
07 DECEMBER 2023

Do you have an “early days” generative AI strategy?

Organizations at the forefront of generative AI adoption address six key priorities to set the stage for success.

by Scott Likens and Nicole Wakefield

Scott Likens

is PwC's Global AI and Innovation Technology Leader. He is a principal with PwC US.

Nicole Wakefield

is PwC's Global Advisory Clients and Markets Leader. She is a partner with PwC UK.

In a recent conversation about generative AI with one of our colleagues, the CIO of a major healthcare company laid out a wide range of issues that concerned her: risk protocols, use case development, cybersecurity, ethics and bias, training and development, and many more. After a few minutes, our colleague asked the client to take a step back: “How clear are you on what you are trying to accomplish, and why? In other words, do you have a strategy?” These questions stopped the CIO, leading her to call a series of meetings with key leaders, and ultimately the board, to create a sharper set of objectives. What emerged was a group of priorities that collectively formed what might be termed an “early days” AI strategy.

Early days, because—let’s face it—that’s exactly where we are with generative AI. It was only in November 2022 that the consumer release of ChatGPT captured the world’s imagination. Since then, organizations have been struggling to keep up with the pace and potential they see in this new, general-purpose technology application. Some organizations are doing better than others, and it’s not too soon to start taking stock of early leaders that are leveraging generative AI to capture value and pull ahead. Across industries, we’re seeing these leaders tackling a number of critical priorities:

- They’re navigating tensions between the need for prudence and risk mitigation, and the importance of moving quickly to grab emerging opportunities.
- They’re aligning their new generative AI strategy with their existing digital and AI strategies, building on these foundations to guide their thinking

rather than starting from scratch.

- They're thinking big—encouraging experimentation across their organizations, with a focus on identifying use cases that can scale.
- Rather than simply looking for ways to improve productivity, they're looking strategically at their options for putting productivity gains to use.
- Relatedly, they're considering impacts on workers, roles, and skills-building, determining how best to both prepare employees to take advantage of the new tools available and include employees in shaping the company's generative AI journey.
- They've realized that with such a potentially disruptive technology, teaming up and collaborating with their ecosystems can be a truly transformative route to a radical rethink of their value chains and business models.

In many cases, these priorities are emergent rather than planned, which is appropriate for this stage of the generative AI adoption cycle. Leaders and organizations are learning as they go.

Priority 1: Manage the AI risk/reward tug-of-war

There's a fascinating parallel between the excitement and anxiety generated by AI in the global business environment writ large, and in individual organizations. At the same time that surging market capitalizations for early AI leaders are providing financial evidence of the opportunity investors and markets see in generative AI, a number of experts in the field are voicing existential angst about the potentially significant unintended consequences that could emerge as the reach of AI grows. Similarly, in many companies we know, there's a tug-of-war going on between the executives and managers seeking to rapidly tap the potential of generative AI for competitive advantage and the technical, legal, and other leaders striving to mitigate potential risks. Although such tension, when managed effectively, can be healthy, we've also seen the opposite—disagreement, leading in some cases to paralysis and in others to carelessness, with large potential costs.

Achieving healthy tension often starts with a framework for **adopting AI responsibly**. At PwC, we developed such an approach several years ago, and we

continue evolving it with the changing nature of AI opportunities and risks. Practical safeguards and guidelines help organizations move forward faster, and with more confidence. Open-minded, agile leadership also is critical: risk-minded leaders deliver better, faster guidance as they internalize the momentous significance of the generative AI revolution. Opportunity-seekers are well-served by spending time immersing themselves in what can go wrong to avoid costly mistakes. And both groups need a healthy dose of appreciation for the priorities and concerns of the other.

One company we know recognized it needed to validate, root out bias, and ensure fairness in the output of a suite of AI applications and data models that was designed to generate customer and market insights. Given the complexity and novelty of this technology and its reliance on training data, the only internal team with the expertise needed to test and validate these models was the same team that had built them, which the company saw as an unacceptable conflict of interest. The near-term result was stasis.

Another company made more rapid progress, in no small part because of early, board-level emphasis on the need for enterprise-wide consistency, risk-appetite alignment, approvals, and transparency with respect to generative AI. This intervention led to the creation of a cross-functional leadership team tasked with thinking through what responsible AI meant for them and what it required. The result was a set of policies designed to address that gap, which included a core set of ethical AI principles; a framework and governance model for responsible AI aligned to the enterprise strategy; ethical foundations for the technical robustness, compliance, and human-centricity of AI; and governance controls and an execution road map for embedding AI into operational processes.

