


The symbiotic enterprise

How cognitive and physical AI are reinventing enterprise execution

Christian Jansen
Daniele Chiarella
Kim Baroudy
Lari Hämäläinen
Lieven Van der Veken
Ruben Schaubroeck
Sébastien Lacroix
Stéphane Bout
Guillaume Dagorret

June 2026



Contents

Executive summary 2

Chapter 1: AI is no longer just a tool—
it is becoming a workforce 4

Chapter 2: The step change comes not from
augmentation but from reinvention 11

Chapter 3: The symbiotic enterprise: Reinventing
how work is executed 17

Chapter 4: Navigating the symbiotic transformation:
Reinvention without destabilization 25

Conclusion 30

Executive summary

AI is no longer just a tool. It is becoming a workforce. Reasoning models and agentic skills now enable AI agents to execute complex cognitive tasks with limited supervision, while physical AI extends automation into the physical world. Together, these advances make close to 60 percent of work hours theoretically automatable.¹

Yet the promise remains largely unrealized. Despite widespread adoption, very few companies report meaningful P&L impact. In most cases, AI remains embedded within existing workflows, generating only incremental gains.

The real breakthrough comes from reinventing execution, not augmenting it. Early human–AI systems already deliver step change improvements when cognitive and physical workflows are redesigned from first principles.

At scale, this transformation gives rise to a new enterprise model: the symbiotic enterprise, in which humans, AI agents, and intelligent robots each contribute according to their respective strengths within flatter organizations and under a new economic model, with technology becoming a primary cost driver. Beyond productivity, the symbiotic enterprise fundamentally changes the economics of growth by enabling organizations to innovate faster, adapt continuously, unlock new revenue opportunities, and scale through software rather than labor.

Traditional advantages such as expertise, workforce scale, coordination complexity, and market frictions erode, lowering barriers to entry and enabling customer re-insourcing and AI-native competitors to challenge incumbents. As AI capabilities commoditize and productivity gains diffuse across industries, durable competitive advantage shifts away from access to AI itself toward assets and capabilities that compound over time: proprietary intelligence built through data, agent skills, and learning loops; customer and ecosystem control points; and mastery of intelligence architecture and execution at scale. At the same time, a new strategic dependency on AI providers emerges, introducing the risk of a “cognitive tax” on enterprise execution.

Transitioning toward the symbiotic enterprise is neither a technology deployment nor a productivity program; it is a strategic transformation. Leaders must avoid two failure modes. The first is incrementalism—optimizing a pre-AI operating model until AI-native competitors erode its economics. The second is overreach—deploying autonomous systems faster than the organization can absorb, adapt, or govern.

Success will require four conditions:

- ***A bold, value-driven North Star.*** Organizations must define a strategic vision of the target state, grounded in future sources of differentiation, control points, and value creation, to guide investment decisions, workforce transformation, and operating-model redesign.

¹ “Agents, robots, and us: Skill partnerships in the age of AI,” McKinsey Global Institute, November 25, 2025.

- ***A dual transformation journey balancing value realization and workforce adaptation.*** Reinvention should take place at the domain level, while targeted augmentation initiatives continue to enhance individual productivity and accelerate organizational learning. The pace of the journey should balance value realization, workforce adaptation, and technological maturity.
- ***Scalable foundations.*** The symbiotic enterprise requires API-enabled legacy systems, reusable data products, a modular and vendor-agnostic orchestration platform, and new disciplines governing AI reliability, behavioral control, security, and economics.
- ***Extended executive leadership.*** This transformation cannot be delegated. The CEO defines the ambition and arbitrates strategic trade-offs. The chief human resources officer leads workforce transition and reskilling, the chief transformation officer orchestrates enterprise-wide execution, and the chief technology officer builds the technological foundations required to scale intelligent execution. Success ultimately depends on mobilizing the full execution muscle of the enterprise, combining disciplined delivery, organizational change, and continuous adaptation over a multiyear journey.

The winners will not be the organizations that deploy the most AI, but those that reinvent their operating models fastest, build compounding advantages around intelligence, and secure the control points where intelligent execution creates value. Beyond enterprise performance, the broader societal challenge will be ensuring that the gains created by human–AI collaboration translate into expanded opportunity, employability, and inclusive economic growth.

About QuantumBlack, AI by McKinsey

QuantumBlack, McKinsey's AI arm, has been helping businesses create value from AI since 2009, expanding on McKinsey's technology work over the past 30 years.

QuantumBlack combines an industry-leading tech stack with the strength of McKinsey's 7,000 technologists, designers, and product managers serving clients in more

than 50 countries. With innovations fueled by QuantumBlack Labs—its center for R&D and software development—QuantumBlack delivers the organizational rewiring that businesses need to build, adopt, and scale AI capabilities.



**AI is no longer
just a tool—
it is becoming
a workforce**

For decades, software transformed how enterprises operated, but its use largely remained limited to supporting human work rather than executing it autonomously. In parallel, physical automation has delivered massive productivity gains, but only inside tightly controlled environments where variability has been engineered out. As a result, most cognitive and physical activities have remained beyond the reach of large-scale automation.

That boundary is shifting rapidly. Recent advances in reasoning models, agentic systems, and intelligent robotics are transforming AI from a tool that supports work into systems capable of executing tasks with limited supervision. For the first time, both cognitive and physical work are becoming partially automatable at scale, potentially reshaping how enterprises operate, compete, and create value.

Latest reasoning models and new agentic skills enable autonomous execution of complex cognitive tasks

Since the launch of ChatGPT, a sequence of technological advances has progressively transformed gen AI from an assistive tool into a professional-grade autonomous-execution capability. Consider how rapidly technology has developed in just four short years:

- ***The gen AI breakthrough: From data to intelligence.*** The first generation of gen AI applications, which emerged in late 2022, represented a genuine breakthrough. For the first time, AI could produce fluent, contextually relevant text through a natural-language interface, dramatically expanding accessibility and unlocking a wide range of productivity use cases across text generation, summarization, and translation.

However, these early systems remained fundamentally limited. Interactions were stateless, with no memory or continuity across sessions. Systems couldn't plan, execute tasks, or interact with external tools. Their reasoning capabilities struggled with complex, multistep problems, and frequent hallucinations required systematic human validation. Users remained fully responsible for all decisions and outcomes.

At this stage, gen AI could inform but not act.

- ***The agentic shift: From responding to executing.*** In 2024, the emergence of agentic AI marked a decisive inflection point. Gen AI systems evolved from passive responders into active agents capable of executing sequences of tasks, maintaining context over time, and interacting with enterprise systems. This shift enabled the automation of cognitive tasks and unlocked materially higher productivity gains.

Yet these first-generation agents still faced structural limitations. Their reasoning capabilities remained constrained, particularly in complex or ambiguous situations. More critically, they lacked access to enterprise-specific know-how: the decision rules, the internal processes, and the constraints that govern how work is actually performed within organizations. In essence, agents could operate as junior contributors but not at a professional standard.

The missing element was not execution but enterprise-grade judgment and know-how.

- ***The cognitive shift: From execution to professional-grade autonomy.*** Since 2025, two complementary advances have jointly closed the gap: a step change in reasoning capabilities and the emergence of agentic skills. Together, these advances transform generic agents into professional-grade digital workers.

