




The Agentic Enterprise

*Elevating Workforce Productivity
Through Human-AI Collaboration*

HBR
Analytic
Services
**BRIEFING
PAPER**

Sponsored by



Elevating Work in the Age of Enterprise AI

As organizations accelerate their adoption of artificial intelligence (AI), one truth is becoming increasingly clear: Transformative results do not come from technology alone. They emerge when enterprises rethink how people, processes, and information work together. It is this shift, from experimental tools to embedded, governed, and collaborative systems, that defines the agentic era of AI. Hyland is pleased to sponsor this paper in association with Harvard Business Review Analytic Services to explore what it takes for organizations to successfully navigate this moment.

Across industries, leaders see extraordinary potential in AI systems capable of acting alongside people, not simply providing suggestions or generating content. Yet many organizations still face structural barriers that limit impact: fragmented information, disconnected workflows, uneven readiness, and the absence of clear governance models. The move toward agentic AI requires resolving these foundational issues so intelligent systems can operate with context, transparency, and accountability.

This research underscores an essential point: AI's value increases when it becomes part of the fabric of work itself. That means grounding systems in an organization's own content and

knowledge, designing workflows where human expertise and agentic capabilities complement each other, and building oversight structures that ensure responsible operation. When these elements come together, employees can focus on higher-value tasks, decision making becomes more consistent, and organizations gain the agility required to succeed in a rapidly evolving landscape.

The transition to an agentic enterprise is not simply a technical upgrade; it is an organizational transformation. It invites leaders to reimagine roles, redefine how decisions move through the business, and elevate the human skill sets that matter most in an AI-enabled world. It also requires the courage to embrace new operating models, the discipline to build strong information foundations, and the commitment to ensure AI augments rather than replaces human judgment.

Hyland's long-standing work at the intersection of enterprise content, operational processes, and application intelligence gives us a unique perspective on this evolution. We believe the organizations that will realize the greatest value from AI are those that align technology with trusted information, governed workflows, and clearly defined responsibilities. By sponsoring this paper, our aim is to help leaders



Jitesh S. Ghai
CEO, Hyland

understand the opportunities and challenges ahead and to spark a thoughtful, practical dialogue about building the next generation of digitally enabled enterprises.

The agentic enterprise is emerging. With the right foundations, it has the potential to lift productivity, enhance decision quality, and expand the capacity of people and organizations to do their best work.

The Agentic Enterprise

Elevating Workforce Productivity Through Human-AI Collaboration

Despite heavy investment, many artificial intelligence (AI) initiatives fall short of expectations—not because organizations lack data or tools but because they lack a cohesive foundation that allows AI agents to understand enterprise context. Fragmented operational processes, unstructured and hidden content, and disconnected systems prevent these agents from recognizing how documents, workflows, and decisions relate to one another.

Without this shared context, agentic AI remains difficult to scale beyond isolated use cases. According to a November 2025 report by New York consulting firm McKinsey & Co., nearly one in four organizations have begun implementing agentic AI systems and another 39% are experimenting with them, yet when it comes to most business functions, fewer than 10% report having scaled deployments of agentic AI systems.¹

As enterprises move beyond early AI experimentation, successful initiatives generate positive results by deploying context-aware autonomous agents inside essential business processes. These agentic AI systems can act and collaborate with humans. When grounded in an organization's

own knowledge and content ecosystem, they can surface insights, automate tasks, and free employees to focus on higher-value activities.

In content- and workflow-intensive environments, this movement from isolated experimentation to embedded execution is particularly significant. Enterprises are redesigning document processing, case management, compliance oversight, and knowledge-driven operations as AI agents move from peripheral tools to governed participants in core business processes.

Realizing this higher value, however, requires more than introducing new technologies. It demands broad AI readiness, such as rethinking how work gets done; preparing content,

HIGHLIGHTS

As enterprises move beyond early artificial intelligence (AI) experimentation, **successful initiatives generate positive results by deploying context-aware autonomous agents** inside essential business processes.

