McKinsey Quarterly

QuantumBlack, AI by McKinsey

The change agent: Goals, decisions, and implications for CEOs in the agentic age

Companies are feeling agentic AI growing pains. Here's what CEOs can do to move past them and position their companies to succeed.

This article is a collaborative effort by Alex Singla, Alexander Sukharevsky, Lari Hämäläinen, Oana Cheta, Olli Salo, Pallav Jain, Raghav Raghunathan, Sandra Durth, Stéphane Bout, and Vito Di Leo, representing views from QuantumBlack, AI by McKinsey; McKinsey Technology; and McKinsey's People & Organizational Performance and Technology, Media & Telecommunications Practices.

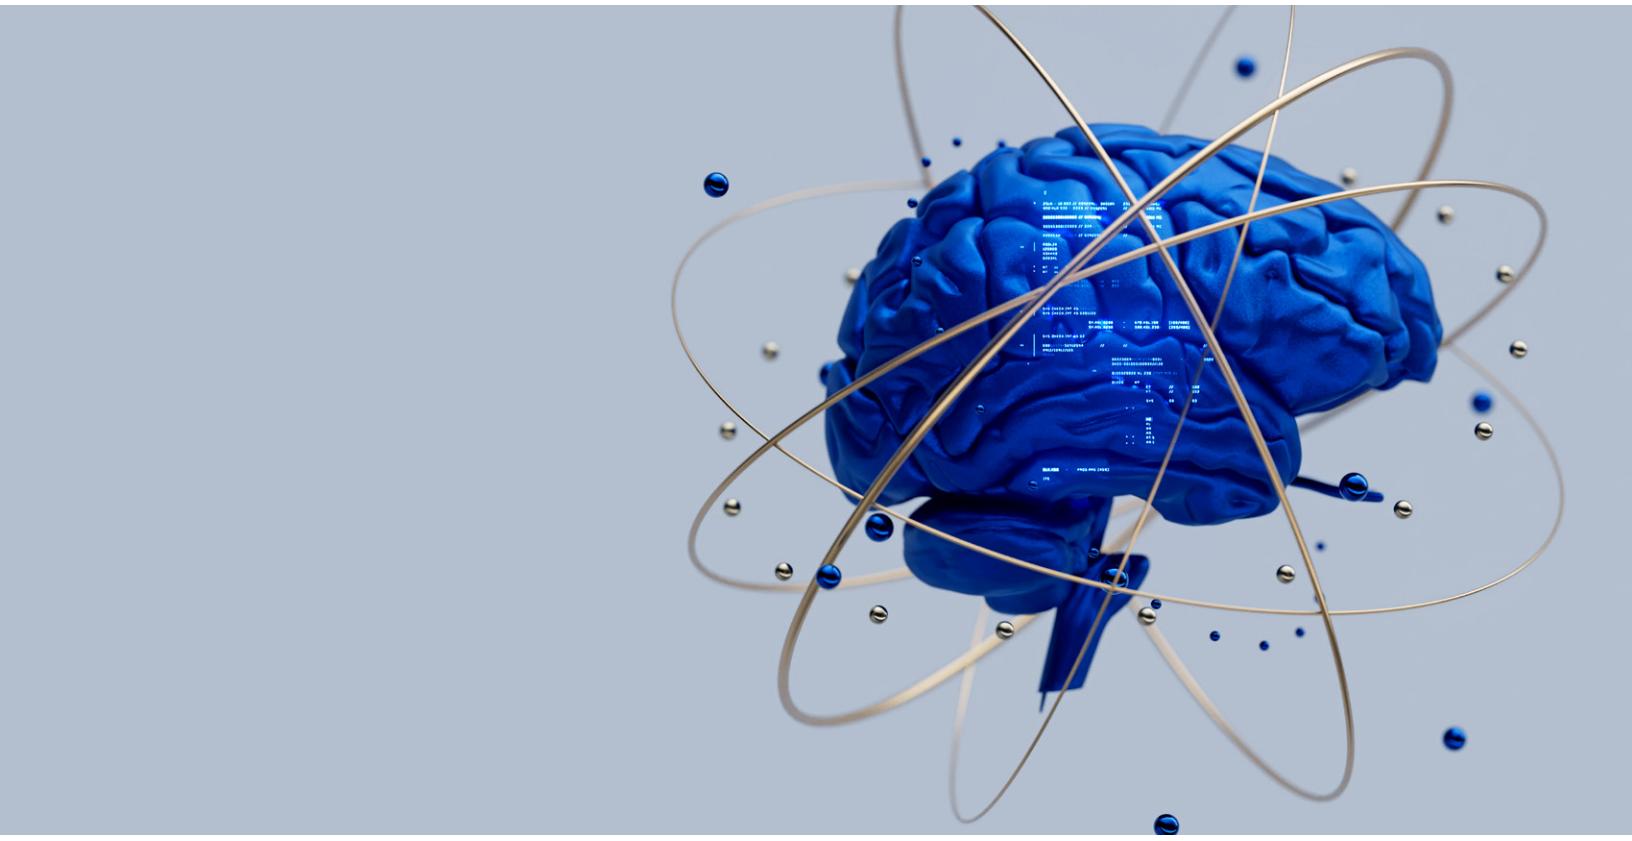

October 2025

Executives are fond of quoting hockey great Wayne Gretzky, who is credited with saying: "I skate to where the puck is going to be, not where it has been." This is sound business advice at one level. But that puck is moving a whole lot faster than it used to as agentic AI rapidly evolves.

The call to move faster may seem tone deaf as CEOs and their senior teams struggle to see bottom-line value from early gen AI investments. Developing and scaling gen AI use cases have proven frustratingly challenging. Some executives remain unconvinced that AI agents will have a significant impact—at least in the short term—and have stepped back from their investments.[1]

As CEOs navigate the uncertainty, it is worth acknowledging both the pace and potential scope of the change that is happening. AI agents—software systems built with gen AI that have the ability to plan, act, remember, and learn to achieve predefined outcomes autonomously—are evolving quickly and, as they mature, could completely change how companies are run and how they generate value (see sidebar "Key trends shaping gen AI and agents"). In fact, this "trough of disillusionment" period, as John Lovelock of Gartner recently called it,[2] is an opportunity for executives to jump ahead of their competitors.

How CEOs manage this change will determine how well they can capture the benefits. Although AI agents are in their infancy, early lessons and experiences highlight four mindsets and actions that can position CEOs to prosper:

— *Reimagine what's possible.* Much of the thinking around agentic AI today is still focused on automating basic tasks or augmenting knowledge. The real win, however, will come from much bolder aspirations of rearchitected workflows and organizations built around agent-first systems.

— *Act with urgency and start the learning.* The rapid rate of improvement of gen AI agents means that a wait-and-see approach is potentially a high-risk move (see sidebar "Breakthroughs in gen AI and agents"). Early practical learnings are invaluable in quickly building a competitive advantage as the technology matures.

