

Special Collection

DATA, AI, & MACHINE LEARNING

The Importance of Enterprise AI Roadmaps

Artificial intelligence works better when it's supported with objective, contextual knowledge about how things are actually done at your enterprise.

The very best of *MIT SMR*, helping readers explore ideas that matter for professional and organizational success.

sloanreview.mit.edu

Sponsored by:



Contents

The Importance of Enterprise AI Roadmaps

1 Why AI Demands a New Breed of Leaders

By Faisal Hoque, Thomas H. Davenport, and Erik Nelson

5 How to Manage Tech Debt in the AI Era

By Koenraad Schelfaut and Prashant P. Shukla

9 Generate Value From GenAI With ‘Small t’ Transformations

By Melissa Webster and George Westerman

16 Every enterprise AI roadmap must start with AI understanding how the business works

Sponsor’s Viewpoint from Celonis

Why AI Demands a New Breed of Leaders

Many CIOs lack the bandwidth and authority to solve the tough cultural and organizational change challenges that can block AI success. It's time for an expanded leadership role.

By Faisal Hoque, Thomas H. Davenport, and Erik Nelson

ARTIFICIAL INTELLIGENCE IS fundamentally transforming how organizations operate, but this transformation extends far beyond technical implementation. Modern AI systems are increasingly taking on roles that previously would have been filled by human workers. People working alongside these AI systems often need reskilling, upskilling, and training in behavioral traits such as critical thinking. To successfully manage this blend of AI tools and humans working together in new ways, leaders need to understand complex human and organizational factors, such as agility and cultural change, personality dynamics, and emotional intelligence.

Yet most organizations continue to treat the implementation of AI as a primarily technical challenge — and current technology leadership roles reflect this mindset. According to [Foundry's 2024 State of the CIO survey](#), 85% of IT leaders say that CIOs are increasingly becoming changemakers in their organizations, but only 28% call leading transformation their top priority. In another [recent survey](#), 91% of large-company data leaders said “cultural challenges/change management” are impeding organizational efforts to become data-driven. Only 9% pointed to technology challenges.

But instead of focusing on the aspects of cultural and organizational change that are relevant to AI, much of the time and effort of IT and data leaders is spent on operational functions, long seen as the bread and butter of these roles. This operational focus seems to be increasing: Sixty-one percent of CIOs in the Foundry survey reported having less time available for strategic responsibilities over the past year than in previous years.

Although AI-enabled transformation clearly has enormous human and organizational implications, HR leaders have, for the most part, not stepped up to deal with such changes either.

When leaders fail to think through the strategic and organizational consequences of their AI plans, the results can be catastrophic. Zillow's failed attempt to use AI-generated property valuations as the basis for its own homebuying division not only cost the company \$300 million in losses but also saw its stock price fall by more than 20% as investors lost confidence in its ability to navigate the AI transformation. In another example, California State University had a clear strategic vision but failed to account for the human element. In early February 2025, the university announced a plan for integrating AI across all its systems and services, led by its newly formed AI Workforce Acceleration Board comprising representatives from [10 leading AI companies](#). Within a week, the initiative faced fierce opposition from staff members and students who objected to both the goal of the project and its implementation.

The consequences of an excessively technical approach to AI implementation can also be seen on a more granular level. When [Air Canada deployed a generative AI-based chatbot](#) to assist travelers with booking flights, the goal was to create a more efficient and streamlined customer service experience. However, when the chatbot made mistakes about bereavement fares, passengers dealing with the loss of a loved one were faced with additional challenges to manage in an already difficult situation. This case underscores the need for organizations to address the human-facing personality of AI models, their decision-making authority, and the kinds of ethical boundaries and special circumstances in which it might be essential to have a human in the loop.

All of those examples highlight a crucial gap in organizational leadership. While CIOs and CTOs play critical roles in technical implementations and system maintenance, they sometimes lack both the bandwidth and the

Sastry Durvasula, the chief operating, information, and digital officer at TIAA, views workforce transformation as a key part of his role.

mandate to address the broader human and organizational implications of AI transformation. Organizations require a new kind of leader.

The new role we envision as being essential in the age of AI — which might be called the chief innovation and transformation officer (CITO) — combines technical expertise, behavioral insights, and strategic vision with a deep understanding of organizational psychology and culture change. This combination of skills ensures that organizations are properly equipped to manage the profound changes that the age of AI will bring. Companies are still debating the job title, but forward-thinking organizations are already employing, or hiring for, such leaders. Let's explore the reasons.

Why AI Requires a New Leadership Model

The challenge of implementing AI effectively extends far beyond technical prowess. Leaders responsible for AI implementation must be able to do the following:

- **Navigate complex ethical landscapes.** AI deployment requires careful consideration of [ethical implications](#), bias mitigation, and alignment with organizational values.
- **Foster cultural transformation.** Successfully integrating AI — and, particularly, AI agents — means transforming an organization's culture to embrace [new ways of working](#) and thinking.
- **Manage human-AI collaboration.** Leaders must understand both AI capabilities and human skills and psychology if they are to create effective partnerships between human workers and AI systems.
- **Drive cross-functional integration.** AI implementation touches every part of an organization, requiring leaders who can work across traditional silos.
- **Deal with citizen development.** AI is enabling businesspeople to [develop systems and models](#) that could previously only be created by IT professionals — a [work trend](#) requiring simultaneous encouragement and risk management.

- **Ensure responsible innovation.** Leaders must balance the drive for business innovation with careful consideration of the potential risks and societal impacts.

Some people serving in traditional IT and data leadership roles that focus on technical implementation and system maintenance may lack the skills and the bandwidth to address these broader challenges. Even for CIOs who have succeeded at driving revenue, the AI age raises the difficulty level. This misalignment is one reason why so many AI initiatives fail to achieve their goals.

New Hiring Patterns, Expanded Leadership Roles

Recent data shows that businesses are responding to the accelerating pace of technological change by broadening their C-suites to include roles dedicated to innovation, AI, and transformation. Increasing compensation levels and a surge in hiring reflect the strategic importance of innovation leadership, AI expertise, and transformation management.

