported by specialized worker agents, improving processing speed by more than 50%.

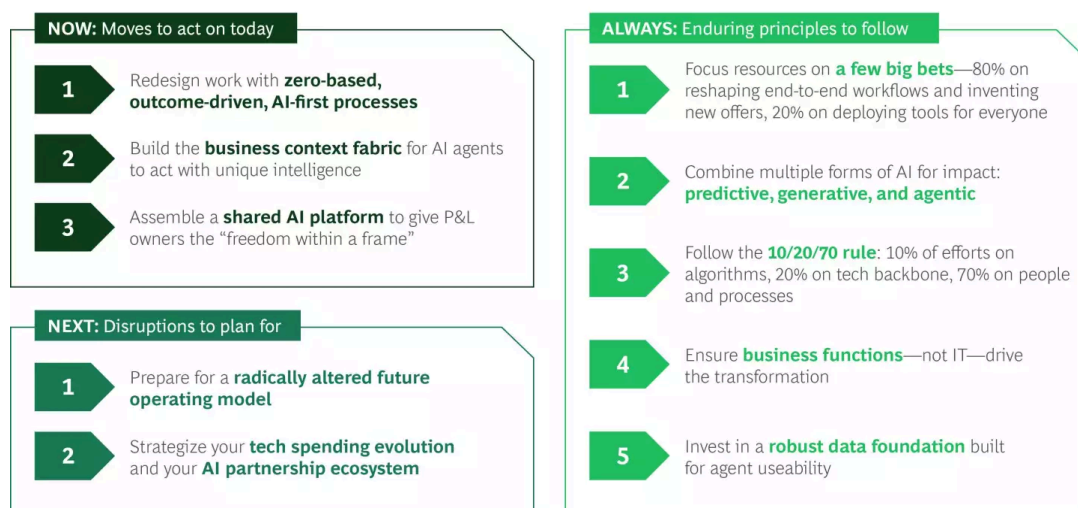
Results like these are spurring rapid investment across industries. In a recent BCG study with MIT Sloan Management Review, 35% of organizations said they are already using agentic AI, and another 44% said they plan to do so soon. That same research found that three-quarters of extensive adopters believe AI is now enabling new sources of value and competitive advantage.

Productivity gains are often the initial benefit of AI adoption, but the real prize in the agentic age is differentiation. And the real challenge is sustaining advantage when AI can enable new, disruptive

competitors and shrink traditional advantages.

The first step in the agentic AI journey is to narrow the agenda to the priorities that create advantage and value. Drawing on BCG's experience with early adopters, we've laid out a roadmap for what CEOs should do now, what they should anticipate next, and the principles they should always follow to achieve the highest returns from AI.

## Now, Next, Always: The CEO's Roadmap for Agentic AI



Source: BCG analysis.

## Now: Three Moves to Make Today

Leaders should use the next year to build their agentic foundation. This will not be completed in 12 months, of course. But making these moves now will help build the capabilities and confidence to scale broadly in the years ahead.

**Redesign work around zero-based, outcome-driven processes.** A common mistake is automating what already exists. Real value comes from a “zero-based” approach, starting with the outcome you want and reinventing how to deliver it. Rethink how tasks flow, where decisions sit, and how roles change as agents take on more execution. Spell out the skills, responsibilities, and the shape of the team you'll need once the system is up and running.

One global consumer goods company did this with its product innovation process. Instead of speeding up individual steps, they started from scratch and rebuilt the workflow around meta-

agents that orchestrated the work and worker agents that generated insights, tested ideas, and created concepts. The unified system cut cycle time by roughly 60% and improved product fit.

**Embed your company's context and proprietary intelligence.** Context drives individual behavior, which shapes overall performance in human-centric organizations. The same is true for organizations leveraging AI agents. They need the same rich context your best workers are guided by, from data and tools to norms, tacit know-how, constraints, and intent. For AI agents to move from generic outputs to distinctive institutional intelligence, you must build a business context fabric. It consists of three components:

- **Objectives:** Define the outcomes (not outputs) the agent must deliver, such as higher conversion or faster compliance.
- **Resources:** Connect the agent to the data, tools, and decision logic your top teams use.
- **Constraints:** Set the guardrails that keep autonomous actions safe and consistent, including regulatory rules, brand tone, and risk appetite.

One global consumer goods company identified its business context fabric by holding expert workshops that helped surface the unwritten principles its creative team followed in developing strong campaigns. They translated that tacit knowledge into the prompts and decision trees their agents follow. The result was brand-consistent recommendations with far less rework.

**Assemble a shared AI platform.** To scale, you cannot have every function building its own infrastructure. You need a shared AI platform that gives P&L owners “freedom within a frame.” Think of it as an enterprise system that everyone builds on.

