McKinsey & Company

Operations Practice

How COOs maximize operational impact from gen AI and agentic AI

By defining the right operating structure, data governance model, and change management approach, COOs can help their companies make the most of their AI investments.

by Curt Mueller, Darryl Piasecki, Marie El Hoyek, and Oana Cheta

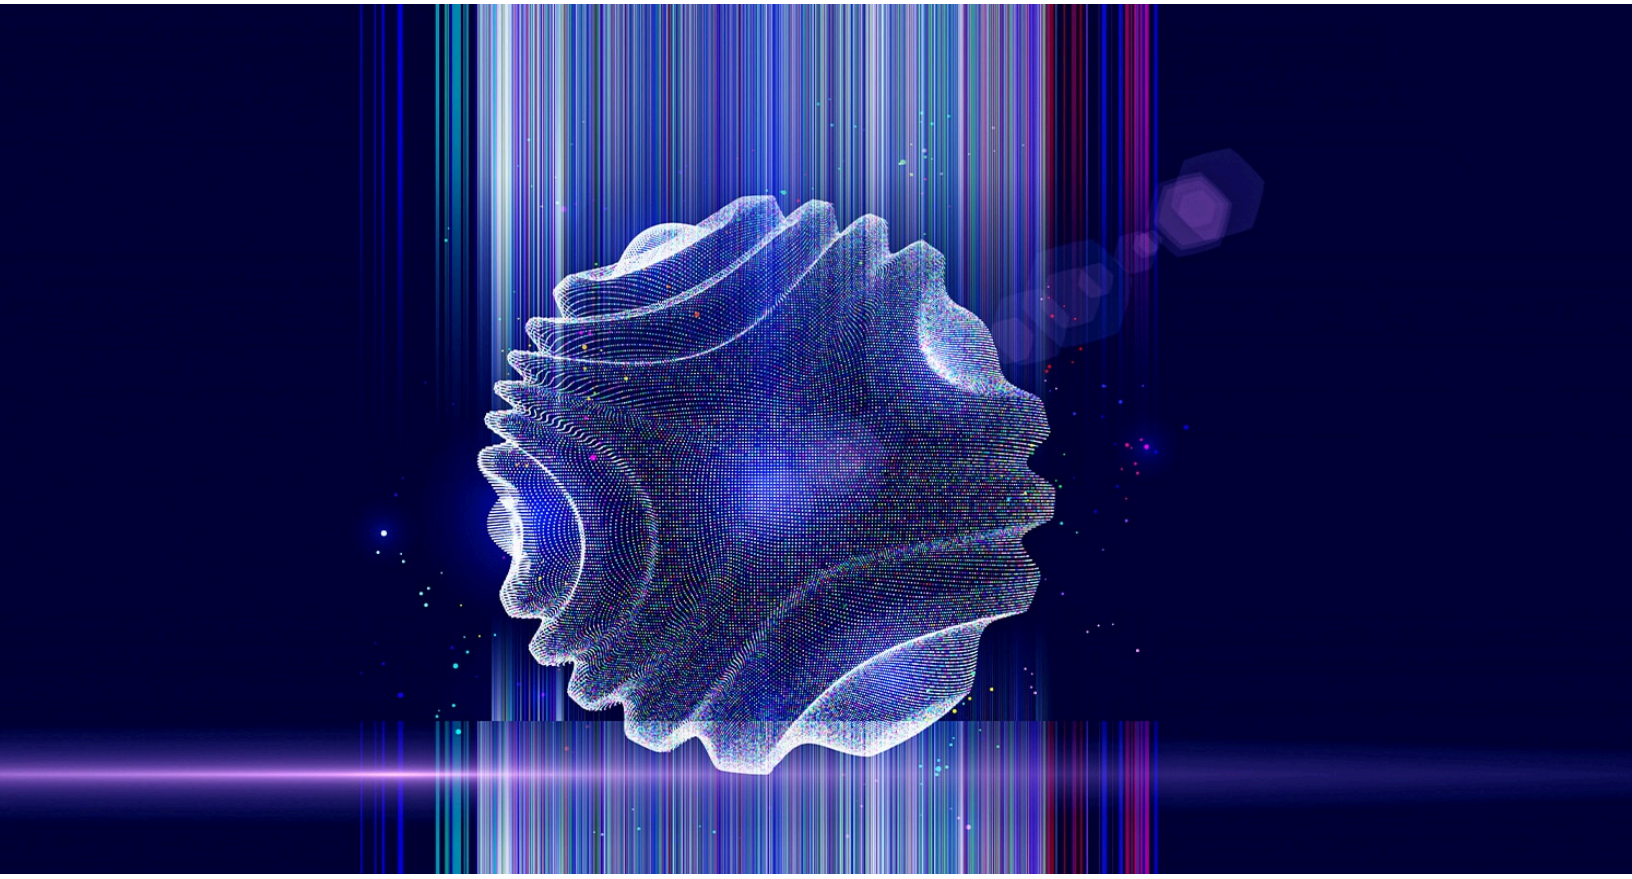

March 2025

Better, faster, easier, cheaper: That's the promise of gen AI. For at least some companies, it's becoming the reality as well, as leaders find new ways for gen AI—and the increasingly capable agents it enables—to automate, augment, and accelerate work across virtually every function. Early adopters are using gen AI to help strengthen supplier negotiations in procurement and improve quality control in equipment maintenance (see sidebar "Gen AI's potential across operations"). One digital marketing platform is even using gen AI to manage "long tail" sales accounts that were previously too labor-intensive to serve, for an annual revenue gain of more than $30 million.

Gen AI's potential across operations

McKinsey estimates that over the long term, gen AI could yield $4.4 trillion in productivity growth potential.[1] That's on top of "traditional" or "analytical" AI, which relies on structured data to solve discrete analytic tasks—such as predictive analytics for optimizing equipment maintenance. Gen AI's deep learning AI models are already helping companies achieve performance breakthroughs across the operations value chain, especially by finding new opportunities to break internal silos. Multiagent systems can achieve even more (exhibit).

— *R&D.* A European automotive supplier initially used gen AI to analyze its existing code base, then deployed targeted gen AI applications to automate new-code generation, quality assurance, and documentation. Together, the efforts across these three subfunctions yielded a *time savings of 20 to 30 percent while improving code quality*.

— *Manufacturing.* A global heavy-industry manufacturer developed a gen-AI-based "technician maintenance copilot" to help frontline workers quickly find nonintuitive root causes for equipment failures. *Frontline workers can now complete tasks that previously required automation experts from other parts of the operation to solve,* reducing maintenance workload by 40 percent while increasing overall equipment effectiveness (OEE) by 3 percent.