Leaders with titles such as chief innovation officer, [chief AI officer](#), and chief transformation officer are becoming increasingly common as companies wrestle with how best to meet these strategic needs. A [study by Boston Consulting Group](#) found that the number of companies hiring chief transformation officers increased by more than 140% from 2019 to 2021, and that those companies experienced a significant increase in total shareholder return in the year after the new appointment. This hiring trend has continued over the past three years.

Some companies are expanding the [scope of existing technology leadership roles](#) to emphasize strategic delivery and value creation versus operational and engineering tasks. Others are [consolidating a variety of tech roles](#) into a single [“supertech” leadership position](#) that focuses on business strategy and vision. Sastry Durvasula, for example, is the chief operating, information, and digital officer at investment company TIAA, where

The most effective leadership roles for managing AI will combine both technical and organizational change responsibilities.

he oversees organizational, technical, and operational change — and 60% of TIAA employees — in all of those domains. He views workforce transformation as a key part of his role and told us that, “since change from AI is imminent at this point, it’s important to manage the upskilling/reskilling and transition of our people into their next job post revolution.”

Regardless of the specific title, the most effective leadership roles for managing AI and related new technologies will combine both technical and organizational change responsibilities. While the latter capability is often overlooked in technology initiatives, our discussions with successful leaders suggest that it must be elevated in importance. AI will touch all parts of an enterprise, including human employees and clients. Harnessing its potential and managing its risks requires the authority to take a view across the whole organization, steering macro-level innovation and managing change that touches on all parts of the business’s identity and culture. The AI evangelist, the strategic innovation lead, and the hands-on transformation expert are at their most effective when brought together in a single office.

Early Examples of CITO-Type Roles

Two major corporations that have appointed individuals to these kinds of expansive roles are PepsiCo and Standard Chartered Bank. In 2020, PepsiCo appointed Athina Kanioura chief strategy and transformation officer, a role that oversees the digitization of the entire business, including the implementation of AI. In 2021, Roel Louwhoff became chief transformation, technology, and operations officer at Standard Chartered Bank. Before taking on that role, Louwhoff was ING Bank’s chief operations and transformation officer, a position that included IT and change management responsibilities.

For people in such positions, there is currently little consistency in title. In some cases, individuals with traditional tech-leader titles have taken on an expanded role that cuts across all of the domains relevant to AI implementation. [JPMorgan Chase’s implementation](#) of its IndexGPT and AI-driven investment analysis systems provides an example of this kind of expanded lead-

ership mandate. Under the direction of global head of data and analytics Teresa Heitsenrether, the bank established a comprehensive AI governance structure that explicitly connects AI implementation to the organization’s purpose and values. It includes both an explicit focus on the ethics of AI interactions and specific protocols for AI decision-making in investment scenarios.

A recent advertisement for a position at State Street Bank exemplifies the breadth of competencies required in these new leadership positions. The responsibilities of the role of chief transformation officer within the bank’s central Global Technology Services unit include “enterprise transformation leadership,” with specific accountabilities for driving “innovation and modernization initiatives, including automation, AI, blockchain, and cloud adoption,” as well as leading “cultural and organizational change efforts to embed agility, efficiency, and a customer-first mindset.” Other responsibilities involve business process optimization, customer experience and digital innovation, and change management. As of mid-March, the job had not yet been filled, and it may be difficult to find such a combination of skills. But the job description perfectly illustrates the responsibilities we envision as coming together in an office of innovation and transformation that is led by the equivalent of a CITO.

The CITO (or equivalent) role encompasses several critical functions. At the strategic level, these leaders align AI initiatives with organizational purpose while developing long-term transformation road maps. They ensure ethical alignment across implementations while creating frameworks for sustainable innovation that balance technical advances with human values.

In terms of cultural leadership, CITOs guide their organization’s evolution through its AI transition while proactively managing resistance to change. They drive innovation across all levels of their organization while maintaining meaningful human connections in automated processes and building psychological safety for both staff members and clients into AI implementations. Technical oversight involves ensuring robust implementation standards, managing risk and security

AI personas can be autonomous or collaborative with human workers, requiring leaders to manage the interplay between human psychology and intelligent machines.

across AI systems, driving integration across platforms and departments, and maintaining data quality and ethical use.

An Emerging Requirement: AI Persona Management

AI persona management is an example of a new responsibility that requires new leadership from a CITO. Generative AI tools already can be asked to take on specific personalities — a teacher, a scientist, a lawyer, and so on. As [agentic AI use](#) grows, companies will need to manage AI agents with increased levels of autonomy and specific attributes and identities.

You may think of AI personas as digital characters or workers with specific traits, priorities, and capabilities that are designed to interact with users and process information in customized ways. They perform defined tasks or serve as interfaces between humans and AI systems. This allows organizations to create distinct roles — such as strategic adviser and customer service agent — tailored to specific use cases. AI personas can be autonomous or collaborate with human workers, requiring leaders to understand and manage the delicate interplay between human psychology and intelligent machines.

Leaders will also need to ensure consistent behaviors from the AI personas and keep outputs aligned with defined business goals. Salesforce, for example, is not only selling agentic AI tools but also using AI personas to guide its own decisions on product updates, training programs, and feature prioritization.

To be strategic, AI persona management will require a collaborative leadership approach that involves a cross-functional team with a combination of technical, ethical, and strategic expertise. Led by a CITO, the ideal leadership team for AI persona management would include senior AI, ethics, technology, and HR executives.

Thinking Differently About Tech Change

As AI continues to reshape our organizations and society, leaders face a clear but complex challenge: determining how to harness the power of these new technologies while preserving and promoting human values. Success will require practical frameworks for responsible innovation, careful balancing of efficiency with human needs, and thoughtful approaches to measuring success that go beyond technical metrics.

The evolution of technology leadership is not merely a matter of adding new roles to the C-suite. Rather, it reflects a growing recognition that implementing AI successfully requires us to think differently about how we manage technological change.

Organizations that adapt their leadership structures to meet the cultural and change management challenges of this new era will be better equipped to make sound decisions about how and when to implement AI. The CITO or equivalent role embodies this shift, providing a structured approach to bridging technical expertise with human insight.