For this company, in short, addressing risk head-on helped maintain momentum, rather than hold it back.

Priority 2: Align your generative AI strategy with your digital strategy (and vice versa)

If you're anything like most leaders we know, you've been striving to digitally transform your organization for a while, and you still have some distance to go.

The rapid improvement and growing accessibility of generative AI capabilities has significant implications for these digital efforts. Generative AI's primary output is digital, after all—digital data, assets, and analytic insights, whose impact is greatest when applied to and used in combination with existing digital tools, tasks, environments, workflows, and datasets. If you can align your generative AI strategy with your overall digital approach, the benefits can be enormous. On the other hand, it's also easy, given the excitement around generative AI and its distributed nature, for experimental efforts to germinate that are disconnected from broader efforts to accelerate digital value creation.

To understand the opportunity, consider the experience of a global consumer packaged goods company that recently began crafting a strategy to deploy generative AI in its customer service operations. Such emphasis has been common among companies. The chatbot-style interface of ChatGPT and other generative AI tools naturally lends itself to customer service applications. And it often harmonizes with existing strategies to digitize, personalize, and automate customer service. In this company's case, the generative AI model fills out service tickets so people don't have to, while providing easy Q&A access to data from reams of documents on the company's immense line of products and services. That all helps service representatives route requests and answer customer questions, boosting both productivity and employee satisfaction.

As the initiative took hold, leaders at the company began wondering whether generative AI could connect with other processes they had been working to digitize, such as procurement, accounts payable, finance, compliance, HR, and supply chain management. It turned out that similar generative AI models, with refinement and tailoring for specific business processes, could fill out forms, as well as provide Q&A access to data and insights in a wide range of functions. The resulting gains, in total, dwarfed those associated with customer service, and were possible only because the company had come up for air and connected its digital strategy and its generative AI strategy. In this case, the alternative would have been a foregone opportunity to turbocharge existing digital efforts. In the extreme, siloed digitization and generative AI efforts might even work at cross-purposes. Given how much companies have already invested in

digitization, and the significance of generative AI's potential, there's no substitute for the hard work of bringing the two together.

A fringe benefit of connecting digital strategies and AI strategies is that the former typically have worked through policy issues such as data security and the use of third-party tools, resulting in clear lines of accountability and decision-making approaches. Such clarity can help mitigate a challenge we've seen in some companies, which is the existence of disconnects between risk and legal functions, which tend to advise caution, and more innovation-oriented parts of businesses. This can lead to mixed messages and disputes over who has the final say in choices about how to leverage generative AI, which can frustrate everyone, cause deteriorating cross-functional relations, and slow down deployment progress. These disconnects are easily avoided, though. At another financial services company we know that was seeking to exploit generative AI in the HR function, the CHRO, the CIO, and the CISO came together quickly to assess the new opportunities against the company's existing data, tech, and cybersecurity policies, providing helpful guidance that maintained momentum.

Priority 3: Experiment with an eye for scaling

The C-suite colleagues at that financial services company also helped extend early experimentation energy from the HR department to the company as a whole. Scaling like this is critical for companies hoping to reap the full benefits of generative AI, and it's challenging for at least two reasons. First, the diversity of potential applications for generative AI often gives rise to a wide range of pilot efforts, which are important for recognizing potential value, but which may lead to a "the whole is less than the sum of the parts" phenomenon. Second, senior leadership engagement is critical for true scaling, because it often requires cross-cutting strategic and organizational perspectives.

Experimentation is valuable with generative AI, because it's a highly versatile tool, akin to a digital Swiss Army knife; it can be deployed in various ways to meet multiple needs. This versatility means that high-value, business-specific applications are likely to be most readily identified by people who are already familiar with the tasks in which those applications would be most useful.

Centralized control of generative AI application development, therefore, is likely to overlook specialized use cases that could, cumulatively, confer significant competitive advantage. Certainly, our experience at PwC—where internal hackathons have identified value creation opportunities comprising 1 to 2% of revenue in some of our service lines—has underscored the importance of engaging individual workers and departments in experimentation and exploration.