In just a few years, large language models (LLMs) have moved from bottom-decile performance on academic benchmarks to top-decile results across most disciplines, now matching top-tier human performance in standardized exams (Exhibit 1). In practical terms, this means models can now solve complex, multistep problems comparable to those professionals face, ranging from legal case analysis and medical reasoning to advanced coding and quantitative problem-solving. At the same time, reliability has improved significantly. While early models exhibited high hallucination rates, recent systems, especially in grounded or enterprise contexts, have reduced these errors by an order of magnitude.

Crucially, these advances are not only quantitative. Modern models can now decompose complex problems, evaluate alternative courses of action, maintain coherence across long reasoning chains, and operate under explicit constraints. This evolution enables gen AI systems to move beyond isolated tasks toward the autonomous execution of increasingly complex workflows.


But reasoning alone is not sufficient. To operate in enterprise environments, AI agents must apply domain-specific rules, processes, data-access patterns, and constraints—the tacit know-how that defines how work is actually performed. Agentic skills provide this missing execution layer, transforming generic agents into professional agentic operators.

An agentic skill defines not only what to do, but how to do it, ensuring execution is consistent, auditable, and aligned with enterprise standards. Once encoded into skills, expertise can be deployed at scale across hundreds or thousands of agents, effectively multiplying the reach of top performers. Then, through continuous feedback loops, these skills can be refined over time, turning operational knowledge into a dynamic and scalable enterprise asset.

Exhibit 1

The latest large language models match or outpace human performance across a range of capabilities.

Large language model performance relative to human baseline, select AI Index technical performance benchmarks, index (human baseline = 100)



Source: Stanford Institute for Human-Centered Artificial Intelligence (HAI), AI Index Report 2026

McKinsey & Company

Skills in practice: From generic agent to professional-grade execution

Code review in software development. A generic AI agent reviewing source code will flag general issues, such as readability, common antipatterns, and basic security risks. This is useful, but not sufficient to produce production-grade code aligned with enterprise standards. Equipped with a company-specific “code review” skill, the same agent applies the organization’s internal standards, such as proprietary naming conventions, mandated authentication patterns for API calls, data encryption rules for sensitive fields, and minimum test coverage thresholds. The skill encodes these rules, exceptions, and thresholds into a reusable component, so rather than guessing, the agent executes against a defined reference framework.

Claims handling in customer operations. Similarly, a generic agent tasked with handling a customer complaint will be able to draft a polite response. Equipped with a company-specific claims-handling skill, however, the same agent processes the claim end to end: applying the correct escalation path based on claim value, customer tenure, and product category; checking regulatory compliance against jurisdiction-specific rules; triggering the appropriate refund or goodwill gesture within preapproved thresholds; and generating a response using legally validated templates. The skill transforms a generic conversational agent into a fully compliant claims handler, aligned with the organization’s decision framework and audit requirements.

Physical AI, a new generation of intelligent robots, is extending the frontier of physical automation

AI advances are no longer confined to cognitive tasks; they are now entering the physical world. For decades, physical automation remained constrained to tightly controlled environments where variability could be engineered out. Industrial robots could execute repetitive tasks with extreme precision, but only under fixed conditions (fixed inputs, fixed trajectories, and predictable outputs). When reality deviated, whether through a misaligned component, a deformed part, or an unexpected obstacle, systems typically failed, stopped, or required human intervention.

The limitation was less mechanical than cognitive. Traditional robots lacked the ability to understand changing environments, adapt actions in real time, and learn from experience.

Three converging technological breakthroughs are now making this possible. Together, these advances allow robots to perceive environments, reason about actions, adapt in real time, and continuously improve through experience.

World models and vision–language–action models: From perception to anticipation

Until recently, robots could recognize objects but could not reason about their relationships or anticipate what would happen if they moved or interacted. World models and vision–language–action (VLA) models are changing this paradigm.

World models give machines an internal, dynamic representation of their physical environment. Trained on massive data sets of video and physics simulations, they enable robots to anticipate the consequences of their actions before executing them, predicting how objects will move, how forces will interact, and how environments will respond.

VLA models, by contrast, act as the “unified pilot.” It is a foundation model that processes visual environments and verbal instructions to output robotic actions directly. A VLA model can function independently, enabling a robot to perform tasks by mimicking learned behaviors. However, a stand-alone VLA model remains essentially reactive: It “knows” what movement to make because it has seen it before, but it does not “understand” the underlying physics.

The breakthrough comes from combining VLA models, world models, and reasoning systems. Together, they enable robots not only to execute actions but also to increasingly evaluate alternatives, anticipate outcomes, and select appropriate responses while facing real-world variability.

The robot, in other words, no longer just repeats movements; it begins to anticipate them.

Embedded intelligence: From cloud dependency to real-time autonomy

Even with advanced reasoning, robots are constrained by latency. When perception and decision-making depend on a centralized or cloud-based infrastructure, response times, often measured in tens of milliseconds, are incompatible with physical interactions where adaptations must occur instantaneously.

Embedded intelligence removes this constraint by bringing perception, reasoning, and control directly onto the machine. Advances in edge computing, including specialized GPUs and emerging neuromorphic architectures, enable real-time inference at the point of action. This could potentially allow robots to continuously adapt to unexpected events (such as an obstructed path or an object slipping), providing stable performance without stopping or escalating and extending automation from stable, engineered environments into live operations. Autonomy is no longer tethered to the cloud; it lives on the machine.

Active digital twins: From monitoring assets to generating intelligence

Historically, training physical AI required millions of slow, expensive real-world trials. Today, platforms like MuJoCo are fundamentally changing these economics by approximating complex real-world dynamics. By applying massive domain randomization and intentionally varying friction, lighting, and materials across vast numbers of virtual scenarios, robots can learn to overcome the “reality gap” and accumulate years of robust experience before deployment. This “sim to real” capability redefines the digital twin. Unlike the passive monitoring tools of a decade ago, modern digital twins act as active “training gyms.” They serve as a core execution layer where physical intelligence is synthesized, stress-tested, and optimized in simulation before any physical action is taken.

Consequently, physical systems are acquiring software-like learning dynamics through a closed-loop feedback system. Real-world anomalies are fed back into the virtual environment to compute new solutions, which are then deployed back to the hardware. While strict safety protocols still require careful oversight, this loop allows physical intelligence to iteratively adapt to edge cases postdeployment, increasingly without human intervention.

In effect, physical systems begin to learn at the speed of software, not the speed of reality.

Taken together, these three breakthroughs give rise to physical AI, a new generation of intelligent robots capable of perceiving, reasoning, learning, and acting under real-world uncertainty.

These robots fundamentally expand the scope of automation, moving from repetitive, structured tasks to variable, unstructured, and decision-intensive activities. This enables the automation of tasks that were previously out of reach, both by deepening capabilities within historical strongholds, such as logistics and manufacturing, and by expanding into entirely new sectors, including agriculture, construction, waste management, healthcare, and retail.

Just as enterprise know-how is codified into agentic skills for cognitive AI, motor know-how is codified into physical skills: reusable behaviors that can be trained in simulation, deployed across robots, and improved through fleet learning.

