Deloitte.

Together makes progress

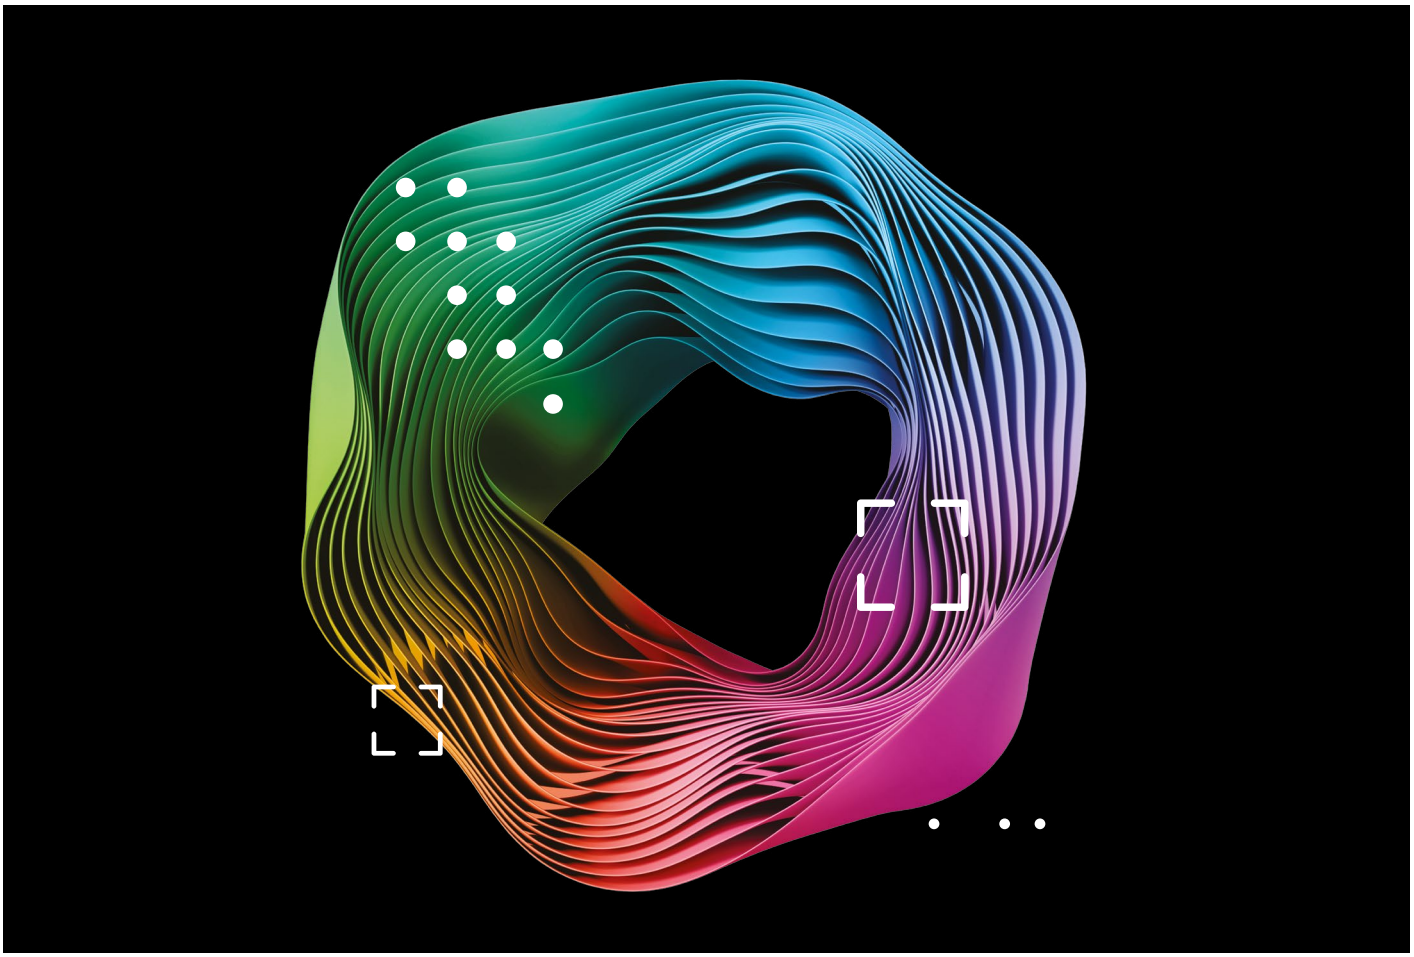

Talk to my agent

Agentic AI Leadership Insights Series – Part 2

Agentic AI is reshaping workflows, shifting cost structures, and challenging long-held assumptions about how work gets done. Layoffs attributed to AI sit alongside claims that these systems will empower people rather than replace them, creating a conflicting narrative. What matters for leaders is cutting through the noise and identifying how leveraging agentic AI will translate into value for their businesses. There will be no one-size-fits-all solution: agents will appear in many forms, from personal assistants to shared semi-automated workflows to deeply embedded enablers of re-engineered processes. They will eliminate some tasks, augment others, and open doors to entirely new forms of work. And because the technology behind them is evolving at extraordinary speed, each successive wave will reset expectations.

Implementation with Purpose

Introducing agentic AI just for the sake of it, to "ride the wave" or "catch up with competition", will almost certainly yield lackluster results. This is no fault of the technology, but of the lack of direction in which it is applied. For agentic AI, or any technology, to be effective, it must be applied deliberately, with purpose, and with clear, measurable goals[1]. The ultimate metric will be return-on-investment (ROI). Organizations must take care to approach ROI with perspective. Limiting the focus to reduced headcount or hours saved would overlook equally valid opportunities for growth and customer loyalty.

[1] Gartner, Gartner Predicts Over 40% of Agentic AI Projects Will Be Canceled by End of 2027, 2025. Gartner attributes projected failures primarily to poor strategic alignment, unclear objectives, and weak governance, reinforcing the need to anchor agentic AI initiatives in well-defined business outcomes rather than hype-driven adoption. https://www.gartner.com/en/newsroom/press-releases/2025-06-25-gartner-predicts-over-40-percent-of-agentic-ai-projects-will-be-canceled-by-end-of-2027

A more meaningful view of ROI requires a disciplined baseline and a broader understanding of value, identifying leading indicators of longer-term ROI that can be felt in business operations now: quality (rework rates), cycle times (throughput), risk reduction (disruptions), and agility (time-to-decision). Without such a structure, conversations around ROI will lack focus and the ability to pinpoint potential for further improvement.

Use Cases Everywhere

