



The agentification of the enterprise

Navigating enterprise
transformation with agentic AI

October 2025



About the Deloitte AI Institute

The Deloitte AI Institute™ helps organizations connect the different dimensions of a robust, highly dynamic and rapidly evolving AI ecosystem. The Institute leads conversations on applied AI innovation across industries and offers cutting-edge insights, all to promote human-machine collaboration in the “Age of With.”

The Deloitte AI Institute aims to promote a dialogue and development of artificial intelligence, stimulate innovation, and examine both challenges to AI implementation and ways to address them. The Institute collaborates with an ecosystem composed of academic research groups, startups, entrepreneurs, innovators, mature AI product leaders and AI visionaries to explore key areas of artificial intelligence including risks, policies, ethics, future of work and talent, and applied AI use cases. Combined with Deloitte's deep knowledge and experience in artificial intelligence applications, the Institute helps make sense of this complex ecosystem, and as a result delivers impactful perspectives to help organizations succeed by making informed AI decisions.

No matter what stage of the AI journey you're in, whether you're a board member or a C-suite leader driving strategy for your organization or a hands-on data scientist bringing an AI strategy to life, the Institute can help you learn more about how organizations across the world are leveraging AI for a competitive advantage. Visit us at the Deloitte AI Institute to access the full body of our work, subscribe to our podcasts and newsletter, and join us at our meetups and live events. Let's explore the future of AI together.

www.deloitte.com/us/AIInstitute



Content

Key takeaways

- Agentic AI is revolutionizing workflows by enabling autonomous decision-making and task execution, freeing human talent to focus on strategic and innovative activities.
- The integration of agentic AI into enterprises is driving significant transformation, enhancing efficiency, productivity, and the ability to scale operations dynamically.
- Organizations must adopt a phased approach to agentification, balancing gradual implementation with bold experimentation to fully realize the potential of digital labor.
- Effective risk management and workforce engagement are critical to the successful deployment of agentic AI, ensuring secure, ethical, and sustainable transformation.

Introduction	4
The value and function of agentic AI	5
A path to agentification	6
Gradual implementation and bold experimentation	
Assessing cost and speed to value	7
Better together	8
Fostering change and adoption in the workforce	
Risk mitigation in the agentic enterprise	9
A bold journey ahead	10
Get in touch & Endnotes	12



Introduction

Organizations across industries are at the dawn of a step change in how work is accomplished: agentic AI. The scale of opportunity and disruption necessitates a reimagination of the enterprise to use technology in a way never before possible.

With agentic AI, workflows are designed with an agent-first approach. Reasoning becomes a joint human-machine endeavor, with humans serving as architects, overseers, and stewards, rather than users. Skilled talent is elevated to tackle new business problems, steer strategic decisions, and reimagine how business gets done.


Alongside traditional value measures such as revenue growth, efficiency, and productivity, the grander value with agentic AI is found in its disruptive, transformational impact on the future of functions and how digital labor is used. What new business models can emerge when digital labor is as essential to operations as people? What differentiators arise when process constraints are removed entirely? What new products and services will be possible when human ingenuity is unleashed because tedious and routine workflows are no longer a part of day-to-day work?

This opportunity calls for bold action.

According to a 2024 Forrester survey, more than 80% of automation leaders expected to accelerate AI agent investments over 2025,¹ and by 2028, a third of enterprise software applications were forecast to include agentic capabilities.² This fast-growing interest is expected to grow the global agentic AI market from US \$3.7 billion in 2023 to US \$103.6 billion by 2032.³

The race to transform has begun, and there are new factors to consider in strategy, cost, adoption, and risk management. What is more, in this heady period of intense innovation, there remains some uncertainty about just what agents can enable and what advancements are still to come.

Still, even without perfect foresight, enterprises can and should begin the agentic journey to work reimaged.