— *Tackle scale and long-term competitiveness issues now.* Critical decisions around technology, trust, governance, what to buy versus what to build, capabilities, and talent are important to drive a wider transformation. While you experiment, start forming your strategy and developing scaling capabilities as soon as possible since execution will take longer than expected due to talent scarcity and organizational complexity.

— *Turn everyone into an agent leader.* As agents and agentic systems take over more of the executional work, everyone in the organization will need to develop agent leadership and supervision skills. The executive team especially needs to role model and champion learning and the evolution of their work habits.

While much is still unknown, building a business for the agentic age will require a fundamental rewiring of how the business operates, innovates, and protects sources of value creation. This article, however, will focus on a few of the most important elements an enterprise CEO should address related to value, scale, and talent. We will outline what a hypothetical two-year agentic journey might look like, what kinds of decisions CEOs should consider, and what the big implications could be for how companies operate.

[1] "Gartner predicts over 40% of agentic AI projects will be canceled by end of 2027," Gartner, June 25, 2025.
[2] "Welcome to the AI trough of disillusionment," *Economist*, May 21, 2025.

Key trends shaping gen AI and agents

AI agents are becoming more human-like in the kinds of tasks they can do and the way people interact with them. These features democratize AI in a way prior technologies haven't and underscore agents' potential to affect a broad set of activities. The increasing possibilities of gen AI—the foundational capability that enables AI agents—are fueled by four mutually reinforcing trends:

— *An acceleration in innovation pace.* Only two new-frontier large language models (LLMs) were announced in 2020;[1] by 2025, the number is in the dozens, even hundreds, depending on counting methodologies.[2] Similarly, the number of new large-scale AI models has grown by 167 percent per year since 2020.[3] The length of tasks AI agents can do (with at least a 50 percent success rate) has been doubling every seven months.[4] At the time of writing, it has been reported that Anthropic's Claude Opus 4 can complete almost as much work as a human can in a day, while a multiagent system outperformed a single-agent Claude Opus 4 by more than 90 percent.[5]

— *Large growth in spend and investments.* The compute used to train state-of-the-art models has been growing roughly four to five times per year.[6] The top three hyperscalers collectively plan to invest more than $250 billion in 2025 on AI and data centers,[7] and in 2023, businesses spent about $15 billion on gen AI solutions, representing roughly 2 percent of the global enterprise software market.[8]