The platform should include data connectors, orchestration frameworks, model management, monitoring, and tiered autonomy with responsible AI controls. These shared components let teams innovate faster without rebuilding the basics or compromising security.

In this new model, IT and the business functions operate as partners:

- **IT maintains the chassis:** They manage the data, infrastructure, and safeguards that ensure systems and AI agents are reliable, compliant, consistent, and up to date.
- **The business drives the design:** Functional leaders and teams use the platform to design agents that deliver outcomes, embedding tech-fluent talent that owns the new agentic-enabled processes as products and enhances them over time.

The impact is significant. In our work with clients, we've observed that shared tools and data can cut costs by up to 30% and lift productivity by 25%. As reuse increases, time to market can improve by 50%.

# Next: Two Disruptions to Plan For

As agents take on execution, CEOs must prepare for two major disruptions.

**A Radically Altered Enterprise Operating Model.** Our research with MIT reveals that as AI agents handle routine execution, the traditional operating model will be radically altered. Among the changes:

- **The workforce will be rebalanced.** Junior talent will spend less time producing “first draft” work and more time directing it. Among heavy adopters, 29% of companies say they’ll need fewer traditional entry-level roles, and 43% expect greater demand for generalists who can manage human-agent teams and hybrid workflows. Early evidence of workforce disruption is already visible—Stanford University research on payroll data found a 16% decline in employment among early-career workers (ages 22 to 25) in AI-exposed roles.<sup>1</sup> The balance across functions and roles also will shift: some workflows will be fully automated, others will stay human-led, and new roles will appear as older ones fade. This makes skills-based strategic workforce planning essential.
- **The organization design will shift.** As humans move away from execution and toward outcome-driven, cross-functional teams, 45% of AI leaders expect to need fewer middle-management layers. Work is moving to diverse teams that combine functional experts, tech and data-science know-how, and relevant legal and human-resource talent. Their work centers on oversight, orchestration, and continual improvement of the agent-driven processes rather than manual delivery.
- **Governance will evolve.** Right now, the vast majority of companies are building most processes with humans in the loop and with ultimate decision authority. But as AI takes on more decision rights, 58% of heavy adopters expect a fundamental shift in governance over the next three years, and one-third believe AI will have more decision-making authority in the same period. These users are drawing clearer lines between human and AI roles. They’re also building cross-functional governance models to monitor and coordinate autonomous activity, develop and enforce standards, and measure outcomes and risks across systems. “Freedom within a frame” must include careful risk management that draws on core responsible AI principles—clear risk taxonomies, rigorous testing, continuous monitoring, and built-in resilience from the start to ensure agents act within boundaries. (See the sidebar “The Trust Protocol: A Framework for Graduated Autonomy.”)

---

## — The Trust Protocol: A Framework for Graduated



# Autonomy

The greatest barrier to scaling agentic AI isn't technology, it's trust. CEOs hesitate for good reason. No one wants a system that hallucinates a discount, offends a client, or misallocates capital.

To move from pilots to scale, leaders should adopt a graduated autonomy framework. Think of it as a promotion path in which agents earn their autonomy through proven performance thresholds and rigorous testing. This framework has four tiers:

- **Tier 1: Shadow Mode (Agent-Assisted).** In this mode, the agent suggests and the human acts. The agent operates as a “digital whisperer,” observing data and drafting options or identifying missing information, much like we use LLMs such as ChatGPT today. The human retains full control and manually executes the final step. In terms of strategic value, this is the training ground. It allows the organization to validate the agent’s logic against the business context without operational risk.
- **Tier 2: Supervised Autonomy (Human in the Loop).** In this tier, the agent acts and the human approves. The agent prepares the action and stages it. Execution is paused until a human provides a digital “thumbs up.” If the agent’s confidence score drops below a set threshold (such as 90%), it reverts to Tier 1. This approach is for high-stakes decisions where accuracy is paramount, such as financial auditing or finalizing complex business proposals.
- **Tier 3: Guided Autonomy (Human on the Loop).** In this tier, the agent acts and the human monitors. The agent executes autonomously within strict guardrails. Humans shift from approvers to exception handlers, intervening only when the system flags an anomaly or during random spot-checks. This approach unlocks speed for routine tasks (such as supply chain restocking) while maintaining a safety net for edge cases.
- **Tier 4: Full Autonomy (Human out of the Loop):** In this tier, the agent operates independently without explicit human oversight for defined workflows. This tier is reserved for highly mature, low-risk environments where the cost of error is negligible compared to the gain in efficiency.

Successful implementations require agents to run in Tier 1 (Shadow Mode) until they prove they align with your organization’s risk appetite. Autonomy is a journey of trust, quantified by accuracy.