Physical AI in practice: From blind repetition to adaptive execution

Order picking in a warehouse. A traditional robotic arm picks items from a fixed bin, at fixed coordinates, at a fixed pace. Introduce variability—a parcel that shifts, a deformed package, or an item stacked at an odd angle—and the system stalls or escalates to a human. A physical-AI-driven picker behaves differently. Its world model anticipates how the stack will settle when a box is removed; its vision–language–action model translates the instruction “pick the fragile item first” directly into motion; and embedded intelligence lets it adjust its grip in real time when a parcel begins to slip, without pausing the line or calling for help. Trained across billions of simulated variations in a digital twin before ever touching a real shelf, it arrives on day one capable of handling edge cases that would have stopped a conventional robot.

Quality inspection on a production line. A conventional vision system flags defects only against preprogrammed templates: It catches the flaws it was explicitly taught to find and misses the rest. A physical AI inspector can draw conclusions about what it sees. It can identify an anomaly it has never encountered, weigh whether the deviation is cosmetic or functional, and decide whether to reject the part, flag it for review, or adjust an upstream parameter, then feed that judgment back into its training loop so the next thousand units improve. Inspection shifts from rigid pattern matching to adaptive judgment in the face of real-world variability.

Nearly 60 percent of work is now theoretically automatable

Thanks to these technological advances, AI can now perform increasingly complex cognitive and physical tasks with limited human supervision, evolving from a tool that supports human work into a workforce operating under human orchestration.

According to McKinsey research, roughly 65 percent of total work hours in Europe and the United States involve nonphysical capabilities (including information processing, analysis, and coordination), and 35 percent require physical capabilities (including operating equipment, handling materials, and performing manual tasks).²

Approximately two-thirds of nonphysical work hours and one-third of physical work hours are now theoretically automatable, representing close to 60 percent of total work hours overall (Exhibit 2). These estimates are based on the capabilities of existing technologies, including those demonstrated in laboratory conditions, with the level of human proficiency required across different categories of work. As AI technologies continue to evolve rapidly, the automation frontier is likely to expand further over time.


Of course, actual adoption will depend on more than technical feasibility. Economic, regulatory, and organizational factors (including policy choices, labor costs, implementation investments, and development timelines) will determine when and where automation is deployed. History suggests that such transformations unfold over decades. Electricity, for example, took more than 30 years to diffuse widely, and industrial robotics followed a similar multidecade trajectory.

² “Agents, robots, and us: Skill partnerships in the age of AI,” McKinsey Global Institute, November 25, 2025; “Agents, robots, and us: How AI reshapes work and skills in Europe,” McKinsey Global Institute, May 12, 2026.

Exhibit 2

With advances in cognitive and physical AI, more than half of all work hours can theoretically be automated.


Distribution of work hours in Europe, by technical automation potential, 2024,¹ %



¹Includes only Czech Republic, Denmark, France, Germany, Italy, Netherlands, Poland, Spain, Sweden, and UK. Automation potential is based on current capabilities of technology to perform human work. Automation potential shown is the late scenario of expert estimates. The early scenario of global technical automation potential ranges from 60% to 70% of current work hours.
Source: National statistical offices; ONET; McKinsey Global Institute analysis ("Agents, robots, and us: Skill partnerships in the age of AI," Nov 25, 2025)

McKinsey & Company

AI agents and intelligent robots are no longer peripheral tools; they are becoming integral participants in how work gets done. The central challenge is no longer whether AI can execute work, but how humans, agents, and intelligent machines can work together at scale inside a new operating model.



**The step change
comes not from
augmentation but
from reinvention**

AI adoption is nearly universal: More than 80 percent of companies deploy AI in at least one function as of 2025, up from 50 percent in 2022.³ But adoption is no longer the differentiator. A radically different level of impact is emerging between organizations using AI to augment existing workflows and those reinventing work around hybrid human–AI execution.


Most organizations still use agentic AI to augment existing workflows, generating only incremental gains

Sixty-two percent of companies are experimenting with AI agents, according to McKinsey’s most recent Global Survey on AI. Yet deployment is limited to one or two functions, with fewer than 10 percent of organizations scaling agents within any given function (Exhibit 3).⁴

Exhibit 3

No more than 10 percent of respondents report scaling AI agents in any individual function.

Phase of AI agent use at respondents’ organizations, by business function,¹ % of respondents (n = 1,933)



Note: Figures may not sum to 100%, because of rounding.
¹Question was asked only of respondents who reported regular use of AI in the respective functions and was rebased to reflect the total sample.
Source: McKinsey Global Survey on the state of AI, June 25–July 29, 2025 (n = 1,993 participants at all levels of the organization)

McKinsey & Company

Among organizations experimenting with agentic AI, most still deploy it within existing workflows, limiting outcomes to incremental productivity gains with little P&L impact. Deployments typically consist of either individual copilots augmenting employees in specific tasks such as writing, research, or translation, or narrowly scoped agents automating isolated workflow fragments such as data consolidation, KPI reporting, and ticket triage. In most cases, humans remain “in the loop” at every critical step: validating outputs, coordinating handoffs, managing exceptions, and driving execution


³ “The state of AI in 2025: Agents, innovation, and transformation,” McKinsey, November 5, 2025.
⁴ “The state of AI in 2025: Agents, innovation, and transformation,” McKinsey, November 5, 2025.

sequentially. AI improves individual tasks, but the overall workflow architecture remains largely unchanged (Exhibit 4).

Exhibit 4

While many organizations are using AI, most are embedding the technology into existing workflows.

Level of AI implementation, % of organizations



Source: McKinsey Global Survey on the state of AI, June 25–July 29, 2025 (n = 1,993 participants at all levels of the organization)

McKinsey & Company

Hybrid human-AI teams operating in reinvented workflows are delivering radically greater gains

When workflows are redesigned from first principles and roles are redistributed between humans and AI agents according to their respective strengths, organizations can achieve step change gains that far exceed those delivered through traditional AI augmentation. The following examples provide a ground-level view of this shift in action.

Software development: The two-shift factory with AI agents

The first generation of AI assistants for developers delivered an initial layer of productivity gains, typically in the range of 5 to 15 percent based on our experience, primarily through code generation support, without changing the development cycle.

With the emergence of next-generation agents, the entire software development process can be reinvented, redefining the distribution of roles between humans and AI agents. This enables a fundamentally new operating model: the two-shift software factory. In this model, software development runs as a continuous cycle. Humans take the day shift, setting direction and enforcing quality, while AI agents take the night shift, executing work continuously and at scale.

During the day shift, human teams define what needs to be built and how. They refine user stories, translate features into detailed specifications, break work into well-scoped, agent-ready tasks, and define acceptance criteria. They provide architectural guidance, clarifying which components can be modified and which must remain stable, and continuously adjust priorities, guardrails, and test coverage. Their role shifts from writing code to directing, decomposing, and controlling the quality of the execution.

As evening begins, the AI agent workforce takes over. A coordinated system of specialized agents executes multistep workflows in parallel across coding, testing, quality assurance, security validation, performance optimization, and documentation.

The entire workflow is managed by an orchestrator agent coordinating handoffs and decision flows across specialized subagents. If tests fail, work is routed for correction; if performance degrades, additional checks are triggered; if policy constraints are violated, execution is halted. By morning, the system has produced a set of ready-to-review pull requests, complete with code, tests, logs, analysis, and structured explanations.