The value and function of agentic AI

Generally, AI applications automate, improve, or support a discrete task or subtask. Agents are different. Agentic AI systems set tasks, make decisions, and act autonomously to achieve defined goals. They are reasoning engines that can understand context, plan workflows, connect to external tools and data, and act with limited or no human intervention.

An LLM-enabled chatbot can help a user plan a trip; an agentic system can plan the trip, book the flight and hotel, and schedule a car—all without being explicitly told how to do so. AI agents achieve this by leveraging a set of technologies: large language models (LLMs), memory, planning and reasoning capabilities, tools to interact with software and data, and feedback loops to improve function.

Agentic systems can be conceived on a trajectory of autonomy.

- **Semi-autonomous**

Agents receive goals set by a user and pursue them by identifying and performing subtasks. To do so, they use short-term memory and context management, and the outputs require frequent human validation. Semiautonomous agents are suited to supporting human work, such as a co-pilot agent making recommendations and automating research and writing. These are the agents primarily used today.

- **Contextually autonomous**

Agents maintain an internal memory to accommodate longer tasks, iteratively adapt, and improve over time. These digital workers, embedded in reimagined processes, require less human intervention, although a human in the loop is still important in decision-making and task execution. With contextual autonomy, agents handle unstructured tasks by interpreting context, planning and changing activities, and drawing from multiple data sources. One example is an AI assistant managing multiple tasks in parallel on behalf of a human.

- **Fully autonomous**

Agents operate independently to plan and execute open-ended tasks and collaborate with other agents or tools. Rather than applying technology for human-designed steps and activities, fully autonomous agents build their own processes optimized against existing systems using reinforcement learning. There is limited or even no human intervention. With this highest degree of reasoning and autonomy, agents orchestrate work, and humans serve as overseers of multiagent operations.

This expected evolution to greater degrees of agent autonomy is dependent in part on technology innovation and maturity. The technologies to enable fully autonomous agents are still being developed and refined. As with any emerging technology, organizations need to experiment, monitor ongoing innovation, and explore where current agent capabilities could enable transformative impact.

Exemplifying the agentic impact

Agents open the door to new products, offerings, and growth vectors. Consider some examples.

- **Tailored offers for customer satisfaction:** A major phone carrier deployed a network agent to track customer data usage and location. When a customer nears their data cap or is roaming abroad, the agent independently creates and presents a bundled offer to purchase more data. The customer can opt in with one click, increasing revenue while also improving customer satisfaction.
- **Agent concierge for a human journey:** A major commercial airline deployed an agentic concierge in the company's app, which anticipates traveler needs and orchestrates actions across a trip, including proactively sending alerts, providing real-time wayfinding at the airport, suggesting rebookings or alternative transport in case of delays or cancellations, and connecting loyalty accounts with other transportation partners.
- **New revenue in aftermarket sales:** A leading automaker rolled out an AI-driven service that continuously monitors vehicle health and offers maintenance and service bundles to drivers before issues worsen. Using real-time vehicle diagnostics, the agent creates micro promotions for service bundles that are adjusted based on mileage, driving patterns, seasonality, and even location. Agents recalibrate bundle pricing and discounts based on predicted customer lifetime value and engagement signals.



A path to agentification

Gradual implementation and bold experimentation

Organizations on the agentic journey should anticipate a multiphased approach. **Taking thoughtful, gradual steps toward greater levels of agent autonomy conditions the organization for more expansive transformations still to come.** It is partly a matter of becoming comfortable, as an organization, with new ways of working.

As an orienting perspective, consider, “If we were to create our organization for the first time, how would we design workflows and assign tasks in light of agentic AI?”

This starting point leads to questions about which workflows may be suited to agentification. Look at where time is spent in the organization, where there is a high variability of results, and where automation would result in financial improvement.