— *Procurement.* At a multinational resources company, upskilled procurement teams worked with the legal department to use gen AI models to review thousands of pages of contracts covering more than $1 billion in spend. The effort revealed more than $15 million in quick-win savings, such as through improved compliance with payment terms—with substantial further savings to come from negotiation insights.

— *Production line.* Future production lines could take these ideas even further by deploying gen AI agents that, underpinned by the most advanced large language models, are increasingly able to take autonomous action to complete complex tasks. Picture a shift turnover in a factory or hospital—a moment notoriously prone to miscommunication, delay, and rework. As the incoming shift supervisor checks in, a gen-AI-based agent could aggregate the prior shift's performance, identify ongoing problems and their resolution status, and even develop remediation plans for the supervisor to review. Once the supervisor gives feedback, the agents can either iterate and refine the plan or start implementation.

Exhibit

With greater reasoning capabilities than single agents, multiagent systems can complete complex workflows quickly.

Benefits of multiagent systems over single agents

Task specialization
Individual agents are optimized for specific tasks, leading to . . .

- higher performance
- greater speed
- fewer hallucinations
- diverse perspectives without confusion in prompts

Independence
Agents execute uninstructed analyses autonomously by . . .

- dividing tasks between agents
- coordinating decision-making
- resolving conflicts

Modularity
Multiagent systems are modular and thereby . . .

- scalable, enabling addition of new agents with new capabilities
- efficient, allowing reuse of smaller components

Robustness
Decentralized control and computational resources distributed across agents reduce likelihood of system failure due to . . .

- resource limitations
- performance bottlenecks
- critical failures

Multiagent systems are more expensive, as they need more large-language-model calls

McKinsey & Company

[1] Hannah Mayer, Lareina Yee, Michael Chui, and Roger Roberts, Superagency in the workplace: Empowering people to unlock AI's full potential, McKinsey, January 28, 2025; The economic potential of generative AI: The next productivity frontier, McKinsey, June 14, 2023.