Faisal Hoque is the founder of Shadoka, NextChapter, and other companies. His latest book is *Transcend: Unlocking Humanity in the Age of AI* (Post Hill Press, 2025). [Thomas H. Davenport](#) is the President's Distinguished Professor of Information Technology and Management at Babson College, the Bodily Bicentennial Professor of Analytics at the University of Virginia Darden School of Business, a fellow of the MIT Initiative on the Digital Economy, and senior adviser to the Deloitte Chief Data and Analytics Officer Program. His latest book is *All Hands on Tech: The AI-Powered Citizen Revolution* (Wiley, 2024). Erik Nelson is a senior vice president at CACI International, responsible for strategic vision and growth in the company's Enterprise IT division.

Reprint 66418. Copyright © Massachusetts Institute of Technology, 2025. All rights reserved.

How to Manage Tech Debt in the AI Era

AI is transforming business — and increasing technical debt in enterprise systems. Consider these four key insights for leaders on evaluating trade-offs and ensuring innovation capacity.

By Koenraad Schelfaut and Prashant P. Shukla

TECHNICAL DEBT IS AN ANCHOR, DRAGGING down business leaders' efforts to run a tight ship. The accumulated costs and effort from IT development shortcuts, outdated applications, and aging infrastructure sap a company's ability to innovate, compete, and grow.

A degree of [technical debt](#) is inevitable. To remain agile, businesses often choose to deploy new technologies as quickly as possible, knowing that they may have to pay to fix their systems later. Addressing this trade-off is becoming more critical as companies adopt artificial intelligence. With over [\\$2.41 trillion in annual costs](#) in the United States alone, tech debt isn't just an IT problem; it's a business liability that requires CEOs' attention.

Outdated infrastructure and applications aren't just slowing companies down. Tech debt prevents organizations from deploying AI solutions that could [reshape how they compete](#). The stark reality is that with AI poised to penetrate every business function, all technical debt is becoming AI technical debt.

In the past, experts recommended injecting a healthy dose of [digital decoupling and modularity](#) into the tech stack to address technical debt — that is, breaking code, infrastructure, and data (for example) into smaller, interchangeable parts that could be replaced more easily when better technology emerged.

To understand how today's business leaders are reinventing their organizations, including the role tech debt plays, Accenture studied 1,500 global companies in 10 countries covering 19 industries and conducted scores of in-depth discussions with C-level leaders. The research found that companies that are well positioned for change have a [reinvention-ready "digital core"](#) — a set of key components such as cloud infrastructure,

data, and AI that can be easily updated. They also typically set aside around [15% of their IT budgets](#) for tech debt remediation.

What the research made clear is that today, addressing tech debt is not about eliminating it but managing it. The key lies in knowing what the debt is, what to fix, what to keep, and how to recognize the tech debt that is boosting your company's innovation capacity.

For example, the snack company Mondelez International, known for popular brands like Oreo, Belvita, and Cadbury, is using new technologies, including generative AI, to keep up with and even get ahead of changing consumer preferences. To enable these applications and prepare for the future, Mondelez upgraded the data and AI portions of its digital core along with the cloud infrastructure, security, and cloud management platforms supporting them.

To prepare for the upgrades, Mondelez needed help with its tech debt. The company was grappling with a highly fragmented IT landscape that included more than 1,000 applications. Some of those applications not only were redundant but also cost time and money to maintain. Together, Mondelez and Accenture assessed each application, evaluating metrics such as software quality, performance, and operating costs. Through this work, Mondelez created a consolidated set of applications that reduced both the total cost of ownership and tech debt through savings on infrastructure, licenses, and application management services.

As a result, Mondelez can make the most of opportunities to use AI. Together with Accenture and Publicis Groupe, [Mondelez launched a new platform](#) that uses generative AI to improve its global marketing capabilities and optimize consumer experiences.

Companies that are well-positioned for change typically set aside around 15% of their IT budgets for tech debt remediation.

Managing Tech Debt Now: Four Key Insights

From our work with companies such as Mondelez, we've identified four key insights that can help business leaders manage technical debt in the era of AI.

1. Zero tech debt is not the goal.

Many leaders aiming for zero debt funnel resources into erasing legacy system inefficiencies. This is a mistake because, as noted earlier, some degree of tech debt is inevitable. Systems start incurring technical debt as soon as they go live. The cost of maintaining and updating these systems grows exponentially with their complexity and age.

Instead, companies need to effectively manage their tech debt by treating it as a necessary cost of doing business. The categories that lenders use to make decisions about their loan portfolios — principal, interest, liabilities, and opportunity costs — offer a way to think about what technical debt to target:

- Principal is the cost of updating outmoded technology.
- Interest is the incremental cost associated with working around outmoded technology.
- Liabilities are the additional issues that occur when outmoded technology is not updated.
- Opportunity cost is the set of opportunities a company cannot realize due to outmoded technology.

Companies should start remediating their tech debt with the principal — that is, at its source. If they effectively manage the principal, as they would financial debt, it will accrue little or no interest and minimize any liabilities or opportunity costs. With these priorities in mind, leaders can make better investment decisions.

For example, companies can focus on the quality of the code in their software. If the code is not well written, debugged, improved, or maintained, technical debt can proliferate. Modern software development tools can help developers minimize errors that will turn into tech debt principal later.

2. Know good technical debt from bad.

Businesses may intentionally incur tech debt to achieve short-term goals. When leaders lack a clear understanding of good versus bad technical debt, their companies can become paralyzed by accumulating inefficiencies.

Fear of debt can also prevent companies from taking strategic risks. The companies Accenture studied differentiate between necessary debt and liabilities they need to address or avoid.

Good technical debt is an investment. It enables experimentation, learning, and faster time to market. One way to ascertain whether a debt is good is to consider whether the opportunity cost of avoiding the debt is greater than the debt itself. This is often the case when companies are deploying emerging technologies.

Consider a company that is deploying a custom generative AI chatbot for internal knowledge management that leaders envision will assist customer service representatives and product designers. Due to knowledge and skill gaps, the company is bound to encounter a learning curve that could result in bugs or poorly written code. Developers may even choose the wrong model as the foundation for the chatbot.

If the prospect of accruing this technical debt gives leaders cold feet, they might pull the plug on the project — and regret it. In that scenario, the company wouldn't have the team or talent in place to take advantage of more mature solutions as they emerged. If a competitor were to use those tools, the company would have lost an opportunity to differentiate its customer service or product design.