— *Sharp gains in model training and inferencing efficiency.* Breakthroughs in architecture and optimization have driven training costs down significantly for a given capability. The inference costs for ChatGPT 3.5 dropped more than 280 times between November 2022 and October 2024.[9] The cost per million input tokens, for example, has decreased about ten times, from about $36.00 in March 2023 to about $3.50 in August 2024.[10] For some models, the cost is less than $1.00.[11]

— *Breakthroughs in model and system capabilities.* New reasoning models deploy "test time compute" thinking during inference ("system-2 thinking"); standardized tool-calling interfaces, such as Anthropic's Model Context Protocol (MCP), let models invoke enterprise APIs safely; vastly larger and more precise short- and long-term memory structures improve both the recall breadth and precision; multiagent orchestration frameworks (for example, LangGraph, AutoGen) enable specialized agents to delegate, monitor, and reconcile their work; and early agent-to-agent protocols (for example, A2A, created by Google and recently donated to the Linux Foundation to maintain as an open-source project) point to a future where agents autonomously discover peers, negotiate roles, and execute workflows.

1 *The 2024 AI Index report*, Stanford University, Human-Centered Artificial Intelligence, 2024.
2 For example, see Anthony Cardillo, "Best 44 large language models (LLMs) in 2025," Exploding Topics, August 28, 2025.
3 Based on Epoch AI data, as cited in Mary Meeker, Jay Simons, Daegwon Chae, and Alexander Krey, *Trends—Artificial Intelligence*, Bond, May 2025.
4 "Measuring AI ability to complete long tasks," Metr, March 19, 2025.
5 Michael Nuñez, "Anthropic overtakes OpenAI: Claude Opus 4 codes seven hours nonstop, sets record SWE-Bench score and reshapes enterprise AI," VB, May 22, 2025; "How we built our multi-agent research system," Anthropic, June 13, 2025.
6 Jaime Seville and Edu Roldán, "Training compute of frontier AI models grows by 4–5x per year," Epoch AI, May 28, 2024.
7 Dan Romanoff, "Is AI investment poised for growth? Top picks and promising applications for 2025," Morningstar, May 27, 2025.
8 Jeremy Schneider, Tejas Shah, and Joshan Cherian Abraham, "Navigating the generative AI disruption in software," McKinsey, June 5, 2024.
9 *The 2025 AI Index report*, Stanford University, Human-Centered Artificial Intelligence, 2025.
10 "Large language models: Performance vs. cost," MentorCruise, January 13, 2025.
11 "API pricing," OpenAI, accessed on September 19, 2025.

Are agents worth it?

Claims about the value of AI agents permeate the internet. But since the technology is still so new, those claims are hard to verify.

Early implementations, however, suggest there is significant value at stake. Our experience with modernizing technology estates indicates that harnessing AI agents can accelerate timelines 40 to 50 percent and reduce costs more than 40 percent while also improving the quality of the outputs.[3]

3 "AI for IT modernization: Faster, cheaper, better," McKinsey, December 2, 2024.

In another case, a leading universal bank faced a major challenge in delivering a high volume of IT projects to drive business outcomes while managing significant technical debt and a shortage of skilled developers. Starting with a small team of three engineers, it built a tech modernization agent factory with 100 agents supervised by just five humans. These agents, under human oversight, executed the entire modernization life cycle—from reverse engineering to designing and building new applications—cutting time and labor costs by more than 50 percent.

Our experience indicates that the initial uses of AI agents to support people and automate tasks can drive 3 to 5 percent in annual productivity improvements at the company level. As teams of AI agents become capable of carrying out more complex workflows, growth could increase by as much as 10 percent or more.

Know your agents: From 'agentic labor' to 'agentic engine'

Executives still tend to have fixed and limited notions of what agents are and what they can do. This confusion can make it difficult for them to understand what decisions related to risk, investment, resource allocation, and change are needed.

While it's tempting to think of agents in human terms, a more objective approach is to consider them as software systems that can do a spectrum of increasingly complex tasks (exhibit). In our experience, this approach sharpens the thinking around what sorts of organizational changes are necessary.

Agentic labor: Agents as tools to help with existing work
Agentic tools can contribute to existing work by facilitating basic tasks undertaken by individuals and by automating workflows.

Individual augmentation. These tools help automate, speed up, or improve tasks that people typically do. Many of the tasks are familiar—drafting research notes, summarizing meetings, generating code, conducting research, or proposing contract clauses.

These tools will likely become, and already are, in areas like programming, a "cost of doing business" similar to employees using email and spreadsheets.