The human team then resumes control, reviewing outputs, validating architectural alignment, refining results where needed, and setting new directions. Over time, the scope of what can be safely automated expands, increasing the system's autonomy and productivity.

This approach unlocks step change productivity improvements across development teams as 24-hour delivery cycles replace multiweek sprints. But the deeper shift is organizational rather than purely operational. Software development moves from a sequential, human-limited process to a continuous system of execution, where humans provide direction and judgment, and AI agents perform the bulk of execution at scale.

While success stories at scale remain relatively scarce, early signs of breakthrough are starting to appear. One large financial services firm, for example, has set up an AI agent factory to develop a greenfield payment system and is improving productivity by 40 percent or more.⁵

Customer operations: From assisted service to autonomous operations

With first-generation gen AI tools, it was possible to assist customer service agents by suggesting responses, summarizing previous interactions, or retrieving relevant knowledge-base articles. This allowed incremental improvements in agent productivity, typically in the range of 10 to 15 percent based on our experience, as well as response quality, but the interaction remained fully human driven. Similarly, first-generation chatbots could address simple questions but not process end-to-end customer requests or claims.

A European utility has recently deployed a next-generation customer service model that relies on hybrid human–AI teams. A multimodal (voice, chat, email) conversational agent interacts directly with customers and orchestrates a series of specialized agents, each responsible for a specific task, such as retrieving customer data, validating account information, applying pricing and policy rules, handling billing adjustments or refunds, and ensuring regulatory compliance. These agents operate as a coordinated system, enabling end-to-end resolution of customer interactions while dynamically adapting to context.

Making this process work required significant architectural investment: a reusable enterprise AI platform integrating CRM, billing, and agent desktop systems; differentiated models optimized for latency versus verification; and dedicated governance metrics to monitor quality, compliance, and guardrails.

In this new model, the role of human agents shifts fundamentally, from handling individual requests to supervising AI-driven interactions and from executing transactions to managing exceptions and high-sensitivity situations.

More than 40 percent of calls are now fully identified and verified by the AI voice agent, reducing the average cost to manage inbound calls by close to 50 percent.⁶ This progression illustrates a broader shift: from AI as a support tool, to AI as a first-line operator, to AI as a fully integrated customer service workforce.

⁵ Eric Lamarre, Kate Smaje, Robert Levin, Alex Singla, and Alexander Sukharevsky, *Rewired: How Leading Companies Win with Technology and AI*, 2nd edition, Wiley, April 2026.

⁶ Eric Lamarre, Kate Smaje, Robert Levin, Alex Singla, and Alexander Sukharevsky, *Rewired: How Leading Companies Win with Technology and AI*, 2nd edition, Wiley, April 2026.

These examples demonstrate that once workflows are reinvented around hybrid human–AI execution, AI no longer behaves as a productivity tool embedded within the operating model. It becomes part of the operating model itself, reshaping how work is executed, coordinated, and scaled across the enterprise.

Early implementations of physical AI are generating disruptive gains when operating models are redesigned

The same pattern—AI moving from tool to execution unit—is now reshaping the physical world. Some pioneering companies have begun deploying physical AI systems in highly structured environments, with thousands of AI-coordinated robots, including autonomous mobile robots (AMRs) and collaborative robots (cobots), operating across warehouses and fulfillment centers. In these environments, physical AI is no longer simply added onto existing workflows. It becomes the core execution system around which operations, workforce roles, inventory flows, and coordination are redesigned.

Amazon, for example, now operates more than one million robots across its global operations. Its Sequoia system identifies and stores inbound inventory up to 75 percent faster than its previous robotic systems, reducing order-processing time by up to 25 percent.⁷ But these gains did not come from simply replacing human operators with robots. Amazon has reinvented the warehouse operating model around AI-driven orchestration. Instead of humans moving through warehouses to locate products, fleets of robots dynamically move inventory across continuously optimized flows.

In June 2025, Amazon introduced DeepFleet, its first foundation model for multirobot coordination. Trained on millions of hours of warehouse data, it dynamically assigns tasks and routes robots around congestion in real time, while supporting the human teams working alongside them. As it does so, human roles shift progressively from repetitive physical execution to supervision, exception handling, and system coordination. The warehouse therefore evolves into a software-orchestrated execution system in which humans, robots, and AI operate as an integrated operational network.⁸

At the online supermarket Ocado, meanwhile, thousands of robots operate simultaneously on a grid to assemble grocery orders, coordinated in real time by a centralized AI layer optimizing routing, inventory access, traffic management, and order sequencing. Here again, the disruptive impact comes not from robotization alone, but from designing the entire fulfillment model around algorithmic orchestration and highly composable execution flows. Ocado reports near-perfect order accuracy and up to 99 percent reliable, on-time dispatch.⁹

Humanoids are being explored to automate physical tasks in less structured, human-centric environments

Most transformative deployments of physical AI to date have occurred in highly structured, purpose-built environments that have been optimized for automation: standardized layouts, predictable flows, flat floors, and tightly controlled operating conditions. In such settings, robotics systems can deliver massive gains in speed, precision, utilization, and scalability when workflows are redesigned around automation.

Yet most physical environments remain far less structured than modern fulfillment centers or automated warehouses. Factories, hospitals, stores, airports, construction sites, and industrial facilities—equipped with stairs, doors, handles, ladders, corridors, shelving, tools, and variable workspaces—were designed around human mobility, dexterity, and interaction. As a result, specialized robots often require costly infrastructure redesign and thus remain confined to isolated tasks embedded within existing workflows.

⁷ Scott Dresser, "Amazon announces 2 new ways it's using robots to assist employees and deliver for customers," Amazon News, October 18, 2023.

⁸ Joey Durham, "Amazon builds first foundation model for multi-robot coordination," Amazon Science, August 11, 2025.

⁹ "Online grocery through the Ocado Smart Platform," Ocado Group, 2026.

Humanoid robots are expected to help address that challenge. By replicating the human form factor, humanoids could potentially automate some tasks inside existing human-designed environments without requiring large-scale infrastructure redesign. Their economic value, therefore, comes not only from targeted labor substitution but also from reducing the need to redesign physical infrastructure around automation.

To date, humanoids remain at an earlier stage of maturity than other categories of robots. Yet early commercial pilots concentrated in manufacturing and logistics environments, such as Renault's pilot of Wandercraft's Calvin 40 humanoid, are delivering early but encouraging results.

Scaling humanoid deployment toward broader use cases will require overcoming four major barriers: advanced cognitive interpretation of dynamic and unpredictable environments; fine motor dexterity beyond transport and inspection tasks; sustained operational uptime well beyond the few hours per battery charge achieved today; and radical cost reduction from current levels (\$150,000 to \$500,000 per unit) toward the \$20,000 to \$50,000 range likely required for large-scale adoption.¹⁰

Meet Calvin, Renault's humanoid robot

At Renault's production lines, Calvin handles tires in assembly workflows where human operators face strict ergonomic regulations limiting load weight and hourly tonnage. The robot, created by Wandercraft, operates within existing factory infrastructure (without requiring layout redesign), and performs tasks beyond human physical constraints.