There are three reasons for the proliferation of agentic AI use cases. First, agentic AI is a versatile concept covering a wide scope of possible applications. Second, successful applications of agents are anchored in specific tasks, datasets, and systems. They are optimized to perform specific work. Even the impressive features of "generic chatbots" can be reduced to a number of specific agents: planners, web-scrapers, interpreters, summarizers, which together produce ever more relevant and reliable results. Third, the "democratization" of AI continues apace, now extending to a developer community that need not master software programming or engineering skills. Basic users can prompt their personal assistants. Citizen developers can build workflows with no- or low-code platforms. Only re-engineered, AI-fueled enterprise processes remain the domain of expert developers. Agentic AI comes in a variety of forms and offers a broad spectrum of implementation possibilities.

Spotlight: The Agentic Evolution of the Chatbot

Chatbots like ChatGPT or Gemini offer a useful example of how agents have substantially enhanced generic Q&A functionality. It has evolved from providing responses based on interpolation across a static knowledge base to richer reasoning, incorporation of the latest information for broader context, and now multi-step operations supported by coordinated agents. Asking a chatbot to perform a basic search, say on the relative pros and cons of no- or low-code platforms, today unleashes a team of dedicated agents to provide the best possible answer, for example:

1. **Interpretation:** understanding the user's request
2. **Planning:** translating the request into a series of steps to deliver a highly relevant and reliable answer, in this case defining more precisely the search criteria and strategy
3. **Search:** translating strategy into key terms for optimized search results and scraping those results
4. **Summarization:** isolating the most relevant components and synthesizing them into a single, coherent response
5. **Citation:** enhancing transparency by inserting links to sources for the major points
6. **Self-examination (iterative):** checking the accuracy of the final result, including the existence of links and the fit between response and original question

This pattern mirrors how enterprise use cases will unfold. Common archetypes include:

- **Complex tasks executed through existing tools:** "Find my Munich contacts, check their availability, and schedule meetings."
- **Analysis with deeper context:** "Validate this dataset by comparing entries to their cohorts, cross-referencing documents against accepted external sources."
- **Operating sophisticated systems:** "Plan the safest and fastest route for an e-vehicle, considering real-time and historical traffic and charging constraints."
- **Investigative and intelligence tasks:** "Correlate CCTV sightings, citizen reports, payment activity, and border data to track down a specific crime suspect."

