

Deloitte.



Work, reworked: What it takes to win in the age of Agentic AI

Agentic AI isn't just a powerful new tool—it's a transformative juggernaut that requires you to reimagine how work gets done.

Executive Summary

Imagine achieving a 10x leap in productivity without hiring a single additional employee.

It sounds impossible, but it's exactly what agentic AI is designed to do. Radical innovation, exponential productivity, next-level growth—they're all brought within reach by these digital workers capable of executing multi-step processes with minimal human supervision, in tight collaboration with humans and other agents. But doing so requires organizations to fundamentally rethink how work gets done in an agentic AI world—a tipping point for the future of work.

Because agentic AI isn't just the latest fleet of chatbots and AI pilots. It is a new approach to work. Workflows, roles, and organizational structures will need to be reassessed and redefined by any organization expecting to gain a competitive advantage from agentic AI.

Enter the Human-Agentic Workforce: Humans and AI agents working in concert.



Introduction

From the ground up

It's often said that if very old cities were to be built today, they would be designed entirely differently—different street grids, public transportation systems and routes, infrastructure layers, and so on. For practical reasons, city governments instead just work around these suboptimal, outdated elements.

In an agentic AI context, most organizations can be viewed as the equivalent of an old city. They are often constrained by the weight of busywork, inefficient work processes, workarounds, outdated organizational structures, and cultural inertia. Add AI agents to this environment and expect to merely replicate the inefficiencies and shortcomings of the organization at a different scale. That's the danger of taking an incremental adoption approach to agentic AI—eventually the organization will simply add an agentic layer to an outdated approach to work itself.

The AI evolution: From tools to transformative force

Traditional automation

Rules-based systems
that handle
repetitive tasks



Generative AI

Tools for individual
productivity
(like content creation)



Agentic AI

Autonomous systems
that orchestrate
multi-step processes
across workflows
through human-agent
collaboration (humans on
the loop) rather than
human micromanaging
(in the loop)


Avoiding this outcome to realize the full, true potential of agentic AI requires organizations to address three core questions:

1. What does an integrated human + AI agent workforce look like?
2. How do we redesign work, roles, and decision-making paradigms to align with this future?
3. How will we ensure our cultural, technological, and operational readiness to enable a new system of work?

Those are big questions with equally significant implications. Are they really necessary? Yes—because agentic AI doesn't just digitize information, it digitizes the entire system of work. Intelligent agents act with autonomy, and human oversight of this technology shifts from execution to orchestration. Humans and AI agents need to work in concert—and that is fundamentally different than any other moment in the history of work.

As leaders in organizations begin to recognize and embrace the radical promise of agentic AI, they can expect to encounter a range of familiar, deep-rooted barriers—starting with structural inertia, cultural resistance, and a lack of strategic readiness. This is the reality of operating in a world at a tipping point—where technological advances consistently outstrip most organizations' ability to adapt and scale. Which is why leaders cannot afford to take the same old incremental approaches to adopting agentic AI, dabbling around the edges, experimenting here and there, and hoping for some wins that can be replicated elsewhere.

Instead, this is the moment for bold leadership. Your next move on agentic AI will drive transformation or delay it. It's time to get human-agentic—from the ground up.



Making the leap: 5 key shifts to enable the Human-Agentic Workforce

Can humans and AI agents really work together as a team? Absolutely—in fact, soon it will be an imperative for competitive advantage, given how quickly agentic AI capabilities are evolving. But it won't happen until organizations reassess, rethink, and reengineer traditional models for how work gets done.

If that sounds like a big deal, it is—but it's also entirely within reach. It will require confronting outdated paradigms of control, capacity, and incremental change. Laying the groundwork for human-agentic change starts with these five shifts:


Technology-led transformation → Human-centric transformation

Most organizations today treat AI as another technology initiative, along the lines of the move to cloud computing. But agentic AI is not an IT program. It's a tool for workforce and business transformation—and that requires a fundamentally different approach. How will your organization use AI to amplify human potential, not just automate processes? The answer requires flipping the standard focus: Human impact is the driver, and technology is the enabler.