Operating in live production environments gives these robots access to continuous real-world data that feeds back into

their AI systems, progressively improving reliability, reducing failure rates, and expanding the range of tasks the robot can perform autonomously. Renault has announced plans to deploy 350 units within 18 months, as part of a broader operational-excellence initiative targeting a 30 percent reduction in production hours per vehicle—a breakthrough ambition in automotive manufacturing. The company plans to industrialize Wandercraft's robots in large series production, leveraging its automotive production expertise.

Across both cognitive and physical domains, the pattern is the same: the largest gains come from redesigning work around the combined capabilities of humans, agents, and robots, not from inserting AI into existing workflows. At scale, this shift transforms AI into a substrate for the enterprise itself, reshaping how work is coordinated, executed, and scaled.

¹⁰ "Humanoid robots: Crossing the chasm from concept to commercial reality," McKinsey, October 15, 2025.



The symbiotic enterprise: Reinventing how work is executed

We’ve seen how the reinvention of individual functions around hybrid human–AI execution (for example, software factories operating continuously across shifts, customer operations handled autonomously at scale, robotic fleets transforming warehouse productivity) can produce dramatic gains. Yet in most organizations, these transformations remain confined to isolated domains while the broader enterprise continues to operate largely unchanged.

The deeper question is: What happens when this reinvention extends across the entire operating model?

When workflows throughout the enterprise are redesigned around hybrid human–AI systems, the outcome is no longer a collection of optimized functions. It is a fundamentally different kind of enterprise, one that transforms not only how work is executed but also how organizations scale, innovate, and compete.


Reimagining workflows at scale gives rise to a new enterprise model

At scale, the reinvention of cognitive and physical workflows gives rise to the symbiotic enterprise—a new enterprise model in which humans, AI agents, and intelligent robots each contribute according to their respective strengths. Humans provide judgment, creativity, supervision, and exception handling. AI systems deliver scalable cognitive execution and coordination, while physical AI systems perform repetitive or operationally constrained tasks. The unit of production is no longer the individual; it is the hybrid human–AI system (Exhibit 5).

Exhibit 5

In the symbiotic enterprise, humans, AI agents, and intelligent robots each contribute according to their respective strengths.

Role allocation in the symbiotic enterprise



McKinsey & Company

As these hybrid systems become the foundational units of execution, human roles, organizational structures, and enterprise economics are fundamentally reshaped.

Roles: From producing work to exercising judgment

Today, enterprise roles are primarily defined by the work humans produce: the analyst builds the model, the lawyer drafts the contract, the engineer writes the code. Humans remain the primary unit of execution. As AI systems progressively take on a growing share of execution tasks, human roles shift from producing work to directing, supervising, and improving systems of execution.

The scale of this transition is not incremental. According to McKinsey research, 75 percent of roles may require significant redesign, including both managerial and frontline positions.¹¹ Three broad role categories are emerging: integrative supervisors orchestrate hybrid human–AI teams and robotized fleets; deep specialists handle edge cases, arbitrate ambiguity, and continuously improve system performance; and AI-enabled frontline workers differentiate through empathy, trust, persuasion, and social and emotional intelligence. As execution shifts to AI systems, strategic thinking, judgment, and creativity become more and more central. Conversely, routine forms of research, analysis, synthesis, and first-draft production progressively decline as tasks humans need to perform.

The magnitude of this shift cannot be addressed through incremental training tied to new tools. Enterprises will need large-scale, multiyear workforce transformation and reskilling programs. They will also need to address two structural risks: the widening divide between AI-empowered and marginalized employees (what we call “the AI inclusiveness challenge”) and the erosion of traditional apprenticeship pathways as junior analytical work becomes automated. Without new mechanisms to develop judgment and operational experience, organizations may ultimately face shortages of experienced talent capable of governing autonomous systems.

Organization: From functional silos and coordination layers to small, outcome-oriented teams orchestrating end-to-end execution

Traditional hierarchical organizations emerged from two structural constraints: the need to distribute specialized expertise across functions and the need to coordinate workflows between them.

As AI agents execute specialized work and coordinate activities autonomously, both of those constraints weaken. Agentic systems operate across functional domains, share context in real time, manage dependencies, escalate exceptions, and coordinate workflows without extensive hierarchical coordination layers. The result is a different organizational model, in which small human teams orchestrate large groups of specialized agents to deliver end-to-end outcomes rather than pass work sequentially between silos. These teams assemble and reconfigure dynamically as priorities evolve.

Two consequences could follow: organizational flattening (as layers of middle management consolidate into smaller orchestration layers) and a significant expansion in managerial reach (as leaders supervise broader portfolios enabled by agentic systems). To be sure, predictions of organizational flattening have had a long history of failing to materialize. What may be different this time is that AI agents are replacing specific coordination functions directly, rather than simply freeing managerial time that is later reabsorbed elsewhere.

The enterprise therefore evolves from a rigid functional hierarchy into a dynamic system of orchestrated workflows.

Economics: From labor to technology as a dominant factor

Historically, operating economics differed sharply across industries. In labor-intensive service and knowledge sectors, costs were driven primarily by workforce size. In capital-intensive industries, such

¹¹ McKinsey analysis based on data from “Agents, robots, and us: Skill partnerships in the age of AI,” McKinsey Global Institute, November 25, 2025.

as manufacturing, logistics, energy, or heavy industry, the operating economics relied more on industrial assets, infrastructure, and machinery.

Cognitive AI shifts costs away from labor toward software, models, orchestration layers, and compute infrastructure. Physical AI extends this shift into operational environments, increasing the role of robotics, embedded intelligence, autonomous systems, and simulation platforms. As a result, technology becomes a growing, and potentially dominant, component of enterprise economics. In knowledge-intensive industries, technology spending could ultimately exceed labor costs altogether.

This creates a new management challenge. Labor costs have been relatively predictable, governed by hiring plans and compensation policies. AI economics behave differently in that costs are dynamic and consumption driven based on compute usage, inference volumes, model selection, and agent activity.

Managing future enterprise economics will require both designing AI systems for cost efficiency from the outset and continuously optimizing their consumption in production through fine-grained orchestration decisions. Much like FinOps emerged for cloud computing, organizations will need to actively manage token consumption, model utilization, and inference costs by routing tasks across models, caching repetitive queries, compressing context, batching inference requests, and balancing workloads across providers (table).

Table

The symbiotic organization represents a fundamental shift in how organizations operate.

	FROM: Traditional enterprise	TO: Symbiotic enterprise
Human role	Planning and executing tasks with AI assistance	Setting strategic priorities and supervising AI systems
Organization	Fixed multilayer structure coordinating vertical functions	Dynamic flat network of outcome-oriented hybrid teams
Economics	Labor or physical capital as dominant cost drivers, fixed and predictable	Technology-intensive cost structure, variable and usage-driven

Beyond productivity: A new model for faster innovation, continuous adaptability, and scalable growth

The symbiotic enterprise unlocks substantial productivity gains by automating cognitive and physical work through hybrid human–AI systems operating at far lower marginal cost. Coordination overhead compresses, execution cycles accelerate, idle time falls, and asset utilization improves across workflows.