Yet, as encouraging as these results are, there's still much to do. In a recent McKinsey survey of 118 US C-suite executives, only 19 percent said that gen AI increased their company's revenue by more than 5 percent. It's a similar picture elsewhere: In mid-2024, just 17 percent of organizations worldwide said that they derive more than 10 percent of EBIT from gen AI.

Not surprisingly, about half of senior executives in that survey describe their organization's development and release of gen AI tools as too slow—despite the fact that three-quarters also say they have at least a draft of their gen AI strategy. Only 12 percent of these organizations have been able to find revenue-generating use cases for gen AI. And while the ultimate goal for these organizations is to achieve gen AI maturity, with gen AI fundamentally changing how work gets done, a mere 1 percent of executives say their organization has reached that point.

That's where the COO plays a critical role, as illustrated by several recent success stories where gen AI and gen-AI-based agents have helped redefine how a company creates value. Specifically, the COO can help build enterprise capabilities for gen AI-based rewiring in three ways: First, they can define the company's operating structure for gen AI, identifying the highest-potential domains for gen AI deployment and building the capabilities needed to scale the technology effectively across the enterprise.

Second, they can shape the organization's data governance, addressing the complex challenges associated with extracting and structuring data from legacy operating systems and minimizing risks associated with inaccuracy. Third, and most important for sustaining gen AI's advantages over time, they can oversee change management initiatives so that people learn, use, and improve the tools and processes gen AI enables.

Getting these three factors right takes work, not just in operations but also in collaboration with other leaders, such as the chief information officer (CIO). But it's how companies' investments in gen AI can pay off: by reshaping how work gets done every day.

Ensuring gen AI creates real business value

Getting gen AI wrong could be costly: not just in wasted investment but also in missed opportunities. Companies that move quickly are already securing major advantages, increasing the stakes.

Sensing gen AI's possibilities, senior leaders of a European equipment maker with more than €10 billion in revenue wanted to avoid one of gen AI's most common pitfalls: fragmentation in development. Too often, individual functions and business units design gen AI tools that optimize their own tasks but fail at the enterprise level—such as a production-scheduling tool that raises factory output higher than the logistics department can absorb.

The company's COO recognized that, in facing the future of operations, he and his team needed much more than a list of potential gen AI use cases. They needed to rethink the entire operating model to see how this new automation could transform people's work.

Rethinking operating structures for gen AI

To an even greater degree than seen in earlier waves of technology-based transformation, gen AI touches virtually every part of a business organization. This expanded scope for coordination makes the operating structure particularly important to get right, both to identify the highest-potential gen AI opportunities at the enterprise level and to see them through to fruition. From the beginning, therefore, the European equipment manufacturer brought together the COO, CIO, chief technology officer (CTO), and heads of manufacturing, procurement, supply chain, and quality control, along with business unit leads responsible for marketing and sales, to undertake a gen-AI-prompted reassessment of its operating assumptions.

Centralization. The equipment manufacturer's leaders recognized that sustaining this sort of centralized approach would be essential, especially as the organization developed foundational capabilities in fields ranging from platform architecture to risk and ethics. At least initially, a center of excellence (COE) or "factory" model, with a steering committee providing executive leadership and an operating committee overseeing day-to-day work, would help keep stakeholders collaborating to generate lasting value (Exhibit 1).

Exhibit 1

A gen AI center of excellence leads capability building and works closely with operational and executive stakeholders.

Example responsibilities for a gen AI center of excellence (illustrative)