Taking an iterative, steppingstone approach may be appropriate for some enterprises and functions. At the outset, evaluate the enterprise’s technology architecture and integrate the necessary tools to prepare for the journey. Depending on technology readiness, leaders can explore pilot demonstrations to test (in a limited, controlled way) whether the agent and reimagined workflow delivers the expected return. Enterprises can also look to layering agents into existing platforms and solutions, injecting

new intelligence into current systems. It sets the stage for increasing degrees of agent autonomy and helps the organization gain experience and confidence in reimagining human-agent workflows. Finding process-specific benefits, the enterprise can consider where else agents could be deployed and which other processes may be suited to agentic AI. Step by step, the organization uses agents to reimagine how work is performed and build toward a widespread use of digital labor.

To be sure, an iterative approach could slip into simply improving how work is performed today, rather than reimagining work against a set of outcomes. For some organizations, a more radical approach may be enticing, and the path ahead is one of bold experimentation and deployment around high-value, differentiating tasks. This approach could accelerate agentification, and as agents learn, they can derive optimized workflows that may not have been conceived through human planning and process mapping. Learning by doing can put agents (and the enterprise) on a faster path to widespread agent use, albeit one that may carry more risk around adoption and process gaps.

Ultimately, enterprises might consider an all-of-the-above approach, taking incremental steps in some areas while also experimenting to discover the fuller transformational impact in others.

Two approaches to agentification



Iterative approach

Navigating stepping stones in a strategic, incremental manner

Level 1

Elevate understanding of AI agent potential and establish foundational AI infrastructure

Explore AI potential

Level 2

Develop and pilot single-task AI agents and begin to integrate with existing enterprise systems

Optimize operational efficiency

Level 3

Integrate multifunctional AI agents across key business functions, optimizing operations and enhancing productivity

Elevate enterprise capabilities

Level 4

Leverage autonomous and multi-agent systems to drive breakthrough innovation, new business models, and automation

Transformative innovation



Radical approach

Taking the leap towards a fully agentified enterprise



Assess

- Identify candidate processes for agentification

Assess tech architecture and embed agent platforms as the backbone of the new design



Action

- Reimagine processes for autonomous agents handling end-to-end workflows

Layer agents into core systems with continuous intelligence



Accelerate

- Deploy agents deliberately across the enterprise

Drive human-agent collaboration via workforce trainings and change management initiatives



Assessing cost and speed to value

There are complex trade-offs between workforce costs, operating costs, and the cost to acquire and maintain the technology itself. At a high level, automating a significant portion of day-to-day operations will lead to lower labor costs. However, there will also be new costs for compute and technology maintenance. There are profit and loss calculations to perform, and organizations should expect operating expenditures to change significantly. Relevant factors include cost-per-token, cloud compute costs, and the ability to reuse agents across processes. As agents take over reimagined end-to-end processes, there may be savings and productivity benefits that can help fund the endeavor.

Organizations will need engineers, process experts, and technical talent to build and sustain agentic systems.


Weighing talent availability against speed to transformation, organizations will need to think through whether agentic AI development should be performed internally or with a vendor. Building internally could afford differentiation by leveraging proprietary data and intellectual property. Alternatively, acquiring agentic capabilities through vendors could accelerate adoption, and marketplace options are proliferating.

The vendor ecosystem is investing heavily in agentic offerings, including embedding agentic functionality in existing platforms and solutions. As such, leaders should engage vendors to learn more about product roadmaps, anticipate when new capabilities will become available, and encourage investment and participation in agentic proofs of concepts and experiments.

Perhaps the most significant ongoing cost is transforming work itself. Reimagining processes across functions is a complex undertaking replete with uncertainties and gaps. The organization needs to understand how its human workforce currently performs roles and tasks, the friction points and broken manual loops, and the disparate structured and unstructured datasets associated with years of work. Process experts with an agent-first mindset are critical in redesigning work. These professionals "should understand process reengineering and bring enough technical depth to think boldly about the most impactful approaches to unite humans, machines, and processes for new ways of working.