Yet productivity is only the first-order effect, and productivity advantages alone are unlikely to remain durable. As AI capabilities become widely available, productivity gains diffuse across industries and are competed away. Over time, they become a competitive necessity rather than a source of differentiation, with customers capturing much of the resulting value through lower prices, improved service, and higher expectations.

The deeper significance of the symbiotic enterprise lies in its ability to fundamentally change how organizations innovate, adapt, grow, and scale.

- ***Accelerating innovation.*** Innovation has historically been constrained by the cost and speed of experimentation. Developing new products, testing new business models, redesigning processes, or exploring new markets typically required significant investments of time, expertise, and resources. Many potentially valuable ideas never progressed beyond the concept stage because experimentation itself was too expensive. The symbiotic enterprise reduces these constraints. AI systems can generate concepts, simulate scenarios, create prototypes, write code, analyze outcomes, and support decision-making at unprecedented speed. As experimentation cycles compress and the cost of iteration falls, organizations can test, refine, and scale innovations more rapidly than before. This effect extends beyond product development. Business models, operating models, customer experiences, pricing mechanisms, and internal processes can all be redesigned and optimized continuously.
- ***Enabling continuous adaptability.*** Historically, organizations have been optimized primarily for efficiency and stability. Processes were designed to minimize variation; organizational structures evolved to support predictable execution and change often occurred through periodic transformation programs. The symbiotic enterprise introduces a more adaptive model. Agentic systems can continuously reallocate work, rebalance resources, adjust workflows, and coordinate activities in response to changing demand, operating conditions, technological developments, or competitive pressures. Hybrid human–AI systems allow organizations to dynamically adjust the mix of human labor, digital labor, and physical automation at the workflow level. As a result, adaptation becomes embedded into daily operations rather than occurring through episodic redesign efforts. In environments characterized by technological disruption, economic volatility, and shifting customer expectations, adaptability may become as important as efficiency.
- ***Expanding growth opportunities.*** The symbiotic enterprise does not merely reduce the cost of growth; it expands the universe of growth opportunities in three ways.¹² First, it enhances existing products and services through real-time optimization of allocation, pricing, personalization, and coordination. Ride-hailing platforms, for example, replace physical taxi queuing with algorithmic dispatch, while financial institutions can tailor offers dynamically based on transaction behavior and risk signals. Second, the symbiotic model expands the product frontier by making economically viable offerings that were previously impossible or prohibitively expensive, such as AI-designed molecules accelerating drug discovery or on-demand delivery platforms extending coordination logic from mobility into food and retail. Third, and potentially most transformative, it enables entirely new business models in which AI becomes the delivery mechanism itself, collapsing cost structures and opening markets that legacy economics made unserviceable, such as AI-native legal platforms delivering services that once required teams of specialists, or AI tutors providing personalized, always-on instruction at a scale and cost previously unattainable.
- ***Scaling through software rather than labor.*** Traditional enterprises expanded through recruiting, training, and coordinating larger workforces. Symbiotic enterprises scale through software, orchestration, and intelligent execution systems. Organizational capacity expands through technology infrastructure, models, orchestration platforms, governance systems, and compute rather than through proportional increases in head count. As a result, growth progressively decouples from labor expansion, enabling organizations to scale output, coordination, and execution with far lower marginal labor intensity. Traditional enterprises scaled at the speed of human coordination, while symbiotic enterprises scale at the speed of software.

¹² Antoine Montard, Dago Diedrich, and Tanguy Catlin, “Where AI will create value—and where it won’t,” McKinsey, April 29, 2026.

The symbiotic enterprise erodes traditional advantages while creating new dependencies

The symbiotic enterprise creates a fundamental strategic tension. The same forces that boost enterprise performance also reshape the foundations of competitive advantage. As cognitive capabilities become widely accessible, transaction costs decline, and execution becomes software driven, several historical sources of differentiation weaken while new dependencies emerge.

Expertise becomes democratized

Specialized expertise has long been a powerful source of competitive advantage because it was scarce and difficult to scale. The symbiotic enterprise changes this equation. AI democratizes access to knowledge, analysis, software development, forecasting, and decision support. Capabilities that once required years of training can now be augmented by intelligent systems. Deep domain knowledge and judgment remain essential, but expertise alone becomes a less durable source of differentiation as access to cognitive capabilities broadens across competitors, customers, and new entrants.

Scale advantages diminish

Historically, scale enabled organizations to spread fixed costs, invest in specialized capabilities, and support complex operations that smaller competitors could not replicate. The symbiotic enterprise lowers many of these barriers. Software development, customer service, analytics, and many operational capabilities can be delivered through intelligent systems rather than large workforces. As a result, the minimum efficient scale required to compete declines, allowing smaller and more agile competitors to challenge incumbents. Scale does not disappear, but advantage shifts away from workforce scale toward data scale, ecosystem scale, and learning scale.

Coordination barriers fall

Many organizations derived advantage from their ability to coordinate large numbers of employees, suppliers, partners, and customers. Managing complexity required extensive hierarchies, formal processes, and significant managerial capabilities. Agentic systems automate much of this coordination work by allocating resources, managing dependencies, monitoring execution, and orchestrating workflows across organizational boundaries. Activities that once required substantial management structures can be coordinated through software, reducing one of the traditional advantages of large organizations.

Friction-based advantages weaken

Many organizations benefited not only from superior capabilities but also from market frictions such as search costs, switching costs, information asymmetries, and coordination inefficiencies. Agentic AI has the potential to reduce many of these transaction costs dramatically. Agents can search, compare, negotiate, coordinate, and transact directly with one another. As transparency increases and switching costs decline, customer mobility rises, intermediary positions become more vulnerable, and industry value chains become more fluid and contestable.

A new strategic dependency emerges

As traditional sources of differentiation erode, a growing share of enterprise capabilities stands to become dependent on foundation models, cloud platforms, orchestration layers, AI tooling, and compute infrastructure controlled by a small number of providers. This dependence extends beyond traditional supplier relationships to the intelligence layer itself. As a result, part of enterprise performance becomes influenced by external model road maps, pricing decisions, infrastructure availability, regulatory constraints, and geopolitical considerations. This creates what could become a new form of economic dependency: a “cognitive tax,” through which a portion of the value created by symbiotic enterprises accrues to the providers of the intelligence infrastructure on which they depend.

Winning in the world of commoditized intelligence

The symbiotic enterprise unlocks powerful new capabilities. Yet as AI technologies diffuse, these capabilities will become available to all organizations. Similarly, new vulnerabilities will affect entire industries, not individual firms.

The critical strategic question is therefore not what AI enables, but where lasting differentiation will accumulate. History suggests that when a transformative technology becomes broadly accessible, advantage shifts away from access to the technology itself and toward the assets, positions, and disciplines that deepen with use.

In the AI era, three sources of durable differentiation stand out.

Building proprietary intelligence that compounds over time

As foundational AI models become commoditized, access to frontier capabilities will no longer differentiate. Instead, what will differentiate is the proprietary intelligence that organizations build on top of these foundations: unique data assets, specialized agent skills, and continuous learning loops fed by operational execution. Proprietary data enriches models trained primarily on public information with company-specific context, while agentic skills codify the organization's unique know-how, processes, and ways of operating. Together, they create forms of intelligence that competitors cannot easily access or replicate.

