McKinsey & Company

Industrials & Electronics Practice

The AI assembly line: Strategic imperatives for CEOs

AI has proven its potential, but value remains hard to scale. To turn experimentation into sustained performance, CEOs must reimagine their organizations.

by Gianmarco Cilento, Steffen Fuchs, and Varun Marya

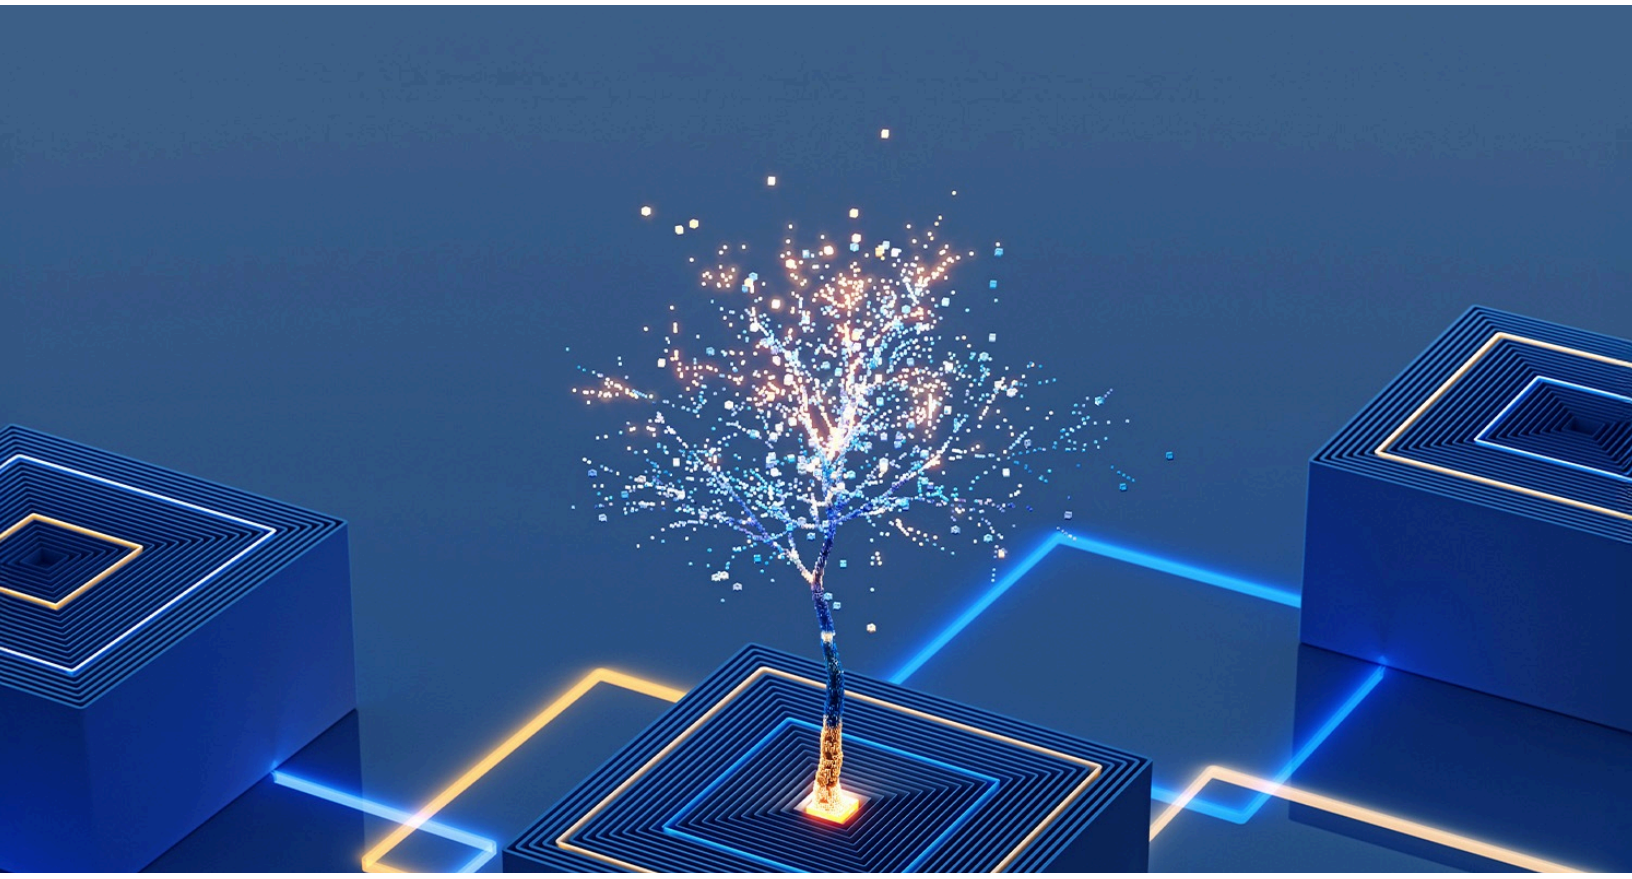

May 2026

In the early 1900s, Henry Ford's assembly line transformed automotive manufacturing. Instead of having skilled workers assemble an entire car, or large portions of it, he divided production into a series of simple repetitive tasks performed by workers at a conveyor belt. The result was not just a manufacturing transformation but a cultural shift and a new operating model. Production time for the Model T fell from more than 12 hours to just 90 minutes, costs dropped sharply, and car ownership moved from a luxury to a mass-market reality. While Ford's transformation happened over 100 years ago, the model he developed still resonates today.

Just as Ford's production line transformed physical labor, agentic AI—systems that can act autonomously rather than just responding to prompts—is now reshaping cognitive work, including engineering design, supply chain planning, and risk assessment. (We will refer to agentic AI simply as "AI" throughout this article.) With AI, companies no longer need to depend solely on the judgment and availability of a small number of experts to make complex decisions or create sophisticated products. Instead, knowledge becomes broadly accessible to anyone with the right AI capabilities, accelerating decision-making, product customization, and other tasks once limited to experts.

Across industries, companies recognize AI's potential and have launched multiple initiatives, often pursuing dozens or even hundreds. Yet the returns remain narrow, with most reporting incremental, function-specific improvements, such as reduced procurement costs or faster invoice processing. Enterprise-wide breakthroughs remain elusive.

Such limited gains are unsurprising. Too often, AI is classified as a technology project and applied to isolated functional domains rather than embraced as a business transformation. Top leaders rarely grapple with the harder questions required to scale impact: How should end-to-end workflows be reimagined with agentic AI? Which decisions should be automated, augmented, or escalated? What happens to roles, spans of control, and performance metrics when cognitive work that once took weeks can be completed in minutes?