Studies show 20 to 30 percent higher personal throughput and sometimes much higher numbers in single-task areas. However, broad horizontal deployments of agentic tooling across the business rarely translate into significant business impact. Furthermore, usage tends to tail off, and retention drops significantly beyond lead users for many tools.

Driving broad adoption of personal support agents requires familiar change management investments, such as embedding the tools in standard operating procedures, integrating the expected outputs and usage monitoring into performance management systems, giving employees proper training to use the tools, and communicating and role modeling the benefits. At the same time, leadership will need to determine how to capture the increased productivity given its spread across several small tasks, often done through budgeting and large-scale efficiency efforts.

Task and workflow automation. The second category focuses on automating existing processes, workflows, and tasks in the organization. An agentic execution layer essentially sits on top of existing processes and systems (with small changes to them).

Exhibit

Agentic systems can work on increasingly complex tasks depending on how well companies can overcome some key constraints.

Agentic systems based on increasing task complexity

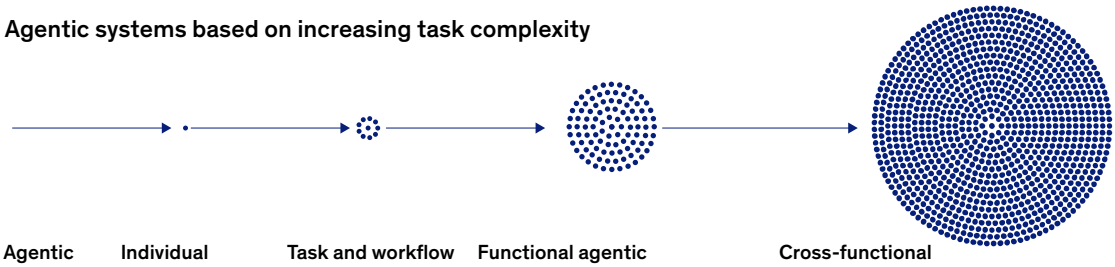

Agentic system	Individual augmentation	Task and workflow automation	Functional agentic workflows	Cross-functional agentic systems
Definition	Agents help automate and improve basic tasks	Automation of existing, low-complexity workflows	Agentic teams working on redesigned workflows	Agent-driven systems work on complex workflows across functions, with high-level decisioning capabilities
Primary benefits	Individual productivity	Cost efficiency, speed, compliance	Improvements in efficiency, speed, and customer satisfaction; greater scale; revenue uplift	Faster production, lower cost per transaction, higher value per customer, accelerated cross-functional processing
Key constraints	High learning curve, continued adoption, significant change management burden	Execution capacity to automate existing tasks at volume and at high quality	Capacity to reimagine processes, engineering advances to manage agent teams at scale	Organizational and operating model redesign

McKinsey & Company

Major technology players are introducing first-generation agentic products while an explosion of new companies is bringing solutions to many functional domains (for example, customer care, financial reporting and monitoring, programming, product development, and procurement).

In our studies, early deployments have delivered 20 to 40 percent faster cycle times or lower handling costs for repetitive, transactional work. In contact centers, certain types of calls (for example, transactional handling of balance checks and address changes) are almost fully automated. Embedding agentic tools into workflows and establishing a continuous improvement approach is a core enabler, but simply giving tools to users won't work.

The impediment to value, however, is that these more domain-specific use cases operate in isolation and rely on other systems and significant human intervention to perform. Furthermore, while model capability improves, companies struggle with execution capacity to automate existing tasks at volume and at high quality.

Agentic engine: Agent-native workflows and operating model

Emerging agentic systems, driven by breakthroughs allowing for teams of agents to work together, offer the most promising opportunities to generate major value. Capturing this value, however, requires rethinking and redesigning workflows to be agent-first, either within a function (for example, customer care) or across them (for example, lead to order).

Functional agentic workflows. In this case, domain-specific workflows (for example, financial planning and reporting) are redesigned to take advantage of teams of AI agents and agentic processes. That means rethinking task order, merging tasks, accessing new data sources, and developing new processes, such as early sensing and resolution of issues before they emerge. Agent-native systems can eliminate the frequent handoffs and fragmented activities that hamper many current processes because teams of agents are orchestrated to seamlessly operate.