Task-based roles → Outcome-based design

Traditionally, workplace roles have been built around fragmented tasks. In an AI-powered system, roles must be redesigned around outcomes, capabilities, and human-agent collaboration—from static jobs to dynamic, skills-based role ecosystems. This requires a system-level redesign of how work is structured and roles are defined.

When digital agents execute workflows at greater speed and scale, humans are free to focus on judgment, empathy, and creativity—all of which contribute to differentiated outcomes in areas like customer value, adaptive operations, and sustainable growth. That only happens when roles are designed in ways that allow both humans and digital agents to thrive.



Digitized information and processes → AI-native business and operating models

Digitizing documents, automating workflows, layering new technologies onto legacy structures—these were all important first steps in enabling the transition to AI. But agentic AI requires more—building business and operating models from the ground up to facilitate deep, seamless human-agent collaboration. It's about designing for what's possible: fluid, adaptive, AI-native systems that scale impact rather than amplify existing inefficiencies.

Capacity-fueled growth from headcount → Scaled performance through redesign

More people = more output? That's no longer a sustainable or effective way for organizations to grow today. In the human-agent era, achieving growth through scale happens as a result of intelligent work redesign rather than workforce expansion. In this model, growth is unlocked by redesigning work, orchestrating tasks in ways that allow humans and agents to drive exponential outcomes that matter most to the organization—*together*.

Cultural inertia → Organizational readiness

Fear of job displacement. Siloed governance. Lack of trust. For agentic AI, these familiar challenges are strategy killers. That makes organizational readiness a core differentiator, enabling the ability to build trust, embrace ambiguity, and scale learning faster than the pace of technological change. Leaders in agentic AI will create the environment and culture needed for humans and agents to co-evolve and thrive together.

These are not simple mindset shifts. Instead, they are systems-level transformations in how enterprises think, lead, organize, and perform—and none can be solved in isolation.



Beyond cross-functional collaboration: Systemic codependency

No single domain can act in isolation in an agentic AI world—not technology, operations, talent, customer, finance, risk, or legal. This means that cross-functional collaboration—typically collaborating across silos—isn't sufficient for enabling successful agentic AI contributions at scale. Systemic codependency, in which progress in one area is only possible through alignment and shared ownership with others, is needed. Each part of the enterprise is a mutually dependent driver of AI-enabled progress.

How does this work in practice?

These three actions provide the most direct path to systemic codependency:

+ Reimagine leadership and governance

CIOs, CHROs, Chief Risk Officers, and business leaders must operate as a connected leadership system, aligning ambitions, pace, and accountability across their domains—because the agentic future can't be achieved through silos. This requires shifting from managing people to managing outcomes, supported by human-agentic teams.

Bold governance intervention—in the form of new structures for overseeing human-agent collaboration and managing trust, risk, and performance across human-agentic workforces—is the foundation for this effort.

+ Redesign work from the ground up

Legacy work models can't scale in an agentic world. Rather than inserting AI into existing roles, organizations must deconstruct and reassemble work to unlock new human-agent teaming models.

In a human-agentic system, performance and growth are powered not by increasing capacity but by reimagining and redesigning the system of work itself: configuring roles, workflows, and organizational structures for fluid human-agent collaboration and meaningful value creation.

In this environment, role-based design and outcome-based structures reflect how work gets done to create value. They are aligned with systems-level codedependencies across roles, domains, and functions to enable performance, agility, and impact at speed and at scale.

+ Build organizational readiness as a strategic differentiator

Organizational readiness is a strategic, dynamic capability that must be built and activated in tandem with innovation. It functions at two levels:

- *Individual:* Workers must feel equipped, empowered, and valued in a world where AI is a teammate.
- *Enterprise:* Infrastructure, policies, leadership, and culture must evolve to absorb and scale agentic AI impact.

Closing the widening gap between technological potential and operational reality requires a dual-speed approach: enabling and accelerating experimentation and learning in targeted areas where readiness exists while simultaneously laying the foundations of trust, governance, and alignment for enterprise-scale transformation.