Addressing those questions requires the same caliber of leadership that reshaped manufacturing in the early 20th century. The model that Ford developed still represents the best strategy for modern enterprises: CEOs must galvanize the leadership team to reimagine and redesign end-to-end workflows, including roles, governance, and work processes, to embed AI across the enterprise. This total rewiring of the business involves applying a framework that focuses on six themes: strategic road mapping, talent, operating model, tech, data, and adoption and scaling. This article focuses on setting a strategy and the underlying technology for AI transformations.

If companies adhere to this framework, they can create an AI assembly line with multiple agents that complete tasks across functions, increasing both efficiency and productivity. The payoff can be substantial, with early McKinsey research showing that CEO-led digital transformations are 1.5 times more likely to succeed than those led primarily by technology teams.

Building the AI assembly line

Although AI can optimize repetitive tasks such as invoice processing and payroll, the most transformative corporate application involves end-to-end decision-making. Companies have traditionally relied on selected leaders, middle managers, and specialists to evaluate options, but it is difficult to share knowledge across large, geographically dispersed organizations. Human judgment is also subject to error, often because of biases, and experts may disagree. Such friction often triggers additional meetings and reviews, delaying action.

These dynamics create cognitive bottlenecks that limit how fast organizations can learn, adapt, and respond. As product complexity increases and data volumes explode, the gap between available expertise and what is required for decision-making widens. AI's true potential lies not merely in automating tasks but in expanding decision-making capacity to a far broader group and accelerating the process (for instance, providing the documentation required to support a decision in minutes rather than days).

Many companies are incorporating AI into decision-making, but initiatives often stall because the teams in charge focus on technology solutions rather than on revising the organization to support an AI assembly line. While chief technology officers (CTOs) or IT groups may manage critical technology decisions, CEOs must serve as the chief architect of the AI transformation.

Strategy: Reconfiguring the organization

In their strategic plan, CEOs must create a compelling vision that sets solid goals and provides concrete examples of the desired impact—for instance, a 50 percent reduction in product development timelines across the entire organization. Because each AI assembly line is custom designed based on a company's needs, there is no set combination of agents that will be deployed in all instances.