Specialist vendors in horizontal and vertical software spaces are building and implementing full-stack, agent-native applications for areas such as customer care, finance, supply chain planning, and software modernization. Deployed correctly, these systems cut end-to-end cycle time, improve resolution times, and drive up customer satisfaction. For call center performance, for example, the estimated impact could be an automated handling of 60 to 80 percent of incoming requests with a comparable or better customer satisfaction score than for current systems (for example, interactive voice response plus first-line support).[4]

Such systems will require a combination of engineering (for example, integrating probabilistic models with more classical, deterministic software) and domain expertise to both build the multiagent systems and redesign the related organization and operating models with sufficient human oversight. It will be critical to instill in these agents governance rules (for example, access rights, decision rights, and quality gates) for targeted workflows (for example, procure to pay, vendor contracting, supplier communication, and policy management) to ensure that supervising humans aren't quickly overwhelmed.

Cross-functional agentic systems. These agent-first systems work on complex workflows (such as end-to-end customer journeys) across functions and have high-level decisioning capabilities. Consider, for example, 24/7 field service operations agents that dispatch technicians, reschedule visits, and order parts autonomously; an insurance team that adjudicates claims; a mortgage that is approved and underwritten in seconds; or a financial cycle with agents handling everything from annual planning to monthly reporting.

These agentic systems can create multidimensional value in terms of, for example, faster time to market, lower cost per transaction, faster issue resolution, and increased revenue through better offer targeting. Early pilots using existing technologies saw up to 70 to 80 percent reductions in cost per transaction for certain labor-heavy processes.

At this level, the key constraints are tied to organizational and operating-model issues. The CEO and board will need to be intimately involved to rearchitect the operating model, including leadership and team responsibilities that historically have lived in siloed corporate functions. Incremental changes won't work; this level of transformation requires a clear break with past practices.

Decisions to make along an agentic journey—and some big implications to consider

To help CEOs visualize the journey and surface some critical decisions along the way, we've laid out a high-level, hypothetical two- to three-year road map. It highlights certain markers to aim for and some of the key decisions that require the CEO to be actively involved. (See sidebar "How start-ups built around AI agents are reshaping business" for outtakes from an interview

[4] "Seizing the agentic AI advantage," McKinsey, June 13, 2025.

with Magnus Grimeland, the cofounder and CEO of the tech venture capital and incubator company Antler.)

The goals and timelines presented in the road map are aggressive, and we acknowledge that much will change over the two-year timeline. In our experience, however, it is crucial for CEOs to set bold aspirations and goals to motivate the business and move with urgency.

Years one and two: Setting the course and creating momentum
Initial goals in year one should include building understanding, creating momentum, and developing the foundations so AI agents can work at scale. The focus should be on driving down operating costs of existing activities in a targeted set of functions and operations (consider a goal of up to 10 percent efficiency gains, for example). First and foremost, though, this phase is about breaking through the inertia, moving with purpose, and learning.

How start-ups built around AI agents are reshaping business

Magnus Grimeland, cofounder and CEO of venture capital firm Antler, talked with McKinsey Senior Partner Lari Hämäläinen about how start-ups are taking advantage of AI agents. The following is an edited excerpt of that conversation.

Lari Hämäläinen: Do you see agentic-first models disrupting business models?

Magnus Grimeland: Absolutely. Some companies are just building more effective cost structures, such as automating the back office or big parts of the development cycle. But some companies are being much more disruptive by completely removing humans from core parts of the business model. One logistics company we backed is looking to completely rearchitect and replace the back-office functions that optimize the entire supply chain around AI agents. Another one is looking to replace process engineers from factories. There are tremendous opportunities across the value chain.

Lari Hämäläinen: Among start-ups that are building around AI agents and AI in general, how are they building up businesses differently?

Magnus Grimeland: That's a broad question, but a few themes have emerged. One is that we're seeing companies grow faster with significantly fewer people than we've ever seen before. Two, they typically have much flatter structures, with leadership being much more active in product development and sales, and the whole team is composed of builders. In the past, when starting up a logistics company, for example, you'd look for people with ten to 15 years of experience in the industry. With AI, it will be more and more about finding the best technology talent and engineers.

And another is that start-ups are really keenly attuned to the speed of change. That means they think in terms of how to continuously upskill and reskill their people, and they constantly scan the market for breakthroughs and potential competitors. AI is developing at a speed that is much faster than previous cycles, and you quickly become irrelevant if you're not on top of new capabilities.

We also see this speed and flexibility playing out, as costs have come down significantly. In the past, when you built an e-commerce site, you needed a tech backbone, which needed a lot of planning and long product development cycles. With gen-AI-first companies, the product cycle is much faster, and there is the ability for people to switch systems. You can build on top of one infrastructure, then switch over to another relatively easily if it's a better option. This means it is much easier, cheaper, and faster to build and experiment.

Lari Hämäläinen: Where is the value going to come from in the AI agent world?