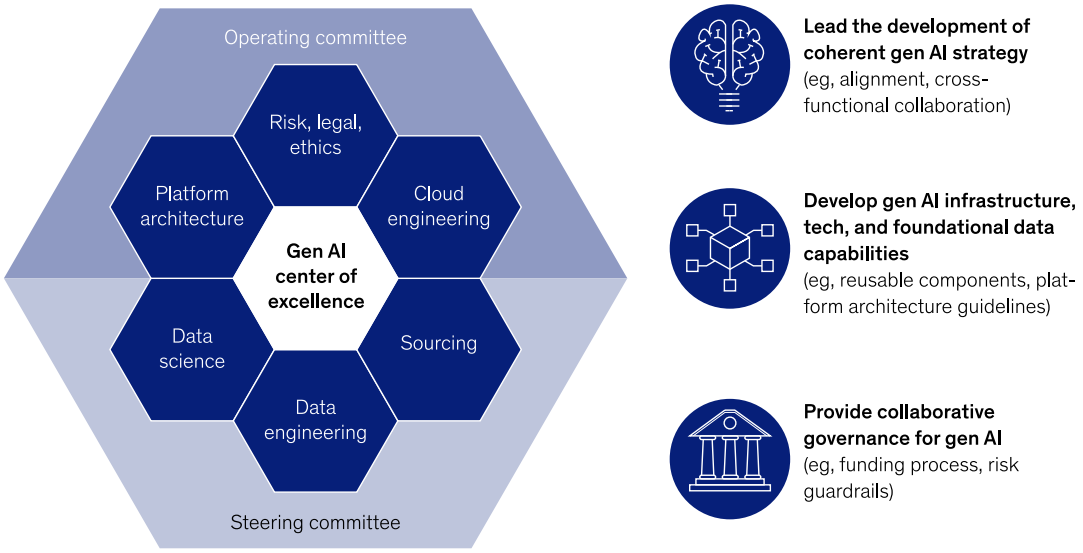

McKinsey & Company

The next question is where the COE should sit within the larger organization. Under the most centralized approach, the COE directs gen AI strategy and reports directly to the CEO, operating in parallel to the business units (Exhibit 2). By enforcing enterprise-wide standards and minimizing the risks of duplication and resource waste, this option is often the most practical one at the very earliest stages of gen AI exploration. For the equipment manufacturer, following this model has so far yielded a prioritized road map of relevant use cases for €300 million in EBITDA improvement.

Exhibit 2

Scaling with a highly centralized operating model enables rapid capability building while optimizing gen AI resources.

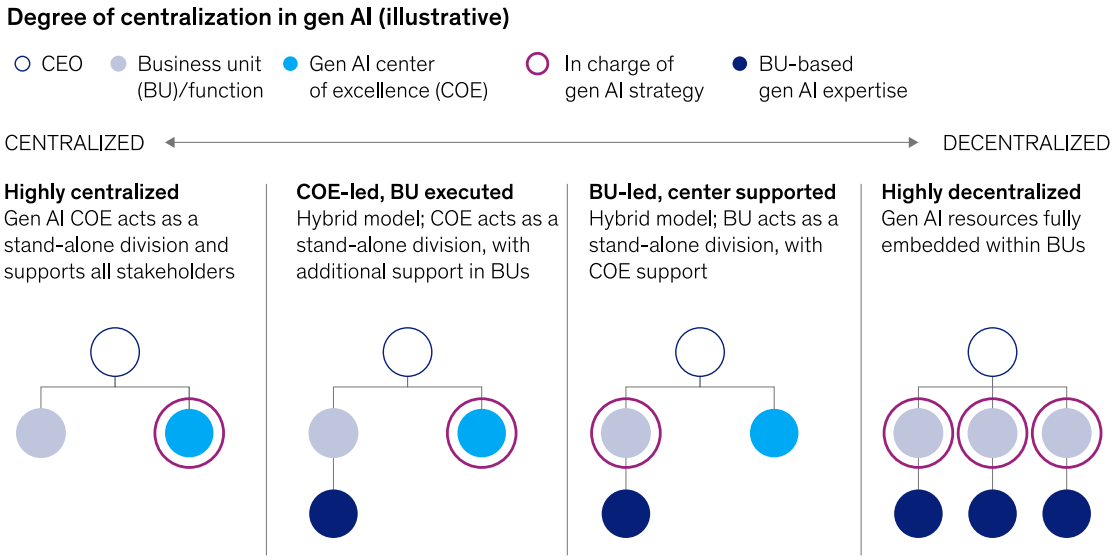

McKinsey & Company

As the company builds more confidence, it could evolve toward one of two middle alternatives in which the business units develop their own gen AI capabilities. In some instances, the COE takes the lead and the business unit executes, while in others, the business unit takes the lead with support from the COE. Only a few organizations have fully decentralized their gen AI function and left it to the business units to run.

Identifying domains. Developing a clear structure helps organizations find the right balance in designing gen-AI-based solutions that are large enough to achieve meaningful end-to-end impact yet small enough to be achievable within a reasonable time frame. Thinking in terms of domains can push gen AI past the "pilot purgatory" stage, in which organizations spend time and resources and incur opportunity costs on developing gen AI tools that have little effect beyond saving workers a few minutes a day.