These are high-level patterns. The actual shape of each agent will vary dramatically by organization.

Upheaval in the Working World

Jobs and human workers will be differently exposed to, and impacted by, agentic-enabled automation. The effects will differ even within the same function. Controllers offer a useful example. Controllers in two organizations may have the same title but face very different day-to-day realities. In mature, high-volume environments, much of their work is already standardized and supported by automation. In bespoke, low-volume settings, controllers still build structures manually, resolve exceptions, and piece together data from scattered sources. Such work is too ad hoc to expect meaningful automation from the rigid, deterministic systems of the past. Two people with the same title, but in different organizations, may find their roles affected by agentic AI in vastly different ways, or to different degrees[2].

Change is in the Air

The "Magnificent 7" dominate headlines about AI productivity, but their hiring and firing patterns have always been cyclical. Some workforce reductions may indeed reflect early automation. Others may reflect a carefully managed narrative more than operational reality. Leaders should be cautious when drawing lessons from these signals. Yet it would also be unwise to dismiss all as merely hype. AI – especially agentic AI – is driving changes everywhere: some gradual, others abrupt and convulsive. The 300-page "Trump America AI Act", besides seeking to replace myriad state laws with a single federal regulation, takes AI-driven job displacement seriously. The proposed federal legislation obliges companies to report AI-related layoffs to the Department of Labor (or face penalties for non-compliance), thereby establishing the first official substantiation on the topic.[3] Low- or non-value-added work will be the first to be automated away, leaving employees freed from such tasks to focus on higher-value work. Results may be, in the case of Controlling, more timely and better-informed decisions for business leaders. They may also require less capacity, or unleash capacity to fuel expansion into new business activities or sectors. Early adopters of the technology may gain a substantial lead over their competition, shifting work from non-agentic to agentic-enabled companies. Competition may come from unanticipated new players, such as start-ups or tech firms, rather than from incumbents. This has happened before with e- and m-commerce, and it will happen again. How change will manifest will vary widely, but change is coming … and has already begun.

Priority moves to consider

1. **Establish clear baselines.**
   Define measurable before-and-after metrics for any agentic deployment, including quality and cycle time, not only cost.
2. **Start with high-variability workflows.**
   Agents excel where traditional automation struggled. Identify processes with many exceptions or manual handoffs.
3. **Build small, testable pilots.**
   Prove value with constrained scopes before scaling. Treat each pilot as a chance to refine your ROI framework.
4. **Tie agents to real decision making.**
   Look for areas where agents can not only execute tasks but also support faster, more informed judgments.
5. **Prepare for uneven impact across roles.**
   Do not assume a uniform effect within functions. Map exposure role by role, not title by title.

[2] Beyond Automation: Redesigning Jobs with LLMs to Enhance Productivity, 2025. The study shows that AI affects work at the task level rather than the job title level, leading to uneven impacts within the same occupation as tasks are selectively automated, augmented, or reallocated. https://arxiv.org/abs/2512.05659

[3] The Republic Unifying Meritocratic Performance Advancing Machine Intelligence by Eliminating Regulatory Interstate Chaos Across American Industry Act (TRUMP AMERICA AI) Act, Senator Marsha Blackburn, 18th March 2026, https://www.blackburn.senate.gov/services/files/C43D3B19-391B-4EB6-84C1-0FC37EEBBA4D

Your contact

David Thogmartin
Partner
aiStudio, AI & Data Analytics,
Trustworthy AI
Tel: +49 151 58072163
dthogmartin@deloitte.de

Deloitte.

Issue 03/2026

MAKING AN IMPACT THAT MATTERS since 1845