As they rewire their organizations to support AI assembly lines, the following actions may help:

— *Breaking down silos to encourage collaboration.* The greatest divide is typically between the tech groups that design AI solutions and the business side that implements them. Ideally, the new CEO-led model will eliminate many communication gaps and mitigate common issues, such as the development of AI solutions that are technologically sophisticated but deliver little business value.

— *Eliminating unnecessary bureaucracy.* CEOs or designated functional leaders must pinpoint the organization's most consequential decision-making bottlenecks. Constraints may be structural, such as layers of hierarchical approvals and complex functional handoffs, or procedural, including an overreliance on manual processes. Once these friction points are clear, leaders can convene a focused working session to eliminate unnecessary steps while ensuring appropriate oversight of AI agents on the assembly line.

— *Standardizing data and using AI tools in combination.* Leaders should ensure that all AI solutions share the same data and operational framework. When companies have a common data foundation, they can use multiple tools in combination, likely creating more value than would be obtained by using each one separately.

— *Setting new goals and expectations for talent.* AI will necessitate some new roles, such as leaders of agentic squads that ensure the tools create the desired impact. In other cases, CEOs may need to revise performance criteria as AI frees employees to handle more advanced tasks or help human resources develop new training programs for new hires (for instance, short courses on AI processes, rather than six months of onboarding).

Technology: Creating a conveyor belt for intelligence

Agentic AI still needs a strong technology and data foundation, with a focus on reusability to gain scale. Underpinning every AI assembly line is an agentic orchestration layer where data, models, and agents interact seamlessly, eliminating the need for bespoke bridges for every new use case (exhibit). This layer can independently plan, reason, and execute complex, multistep tasks with minimal human intervention. Within the AI assembly line, it acts as the conveyor belt by managing the flow of tasks, decisions, and data across the organization.

The agentic orchestration layer contains functional agents that operate within specific domains and enterprise agents that oversee the end-to-end flow of intelligence. Working in sync, functional and enterprise agents provide and disseminate comprehensive insights across the entire organization. In healthcare, for instance, functional agents might analyze diagnostic data, while enterprise agents ensure that insights are delivered to the right clinician at the right time, enabling real-time optimization of patient care.

By managing routine tasks and synthesizing data, the agentic orchestration layer allows humans to focus on higher-order judgments that require creativity, empathy, and strategic thinking. It has enough intelligence to realize its limitations and flag certain tasks for human intervention. For example, a financial institution's agentic orchestration layer might process 90 percent of loan applications but flag the most complex or ambiguous cases for human review. This hybrid model ensures efficiency while preserving accountability and trust.

As companies scale AI, one of the most consequential choices is whether to build critical tech capabilities in-house or outsource them. Technology vendors often cost less and move faster—a big advantage in a field where tools evolve quickly—but in-house creation makes sense when AI capabilities are central to competitive differentiation. For example, a company might develop a proprietary software program that segments customers into highly specific cohorts and generates tailored sales guidance, drawing on unique internal data. Owning the intellectual property ensures that competitors cannot easily replicate the resulting insights. Companies must also decide whether to outsource certain AI-related services, such as tailoring AI systems to suit a company's internal processes or anonymizing data before model input.

Exhibit

The agentic orchestration layer directs workflows.

Agentic workflow, illustrative

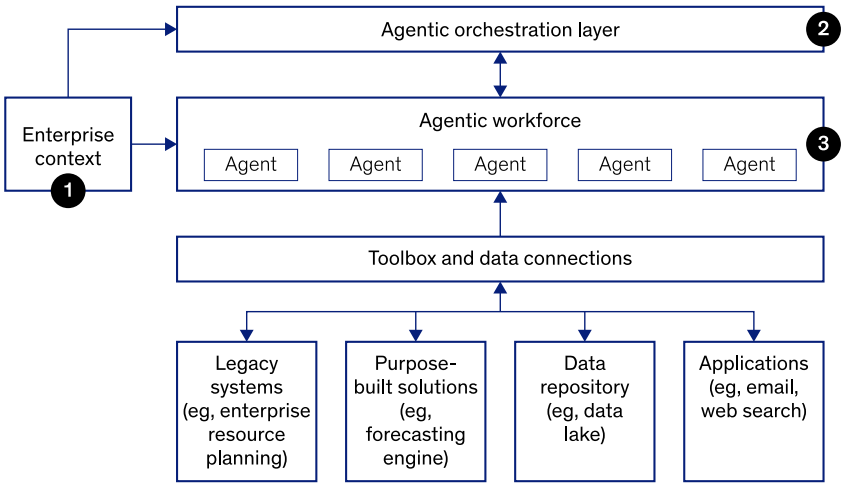

1 Information on the enterprise context is readily available, allowing agents to learn and adapt

2 The agentic orchestration layer routes the right agents to tasks and provides the right context for the task (eg, sales agent must coordinate with the production agent)