Most important is to start by assessing the strategic fit for gen AI, with an expansive view of the art of the possible so that the solution can have a lasting effect. For example, a finance function might start by identifying a pain point—such as analysts being overloaded with simple requests from other managers that would take days to answer. An initial response might be to create a gen AI chatbot that would allow anyone in the company to directly query finance data on their own.

This addresses the initial problem by enabling faster query resolution and freeing up analysts for higher-value work. But a deeper examination would seek the root causes for the frequent queries, and whether a more sophisticated gen AI tool—perhaps an agent or a set of agents—could start to produce certain analyses automatically when certain scenarios occur.

This sort of thinking implies another major question: Can we keep gen AI from destroying value? Automating financial analyses for internal purposes, such as to find lessons relevant to new product launches, tends to be substantially less risky than automating analyses for compiling into quarterly securities reporting. And that leads to a final question: Is gen AI the right solution? For some reporting, simple and (comparatively) inexpensive analytic AI may be completely adequate.

Data governance

Centralization helps operations leaders deal with what 70 percent of gen AI high performers reported as a challenge: managing data (Exhibit 3). With gen AI, the accuracy, availability, and usability of operational data become even more important, yet old challenges persist. A global materials company provides a typical example, with teams in different functions each developing their own unique information about the same products. The R&D department's data focused on safety issues; the application engineering team developed tailored customer solutions; commercialization owned the product descriptions; and customer support assembled a set of highly specific product details to answer user queries. With no single source of truth, conflicts naturally arose in the underlying data, which gen AI models struggled to parse.

To resolve the issue, the company is now following a centralized data management system that harmonizes data from different sources, eliminating discrepancies and ensuring that all teams have access to the same accurate information. Crucially, the system emphasizes human oversight to maintain high data quality and reliability, especially for AI-generated answers. A robust governance structure further validates and regularly updates data.

Exhibit 3

For gen AI high performers, data management is the top challenge in capturing value from the technology.

Elements that have posed challenges in capturing value from generative AI,[1] % of respondents

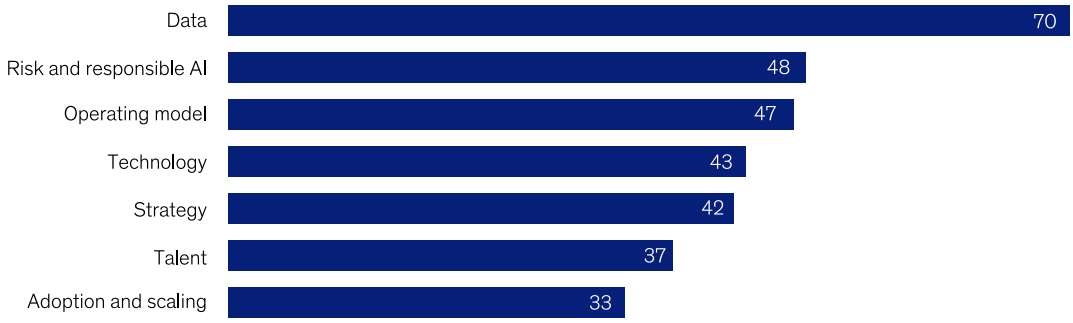

Note: Figures do not sum to 100%, because respondents could choose multiple answer options. Respondents who said "don't know/not applicable" are not shown.
[1]Gen AI high performers are respondents who said that at least 11% of their organizations' 2023 EBIT was attributable to gen AI use. For respondents at gen AI high performers, n = 46.
Source: McKinsey Global Survey on AI, 1,363 participants at all levels of the organization, Feb 22–Mar 5, 2024

McKinsey & Company

Change management

As with earlier waves of digital innovation, gen-AI-based transformations are less about the technology itself and more about rethinking how humans work. If anything, gen AI's potential to enhance creativity and innovation makes change management even more central, particularly as its impact depends on integrating human and gen AI capabilities.

It's a tall order. A gen AI transformation must not only account for the complexities of an evolving technology landscape while yielding clear business results but also address risk concerns (see sidebar "Mitigating risk"), overcome skill gaps, and foster innovation and adaptability. And gen AI itself must keep improving, with AI agents subject to their own performance management systems.