Powerful as pilots like this are for spotting business-specific trees of opportunity, they run the risk of missing the forest (at best) or (at worst) veering toward the “pilot purgatory” state in which many corporate advanced data analytics efforts found themselves a few years ago, with promising glimmers generating more enthusiasm than value. The above-mentioned financial services company could have fallen prey to these challenges in its HR department, as it looked for means of using generative AI to automate and improve job postings and employee onboarding.

Fortunately, the CHRO’s move to involve the CIO and CISO led to more than just policy clarity and a secure, responsible AI approach. It also catalyzed a realization that there were archetypes, or repeatable patterns, to many of the HR processes that were ripe for automation. Those patterns, in turn, gave rise to a lightbulb moment—the realization that many functions beyond HR, and across different businesses, could adapt and scale these approaches—and to broader dialogue with the CEO and CFO. They began thinking bigger about the implications of generative AI for the business model as a whole, and about patterns underlying the potential to develop distinctive intellectual property that could be leveraged in new ways to generate revenue.

This same sort of pattern recognition also was important to scaling at the consumer packaged goods company we mentioned earlier. In that case, it soon became clear that training the generative AI model on company documentation—previously considered hard-to-access, unstructured information—was helpful for customers. This “pattern”—increased accessibility made possible by generative AI processing—could also be used to provide valuable insights to other functions, including HR, compliance, finance, and supply chain management. By identifying the pattern behind the single use case initially envisioned,

the company was able to deploy similar approaches to help many more functions across the business.

As leaders make such moves, they also need to take a hard look at themselves: What skills does the organization need to succeed at scale with AI, and to what extent do those capabilities already reside somewhere in the company? What's the plan for filling skills gaps, and on what time frame? Failure to pose questions like these can lead to problems down the road—and they're much better answered in the context of early experiments than in the abstract.

Priority 4: Develop a productivity plan

Generative AI's ability to find relevant information, perform repetitive pattern tasks quickly, and integrate with existing digital workflows means the increased efficiency and productivity it can deliver can be almost instant, both within individual departments and organization-wide. Such opportunities aren't unique to generative AI, of course; a [2021 s+b article](#) laid out a wide range of AI-enabled opportunities for the pre-ChatGPT world.

Generative AI has boosted the awareness and interest of many leaders in AI-enabled productivity gains, which companies can do three things with:

- Reinvest them to boost the quality, volume, or speed with which goods and services are produced, generating greater output, broadly defined, from the same level of input.
- Keep output constant and reduce labor input to cut costs.
- Pursue a combination of the two.

PwC firms in mainland China and Hong Kong followed the first approach in small-scale pilots that have yielded 30% time savings in systems design, 50% efficiency gains in code generation, and an 80% reduction in time spent on internal translations. When generative AI enables workers to avoid time-consuming, repetitive, and often frustrating tasks, it can boost their job satisfaction. Indeed, a [recent PwC survey](#) found that a majority of workers across sectors are positive about the potential of AI to improve their jobs.

Generative AI's ability to create content—text, images, audio, and video—means the media industry is one of those [most likely to be disrupted](#) by this new

technology. Some media organizations have focused on using the productivity gains of generative AI to improve their offerings. They're using AI tools as an aid to content creators, rather than a replacement for them. Instead of writing an article, AI can help journalists with research—particularly hunting through vast quantities of text and imagery to spot patterns that could lead to interesting stories. Instead of replacing designers and animators, generative AI can help them more rapidly develop prototypes for testing and iterating. Instead of deciding that fewer required person-hours means less need for staff, media organizations can refocus their human knowledge and experience on innovation—perhaps aided by generative AI tools to help identify new ideas.

It's also important to consider that when organizations automate some of the more mundane work, what's left is often the more strategic work that contributes to a greater cognitive load. Many studies show burnout remains a problem among the workforce; for example, 20% of respondents in our [2023 Global Workforce Hopes and Fears Survey](#) reported that their workload over the 12 months prior frequently felt unmanageable. Organizations will want to take their workforce's temperature as they determine how much freed capacity they redeploy versus taking the opportunity to reenergize a previously overstretched employee base in an environment that is still talent-constrained.

Other companies may focus more on cost savings, which can be substantial, but which also carry with them risks—for example, worker unrest (as we saw in Hollywood), or the hollowing out of the capabilities that companies need to differentiate themselves from competitors. Some organizations may decide these risks are worth taking; the right approach will obviously vary from industry to industry, company to company, and even department to department. What's crucial is to have a plan: What is the relative importance of speed, quality, and cost improvements? What time horizon are you solving for? What will you do with employees whose skills have become redundant as a result of new generative AI capabilities? Getting clarity on the answers to questions like these is an important starting point for focusing your plan.