3 Toolbox and data connections allow agents to fetch the required data and queries to perform their tasks

McKinsey & Company

Benefits of the AI assembly line

Ford did not win because he built a better Model T; he won because he changed the physics of the unit cost. Before moving to the assembly line, his factory produced 12,000 cars annually at a cost of $825 each. By 1925, after fully industrializing production, factory output soared to two million cars annually at $260 each.

Today, in traditional industrial firms, each incremental increase in product complexity—whether from additional sensors, embedded software, or new regulatory requirements—drives a disproportionate rise in “artisanal” overhead: the specialized expertise required to design, integrate, and validate these offerings. Cross-functional meetings, emails, and manual tests may consume thousands of hours. Industrial companies, such as those in robotics, semiconductors, and automotive, are especially vulnerable to slow, cumbersome development processes because their products are increasingly software defined. In aerospace, for example, software now accounts for roughly 40 percent of product value, up from just 10 percent in 2010.

If the AI assembly line can handle the cognitive load associated with a tenfold increase in product complexity, a company could expand its margins without greatly increasing head count. While a capable employee might juggle ten complex tasks simultaneously, an AI assembly line can manage 16 million concurrent, individualized logic streams. And unlike a factory floor, which has physical limits, an AI assembly line can scale to billions of "virtual workers" instantly to solve a spike in complexity. If an AI system can implement an automotive safety update over the air in 12 seconds, while rivals require 12 months, it would gain a huge lead-time advantage.

Just as the components of the AI assembly line vary, so will the benefits. Consider an organization that wants to redesign its entire customer journey because it is plagued by inefficiencies, including portfolios with redundant offerings and a large number of SKUs, limited margin leakage management, and low use of operational, credit, and payment data to prioritize collections. Rather than dividing up tasks by stage and assigning employees to specific steps, the company could use an AI assembly to handle multiple tasks simultaneously (table). For instance, agents might handle scheduling discovery calls, triage and routing of inquiries, and pricing configuration.

As the cost of cognitive intelligence drops, new opportunities for value creation may surface across industries. Imagine a global bank that can adjust risk models in real time, or a pharmaceutical company that can personalize treatments for every patient or compress the timeline for developing new drugs by years. Services that are considered luxuries, such as personalized wealth management, might become universally available as they become more affordable. Such changes are not simply incremental improvements; they are paradigm shifts that redefine customer expectations, operational efficiency, and competitive dynamics.

While AI is still a nascent technology, it has already demonstrated real benefits. By accelerating decision-making, AI can help companies innovate or respond to market changes more quickly. Research also shows that one early AI mover achieved 20 percent greater profitability than its peers (see sidebar, "Case examples: Accelerating AI's impact").

The proof of AI's benefits may be most apparent in China's automotive sector. By using AI to expedite R&D and software integration, Chinese manufacturers have collapsed the development cycle for new cars to just 24 months. This industrialized efficiency has lowered the barrier t entry so significantly that over 50 new electric-vehicle brands have entered the Chinese market in the past five years. AI has also enabled Chinese automakers to create more sophisticated features at a lower cost.

By producing high-quality, low-cost vehicles, Chinese brands have captured more than 50 percent of their home market for the first time. Meanwhile, the collective market share for non-Chinese joint ventures has plummeted from 64 percent in 2020 to roughly 43 percent today. If legacy OEMs continue to cling to their traditional artisanal engineering processes, they will face a permanent margin squeeze as Chinese OEMs expand globally.

Table

Agentic AI could transform the customer journey.