This requires redefining meaning, motivation, human performance, and value in a system where both humans and intelligent agents can grow together and thrive.

Together, these three imperatives form the architecture for reinvention. They are not sequential—they are codependent and concurrent.

Knowing what must change is only the first step. Leaders need a plan for turning this shared awareness into decisive, coordinated, action.



The Human-Agentic Workforce in action: Lemonade

When Lemonade, an insurance provider, launched in 2015, it anticipated the AI-powered future of work. Long before agentic AI became a boardroom priority, Lemonade embedded AI agents into the very core of how the business operates. Today, the company's experience is a powerful testament to what's possible when human and agentic systems are designed to perform as one.

Maya & Jim: MVP team members (who happen to be digital agents)

From the start, Lemonade chose not to merely digitize insurance, but to reinvent it, building an AI-native enterprise where digital agents such as "Maya" and "Jim" weren't experiments at the edge, but central to the company's strategy. Maya focuses on creating playful, compelling onboarding and customer experiences. Jim processes claims—often in seconds. Both work autonomously to deliver real-time service and allow humans to focus on higher-order judgment, empathy, and innovation.

Stronger business outcomes without growing headcount

Lemonade didn't bolt AI on to legacy workflows. Nor did it replace people. Instead, the company designed roles and operating models to enable scalable performance through human-agent collaboration by design. Rather than focus on task-based roles and job descriptions, the company organized work around outcomes, skills, and intelligent collaboration. AI agents handle high-volume execution while humans intervene to manage exceptions, ensure quality, and shape strategy. This has allowed Lemonade to scale performance without an equivalent increase in headcount, challenging the traditional belief that more people equals more growth.

Human-driven, AI-accelerated

At Lemonade, this was never a technology-first initiative. The company's leaders approach AI not as a solution to implement, but as a lever to drive speed, transparency, and trust—putting human impact and customer experience at the center with AI as the enabler. Every digital interaction is designed with a human outcome in mind. The result: a radically simplified, customer-centric model in which AI enables (not replaces) human value.

What makes Lemonade even more relevant today is its focus on readiness as a capability, not a checkpoint. It has built in governance, transparency, and continuous learning from the start. It didn't wait for AI to mature—it created conditions where it could thrive responsibly, ahead of the curve.

No more experiments. It's time to scale.

Agentic AI has delivered a defining moment with a clear imperative. Organizations must shift from fragmented experimentation to enterprise-wide transformation—from edge use cases to core reinvention. This requires activating a new model of work that is **human-led, agent-enabled, and outcome-anchored**, beyond merely implementing a new technology. Doing so will require clarity, courage, and focus. Here's how to get started:

1. Start where you stand

Every organization's entry point to the human-agentic workforce is different. To assess your readiness and determine where and how to get started, begin by asking these three questions:

- What is our overarching ambition for AI—what does our business need?
- How much trust and autonomy are we willing to extend to AI systems?
- How ready are we—technologically, operationally, and culturally—to operate in a world of human-agent teaming?

For example, **NVIDIA's** decision to scale Generative AI across heavy industry was rooted in sharp clarity, focusing on three core challenges: data readiness, infrastructure scalability, and compute power infrastructure. Their transformation didn't start with use cases. It started with defining the company's value ambition and conditions for success.

2. Avoid chaos with a dual-speed approach

Transformation cannot wait for perfect conditions—but it also can't outpace trust. Agentic AI requires bold action without compromising the integrity and stability of core operations. This means running a dual-speed system with two codependent tracks—scaling only what works, where it works, with the right safeguards in place:

- *Fast lane*: Safe spaces to test, learn, and rapidly deploy agentic AI in areas primed for innovation.
- *Stable lane*: Guardrails to protect business continuity and operational integrity, ensure trust, and manage risk, through clear oversight, governance, and alignment at scale.

PepsiCo has embraced a dual-speed approach that uses AI pilots in demand forecasting, logistics, and service (fast lane), while simultaneously strengthening infrastructure, data governance, and risk controls across its global operations (stable lane). The result? Swift value generation that doesn't compromise enterprise resilience.