3. Focus on the highest-value fixes.

Many companies lack a structured approach to prioritizing debt remediation. They waste resources on low-value fixes while neglecting critical ones.

We developed the [PAID framework](#) (as shown in the graphic) to help leaders assign actions based on business value and the amount of tech debt that needs remediation. The acronym is based on four key actions:

- **Prioritize:** When both the level of tech debt and its effect on business value are high, the debt needs immediate remediation. Consider, for example, a

health care provider that has a large repository of patient data that is scattered across numerous systems and silos. Such decentralization affects the speed and quality of care the organization can deliver, with potential consequences for patients' well-being.

- **Address:** When the level of tech debt is high but the effect on business value is low, the debt should be scheduled for remediation before the effect on business value increases. A bank might have an outdated customer relationship management system that works efficiently but is limiting the personalization it can offer customers. It will want to upgrade this system to avoid losing customers to competitors.
- **Investigate:** When the level of technical debt is low but the impact on business value is high, neglecting the situation could be risky. Consider a consumer goods company with an outmoded inventory management system: It might notice some difficulty with long-term financial planning, and although the system does not need immediate attention, leaders should plan to monitor it to understand when remediation will be needed and avoid unexpected shocks.
- **Document:** When both the level of tech debt and impact on business value are low, tech debt should be monitored, but remediation should be a low priority. For example, a small HR team may be using obsolete report-generating software, but if the team and its work may soon be merged with another business unit, it will no longer need the application.

The PAID framework is useful whether or not a company is getting ready to deploy AI. However, if a company's leaders intend to leapfrog to the latest AI-enabled services, this framework can facilitate decisions to strategically write off technical debt that falls into the framework's Document category.

CTT - Correios de Portugal, the 500-year-old national postal service of Portugal, prioritized tackling a costly patchwork of on-premises digital systems. They were a challenge to manage and were hindering the development and deployment of new services. CTT reduced IT operating costs by 15% by migrating some of its data center workloads to the cloud. And it is poised to reap dividends from accessing the cloud provider's technological advances immediately and continuously rather than accruing its own technical debt internally.

4. Reframe tech debt as innovation capacity.

Companies burdened by outdated systems struggle to scale, compete, and deliver new value. They are at risk of being overtaken by savvier competitors. When technical debt reaches critical levels, it leaves companies vulnerable to being disrupted or going out of business.

Technical debt was one of several key factors that led to the 2022 [bankruptcy of Sungard Availability Services](#), a provider of on-premises IT disaster recovery services. In the years before the bankruptcy, Sungard began adjusting to a changing business environment. Customers were starting to prefer cloud-based services to owning on-premises data centers for backups. Sungard had outlined its plans to address its technical debt as it shifted its business model accordingly.

But the organization did not move fast enough. When demand for cloud services, including remote working capabilities, intensified during the COVID-19 pandemic, Sungard's substantial technical debt meant it couldn't compete with more cloud-focused competitors.

Now that AI is becoming synonymous with innovation, reframing tech debt remediation as an investment in innovation capacity can shift the conversation about remediation to one about the upside of AI-enabled opportunities. For example, thanks to CTT's cloud environment, CTT can get to work on a new AI solution in a matter of hours, not weeks, when customer preferences or the business environment changes.

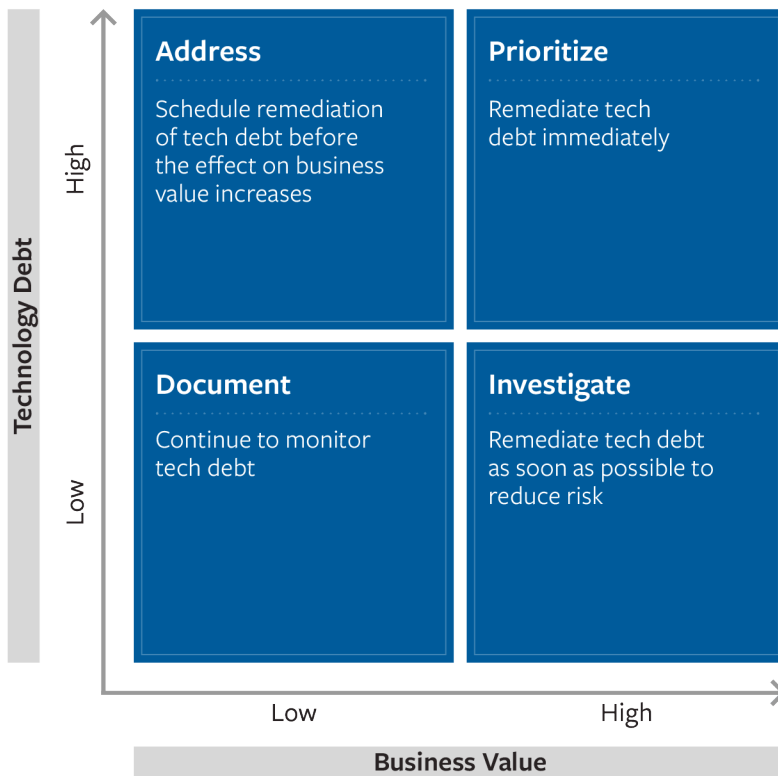
Tech debt is no longer an IT-only issue; it's a strategic challenge that requires CEO and leadership oversight. Leaders who manage tech debt effectively can transform their IT systems into a tool for growth, balancing the trade-offs between agility and efficiency. By embracing continuous management of technical debt, distinguishing good debt from bad, prioritizing high-value remediations, and reframing tech debt remediation as an investment in innovation capacity, leaders can ensure that their organizations not only survive but also thrive in an AI-first world.

Koenraad Schelfaut is senior managing director and lead of Technology Strategy & Advisory at Accenture. Prashant P. Shukla is principal director at Accenture Research. The authors would like to thank Laurence Mackin, Krish Jhaveri, and Gargi Chakrabarty for their contributions to this article.

Reprint 66345. Copyright © Massachusetts Institute of Technology, 2025. All rights reserved.

Managing Technical Debt: The PAID Framework

Decisions about whether — and when — to remediate a particular technical debt should be based on the level of debt and how much the debt affects business value.



Source: Accenture

Generate Value From GenAI With ‘Small t’ Transformations

Business leaders are getting real value from large language models by working their way up the risk slope and building the foundation for larger, future transformations.