Magnus Grimeland: What value means and where it comes from has to switch. Many services today are essentially adding an "AI agent" button onto their existing services and then charging 30 to 40 percent more for it. That model won't work very well for long. All companies have to think about rebuilding AI first. That means rethinking how to engage with the user, how to access information, and how to serve customers differently. This will eventually be about replacing services rather than just creating incremental revenue streams.

In years one and two, look for the following agentic business markers:

— *Agent "fluency" grows quickly.* Being able to use AI agents productively is a requirement for all workers. While the value to the business is low, building up this capability among employees is the "cost of doing business." The goal should be to have more than 25 to 50 percent of employees working with enterprise agents and AI tools regularly. All employees should be interrogating data through "chats" with agents rather than just reading reports.

— *Agents are automating a broad range of existing processes with first-generation tools.* This includes critical processes such as financial filings and broader document authoring, approvals in existing processes, and so on. Concrete benefits, like significantly faster lead times and lower transaction costs, should be clearly evident. For example, in targeted cases, such as correcting simple, well-structured data quality issues, agentic AI can resolve 90 to 95 percent of the issues.

— *First agentic software systems are integrated into key systems.* Key systems are changing their interfaces toward prompt-based querying rather than static commands. Agents are in place in key systems to automatically create insights, execute tasks, and coordinate activities. Automated capabilities such as planning or information gathering are increasingly the norm, and supporting systems (such as reporting) are changing fast.

— *A front-runner team launches a lighthouse to reimagine a complete end-to-end process.* The team designs a 24-month target state vision for a complete process (for example, order to cash, record to report, automated loan acceptance and provisioning) and starts by releasing a series of minimum viable products to test and expand capabilities. The goals should be bold. For an order-to-cash process, for example, the goal could be to automate greater than 70 percent of transactions across all channels.

— *The demand for certain roles may lessen as productivity increases.* AI agents have reliably and efficiently taken over simple coding tasks, reducing the burden on some existing roles. The latest coding agents (especially front-end code execution), for example, significantly increase productivity by as much as 50 to 100 percent.

Enabling the business to hit these markers will require CEOs to address some core business areas:

Architecting the transformation for value. Too much of the current focus is on individual productivity; while useful, this isn't where the greatest pools of value are. CEOs need to aspire to transformational value, which will come from rearchitecting and redesigning entire workflows with agents.

The CEO will need to ensure teams move from working on isolated use cases to focusing on cross-functional priority workflows. This will necessitate an organizational shift away from siloed AI teams toward cross-functional agentic teams that include AI, data, IT, technology, and functional experts from relevant domains.

As these teams work, it will be crucial to put a premium on the learning that is generated. This means ensuring that enterprise-wide learnings are centrally captured and reused across the organization. It also means codifying an agent-first workflow redesign playbook, including ROI criteria, multiagent orchestration patterns, tech and data integration best practices, controls and evaluations, and when or when not to apply agents.

To lead this effort, organizations will need a central team ("agentic factory") responsible for identifying the workflows, managing the redesign, and scaling the redesign.

Scaling the transformation. In 2022, Amazon's founder, Jeff Bezos, mandated that developers' code include APIs that could be exposed to third parties. This mandate was Amazon's way to enable scale. CEOs need their own way to establish—and enforce—the same design principles for AI agents. By developing composable agents, rather than monolithic flows, elements can be reused and reconstituted for other tasks. This should be an organizational mandate, not a tech preference.

The agent factory is critical for industrializing this scaling capability. It is made up of dedicated teams that build and deploy agents while developing systems and standards to ensure they scale. Examples include developing a reusable library of blueprints for operational processes, risk checks and guardrails, systematic technical measures to handle agent performance evaluation, and standardized top-down KPIs.

This factory needs to have a strong governance role to avoid uncontrolled agent sprawl across the organization and to check agent compliance with company and country regulations.

Key shifts to enable the agent factory to operate effectively include:

— *Tech and data.* To ensure coordinated agent development and management, CEOs will need to prioritize the creation of appropriate architecture. In some cases, this can be provided through agentic service providers. As a counterbalance, remaining vendor agnostic could avoid lock-in and allow for combining custom and off-the-shelf agents across multiple tech platforms.

  Companies with a higher investment appetite should consider an agentic AI mesh architecture, a set of patterns, practices, and principles that orchestrate workloads. Similarly, agents will only be able to scale if they can access curated, dynamic, and structured data. This effort is not trivial, and it requires senior leadership to prioritize strong data foundations for agents.

— *Platforming.* Companies will have hundreds, if not thousands, of agents, which can quickly become a management nightmare. CEOs should prioritize the development or acquisition of a platform that automates agent management and observability. This should include protocols to manage the cost and performance of agents, onboarding and retiring agents, access rights (for example, determining which agents can access data based on local regulations), privacy rights, and security.