3. Pivot from vision to scaled impact

Your ability to develop a compelling AI vision is only as valuable as your organization's ability to deliver on it. Organizations must move beyond abstract ambitions and act across three interconnected moves to deliver outcomes and impact at scale:

- **Identify your baseline requirements for realizing value from agentic AI**
 - Strategic alignment: Link the deployment of agents to clear, measurable business outcomes in areas such as growth, innovation, productivity, efficiency, agility, and risk reduction.
 - Data and workflow readiness: You can't automate chaos. Tasks, processes, workflows, and data must be structured and ready for execution by AI agents.
 - Human-agent autonomy thresholds: Define the rules of engagement. When should agents act alone? When are humans required to validate? When is collaboration required? These types of fundamental questions should be on the agenda in every planning meeting.
- **Create momentum with high-impact quick wins**
 - Use tools like work analyzers to identify high-impact opportunities for delivering immediate value with AI.
 - Embed agents into the flow of daily work—not as side experiments, but as core contributors.
 - Focus on momentum and speed, not perfection—because early wins generate proof, belief, and pull.
- **Avoid endless AI pilots without work redesign**
 - Spend your time redesigning work to enable human-agent collaboration and codependency rather than in more tech demos.
 - AI layered on top of legacy roles and workflows delivers marginal returns. Instead, intentionally redesign work with the human-agent system in mind, focused on value outcomes—from how roles are structured to how decisions are made and performance is defined.
 - Scale what works across the enterprise, not just within teams.

For example, **Volkswagen** started with small pilots using digital workers to streamline production processes and reduce inefficiencies on the factory floor. But the company didn't stop at isolated wins. These early successes became a launchpad for scaling transformation across manufacturing and customer operations, supported by human-centered roles redesigned for human-agent collaboration. The result: a clear shift from experimentation in pockets of the organization to enterprise-wide impact.