Melissa Webster and George Westerman

LESS THAN TWO YEARS AGO, GENERATIVE AI made headlines with its amazing new capabilities: It could engage in conversations; interpret massive amounts of text, audio, or imagery; and even create new documents and artwork. After the fastest technology adoption in history — with over 100 million users in the first two months — businesses in every industry began experimenting with it. Yet, despite two years of broad managerial attention and extensive experimentation, we are not seeing the large-scale GenAI-powered business transformations that many people initially envisioned.

What happened? Has the technology failed to live up to its promise? Were experts wrong in calling for giant transformations? Have companies been too cautious? The

answer to each of those questions is both yes and no. Generative AI is already being used in transformative ways in many companies, just not yet as the driver of a wholesale redesign of major business functions. Business leaders are finding ways to derive real value from large language models (LLMs) without complete replacements of existing business processes. They’re pursuing “small t” transformation, even as they build the foundation for larger transformations to come. In this article, we’ll show how they’re doing this today and what you can do to generate value with generative AI.

How Businesses Are Transforming With GenAI

Our project team interviewed the senior managers of



Neil Webb

various functions, including artificial intelligence, data science, innovation, operations, and sales, at 21 large companies. We focused on understanding what organizations with relatively early and broad GenAI adoption are doing and why. We also reviewed public information about companies similar to those we studied.

To start, we needed a definition of what *digital transformation* means. An early definition is “the use of technology to radically improve the performance or reach of an organization.”¹ More recently, OpenAI’s ChatGPT synthesized this definition: “a comprehensive integration of digital technologies that fundamentally reimagines business models and processes, contrasting with incremental change, which focuses on gradual improvements.” Digital transformations, in general, consist of numerous technology-enabled improvements, often assembled over time, to create broader change in how a company operates. They are driven not by a single technology but rather by using the right technologies for the right tasks to deliver a new way of doing business.

Our research shows that most companies are following a more targeted approach to transforming with generative AI. While GenAI can potentially increase the speed and quality of many tasks, it also comes with a variety of risks around accuracy, security, and intellectual property management. The leaders we interviewed tend to apply the logic of a risk slope when making their decisions, attaching a higher risk to customer-facing processes than to internal ones. Leaders in industries like medicine and financial services also see these risks through the lens of regulatory compliance.

Some leaders are thinking beyond these highly publicized GenAI risks to also consider the costs and risks of preparing the organization for large-scale implementations. They seek to reduce investment in software licenses and hiring skilled people until the returns are clearer. They also see risks in scaling AI transformation pilots to the enterprise level without first cleaning up the data and back-end systems that feed into them.

At the lower end of the risk slope are discrete uses that can deliver immediate value at relatively low risk. More extensive transformations may provide significant value, but they also have higher risk. As a result, many IT and digital leaders are investing first in use cases at the bottom of the risk slope, aiming to deliver early value while also developing capabilities that can partially de-risk implementations that are higher in cost, risk, and potential return.

Three categories of transformation represent different areas of the risk slope, starting with low-risk individual uses, then moving to role- and team-specific tasks, and finally to products and customer-facing experiences.

1. GenAI in Tasks That Are Common to Individuals in Many Roles. At the lower end of the risk slope, employees are using LLMs in ways that are useful to many roles, such as writing, synthesizing information, generating imagery, and documenting meetings. The near-ubiquitous nature of GenAI can have real impact within the organization. As early adopters considered what tools to provide to which employees, those employees began using public tools such as ChatGPT or Gemini without asking permission. Wary of privacy and accuracy concerns but mindful of the added cost, many companies have begun to offer de-risked GenAI tools to at least some of their workers. Some are now starting to buy or create integrated tool sets that link generative AI to other functions that employees typically perform. Benefits vary by use and user, with individual initiative-taking and prompting skills influencing the value they derive. Consider the following approaches.

Company-specific LLM instances: One way to address privacy and security concerns is to license private instances of major LLMs, such as ChatGPT or Anthropic’s Claude, which can be accessed through secure cloud platforms. Employees can use these stand-alone GenAI tools for synthesis, writing, and other content-generation tasks without leaking confidential information. Beyond generation, some are using the tools for learning. For example, a technical product manager we spoke with at a large U.S. technology company uses an internal tool that simulates feedback from an executive on draft presentations and reports.

Built-in integration with common office productivity tools: Use cases and productivity gains expand when an organization can integrate an LLM with company information and desktop tools, such as Copilot and the Microsoft 365 suite or Gemini and Google Workspace. They create what one manager called “a super search engine,” able to identify useful documents both within a team and across a global company. Integration allows employees to pull content from different sources, such as emails, meeting transcripts, and internal documents. It also makes it possible for a user to request specific information from their inbox, such as a list of open items from the past week, and to prepare spreadsheets or presentations without having to start from scratch, given the LLM’s access to company information.

Custom integration: Some companies are going beyond basic desktop integration to add company-specific intelligence, by training models on terminology and information that are proprietary to the company. Global consulting firm McKinsey built Lilli, a platform that links generative AI to its intellectual property from over 40 internal sources. The effort involved significant technical

hurdles — for example, the tool needed to be modified to read PowerPoint slides, which are one of the company’s main ways of communicating project information — but the platform is providing value to the company. If a consultant has a question about green energy business models in less-developed economies, for instance, Lilli can quickly find and synthesize information from projects that have already studied the problem somewhere in the world. The platform’s capabilities, combined with robust employee education, led to about 75% of employees actively using Lilli in less than a year, time savings of up to 30%, and substantially improved quality, McKinsey has reported.

McKinsey is not alone in developing these specialized models for the general workforce. Another company we studied has extended its specialized LLM to not only find information but also autogenerate initial drafts of proposal text and slide decks.

External tools: Company-specific tools can be useful and secure, but some people will look elsewhere for functionality that they cannot find internally. One program manager in a large tech company told us that he uses ChatGPT for tasks involving nonconfidential information, such as writing specifications or structuring documents, while using DALL-E to create simple visuals easily instead of using more cumbersome tools he has access to. A coworker in product management turns to Superwhisper — a voice-to-text/text-to-voice tool that keeps all information on his device — to dictate, summarize, and clean up his thoughts during performance reviews. He also uses tools like Perplexity to fact-check and cross-reference information in his writing tasks.