Here too enterprises may look to a vendor for value-based approaches to transformation. Organizations may consider temporarily handing off processes to a vendor for agentification or outsourcing the effort long term.

The **speed to value** will depend on the **complexity of the processes** being automated and the **maturity** of the organization's existing technology stack.





Better together

Fostering change and adoption in the workforce


In a traditional technology-led initiative, workforce adoption is sometimes given lesser importance, treated via limited, short-term change management activities. That approach is not viable with agentic AI. The workforce will need to engage, over time, with new tools and processes.

In the agentic enterprise, human workers will need to participate in the reimagined workflows and not regress to old ways of working. The challenge is one of changing behaviors. Adoption grows out of an enterprise culture that embraces disruption and celebrates the transformation of work. The organization needs to be able to articulate to the workforce how their work and roles will change and why. The value narrative should help employees see how the shift in work impacts and benefits them.

One reason workforce buy-in is important is that human expertise is necessary to prepare agents for a reimagined process. Agent design and training is based in part on how humans reason and

act, and knowledge-based workers will need to be incentivized to impart their subject matter and process expertise for agent development. **In doing so, employees should have a sense of how digital labor creates a more rewarding future of work, whether it is increased capacity for higher-value tasks or even the opportunity to take on a new role entirely.**

Upskilling and continuous learning are also factors. Because agentic AI will come to touch every part of the enterprise, organizations will need to embark on broad-based upskilling across the workforce. At the same time, new capabilities and opportunities will emerge as agentic AI matures, and human workers will need ongoing AI learning to keep up with developments. With up-to-date knowledge, workers understand agent capabilities and risks, which can lead to creative thinking about where else agents might solve a business problem. It might also inspire workers to consult their human ingenuity and expertise to solve a problem AI cannot.





Risk mitigation in the agentic enterprise

Agents introduce a new class of enterprise risks that scale nonlinearly across decision chains and interfaces. Some risk owes to how agentic AI functions. For example, cybersecurity becomes more complex. Connecting more data and systems with on-premises, cloud, and third-party technology creates new potential vulnerabilities. In addition, long-context agents create persistent memory exposure, yielding a larger attack surface for cyberthreats.

When Generative AI emerged, it brought with it the risk of hallucination, whereby an application computes an inaccurate output with total confidence. **With agentic AI, operational risk is elevated because the hallucination risk compounds.**

An inaccuracy at some point in the system can percolate throughout the process, propagating and expanding errors. Agents granted greater autonomy could inappropriately trigger disruptive or even destructive actions due to inaccurate computations. In addition, as agents come to take on more processes, poor coordination between agent swarms can create message fidelity gaps, unseen conflicts, and compounding errors.

To help manage these risks, the human workforce needs a mindset shift from AI users to AI architects and stewards. Workers will have a larger oversight role in output validation, while also monitoring for unintended bias, security vulnerabilities, and other risks.

Throughout, decision rights are how the enterprise controls the level of autonomy in a process, and it becomes more complex as the number of agents grows. Early in the agentic journey and for each agent along the way, the enterprise needs to define where human judgment is required. It informs how the agent is designed, tested, and governed. Reimagined workflows should include controls to dictate when an output is discarded or accepted, with thresholds for when the agent defers to a human.

Overall, agentic risk management requires a holistic approach, including:

- Establishing a governance framework for AI agents
- Incorporating standards for agent communication, context management, and secure containerization
- Testing and validating agent decision logic and reasoning
- Proactive, real-time monitoring for system anomalies, security threats, and unexpected behaviors
- Maintaining a human-in-the-loop with the authority to override decisions and perform dynamic audits of agent behavior



A bold journey ahead

There are many entry points to the agentic journey. Opportunities to embark on widespread change arise when refreshing business strategy, performing a technology modernization initiative, reevaluating vendors, or moving forward after an M&A or large-scale reorganization. Yet, the true catalyst for the agentification of the enterprise is a bold strategy and eagerness to change.

