



The age of agents

Agentic AI Leadership Insights Series – Part 1

Agentic AI is advancing at a pace that has caught many leaders off guard. Adoption is already widespread, and by 2028 a significant share of enterprise software is expected to embed agentic capabilities. [Deloitte's latest 2026 State of GenAI report](#)¹ shows that 23 percent of companies are already using agentic AI at least moderately today, with 74% expecting to reach that level within the next two years. Predictions such as [Gartner's](#)², that roughly 15 percent of today's daily work will be automated by

agents, are no longer theoretical. They reflect a shift that is already underway. The central question: are organizations prepared for systems that not only inform and advise, but also act?

The first step toward answering this question is to define what constitutes agentic AI and how it differs from the AI that came before it. Doing so helps cut through the hype, revealing that agentic AI is a much greater prospect than many realize. It is a

pathway to a more broadly capable form of AI that holds the power to transform processes and businesses.

AI agents don't talk, they do

Agentic AI is about action. Traditional AI systems, including chatbots, make predictions and provide answers. Agents go further: once given an objective, they trigger workflows, call tools, and interact with systems to complete tasks end to end. ➔

¹ Deloitte, State of AI in the Enterprise, 2026. The report highlights accelerating enterprise adoption of advanced AI capabilities, including the transition from experimental use to broader, operational deployment across business functions.

<https://www.deloitte.com/us/en/what-we-do/capabilities/applied-artificial-intelligence/content/state-of-ai-in-the-enterprise.html>

² Gartner, Intelligent Agents in AI Really Can Work Alone. Here's How, 2024. The article describes agentic AI as systems capable of planning and executing tasks autonomously across tools and workflows, marking a shift from assistive AI to operational automation.

<https://www.gartner.com/en/articles/intelligent-agent-in-ai>

A new age of generic agents are upon us, from Anthropic's highly polished & powerful "Claude Cowork" to the unhinged, open-source "OpenClaw" (previously ClawdBot, MoltBot). Their differentiating factor: they execute, they find and address related topics, they don't necessarily ask for permission at every step. OpenClaw – despite many reported unintended outcomes – has gained a sizable following, especially in China, even by non-technical users. At its Global Technology Conference (GTC) 16th–19th March 20206, NVIDIA announced "NemoClaw", a more enterprise-friendly variant of OpenClaw based on its Nemotron LLM and including numerous security guardrails.

Agentic AI is also about collaboration. What appears to the user to be a single agent is most often a coordinated group of dedicated sub-agents, tools and micro-services, each focused on specific tasks that, together, solve a larger problem. These multi-agent systems combine modularity and interactivity, enabling them to take on work that would be either too complex or too abstractly defined for a single model to handle.

Agents are already extending the capabilities of chatbots – and enriching the experiences of their human users. The true potential is apparent when agents are

embedded into workflows, enabling automation at scale or even veritable "digital co-workers." Where chatbots are limited to responding to individual user questions, leaving a process otherwise unchanged, agents can execute full process chains with minimal human intervention. From a user perspective, they effectively shorten these process chains from, say, twenty steps down to two: (1) input data, (2) check quality of output. This promises nothing short of a revolution in productivity.

Example: A procurement manager asks an AI agent to consolidate all open contracts, check renewal dates, benchmark pricing against historical data, and draft negotiation points. A traditional system might retrieve documents. An agent completes the workflow: gathers the data, runs comparisons, drafts recommendations, and prepares a ready-to-send email package. The difference is execution, not just assistance.

Agentic differs from earlier forms of AI

As with earlier forms of AI, agentic systems are probabilistic: they can better deal with ambiguity or loosely-defined objectives better than hard-coded, deterministic systems. Agents are powered by "reasoning engines" (specialized GenAI Large Language Models), giving them the ability to interpret. The difference vs GenAI lies in what agentic AI does with that interpretation.

A well-designed multi-agentic system can:

- break down a problem,
- design a sequence of actions,
- identify which data, documents, or tools it needs,
- select or even generate the right tool for the task,
- execute the steps,
- and assemble the final output.

This is possible because the underlying models are tuned not only to generate text but to reason, plan, and coordinate. In most architectures, these capabilities sit with the orchestrator, the component that behaves like a virtual project manager. Humans interact with the orchestrator, which then assigns work across other agents and supporting tools.