2. Specialized GenAI in Specific Roles and Tasks.

Companies working their way up the risk slope are developing generative AI capabilities that will improve productivity and quality in specific job roles or business processes. Here, there is less tolerance for unacceptable output, though not yet to the same degree as with customer-facing applications. These solutions typically maintain a human in the loop, where employees interact with the tools and review the outputs rather than allowing the GenAI tools to automatically make decisions or produce outputs. Consider the following use cases.

Coding and data science: This is one of the earliest and most common GenAI-assisted tasks in every industry. Rigorous studies have shown major productivity gains for software engineers who use coding copilot tools to speed up tasks like writing code, finding useful libraries, or conducting code reviews.² Yet another productivity gain comes from generating sample code or data for training purposes. Data scientists, meanwhile, are using generative AI tools to conduct data analysis or

Copilot tools can not only help experts work more efficiently but also help novices improve their skills.

produce scripts for analytic tools. By creating readable documentation, GenAI offers significant time savings to coders, data scientists, and their managers in performing what many consider a tedious task. Copilot tools can not only help experts work more efficiently but also help novices improve their skills. However, these benefits are not automatic: People often need training to get the most benefit from the many features available.

Support of customer-facing individuals (with a human in the loop): One of the first uses of role-specific GenAI for customer-facing applications is in customer service. The tools can help agents find information quickly and also suggest actions in real time. Some can even coach the agent later, synthesizing numerous calls to identify patterns and opportunities for improvement.

In one study of a tool that provides GenAI-based coaching for call center representatives, MIT researchers found that access to the tool increased productivity, as measured by the number of issues resolved per hour, by 14% on average, with a 34% improvement for novice and low-skilled workers.³

Current examples of how generative AI can assist with customer service include:

- Amazon Pharmacy’s internal chatbot to support customer care representatives can retrieve answers from the help center knowledge base and summarize the information for the representatives, allowing the reps to answer customer questions in less time, according to the company.
- Morgan Stanley found that its knowledge assistant tool, trained on more than a million pages of internal documents, speeds financial advisers’ process of finding information, allowing them to spend more time focusing on customer needs. A new tool summarizes customer video meetings and drafts tailored follow-up emails. This kind of interaction is not limited to finance.
- Sysco, the world’s largest wholesale food distributor, is using GenAI in tasks ranging from making menu recommendations for online

customers to generating personalized scripts for sales calls, based on customer-specific data.

Online content generation: CarMax, the largest omnichannel used-car retailer in the U.S., was an early adopter of OpenAI's generative tools. CarMax uses AI to produce text for its car research pages, which help customers make a purchase decision, and to embed keywords and organize content in order to boost a web page's search ranking. It summarizes customer reviews of a car's model into a few sentences, saving customers from wading through hundreds of individual reviews. Summarizing over 5,000 car pages manually would take multiple humans 11 years, according to the company. With generative AI, it runs this process regularly, and it takes only a few hours. The quality has been even better than anticipated since it did a little fine-tuning, with an 80% editorial review approval rate. CarMax has since expanded its use of GenAI to include marketing design, chatbots for customers, and tools for internal associates.

Creative processes: Dentsu, one of the world's largest creative agencies, uses generative AI in all stages of the creative process, from proposal to project planning to creative ideation. Employees can use it to turn a few lines of copy into a proposal, manage complex budgeting spreadsheets, or make sense of notes from numerous planning meetings. In creative sessions with a client, instead of gathering ideas and making the client wait days or weeks for visual concepts, the team can iterate with the client in real time. "They're getting less time on mundane, tedious tasks," said Kate Slade, Dentsu's emerging technology enablement director. "They can be creative and create higher-quality content with less effort." Dentsu and other companies can use foundational LLMs to generate product image mock-ups in the conceptual phase or use specialty tools like Flair.ai to create polished product photos, including for clothing and accessories, which are displayed on AI-generated models.

Finance and regulatory: Multiple surveys have shown that finance teams are relatively late adopters of new technologies, with CFOs citing technology gaps, data concerns, and competing priorities as reasons for that lag.⁴ However, some companies are innovating within this business function. One international energy company we studied created a tool using a mix of GenAI, traditional AI, and other algorithms that can suggest mitigations or help rewrite an audit report. Other companies are using generative AI to assist in drafting reports for audits or regulatory compliance. At Amazon, the finance function uses a mix of rules-based AI, machine learning, and LLMs to address tasks in fraud detection, contract review, financial forecasting, personal productivity, interpretation of rules and regulations, and tax-related work.

Managers have reported improved performance on those tasks and that employees have been able to shift their efforts away from repetitive tasks to instead focus more on work that involves critical thinking.

While productivity gains are the expected and common benefits of applying GenAI to specialized roles and tasks, the technology's true impact extends further: Generative AI is fundamentally transforming what professionals can achieve across industries. By not only enhancing efficiency but also expanding the realm of possibilities within various functions, GenAI is enabling innovations and reshaping traditional processes.

3. GenAI in Products and Customer-Facing Interactions. When people ask about GenAI-enabled business transformation, they often mean changes in products and other customer experiences. However, these changes are often higher on a company's risk slope, and judicious action is warranted even as companies strategically consider more extensive future applications. Traditional companies are starting to implement GenAI-enabled customer service to answer simple queries and to GenAI-enable the sales process. Meanwhile, major software companies are already incorporating GenAI-powered functionality into their products. Consider the following use cases:

Direct customer service interactions: GenAI is taking the traditional phone menu or robotic process automation-enabled chatbot to a new level of sophistication. It offers natural language interactions and flexibility not possible with rules-based AI, plus it adds multilingual capabilities. For example, life insurance company John Hancock built chatbot assistants to handle common queries, freeing up human agents to handle more complex issues. This reportedly reduces company costs, customer wait times, and employee time spent on simple tasks. Now GenAI is expanding to voice interactions in organizations like Starbucks, Domino's, CVS, and banks. It's only a matter of time before these tools expand to include video.

Personalized shopping experiences: Customers are accustomed to getting e-commerce product suggestions based on what they (or others) have bought, or banner ads based on what they have viewed. Companies now are using GenAI to tailor the shopping experience throughout the customer journey.