**A Recalibrated Tech Strategy.** As labor costs rise and digital costs fall, companies are shifting their spending from people to technology. In that context, higher technology spending (for example, 35% to 45% of operating cost in retail banking, up from 20% to 30% today) reflects a rebalanced cost mix, not a budget hike. This shift makes tech strategy inseparable from competitive strategy.

As the technology ecosystem consolidates around a few major platforms and a long tail of specialists, leaders need to keep a steady hand on the increasingly complex partnership environment. Cloud and model providers are existential dependencies that must be treated as strategic alliances. Leaders also will need to keep systems modular to avoid lock-in and protect the proprietary value of their data.

## Always: Five Enduring Principles

Technology changes, but the fundamentals of AI transformation remain constant. Here are some principles CEOs should follow now and in the years to come.

**Focus on a few “reshape and invent” big bets.** Don't spread AI across the organization to achieve a myriad of small productivity gains. To create real value and distinctive advantage, concentrate resources on reshaping end-to-end workflows and inventing new, revenue-generating business models. Get the early wins from these big bets in critical areas to build the confidence and muscle to provide the foundation to expand.

**Combine multiple forms of AI and technology for impact.** The most powerful systems combine predictive AI to optimize decisions, generative AI to create content, and now agentic AI to orchestrate execution. To build them, leverage what already exists, integrating current systems and data rather than treating them as silos.

**Follow the 10/20/70 rule.** Algorithms account for 10% of the work in AI transformation, the tech backbone 20%. The remaining 70% comes from people and processes. AI agents fundamentally change how work gets done and by whom; as such, the lion's share of the leadership team's effort must go into redesigning roles, managing change, and training your workforce to provide the right oversight and guidance to AI agents—and to reinforce expertise where humans make a difference.



**Make the transformation business-led.** The IT function is critical to building the infrastructure, but they cannot design the work on your behalf. Business leaders understand the nuance of the workflows, where judgment matters, and where AI agents can truly move the needle. They must lead the transformation, in close partnership with IT, to ensure agents are solving the right business problems and being given the right context, data, and constraints to ensure high quality and responsible outcomes.

**Invest in a robust data foundation.** AI agents rely on high-quality, connected data to reason effectively. Don't wait for a perfect data lake. Invest incrementally and use each new agent implementation to expose and close specific data gaps, so that the core foundation gets stronger with every step.

## The CEO's Next Step

AI agents mark a turning point in how organizations create value and compete. With AI-native firms already achieving 25 to 35 times more revenue per employee compared to traditional peers, the potential is clear.

The models are ready. The challenge is everything around them—the systems, data, and workflows they depend on and the organizational capabilities needed to use them at scale. As basic productivity becomes universal, advantage will shift to those that master integration and use agents not only to drive efficiency, but to unlock proprietary intelligence and create new revenue streams.

In this landscape, strategy takes on greater weight. CEOs need a clear sense of where to start, how to secure early wins that strengthen capability and confidence, and how to extend that momentum across the business. The “Now, Next, and Always” framework helps leaders make those choices—while laying the foundation for enduring advantage.

# Authors



Vinciane  
Beauchene

Managing Director & Partner  
Paris



Nicolas de  
Bellefonds

Managing Director & Senior  
Partner  
Paris



Djon Kleine

Managing Director & Partner  
San Francisco - Bay Area



Rich Lesser

Global Chair  
New York



Amanda Luther

Managing Director & Senior  
Partner  
Austin



Tom Martin

Platinion Director  
London



Daniel Sack

Managing Director & Partner  
Stockholm



## ABOUT BOSTON CONSULTING GROUP

Boston Consulting Group partners with leaders in business and society to tackle their most important challenges and capture their greatest opportunities. BCG was the pioneer in business strategy when it was founded in 1963. Today, we work closely with clients to embrace a transformational approach aimed at benefiting all stakeholders—empowering organizations to grow, build sustainable competitive advantage, and drive positive societal impact.

Our diverse, global teams bring deep industry and functional expertise and a range of perspectives that question the status quo and spark change. BCG delivers solutions through leading-edge management consulting, technology and design, and corporate and digital ventures. We work in a uniquely collaborative model across the firm and throughout all levels of the client organization, fueled by the goal of helping our clients thrive and enabling them to make the world a better place.

© Boston Consulting Group 2025. All rights reserved.

For information or permission to reprint, please contact BCG at [permissions@bcg.com](mailto:permissions@bcg.com). To find the latest BCG content and register to receive e-alerts on this topic or others, please visit [bcg.com](https://bcg.com). Follow Boston Consulting Group on [Facebook](#) and [X \(formerly Twitter\)](#).

<sup>1</sup> Erik Brynjolfsson, Bharat Chandar, and Ruyu Chen,

*Canaries in the Coal Mine? Six Facts about the Recent Employment Effects of Artificial Intelligence,*

Stanford University,  
November 13, 2025