Every customer interaction, every transaction processed, and every decision made within the symbiotic enterprise generates data that, when captured and fed back into AI systems, improves their performance over time. Organizations that systematically capture, structure, and reinvest this operational knowledge create compounding advantages: the more they operate, the better their systems become, and the wider the performance gap grows. As foundation models commoditize, competitive advantage shifts from access to intelligence toward ownership of proprietary intelligence.

Owning customer and ecosystem control points

As coordination frictions decline in AI-mediated markets, structural advantage increasingly accrues to organizations that control customer relationships and ecosystem positions.

The first source of advantage is deep integration into customer workflows. Organizations that embed intelligent systems directly into customer operations through agent-to-agent interactions, real-time data exchanges, and continuous learning from customer-specific contexts create relationships that become progressively more difficult to displace.

The second is ecosystem positioning. As AI agents discover (for example, selecting suppliers, shortlisting options), compare (for example, benchmarking prices, evaluating service levels), and transact (for example, switching providers, routing purchases) on behalf of their principals, the organizations that control the interfaces through which these interactions occur will capture a disproportionate share of demand. The strategic imperative is one of positioning and timing; organizations that position themselves at these new control points early can define the rules before market structures stabilize.

Mastering intelligence architecture and execution

The symbiotic enterprise introduces an entirely new operational layer that must be designed, governed, optimized, and scaled. While AI capabilities may become broadly available, the ability to architect and operate intelligent execution at scale will remain a powerful source of differentiation. Doing so requires mastery of five disciplines.

The first is intelligence architecture engineering. Organizations must design the cognitive architecture of the enterprise itself: determining how work, decision-making authority, and accountability are distributed between humans and AI agents; defining the optimal mix of specialized and general-purpose agents; establishing coordination mechanisms between agents and humans, whether hierarchical, process driven, or network based; determining the appropriate balance between centralized and distributed intelligence; and defining performance criteria for hybrid human–AI teams. As intelligent execution becomes a core production capability, the quality of this architecture increasingly determines enterprise performance. As intelligent execution becomes a core production capability, the quality of this architecture increasingly determines enterprise performance.

Second is intelligence industrialization. Organizations must develop the capabilities required to industrialize the creation and deployment of intelligent capabilities across the enterprise. This includes creating reusable agent skills, maintaining libraries of proven components, developing evaluation harnesses and simulation environments, and establishing engineering practices that maximize reliability, reuse, and deployment speed. As intelligent execution expands across the enterprise, the ability to rapidly replicate, adapt, and scale proven capabilities becomes a critical source of productivity and competitive advantage.


Third is orchestration at scale. Enterprises must coordinate thousands of agents, workflows, and human–AI handoffs across the organization with the reliability, resilience, and auditability required in production environments. Operating a handful of agents is relatively straightforward; orchestrating large-scale intelligent execution is substantially more complex.

The fourth component is behavioral governance and security. Organizations must control what autonomous systems can and cannot do, ensure compliance, manage risk, maintain trust, and protect against failures, misuse, and cyberthreats as the scope of autonomous execution expands.

Finally, organizations must rigorously control AI economics. As AI becomes a major component of enterprise cost structures, organizations must manage intelligent systems with the same rigor applied to financial and operational assets. This includes monitoring and optimizing inference costs, token consumption, model utilization, and cost per transaction while continuously improving the economics of intelligent execution.

As intelligent execution scales, the gap between organizations that master these disciplines and those that do not becomes a structural source of advantage. Not every organization will be able to design, govern, and operate thousands of intelligent agents safely, reliably, and economically. Those that can will benefit from superior execution, faster learning, lower costs, and greater confidence in autonomous operations.

The symbiotic enterprise is not a technology deployment or a productivity program. It is a strategic transformation that reshapes operating models, competitive dynamics, and industry value chains. As AI disrupts existing sources of advantage, organizations face a dual imperative. Offensively, they must build the proprietary intelligence, ecosystem positions, and operational disciplines that will define competitive advantage in the AI era. Defensively, they must protect against the commoditization of cognitive capabilities, the collapse of friction-based advantages, and excessive dependence on external AI ecosystems. The organizations that will lead will not be those with access to the best models, but those that combine both imperatives into a coherent transformation.



Navigating the symbiotic transformation: Reinvention without destabilization

As we have shown, recent advances in cognitive and physical AI are making the symbiotic enterprise increasingly achievable. Early deployments of hybrid human–AI systems already demonstrate disruptive gains when workflows are redesigned from first principles. The challenge is no longer whether to move, but how fast and how far. Leaders must navigate between two failures: incrementalism, which leaves organizations trapped in a pre-AI operating model, and overreach, which deploys autonomous systems faster than they can be absorbed and governed. Future competitiveness will depend on striking the right balance.

Achieving success requires a bold, value-driven target operating model; a dual transformation path balancing value realization and workforce transition; scalable AI foundations; and extended executive leadership.

Define a bold, value-driven target state before you move

The transformation to a symbiotic enterprise model is too large and too uncertain to execute as a big bang. But progressing without a clear destination is equally risky. Enterprises need to define the strategic direction of the future operating model up front, even if the target state evolves over time.

A clearly articulated ambition is essential to size the opportunity, calibrate investment, guide workforce transformation, and ensure consistency in decision-making throughout the journey.

The target state cannot be built from the bottom up based on isolated use cases focused solely on productivity. Because AI reshapes competitive dynamics and economics, it must be derived top down with a three-step approach:

- ***Identify future profit pools and sources of differentiation.*** Two questions need to be addressed. The first is defensive: For each revenue stream, does the current offering risk commoditization, customer re-insourcing, disintermediation, or disruption from AI-native entrants operating with structurally lower costs and faster learning cycles? Enterprises need to identify where differentiation can still be sustained or rebuilt and which activities or positions may become structurally indefensible. The second question is offensive: Where can AI create entirely new value through intelligent products, AI-enabled services, embedded capabilities, autonomous interactions, hyperpersonalized experiences, or new forms of monetization?
- ***Define a macro vision of the target operating model.*** This second step aims to define how the enterprise should operate in the AI era to support future profit pools. The enterprise should notably determine the distribution of roles between humans, agents, and robots in the target state.
- ***Assess workforce transformation needs and value at stake.*** The target operating model provides the foundation to build a multiyear workforce transformation plan by comparing future workforce requirements with the current workforce base. The objective is not only to estimate head count evolution but also to identify which populations will be affected, the magnitude of change required for each role, and the pace at which the workforce can realistically transition. Based on the expected future revenue pools and an estimate of the future running costs of the target operating model, transformation value at stake can be estimated.

From enterprise software to digital workers: The reinvention of IFS

Industries like manufacturing, energy, utilities, aerospace and defense, telecommunications, and construction succeed or fail on the performance of their physical assets. The Sweden-based enterprise software company IFS is reinventing itself to be a leading AI-native software firm for industrial and asset-heavy companies, helping its customers perform with less unplanned downtime, faster service, lower operating costs, and higher output. IFS is now delivering those outcomes in a fundamentally new way, having embarked on a fundamental reinvention of its product portfolio, business model, operating model, and culture.