Mitigating risk

McKinsey research has identified eight categories of risk associated with deploying gen AI, which companies can systematically assess as they adopt gen AI across their organizations (exhibit).

- *data privacy and quality* due to unauthorized use or disclosure of personal, sensitive, or other legally protected data, or use of incomplete or inaccurate data in model training
- *impaired fairness* due to bias in training data, model performance, or misrepresentation of model outputs
- *intellectual property (IP) infringement* due to unauthorized use of proprietary content or inadvertent "leakage" of IP into the public domain
- *malicious use* due to heightened potential for bad actors to create harmful AI-generated content (such as deepfakes, hate speech, or phishing scams)
- *poor performance or explainability* due to model limitations that result in inaccurate, outdated, or completely fabricated output ("hallucinations")
- *security threats* due to AI system vulnerabilities that could allow for bypass of existing filters or manipulation of open-source models
- *strategic missteps* due to noncompliance with standards, regulations, or social norms, resulting in reputational or societal risk
- *third-party risks* due to misunderstanding where and how outside actors may be using gen AI, creating potential unknown exposures from, for example, use of proprietary data in public models

Exhibit

Organizations that deploy generative AI use cases can create a heat map ranking the potential severity of various categories of risk.

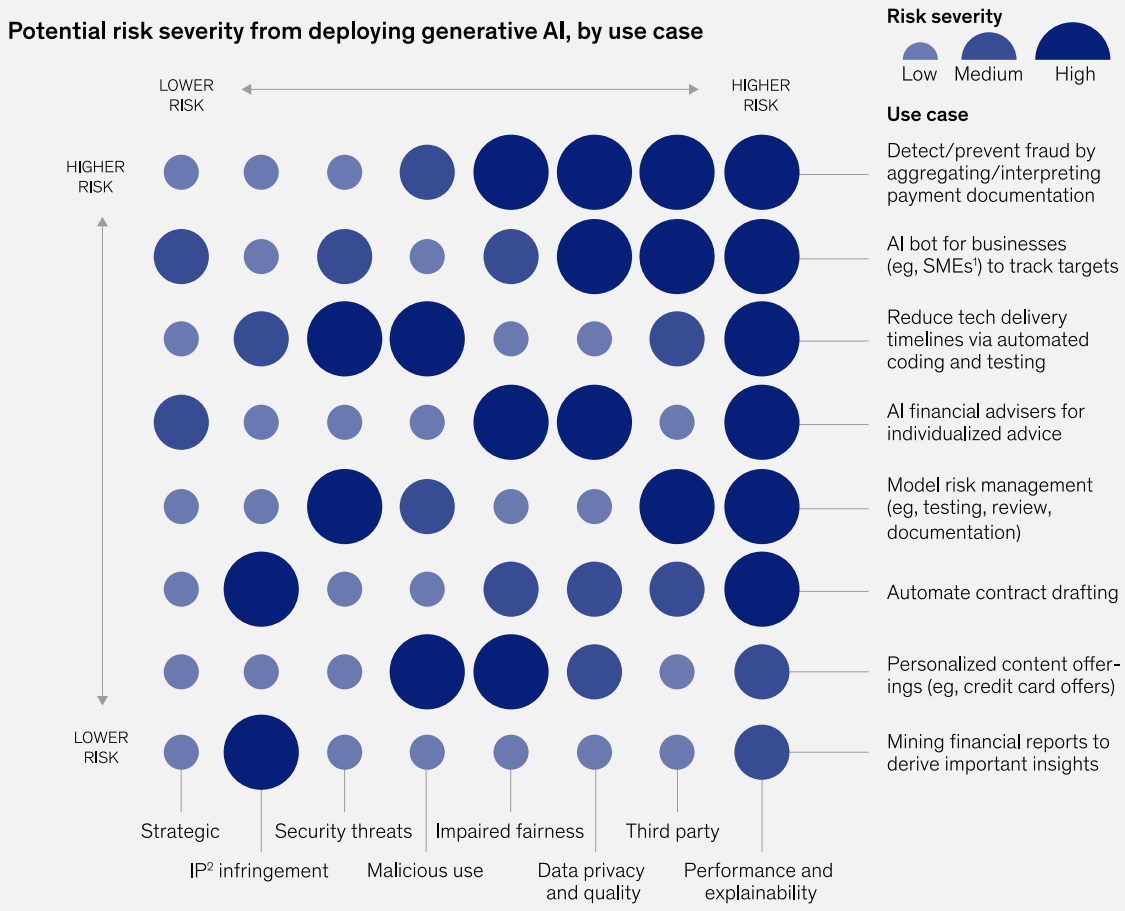

[1]Small and medium-size enterprises.
[2]Intellectual property.