The road ahead with agentic AI is in flux and subject to innovation. Still, with a bold reimagination of the business, organizations can take up the banner of change and pursue the transformed future of work.

Where **to begin** with agents?

- **Imagine an agent-first future state:** Define how a process could work using agents and without being constrained by the current state. Align strategic goals, identify opportunities for agents, and prioritize where to start based on expected impact.
- **Articulate the agent-enabled process:** In the imagined future state, document the tasks or workflows to be completed by humans and agents, the handoffs and coordination points, and the process KPIs to evidence value.
- **Experiment, pilot, and learn:** Stand up agent prototypes and run controlled pilots to explore agentic capabilities, validate technical fit, begin engineering human-agent workflows, and assess outcomes.
- **Shift into agentification at scale:** Move from pilot to production, deploying agents in redesigned workflows. Focus on workforce buy-in, upskilling and continuous learning, and creating the governance and controls to mitigate risks.



For a deeper dive on AI agents, multiagent AI systems and the future of work, download and read our previous reports:

[“Prompting for action: How AI agents are reshaping the future of work,”](#)

Deloitte, November 2024.

[“The cognitive leap: How to reimagine work with AI agents,”](#)

Deloitte, January 2025.

[“The rise of AI agents and collaborative automation,”](#)

Deloitte, July 2025.



Get in touch



Jim Rowan

Principal,
US Head of AI
Deloitte Consulting LLP
jimrowan@deloitte.com



Lou DiLorenzo Jr

Principal,
Tech, AI & Data Strategy Leader
Deloitte Consulting LLP
ldilorenzojr@deloitte.com



Tim Smith

Principal,
Monitor Deloitte Strategy Leader
Deloitte Consulting LLP
timsmith6@deloitte.com



Edward Van Buren

Principal,
GPS Applied AI Leader
Deloitte Consulting LLP
emvanburen@deloitte.com



Gregory Vert

Principal,
Deloitte Consulting LLP
gvert@deloitte.com



Baris Sarer

Principal,
Global AI Leader for TMT
Deloitte Consulting LLP
bsarer@deloitte.com



Mark Szarka

Senior Manager,
Deloitte Consulting LLP
mszarka@deloitte.com

Endnotes

1. Jessie Johnson, [AI Agents Are Your New Coworkers: Get To Know Them Better At B2B Summit](#), Forrester, December 18, 2024.
2. Tom Coshow, [Intelligent Agents in AI Really Can Work Alone. Here's How.](#), Gartner, October 1, 2024.
3. S&S Insider, [AI Agents market Report Scope & Overview](#), October 2024; full report report code SNS/ICT/4611, gated.



Deloitte.

About Deloitte

As used in this document, “Deloitte” means Deloitte Consulting LLP, a subsidiary of Deloitte LLP. Please see www.deloitte.com/us/about for a detailed description of our legal structure. Certain services may not be available to attest clients under the rules and regulations of public accounting.

Deloitte provides industry-leading audit and assurance, tax and related services, consulting, financial advisory, and risk advisory services to nearly 90% of the Fortune Global 500® and thousands of private companies. Our people deliver measurable and lasting results that help reinforce public trust in capital markets, enable clients to transform and thrive, and lead the way toward a stronger economy, a more equitable society, and a sustainable world. Building on its 175-plus year history, Deloitte spans more than 150 countries and territories. Learn how Deloitte’s approximately 457,000 people worldwide make an impact that matters at www.deloitte.com.

This publication contains general information only and Deloitte is not, by means of this publication, rendering accounting, business, financial, investment, legal, tax, or other professional advice or services. This publication is not a substitute for such professional advice or services, nor should it be used as a basis for any decision or action that may affect your business. Before making any decision or taking any action that may affect your business, you should consult a qualified professional advisor. Deloitte shall not be responsible for any loss sustained by any person who relies on this publication.

Copyright © 2025 Deloitte Development LLC. All rights reserved.