Productivity gains emerge not from wholesale automation, but from carefully redesigning workflows so that AI handles structured, repetitive tasks while **professionals focus on framing problems, interpreting edge cases, and exercising judgment.**

When grounded in a structured information architecture, governed by clear accountability, and embedded in operational workflows, **AI agents can augment human capability** rather than displace it.

One fundamental impact of AI lies in its ability to restructure the way an enterprise coordinates its work rather than just accelerating existing tasks.

data, and workflows so they can be used effectively by AI systems; and establishing strong governance procedures. As Thomas W. Malone, the Patrick J. McGovern Professor of Management at the MIT Sloan School of Management and founding director of the MIT Center for Collective Intelligence in Cambridge, Mass., has noted, one fundamental impact of AI lies in its ability to restructure the way an enterprise coordinates its work rather than just accelerating existing tasks. In this sense, agentic AI represents not simply a technical upgrade but also a structural shift in how enterprises coordinate work across content, systems, and teams.

The stakes of this structural shift are especially evident in highly regulated, information-intensive environments such as health care. At Great Ormond Street Hospital (GOSH) in London, clinicians manage some of the most complex pediatric cases in the United Kingdom. Patient journeys often span many years, multiple specialties, and numerous digital systems. Electronic health records, imaging repositories, pathology platforms, genomic data, referral letters, and multidisciplinary notes all form part of the clinical picture. The constraint, however, is not simply the fragmentation of data. Nor is it a shortage of information.

“Health care does not have a data scarcity problem, and fragmentation is only part of the issue,” says Dr. Shankar Sridharan, chief clinical information officer at GOSH and national clinical lead for AI for the U.K. National Health Service. “The real challenge is that we have focused AI almost exclusively on clinical decision support, while our assurance and regulatory frameworks are still catching up with technologies that are partially autonomous and adaptive.”

This paper examines how organizations are addressing the challenges facing AI stakeholders and achieving measurable results in information-rich domains such as case management, claims processing, contract review, and clinical support. It explores how enterprises can embed agentic AI into everyday work, enabling meaningful human-AI collaboration inside defined roles and governed workflows while upholding accountability and preserving human judgment. Ultimately, the paper will prove that agentic AI is less a

technology story than an organizational one hinging on how effectively enterprises align intelligent agents with content and workflows and the people who depend on them.

Information-Intensive Challenges

GOSH operates in an information-intensive and risk-sensitive environment, where clinicians must synthesize longitudinal patient records, imaging results, laboratory data, genomic information, and specialty notes under strict regulatory requirements and often time-critical conditions. Even short delays can affect patient treatment choices, prolong time to intervention, and negatively influence outcomes. When these health care workers cannot easily locate or understand the right information at the right moment, productivity suffers, decisions take longer, and teams produce inconsistent results.

As Sridharan points out, for much of the past five years the health care sector has prioritized clinical decision-support systems using models designed to predict deterioration, improve diagnostic precision, and recommend interventions. “These are important developments, but they sit at the highest regulatory threshold and depend upon assurance frameworks that are still evolving,” he explains. “At the same time, the operational machinery that underpins care delivery, such as referral triage, case preparation, documentation, [and] coordination across teams and sites, continues to absorb significant clinical time and cognitive effort.”

The problem is not unique to health care. Across industries, enterprises confront the same structural barrier that a more collaborative human-AI approach may help address. A December 2025 Harvard Business Review Analytic Services survey of 325 respondents from the *Harvard Business Review* audience reveals the current state of readiness regarding agentic AI. When respondents from organizations that are actively using, piloting, or exploring AI were asked the extent to which the AI tools used by their organization are embedded within or separate from employees’ workflows, the

largest share (39%) say most of their organization's AI-enabled workflows rely on separate, standalone tools. Only 12% report that AI is embedded within the flow of the task to automatically assist employees. **FIGURE 1**

In many organizations, AI remains a separate capability—used for drafting, summarizing, or experimenting—rather than integrated into end-to-end governed business processes. As a result, many organizations are rethinking not just their AI tools but also the foundations that support them. This effort entails improving data quality, standardizing metadata, connecting previously siloed data repositories, and ensuring that content is organized so AI systems can interpret and act on it reliably.