Key process activities for performance improvement, nonexhaustive

	Traditional	Automation potential with agentic AI
Lead to quote	• Reliance on data enrichment vendors • Rule-based lead scoring • Standardized outbound email sequences	• Autonomous triage and routing • Self-optimizing nurture cadence engine • Autonomous scheduling for discovery calls
Quote to order	• Mature configure-price-quote rule set • Discount governance dashboards • Central clause templates	• Fully generated configuration pricing • Automatic multirole approval routing • One-click quote-to-order conversion flow
Order to delivery	• Kickoff playbook templates • Scripted provisioning pipelines • Migration progress dashboards for tracking	• Self-provisioning multicloud tenancy setup • Continuous telemetry health monitoring • Proactive anomaly remediation workflows
Delivery to cash	• Scheduled enterprise resource planning invoice batches • Structured dunning cadences conducted manually • Collections escalation playbooks	• Closed-loop usage-based billing cycle • Adaptive multichannel dunning cadence • Autonomous credit memo issuance

Case examples: Accelerating AI's impact

McKinsey has assisted with many CEO-led digital transformations, some of which are described in the second edition of *Rewired: How Leading Companies Win with Technology and AI*, which was released in April 2026.

In one recent engagement, McKinsey worked with Toshiba Tec to help retailers unlock value from the billions of data points generated by its point-of-sale terminals. While retailers had long relied on retrospective analyses of promotions and sales trends, real-time personalization and dynamic profit optimization remained out of reach.

To address this opportunity, Toshiba Tec launched Gyainamics, a new software engineering company within its business-building factory. In collaboration with McKinsey and its AI arm, QuantumBlack, the company deployed Nvidia's AI infrastructure and software to power a next-generation analytics platform that could process retail data in real time—a task that had previously required hours.

Compared with traditional approaches, the transformer model, enabled by a graphics processing unit, dramatically expanded the coverage of retailers' long-tail catalogs—to nearly 100 percent, from roughly 8 percent—and improved personalization performance sevenfold. Training times for large data sets fell by more than 80 percent, while inference time for 100,000 customers dropped from hours to under a minute. By dynamically adjusting promotions in real time, retailers increased average transaction value by 5 percent, expanded SKU coverage to 99.9 percent, and boosted long-term customer value by up to 7 percent compared with prior manual segmentation methods.

In another engagement, McKinsey partnered with the US division of a major automotive OEM. In the aftermath of the COVID-19 pandemic, the company struggled to meet customer demand, resulting in extended vehicle wait times and lost sales. Business leaders spearheaded the transformation, working closely with technology teams to develop a demand-forecasting model capable of generating detailed insights into customer preferences.

From the outset, the OEM's solution development team involved subject matter experts (SMEs) to ensure both relevance and adoption. SMEs recommended tool features, reviewed model outputs, and provided continuous feedback. To build trust in the AI, the team also held working sessions to explain how the models learned from data and generated predictions, increasing transparency and confidence in the outputs.

After multiple iterations, the resulting model could analyze approximately 5,000 possible vehicle configurations for the millions of vehicles sold annually. It incorporated predictive signals, including customer behavior patterns, market conditions, and competitor sales data. When forecasting demand for specific configurations at the level of individual dealer markets, the model achieved 85 percent accuracy. This improvement enabled the OEM to shorten sales cycles, reduce purchase incentives, and lower working capital costs by 20 percent, resulting in a $200 million increase in profit in the first year alone.

The OEM launched a related initiative to embed AI into a digital twin that connected demand forecasting with manufacturing. The digital twin integrated inputs from demand models and real-time manufacturing constraints while also accounting for enterprise objectives, such as profitability targets, to optimize production and sales decisions. Where traditional planning models typically provide only 13 weeks of visibility, the digital twin extended the planning horizon to three years. Supply chain planners and manufacturing teams used the platform collaboratively to resolve bottlenecks and set production priorities in real time.

The digital twin operated within an agentic orchestration layer that made insights accessible across the enterprise. Planners could assess competitive dynamics, regulatory considerations, and the status of ongoing initiatives in real time, and run alternative scenarios collaboratively with other teams. Overall, the solution automated approximately 80 percent of core planning activities, improving efficiency and reducing errors.

The AI assembly line will not simply enhance decision-making; it will industrialize the process, using customized agents specifically designed for each company. When the cost of judgment falls and its availability scales, personalization becomes standard, forecasting becomes continuous, and organizations can pivot in near real time. This shift requires more than technology. It demands a CEO-led redesign of workflows, roles, and governance to ensure that the speed of operations matches the speed of intelligence.

Gianmarco Cilento is a partner in McKinsey's Miami office, **Steffen Fuchs** is a senior partner in the Dallas office, and **Varun Marya** is a senior partner in the Bay Area office.

This article was edited by Eileen Hannigan, an executive editor in the Boston office.