Agents are smart, but not intelligent

It is tempting to call these systems intelligent. They are methodical, even creative, and can appear self-directed. Yet the so-called “reasoning engines” are not truly, generally intelligent. Instead, they work by simulating reasoning, effectively “thinking out loud” through a process. Beneath that behavior, they are still performing a sophisticated form of next-step prediction. They can produce impressive results, yet they can also miss context, rely on faulty assumptions, or generate outcomes that look plausible but are fundamentally wrong. Accuracy can be improved through various techniques, such as larger context windows (bigger models), incorporating “agent skills” and adding back-end quality controls, such as using a secondary model as a reviewer. As an ensemble, these can greatly increase robustness, although they do not entirely eliminate inaccuracy risk.

While agent teams can handle more complex tasks, not all architectures are equally effective. Investing too many “agent skills” into an orchestrator will create a top-heavy, confused agentic system. The converse is also true: research by Berdoz, Rugli, and Wattenhofer at ETH Zurich finds arbitrary agent “swarms” to be ineffective.³ Nobody has all the answers: another piece by Gloaguen/Mündler/Müller/ Raychev/ Vechev at ETH Zurich challenges the “logging” method (AGENT.md) to improve performance, championed by Anthropic (CLAUDE.md).⁴

Perhaps the greater danger is the false confidence that precise and thorough (although potentially inaccurate) results can instill in the human user. The speed of execution becomes the speed of error

when governance is weak or oversight is absent. Automation removes humans from the details, but it cannot remove accountability. Users remain responsible for the quality of their work, whether produced with or without agentic AI. Leaders remain accountable for their organizations, ensuring an even greater degree of human oversight than for employees exercising human judgement, especially when it comes to work of material impact.

Agentic AI can and must be implemented responsibly

To realize value without creating unnecessary risk, organizations should anchor their efforts in a few practical principles:

- Set clear, measurable outcomes. Without them, ROI becomes guesswork.
- Re-imagine processes, not just automate them. Agentic AI performs best when built into redesigned workflows, not bolted onto legacy ones.
- Invest as much in governance as in automation. Guardrails, review mechanisms, and escalation paths matter as much as the agents themselves.



Priority moves to consider

A practical next step depends on your role:

- CEOs should identify two or three enterprise workflows where end-to-end automation could unlock immediate value.
- CIOs should assess which existing systems can support agent orchestration and where integration work is needed.
- COOs should map processes constrained by manual effort and validate which stages are ready for autonomous execution.
- CHROs should plan for capability uplift, defining where human control remains essential and where agents augment capacity.

³ “Can AI Agents Agree?” Frédéric Berdoz, Leonardo Rugli, Roger Wattenhofer, ETH Zurich, 12th March 2026, <https://arxiv.org/pdf/2603.01213>

⁴ Evaluating AGENTS.md: Are Repository-Level Context Files Helpful for Coding Agents?, Thibaud Gloaguen/Niels Mündler/Mark Müller/Veselin Raychev/Martin Vechev, ETH Zurich, 12th February 2026, <https://arxiv.org/pdf/2602.11988>

Your contact



David Thogmartin

Partner

aiStudio, AI & Data Analytics,

Trustworthy AI

Tel: +49 151 58072163

dthogmartin@deloitte.de

Deloitte.

Deloitte refers to one or more of Deloitte Touche Tohmatsu Limited (DTTL), its global network of member firms, and their related entities (collectively, the "Deloitte organization"). DTTL (also referred to as "Deloitte Global") and each of its member firms and related entities are legally separate and independent entities, which cannot obligate or bind each other in respect of third parties. DTTL and each DTTL member firm and related entity is liable only for its own acts and omissions, and not those of each other. DTTL does not provide services to clients. Please see www.deloitte.com/de/UeberUns to learn more.

Deloitte provides leading professional services to nearly 90% of the Fortune Global 500® and thousands of private companies. Legal advisory services in Germany are provided by Deloitte Legal. Our people deliver measurable and lasting results that help reinforce public trust in capital markets and enable clients to transform and thrive. Building on its 180+-year history, Deloitte spans more than 150 countries and territories. Learn how Deloitte's over 470,000 people worldwide work together every day to make an impact that matters at www.deloitte.com/de.

This communication contains general information only, and none of Deloitte GmbH Wirtschaftsprüfungsgesellschaft or Deloitte Touche Tohmatsu Limited (DTTL), its global network of member firms or their related entities (collectively, the "Deloitte organization") is, by means of this communication, rendering professional advice or services. Before making any decision or taking any action that may affect your finances or your business, you should consult a qualified professional adviser.

No representations, warranties or undertakings (express or implied) are given as to the accuracy or completeness of the information in this communication, and none of DTTL, its member firms, related entities, employees or agents shall be liable or responsible for any loss or damage whatsoever arising directly or indirectly in connection with any person relying on this communication. DTTL and each of its member firms, and their related entities, are legally separate and independent entities.