With a long heritage in enterprise resource planning, enterprise asset management, and field service management solutions, IFS is evolving into a leader in industrial AI. Its approach rests on a complete industrial AI stack. It starts with data: decades of asset and operational data drawn from real-world operations and difficult to replicate at scale. On top of that data sits a layer of intelligence and analytics that turns raw operational signals into decisions. Acting on those decisions are AI-enabled digital workers that execute complex workflows autonomously across maintenance, service management, and supply chain. Rather than simply supporting human work, these digital workers increasingly participate in operational execution, augmenting teams and automating complex workflows.

To accelerate this transformation, IFS created Nexus Black, an industrial AI accelerator that embeds AI engineers, product leaders, and industry experts directly with customers to codevelop solutions, prove the business case through

rapid experimentation, and industrialize what works across the platform. To deploy these innovations at scale, it acquired TheLoops, an agentic platform that orchestrates digital workers across customer environments.

The commercial model is changing to match this reinvention. When digital workers perform the work, and value no longer tracks the number of human users, the old model of seat-based licensing stops making sense. IFS is instead moving to asset-based pricing, tied directly to the assets it helps manage, optimize, and keep running. This marks a shift in enterprise software: from paying for access to applications to paying for outcomes delivered.

Staying close to customer outcomes also demands a new organizational structure. IFS has replaced its market-and-region structure with one built around value chains, placing developers alongside sales and customer success, ensuring product relevance in a context of disruption. This reinvention is enabled by a culture of self-disruption, where every function is expected to continuously question and re-earn its own relevance.

The transformation of IFS illustrates how an enterprise software company can reinvent its products, platform, operating model, and commercial model around AI capabilities—moving from providing software that supports work to intelligent systems that increasingly participate in its execution.

Adopt a dual transformation journey balancing value realization and workforce adaptation

Realizing the full potential of the symbiotic enterprise requires the coordinated transformation of the operating model, technology, data, and talent, guided by a clear North Star and enabled by strong change management. Because these dimensions are deeply interconnected, transformation must be managed as an integrated enterprise effort rather than a collection of isolated initiatives.

The right unit of transformation is the operational domain, such as product development, customer service, or order to cash, where workflows, roles, coordination mechanisms, and human–machine collaboration can be reinvented at sufficient scale to justify transformational investment and meaningful organizational change.

Yet enterprises should not wait for full domain-level reinvention before deploying AI broadly. Tactical augmentation use cases can improve individual productivity, accelerate organizational learning, and build momentum for larger-scale transformation.

A dual transformation path is therefore required. Tactical augmentation delivers near-term value and accelerates learning, while domain-level reinvention fundamentally redesigns operating models from first principles to unlock enterprise-wide gains in speed, scalability, coordination, and economics.

Sequencing should balance value potential and feasibility, taking into account technological readiness, data availability, organizational complexity, workforce impact, and competitive pressure. Domains already operated through outsourcing or offshoring often provide attractive starting points, as they can typically be redesigned faster and with lower organizational disruption. At the same time, workforce adaptation must progress in parallel, ensuring that employees have sufficient time and support to develop the skills required to operate effectively alongside increasingly capable AI systems.

Build scalable foundations

Scaling sustainably requires a new operational foundation that orchestrates intelligent execution while preserving control, resilience, security, economic discipline, and human oversight. Three foundational elements are critical:

- ***AI-ready enterprise systems and data.*** Legacy applications must expose their capabilities through clean APIs, and enterprise data must become agent readable and agent actionable, that is, structured, documented, governed, and accessible in real time.
- ***A modular, vendor-agnostic orchestration platform.*** Relying initially on integrated AI stacks from major technology providers may accelerate adoption. But to scale AI sustainably, enterprises need modular and vendor-agnostic architectures capable of preserving portability, diversifying dependencies across multiple providers, and avoiding excessive concentration around any single technology ecosystem.
- ***Intelligence architecture and execution of new disciplines.*** As intelligent execution becomes a core production capability, competitive advantage will increasingly depend on mastery of five disciplines: intelligence architecture engineering, intelligence industrialization, orchestration at scale, behavioral governance and security, and AI economics management. Together, these capabilities enable organizations to design, build, operate, govern, and economically optimize human–AI systems at scale.

Drive the transformation through executive leadership

The symbiotic transformation cannot be delegated to a technology function or managed as a collection of isolated initiatives. Given its strategic importance and enterprise-wide impact, it requires direct CEO leadership to define the ambition, articulate the target state, arbitrate trade-offs, allocate resources, and set the pace of transformation. Yet according to McKinsey research, only 30 percent of CEOs today actively oversee their organization's AI agenda.¹³

¹³ "The state of AI: How organizations are rewiring to capture value," McKinsey, March 12, 2025.

At the same time, the CEO cannot lead this transformation alone. Three executive roles become particularly critical. The chief human resources officer should lead the workforce transition by redesigning roles, orchestrating reskilling and redeployment, and ensuring that incentives, career paths, and performance systems reinforce AI-enabled ways of working. The chief transformation officer (or equivalent transformation leadership) should orchestrate execution across domains by translating strategic ambition into coordinated programs, aligning business and functional leaders around shared priorities, reallocating resources as the road map evolves, and ensuring accountability for outcomes. And the chief technology officer (or chief information officer, depending on the organization) should build the foundations of the symbiotic enterprise, such as modernizing legacy systems, establishing reusable data and AI capabilities, deploying the agentic execution platform, and ensuring the scalability, security, reliability, interoperability, and economic sustainability of AI-enabled operations.

Together, these leaders form the executive backbone of the transformation. The organizations most likely to succeed will not be those that deploy the most AI, but those that align leadership, workforce, operating model, and technology around a shared vision of how work will be executed in the future.

Turning that vision into reality requires mobilizing the full execution muscle of the enterprise, combining disciplined delivery, organizational change, and continuous adaptation over a multiyear journey.

Conclusion: A new era of reinvention

Enterprises have reinvented themselves before—through industrialization, lean management, digitization, and cloud computing—but never before have they had to redesign an enterprise where humans are no longer the sole engine of execution, where autonomous systems make consequential decisions at machine speed, and where the boundaries of the firm, the economics of work, and the sources of competitive advantage are all being reshaped simultaneously.

The organizations that lead in this new era will not necessarily be those with access to the best models. They will be those capable of defining the clearest future sources of differentiation, sequencing the transition most effectively, building the operational foundations required to govern intelligent execution at scale, and leading the transformation coherently from the top.

The symbiotic enterprise is not simply a technology shift. It is a fundamental redesign of how enterprises create value, organize work, and compete. And for leaders capable of navigating this transition deliberately, it may become one of the largest sources of competitive advantage of the coming decade.

Christian Jansen is a partner in McKinsey's Hamburg office; **Daniele Chiarella** and **Sébastien Lacroix** are senior partners in the Paris office, where **Guillaume Dagorret** is an associate partner; **Kim Baroudy** is a senior partner in the Copenhagen office; **Lari Hämäläinen** is a senior partner in the Seattle office; **Lieven Van der Veken** and **Stéphane Bout** are senior partners in the Lyon office; and **Ruben Schaubroeck** is a senior partner in the London office.

This article was edited by Larry Kanter, a senior editor in the New York office.

Copyright © 2026 McKinsey & Company. All rights reserved.

June 2026

Images: © Getty Images

Copyright © McKinsey & Company

www.McKinsey.com

 @McKinsey

 @McKinsey