Not next quarter or next year—*now*

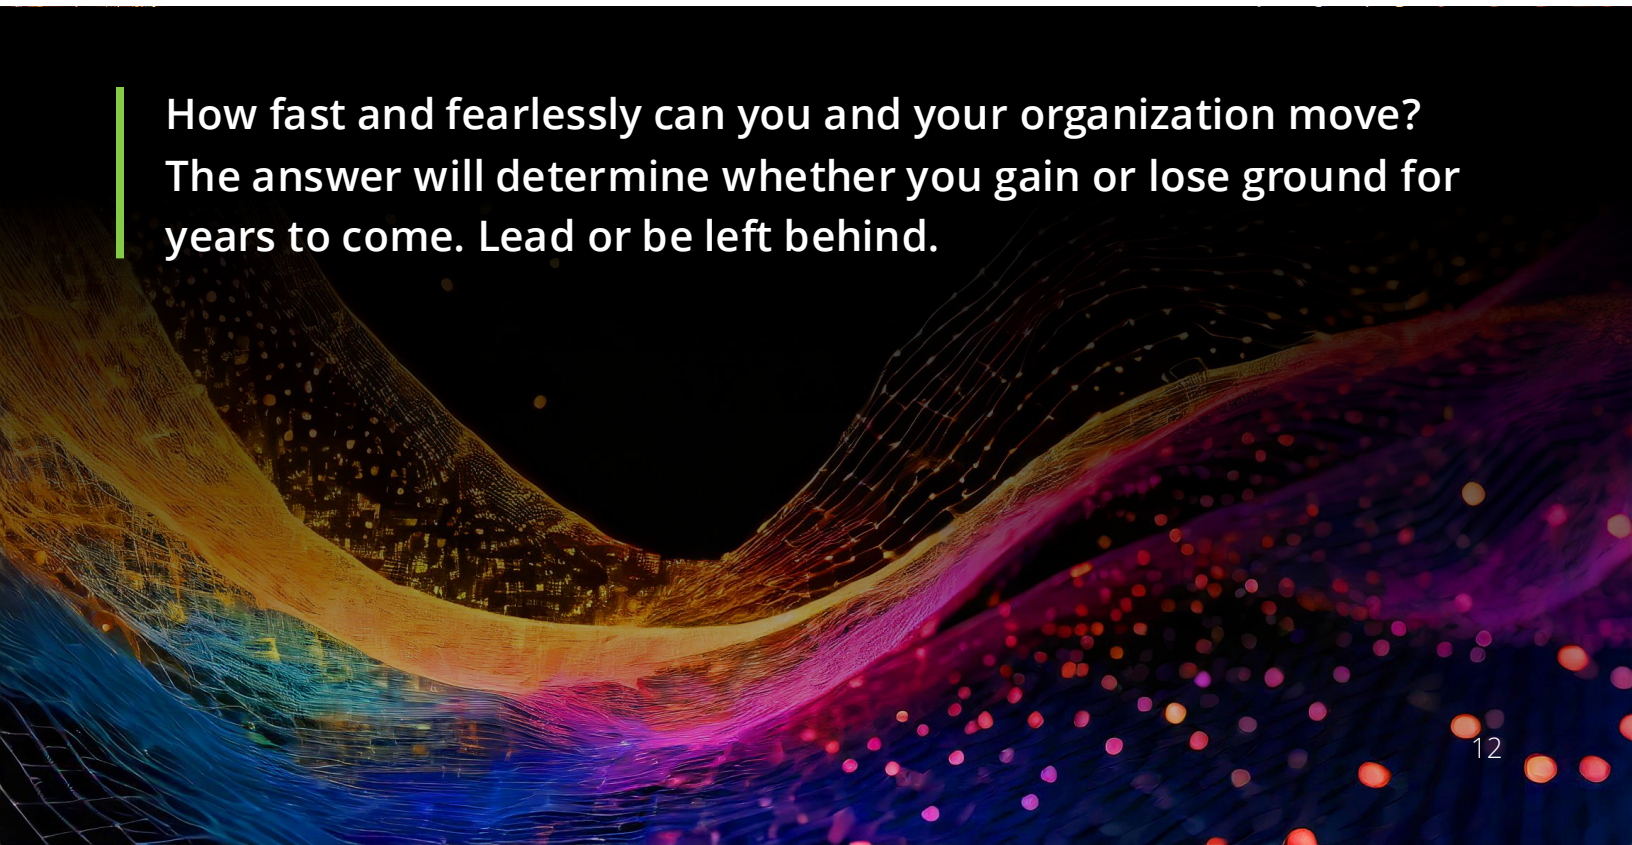
Consider how quickly your own thinking about AI has changed in only the past year or so. AI went from “someday” to “right now” in the blink of an eye—and it shows no signs of slowing momentum. By the time you land “Agentic AI strategy discussion” on the calendar in a few weeks, your competitors may have already had that meeting and started executing their strategy.

All of which makes this the moment for bold decision making. The choices you make now will determine whether AI merely becomes another layer of complexity for your organization to navigate, or a way to unlock a more human, intelligent, productive system of work. And that has big implications for your organization’s ability to perform, innovate, and grow.

This is the time to act at the tipping point—to adapt decisively, not incrementally.

To shift from pilot paralysis to enterprise reinvention.

To architect your human-agentic future—where humans and agents together create outcomes we could never achieve alone.



**How fast and fearlessly can you and your organization move?
The answer will determine whether you gain or lose ground for
years to come. Lead or be left behind.**

Authors

Steve Elliott

Work Redesign Leader
Partner, Deloitte Australia
selliott@deloitte.com.au

Karen Pastakia

Deloitte Global Human Capital Leader
kapastakia@deloitte.ca

Stuart Scotis

Deloitte Global Agentic AI Leader
sscotis@deloitte.com.au

Meagan Lettau

Agentic AI Change Leader
Partner, Deloitte Canada
mlettau@deloitte.ca

Nic Scoble-Williams

Deloitte Global Future of Work and
Agentic Workforce Leader
nsoble-williams@tohatsu.co.jp

About 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/about to learn more.

Deloitte provides leading professional 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 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 approximately 460,000 people worldwide make an impact that matters at www.deloitte.com.

This communication contains general information only, and none of 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.

© 2025. For information, contact Deloitte Global.