McKinsey & Company

Setting a bold aspiration for enterprise-wide impact. These obstacles are all too familiar to the typical COO, who is charged with leading the continuous-improvement efforts that sit at the core of next-generation operational excellence. They were the starting point for a tech industry COO who recognized gen AI's potential to break long-standing operational logjams—and understood that success would depend on how well people embraced gen AI solutions.

The tech company's work with gen AI started by tackling one of its thorniest cross-functional problems, where complex coordination led to frequent delays in generating highly tailored statements of work that outlined the details of the technology services each client would buy. Assembling a statement of work required the relationship manager to collect input from experts in internal functions ranging from finance and legal to data security, as well as from the delivery managers and solution architects leading the day-to-day work—and the client, too. Rework and errors were a fact of life, slowing response times to such a degree that relationship managers missed deadlines for important requests for proposals.

To build a tool that could generate statements of work for more than a dozen product lines, the company needed to scale quickly. The answer for this organization was to centralize. Leaders created a single working group comprising three main teams: one for engineering, one for business and data requirements, and one for change management.

The three teams collaborated extensively, particularly in reimagining workflows that would take full advantage of gen AI's efficiencies. Previously, for example, creating a statement of work involved elaborate rounds of requirements gathering, feasibility analysis, and risk assessment—inevitably generating rework as later reviews identified issues that affected earlier decisions. By analyzing thousands of earlier statements, the new tool developed templates that highlight the most frequent potential problems up front. Specialist experts in legal, compliance, or related functions can instead focus their efforts on problems that don't have a clear precedent.

Increasing employees' confidence in a gen AI solution. The change management team's involvement proved crucial not only in building the tool but also in ensuring uptake once it was deployed. Following the core principles of the influence model, leaders ensured that each product line had its own dedicated change champion, who served as an intermediary between users and the working group to develop and adapt statement-of-work templates that would meet user needs. The change champion would then help communicate with users and build their skills both in using the tool and in improving its capabilities.

The ultimate result is a templatized statement of work that replaces hundreds of document variations, each taking days to produce, with just five that now require only hours to build. This has eliminated thousands of hours of repetitive labor, freeing experienced employees to focus more on high-value work.

Strengthening COO–CIO collaboration

These examples illustrate how using AI to rethink a stream of value can yield much more

improvement than simply automating a few tasks. It also requires a much closer integration between the COO and CIO, whose traditional incentives have often been in tension.

COOs charged with modernizing complex, legacy operations have often found off-the-shelf IT solutions to be a difficult fit at best. Yet the cost and complexity of bespoke technology can create substantial burdens for the IT function and the CIO. Some of the friction has dissipated as newer technologies, such as edge computing and standardized industrial communications protocols, have taken hold—along with modular IT architecture and more flexible development practices. But there's more to be done.

AI's short innovation cycles and high resource needs have raised the pressure for technology investments to yield their projected returns on schedule, if not sooner. When COOs and CIOs collaborate more effectively, troves of data can become usable insights for revamping operations and creating entirely new sources of value.

The technology company shows how this collaboration can produce results. The COO of the business took the lead in identifying the transformation opportunity and developing it so that it met operational requirements. The CIO's involvement expanded the vision of what was possible, such as by finding new opportunities to adapt enterprise-wide gen AI investments for the specific data needs of creating statements of work. Along the way, the CIO's team became more agile in working with the operations team so that the entire project could meet milestones.

COOs already know that dozens of narrow gen AI use cases are unlikely to add up to lasting operational improvement. Instead, gen AI's potential comes from how it helps leaders rethink entire value chains. This is at the heart of the COO's role, and its future.

Curt Mueller is a senior partner in McKinsey's Chicago office, where **Oana Cheta** is a partner; **Darryl Piasecki** is a partner in the St. Louis office; and **Marie El Hoyek** is a partner in the London office.

The authors wish to thank Josselyn Frankiewicz Cavnar and Melissa Knelly for their contributions to this article.

This article was edited by Christian Johnson, an executive editor in the Washington, DC, office.