For example, Tapestry, the parent company of brands such as Coach and Kate Spade, uses real-time language modifications to personalize the online experience to individual shoppers as they are moving through a retail site. This includes injecting a conversational tone that mimics the experience of engaging with a store associate. Tapestry claims to have seen an e-commerce revenue increase

of at least 3% due to these personalization changes.

Amazon is also personalizing the customer journey, by offering product recommendations and descriptions that fit a customer's holiday, sport, or diet preferences, or their household size. In addition, it is executing a phased rollout of tools to help vendors improve the shopping experience that they provide.

GenAI video generation and a chatbot for insights and recommendations on their business performance are currently in beta mode with some U.S. retailers.

Enhancement of existing software products: Even as traditional companies experiment with GenAI in customer interactions, leading software companies have already begun integrating generative AI capabilities into their products, whether to improve existing features or add new ones. Generative AI's natural language processing capability smooths the path for a user of Lucidchart, for example, to create a flowchart just by writing what they want rather than having to go through the steps and menus manually, resulting in a draft that is editable and sharable.

Canva, another visual communication tool, uses ChatGPT to ease the process of creating and modifying slides, images, videos, presentations, and social media posts. This increases productivity for any user, with the added benefit of decreasing frustration for inexperienced users.

Similarly, Adobe has embedded GenAI features throughout its product suite so users can easily create and modify images, adjust them for brand style, or "chat with a PDF." In marketing campaigns, Adobe's GenAI features help to track consumer behavior, personalize content, and improve performance measurement.

As these and other companies integrate generative AI into their product suites, their corporate customers may choose to wait for features rather than build the functionality themselves.

Such integrations into products and customer-facing interactions are enhancing customer experiences without a need for employees in the loop. This strategy serves a dual purpose: It personalizes interactions to increase customer engagement and sales, and it empowers users to achieve new levels of productivity and capabilities on their own.

Generating Transformation

What is the right kind of transformation with generative AI? The answer is, anything that a company can do to change its performance or reach using the technology, in concert with other technologies and, often, human action. Our research uncovered examples in industries from fashion to finance and roles from auditing to mar-

Bring your IT and data teams to the table, as well as employees from domains already adopting GenAI.

keting. The companies we studied are being careful as they work their way up the risk ramp. They are pursuing small-t transformation, often with a human in the loop, as they build capabilities that can enable the development of applications that have higher value and risk.

Our research suggests a number of actions leaders can take to generate transformation with generative AI.

Identify key pioneers in your organization, from decision makers and stakeholders to power users, and develop your message for them. With generative AI, innovation often comes from "cyborgs" — early adopters who integrate the technology into their work and are motivated to use it to solve a problem for themselves or for customers.⁵ However, these workers, as well as the later adopters, may be concerned about a negative reaction from their employer or about GenAI replacing their jobs, hence the need to communicate your innovation vision. Bring your IT and data teams to the table, as well as employees from domains already adopting GenAI or who are the most interested.

Assess where your company is now on the risk slope relative to the companies we've described. What are you already doing, and what would be the next level of complexity and reward? Look at the opportunities in the areas of individual productivity, role-specific enhancements, and innovations in product or customer engagement. Keep in mind that while companies can develop in all three simultaneously, the maturity levels likely will vary.

The companies we studied consider security issues, integration with other systems, and output sensitivity in deciding where to begin and where to go next. CarMax's GenAI solutions currently have some human-in-the-loop elements, according to Shamim Mohammad, the company's executive vice president and chief information and technology officer. "As GenAI matures and becomes ever more sophisticated, transformational or game-changing use cases will emerge," he said. "But in my view, it will take some time before organizations can deploy GenAI solutions without human intervention and

supervision comfortably.” As leaders consider moving beyond the low-hanging fruit that many have pursued so far, they face challenges that must be cleared away, including scalability, management buy-in, and the need for foundational capabilities.

Consider scalability. Numerous research participants described how the process of going from pilot to scale is nontrivial. “It’s easy to do the proofs of concept, but bringing it to the right level of trust among a large group of users is much more difficult,” according to Pentti Tofte, staff senior vice president of data analytics at FM, a large commercial insurer. Data outside of the pilot environment is messier and less connected than it is inside of it. Low-probability events appear more often when working with large volumes of data or customers. According to the head of AI at a large bank we spoke with, “The more stuff you do, the more stuff you find to do.”

Beyond the technical questions of scaling specific pilots, leaders are asking strategic questions about whether it is better to move quickly — even if it means training and integrating custom models and reworking them as technologies change — or to wait until software vendors incorporate new features into the systems they already use.

Secure management buy-in. Management buy-in is essential for larger projects, since managers have heard about the risks of generative AI and may have learned to be skeptical about the promise of new technologies. An executive at a medium-sized tech company in New England reported that their GenAI innovation stagnated until the CEO saw its potential, allocated resources, and communicated how GenAI would be expansive for employees as well as the company. Small-t innovations can help to make the value story real and make the case for investments that can reduce the perceived risk of larger opportunities.

Investigate foundational investments that can improve the risk-return ratio higher up the risk ramp. Some of the boldest use cases will require extensive investment in data cleansing, model training, and integration before they can be ready for a real-world test. “The most important thing with GenAI is, you have to make sure your data is correct because good data will give you good results from AI — GenAI especially,” Mohammad said. “You also have to have the right AI governance in place to ensure that the AI is being deployed responsibly.”

Two notable examples where this is the case are know-your-customer applications in financial services, and regulatory compliance in financial services and health care. Large banks and insurers may have thousands of people doing these tasks, and much of the work is about inte-

Small-t innovations can help to make the value story real and make the case for investments that can reduce the perceived risk of larger opportunities.

grating and interpreting large amounts of unstructured information.