— *Trust.* People need to understand how AI agents operate in order to trust them. Without that trust, adoption will suffer. To help develop that trust from the beginning, CEOs should prioritize systems to help users understand how agents make decisions and where biases are.

Shifting the talent and operating model. Many organizations have rushed to develop training on how to use gen AI and AI agents as "tools." This has been useful in providing a baseline of skills. More significant, however, is that everyone will need to be capable of developing and supervising agents. Agentic will not be most effective as a "tool on top" of regular work; it needs to be built into how every person works.

That shift will require people leaders to start focusing on building a new human–agent hybrid operating model, which should include embedding new skills such as how to build and apply agents effectively, how to train them, how to set tasks for them, how to track and correct their work, and how to string together a series of agents to perform more complex tasks. This change in operating model—essentially how work is done in the future—is a key element of the sustainability of the rewired processes and a critical element of change that needs to be implemented as part of deploying the agentic solution.

Part of embedding the change is that performance management and reviews will need to adjust to focus on "agent management." Supporting this change will require changing KPIs so they become tied to how well people work with agents (for example, number of tasks completed with agents, quality of output, et cetera).

Years two and three: Scaling across the organization

The second year of the agent journey is all about scaling early learnings, with a shift from agent activity measures to P&L impact. This is the period when companies should start to "agentify" key customer journeys and workflows, not just integrate agents into existing processes and flows.

In years two and three, look for the following agentic business markers:

— *The first agentic "lighthouse" hits the tipping point.* The first reimagined customer journey lighthouse, established in year one, gets to scale with greater than 90 percent automation. Agentic interactions have high performance and satisfaction rates for standard processes and effective escalation to people for exceptions.

— *Agentic automation scales in key value streams.* Agentic automation is driven in 90 percent of key value streams across the company. Agentic systems become the standard option not only as a system upgrade but also as a new operating model.

— *Traditional operating models built around static functions and linear handoffs begin to erode.* At least five priority customer journeys are primarily managed by teams of AI agents with human oversight and cut across traditional functions (for example, a single AI agent team analyzes customers, develops and sends personalized marketing offers, completes the sale, and manages customer service). Similarly, agent-to-agent interactions crossing the functional boundaries of the enterprise (including with customer representative agents or sourcing representative agents) become common.

— *Employees are becoming "agent team leaders."* Use of AI agents is becoming the new normal, with adoption north of 75 percent in all key functions. Almost every professional has at least one AI agent and often three to five agents working "for" them, with a few advanced users juggling dozens of them.

— *The ratio of agents to full-time equivalents (FTEs) has significantly shifted.* Agent reliability for basic and medium tasks is high, though human oversight and review are important. Functional modernization happens in a number of domains. FTE needs in the software development life cycle, for example, could drop by as much as 30 to 40 percent while the workload for your financial planning and reporting organization is reduced by 75 percent.

As a company's agentic capabilities mature, the CEO's focus will need to shift to more organizational and operational changes.

Reimagining for value and architecting the transformation. Initial programs to cut costs and learn are important, but CEOs should focus on productivity and value improvements well north of 50 percent and have an appetite for questioning everything the business does. At this stage, CEOs should meaningfully reimagine the art of the possible for creating value and managing the business.

Specifically, they should be challenging what elements of the business model are defensible in an agentic world and what agentic innovations and experiences could create a source of sustainable differentiation. A sharp and clear perspective on future sources of differentiation is critical to inform choices to build versus buy versus partner, as well as to inform the data, intellectual property, and competitive advantage the organizations should protect.

As part of this accelerated reimagining program, CEOs who have the mindset to use agents to test, learn, and adapt at pace could gain important strategic benefits. Agents operating in digital-twin environments, for example, will be able to test broad ranges of processes, applications, and offers at scale, cheaply and quickly. Successful tests can be introduced almost immediately as agents affect the changes directly, with change management costs dropping significantly for highly automated workflows.

As part of this evolution, CEOs will need to rethink budgets, business models, and investments in line with shifts driven by agentic economics. Some of those dynamics will include significant changes in the ratio of employees to agents, with investments shifting meaningfully from labor to technology. CEOs, working with their CFOs and chief human resource officers (CHROs), will need to develop a financial and workforce capability model to better budget for, track, and allocate technology and people.

Scaling the transformation. As end-to-end agentic workflows and customer journeys cross functions, traditional organizational models will become less meaningful. Holding on to those constructs and building agentic systems that mirror them will choke off scale.