Priority 5: Put people at the heart of your generative AI strategy

Regardless of the productivity path you choose to pursue, considering its impact on your workforce and addressing it from the start will make or break the success of your initiatives.

Our [26th Annual Global CEO Survey](#) found that 69% of leaders planned to invest in technologies such as AI this year. Yet our [2023 Global Workforce Hopes and Fears Survey](#) of nearly 54,000 workers in 46 countries and territories highlights that many employees are either uncertain or unaware of these technologies' potential impact on them. For example, few workers (less than 30% of the workforce) believe that AI will create new job or skills development opportunities for them. This gap, as well as numerous studies that have shown that workers are more likely to adopt what they co-create, highlights the need to put people at the core of a generative AI strategy.

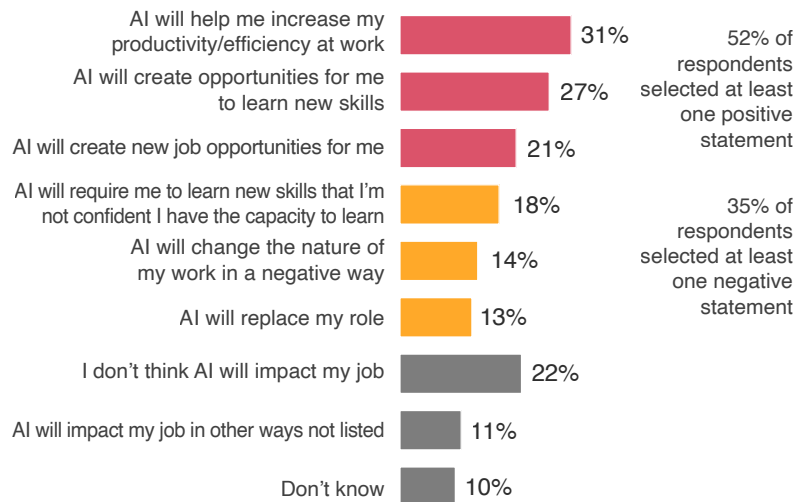
Companies are investing in AI, but most workers aren't sure what that means for them

69% of CEOs say their company is investing in advanced technologies such as AI



Impact workers expect AI to have on their career in the next five years

Positive Negative Neutral



Sources: PwC's 26th Annual Global CEO Survey and Global Workforce Hopes and Fears Survey 2023

To ensure your organization is positioned to capitalize on the promise of generative AI, prioritize steps to engage employees in the creation and selection of AI tools, invest in AI education and training, foster a culture that embraces human–AI collaboration and data-driven decision-making, and support innovation. To this end, we suggest several key strategies:

- **Engage your people early and often.** Continually communicate why AI is important and how it fits into the company’s goals. Explain how AI can make employees’ jobs better and not replace them, and highlight that amassing AI skills will be critical for workers to succeed in their careers going forward.

But remember that communication should be a two-way street. Provide mechanisms to gather feedback from employees about their AI experiences, and use it to refine tools and training programs and address any concerns or challenges.

- **Offer customized training and upskilling.** Assess your employees’ current AI skills and knowledge, and provide role-specific training programs, learning resources, and certifications to address the gaps. Consider teaming up with educational institutions or AI training providers to offer these programs. Create mentorship opportunities that give employees guidance on their AI journey, and provide a way for them to get advice and feedback from AI experts within your company.

And although it’s still difficult to predict many of the new roles that generative AI could give rise to, we know they’ll materialize. Preparing employees for these roles and highlighting the opportunities can energize those looking for career growth and tamp down workers’ fears of replacement. Prompt engineering is a much-discussed role, though it may prove to be a short-term one as generative tools advance. Many other emerging roles involving AI ethics and training will become more prevalent, along with unforeseen roles.

- **Promote a growth mindset.** Create a workplace where learning and trying new things with AI is encouraged by recognizing and rewarding those who do so. And, importantly, make it clear that, with proper guardrails and protections in place, failures mark innovation and are expected,

and even celebrated. One financial services firm we know, for example, highlights at least one instance of failure on a weekly stand-up call among its designers to make visible that these occurrences are acceptable and incur no punitive measures. Unfortunately, this organization remains in the minority—in our 2023 Annual Global CEO Survey, 53% of respondents said leaders in their company don't often tolerate small-scale failures (and employees think that figure is closer to two-thirds).