But getting GenAI to produce accurate results in a tightly regulated environment is very difficult. “In this space, the fruit is not so low,” said an executive from the financial services industry. What seems like a perfect application for GenAI is slowed by the need to build the right internal foundation of data and process before proceeding in these areas, where inaccurate decisions have high costs. If generative AI is going to create an enduring transformation of the business, “it doesn’t go hand in hand with the gold-rush mentality that we have to mine it now,” said Prem Natarajan, chief scientist and head of enterprise AI at Capital One. Instead, companies should think this through, he advised: “Figure out how to do it thoughtfully and responsibly. Build the scaffolding to bring everybody along.”⁶

Maintain a long-term perspective. The gold-rush mentality is real, but costs and uncertainty are too. “The transformative cases take longer to build the business case, test the models, change behaviors, etc.,” said Chris Bedi, chief customer officer at software company ServiceNow. “The challenge is not only technical but also leaders taking time to reimagine their future with big ideas.” Leaders are trying to separate hype from reality while also understanding that they may need to invest in data and technical foundations before they can get the returns they seek. Applications such as customer service and personalized shopping can show real innovations with material returns, but more complex work is needed to make broader change happen.

Companies are already using GenAI to pursue small-t transformation nearer to the bottom of the risk slope. For larger transformations, GenAI will be one of many pieces in the puzzle. “Instead of one transformative thing, we’ll stitch together many technologies, including AI, to rein-

vent a whole process,” said FM’s Tofte. Although it may take time before your company feels ready to launch transformations higher on the risk slope, you need not wait to make progress. You can experiment on some tasks while making foundational investments in data and integration that will make larger transformations possible over time. Choose the areas where you want to invest, both in the short and long terms. Then invest in building awareness and cross-cutting capabilities that can make you faster and more efficient in the future. ■

Melissa Webster is a lecturer in managerial communication at the MIT Sloan School of Management. **George Westerman** is a senior lecturer at the MIT Sloan School of Management and a coauthor of the award-winning book *Leading Digital: Turning Technology Into Business Transformation* (Harvard Business Review Press, 2014).

REFERENCES

1. G. Westerman, D. Bonnet, and A. McAfee, “Leading Digital: Turning Technology Into Business Transformation” (Boston: Harvard Business Review Press, 2014).
2. Z. Cui, M. Demirer, S. Jaffe, et al., “The Effects of Generative AI on High Skilled Work: Evidence From Three Field Experiments With Software Developers,” SSRN, published Sept. 5, 2024. <https://ssrn.com>.
3. E. Brynjolfsson, D. Li, and L.R. Raymond, “Generative AI at Work,” working paper 31161, National Bureau of Economic Research, Cambridge, Massachusetts, April 2023.
4. “CFO Signals,” first quarter 2024, PDF file (London: Deloitte Development, 2024), <https://www2.deloitte.com>.
5. S. Randazzo, H. Lifshitz-Assaf, K. Kellogg, et al., “Cyborgs, Centaurs and Self Automators: Human-Genai Fused, Directed and Abdicated Knowledge Co-Creation Processes and Their Implications for Skilling,” SSRN, Aug. 8, 2024, <https://ssrn.com>.
6. L. Gevelber, M. Eijsackers, and P. Natarajan, “Skills, Jobs, and LLMs — How Companies Are Redesigning Work as They Deploy AI,” panel discussion moderated by B. Armstrong, Sept. 17, 2024, at the 2024 MIT Digital Technology and Strategy Conference, Cambridge Massachusetts, video, 46 min., 21 sec., <https://ilp.mit.edu>.

Reprint 66301. Copyright © Massachusetts Institute of Technology, 2025. All rights reserved.



Every enterprise AI roadmap must start with AI understanding how the business works

By Carsten Thoma, President, Celonis



About the author:

Carsten Thoma is President of Celonis, partnering with co-CEOs Alex Rinke and Bastian Nominacher. An early investor, board member, and advisor since 2016, he joined in 2022, bringing over 25 years of industry experience. He leads strategy, marketing, ecosystem, and culture, while advancing innovation through Celonis Labs.

It's difficult to overestimate the impact of AI on business. The predictions — from trillion-dollar GDP boosts to radically reshaped workforces — are both mind-blowing and convincing. On the scale of the Industrial Revolution and digitization in centuries past, AI will define business in the twenty-first.

But AI's potential is not inevitable. Only one in twenty enterprise AI pilots is achieving measurable P&L impact, reports [MIT NANDA](#). There's a huge gap between AI's hype and real-world execution.

At its heart, AI is a decision engine. Give AI the data and context it needs to know how the business runs, and you empower it to autonomously make smart decisions that produce value. Give AI insufficient data, and it will make flawed decisions.

The data that AI needs lives in the systems that underpin enterprise processes: in systems of record like ERPs, data lakes, emails, spreadsheets, and so on. Hidden in all of this is the blueprint for how the business works.

Free the process, free AI's potential

In the past few decades, “digital transformation” and “enterprise modernization” have driven companies to continuously bring in more systems, upgrade them, and migrate to new ones. There was a time when the best-run businesses could benefit from processes in a rigid system. But that was before the internet, and long before the cloud, automation, and AI.

Now, processes and value chains, like the supply chain, run across dozens or hundreds of pieces of technology. Eventually, the systems started to dictate the process. Suddenly, companies had to settle because of system limitations.

How do businesses expect to deploy AI on top of this? We dream of autonomous agents in our supply chain that predict and react to disruptions before they happen. But if we ever hope to build them, businesses need to free their processes from their systems. In fact, 58% of business leaders are concerned the [current state of their processes](#) may limit what they can do with AI.

It all comes back to processes

Once enterprises solve this problem and are able to build AI agents on business understanding, that is step one to achieving the much sought-after agentic, autonomous enterprise.

Then, orchestration will be critical: the ability to monitor, fine-tune, and steer a workforce of agents at scale. Just as HR has developed as a practice for developing and supporting human labor, AI orchestration will do the same for agentic workers.

The promise of AI is real but not guaranteed. Mastering AI depends fundamentally on mastering business processes and their context. The companies that do this successfully are primed to enlighten enterprise AI and unlock its full value.

Just like the great enterprise revolutions of centuries past, those who succeed will be the ones who truly know their business, and therefore how to transform it for the future.

PDFs · Reprints · Permission to Copy · Back Issues

Articles published in *MIT Sloan Management Review* are copyrighted by the Massachusetts Institute of Technology unless otherwise specified.

MIT Sloan Management Review articles, permissions, and back issues can be purchased on our website, **shop.sloanreview.mit.edu**.

Reproducing or distributing one or more *MIT Sloan Management Review* articles **requires written permission**.

To request permission, use our website **sloanreview.mit.edu/store/faq** or email **smr-help@mit.edu**.