Instead, it will be critical to reorient the business to outcomes and design the agentic organization around them. One model could be to organize human–agentic teams around value streams, such as customer onboarding or new-customer sales. The transition to this model will require significant and often difficult trade-off decisions around governance (such as who will manage and be responsible for the value streams) as well as investments in mechanisms to track agent performance.

These areas will require CEOs to ensure that the promulgation of an agent-first organization doesn't lead to chaos. For example, if different departments or teams develop competing agents, it could create confusion and open up the business to new risks.

Talent and operating model. McKinsey research postulates that up to 30 percent of current worked hours could be automated by 2030, while the potential is much larger. This transformation will fundamentally alter the ratio of humans to technology in organizations and create new models of agent–agent and human–agent workplace dynamics.

Working closely with the CHRO, the CEO will need to determine what the agentic workforce will look like, how the skill-building capability will need to evolve as agents mature, and how to redeploy freed-up capacity.

Given the pace of change, HR systems will need to be significantly recalibrated to regularly review which profiles need to be recast, created, or removed and to oversee the increasing pace of up- and reskilling. People allocation, career structures, performance management, incentives, and reporting systems will need to be redesigned to enable this agentic operating system.

This will include, for example, embedding agent management KPIs into performance reviews and defining new role archetypes, such as "agent orchestrator" and "agent trainer." Career progression should be linked to proficiency in these roles and the ability to manage teams of agents and people.

As agents proliferate across the enterprise, they have the potential to profoundly affect every part of the business. This will require every CEO and board to operate at two speeds by driving transformation in the short term while thinking through longer-term consequences. (See sidebar "The CEO agenda: An overview" for more on how CEOs can facilitate agentic change in the organization).

The following questions can help CEOs and their boards start to think through the long-term implications of an agentic business:

— How will agentic AI affect my business model, challenge our existing sources of differentiation, and create new ones? What strategies can we implement to defend our markets and expand into new ones?

— How could agentic AI disrupt the core business and relationships with customers, suppliers, and partners (for example, agent-produced disintermediation)?

— How do I prepare to shape and manage the "agentic workforce" while maintaining the values and culture of the company?

— How should we manage the transition to a hybrid human–agent operating model as workflows seamlessly cross traditional functional boundaries?

— What is my talent strategy, and how should it inform the ratio of in-house talent to outsourced capabilities? What is the optimal balance of open-source, multivendor, and single-platform technology options to provide maximum flexibility and operational value?

— What should my transformation and investment road map look like to both meet near-term business goals and establish the right foundations to develop and scale transformational change?

The CEO agenda: An overview

To drive agentic change in their organizations, CEOs will need to think through key steps at each of the five stages of the transformation (exhibit).

Exhibit

CEOs need to consider key steps along the agentic AI transformation journey.

Key steps for each stage of the transformation

Develop the North Star and road map

Clarify and define the long-term CEO-led vision and strategy, including a concrete implementation road map, business case, and the related objectives and milestones for the journey

Design foundations

Design the "from-to" changes in your organization, talent base, enterprise technology, and data to execute the vision; ramp up your initial agentic factory to drive implementation

Drive the 2-speed change

First, incubate bold, high-velocity pilots to automate work (bottom up); second, initiate two to three agentic lighthouse workflows to start the change (top down); and third, establish "fast lane" governance to unbottleneck early efforts and enable scaling

Commit and personally role model the change

Lead your lighthouses personally—take interest; commit the top team to drive the AI transformation as role models; establish the incentives and support for the organization; communicate to set expectations and highlight early successes

Invest in AI-ready leaders

Acquire, invest in, upskill, and elevate your leaders to drive the agentic change in the company

McKinsey & Company

The "pace of change" may be one of the most overused phrases in business, but agentic AI has injected it with new relevance—and urgency. The shroud of uncertainty that still veils AI agents shouldn't hold leaders back from moving quickly. In fact, decisive and thoughtful action is the only way for leaders to strip away the uncertainty and develop agentic businesses that open new opportunities for productivity and growth.

Alex Singla is a senior partner in McKinsey's Chicago office, where **Oana Cheta** is a partner; **Alexander Sukharevsky** is a senior partner in the London office; **Lari Hämäläinen** is a senior partner in the Seattle office; **Olli Salo** is a partner in the Helsinki office; **Pallav Jain** is a senior partner in the Atlanta office; **Raghav Raghunathan** is a partner in the Bay Area office; **Sandra Durth** is a partner in the Cologne office; **Stéphane Bout** is a senior partner in the Lyon office; and **Vito Di Leo** is a partner in the Zurich office.

This article was edited by Barr Seitz, an editorial director in the New York office.

Designed by McKinsey Global Publishing