Fostering a growth culture also includes encouraging employees to share their learnings with each other as they begin working with these tools. Some companies we know are establishing prompt libraries, for example.

- **Advocate and enable ethical AI use.** Provide clear guidelines that articulate how your organization defines the ethical use of generative AI, and ensure that employees understand the importance of fairness, transparency, and **responsible AI practices**. At PwC, for example, we've created an internal microsite articulating the generative AI tools approved for employee use, acceptable business use cases, restrictions on the nature of information employees can input into these tools, requirements for human oversight and quality checks, and more.
- **Measure impact.** Knowing what's working and what isn't requires not only worker feedback but also measurement. Implement key performance indicators to assess the impact of AI on productivity, innovation, and customer satisfaction; and actively promote the results. Some companies we know are conducting controlled experiments, such as by having software engineers use coding assistants, to measure productivity improvements.

By following these strategies, organizations can systematically equip and empower their workforce to position themselves, and the organization, for success in an AI-driven world.

Priority 6: Work with your ecosystem to unlock even bigger benefits

Recent PwC analysis has found that companies with **a clear ecosystem strategy**

are significantly more likely to outperform those without one. It's important, as you experiment with AI, to look outside the four walls of your company: Do you know how your suppliers, service providers, customers, and other partners are planning to leverage this technology to improve their service proposition? What implications does their use of AI have for your early days strategy? Will it impose new conditions and demands? Could closer collaboration on AI lead to fresh opportunities to develop stronger propositions?

The holy grail of healthcare and pharmaceutical firms, for instance, is the ability to access patient records at scale and identify patterns that could uncover routes to more effective treatments. Yet information sharing between organizations has long been restricted by privacy issues, local regulations, the lack of digitized records, and concerns about protecting intellectual property—all of which limit the scope and power of ecosystem collaboration.

Meanwhile, the use of AI has already become widespread across the industry. Medical institutions are experimenting with leveraging computer vision and specially trained generative AI models to detect cancers in medical scans. Biotech researchers have been exploring generative AI's ability to help identify potential solutions to specific needs via inverse design—presenting the AI with a challenge and asking it to find a solution. This AI-supported treatment discovery approach is already being used for both precision medicine (via genetic and healthcare record analysis to identify the best treatments given an individual's specific circumstances) and drug development (via protein and chemical model synthesis that can create custom antibodies).

Until recently, the true potential of AI in life sciences was constrained by the confinement of advances within individual organizations. Today, organizations can combine generative AI's ability to help create and manage records with its capacity for creating statistically reliable, yet fully anonymized, synthetic datasets to enable safe, secure, large-scale data-sharing and data-pooling among healthcare organizations and their partners. That larger pool of information increases the opportunity for medical breakthroughs by helping researchers identify commonalities that can reveal more effective treatments—as well as new opportunities for collaboration between organizations, new business models,

and new ways to capture value along with improved patient outcomes.

Use cases have come up several times as we've described these priorities. That makes sense, because generative AI is a general-purpose technology, suitable for an enormous range of business activities; it's hardly surprising that emerging leaders are emphasizing the search for smart, targeted applications. Here again, though, it's important to underscore that it's still early days. To understand how early, consider another general-purpose technology: electricity. Beginning with lighting in the 1870s, electricity began permeating a range of industrial settings and applications, bringing with it a variety of productivity improvements in the decades that followed. Electricity was the force behind a key feature of Henry Ford's automated assembly line—the overhead monorail conveyor system that made it possible to move parts and materials smoothly throughout the plant.

Looking back, no one talks about Ford's "electricity strategy." Rather, the focus is on the moving assembly line. We suspect the same will be true with generative AI, which will give rise to revolutionary business innovations that are beyond our imagination today. That makes early days AI strategies and priorities like the ones we've described even more important. They won't just yield near-term business benefits; they'll also build muscle and generate valuable experience that sets up today's leaders to achieve much bigger breakthroughs to make product, process, and service innovations that represent the assembly lines of the future.

The authors would like to thank Lois Geraldo and Julia Lamm for their contributions to this article.

strategy+business

a **pwc** publication

- strategy-business.com
- pwc.com/strategy-business
- facebook.com/strategybusiness
- linkedin.com/company/strategy-business
- twitter.com/stratandbiz