Rather than deploying isolated AI tools, GOSH pursued a deliberate foundation-first strategy through its Data Research, Innovation and Virtual Environment and Digital Research Environment efforts, which together form a secure analytics and research platform that links clinical, operational, and research data across systems. The goal was not to bolt AI tools onto disconnected systems but to build a unified foundation in which data standards, governance rules, and workflow logic were integrated from the outset. “We deliberately avoided starting with headline-grabbing pilots,” Sridharan explains. “Before deploying agents capable of acting across systems, we needed governance maturity. Agentic AI without bounded autonomy increases risk rather than productivity.”

From Aspiration to Operational Reality


The experience at GOSH highlights what many enterprises are only beginning to confront: Enthusiasm for human-AI collaboration is rising faster than the operational structures required to support it. The survey reveals a striking contrast: Eighty-seven percent of respondents say their organization sees potential in having employees and AI collaborate on tasks, yet only 5% report that their organization has little or no interest in human-AI collaboration.

FIGURE 1

Organizational Use of Artificial Intelligence

Most respondents say it's through standalone AI tools to help with a task


Overall, to what extent are the AI tools used by your organization embedded within or separate from employees' workflows? (Among those at organizations exploring, piloting, or actively using artificial intelligence). *Select one.*



Source: Harvard Business Review Analytic Services survey, December 2025

This gap between aspiration and operational reality sets the stage for the deeper economic question of when, and under what conditions, does AI actually translate into valuable performance improvements? Erik Brynjolfsson, director of the Stanford Digital Economy Lab in Stanford, Calif., and a leading scholar of productivity economics, argues that, with AI initiatives, early enthusiasm often precedes measurable productivity gains. Organizations frequently experience what he calls a “productivity J-curve,” in which performance dips before it rises as organizations invest in redesigned workflows, governance, and training. “People underestimate some of the complexity of getting these systems to really work effectively,” he says. “The real gains will come not just from faster processing but from fundamentally reimagining how knowledge is synthesized and applied across the enterprise.”

Through his work at Workhelix, a Beaverton, Ore.-based company Brynjolfsson cofounded to analyze how AI reshapes jobs at the task level, he has examined how organizations decompose roles into individual activities to determine which tasks can be automated, which can be augmented, and which should remain human-led. His findings reinforce a central insight: Productivity gains emerge not



“The real gains will come not just from faster processing but from fundamentally reimagining how knowledge is synthesized and applied across the enterprise.”

Erik Brynjolfsson
Director
Stanford Digital Economy Lab

“Agentic systems expand capability and risk simultaneously. They increase the potential cyber surface area and introduce new governance complexity because they can act, not just analyze.”

Dr. Shankar Sridharan, chief clinical information officer,
Great Ormond Street Hospital

from wholesale automation, he explains, but from carefully redesigning workflows so that AI handles structured, repetitive tasks while professionals focus on framing problems, interpreting edge cases, and exercising judgment.

The most valuable human skill in the age of AI, Brynjolfsson suggests, may be the ability to ask better questions and frame problems correctly. “I believe the ‘chief question officer’ will become a critical role in the future organization,” he predicts. “Teams will evolve from executing processes to managing and auditing the agents that execute them.”

In research on generative AI in customer support environments, Brynjolfsson and his colleagues documented productivity gains of approximately 14% when AI tools were deployed in this way. The real gains from AI come not from isolated automation, he notes, but from complementary investments in workflow redesign, skills, and organizational change.

Workflow automation is particularly important in health care settings. According to GOSH’s Sridharan, large language models now assist clinicians by summarizing clinical histories, highlighting laboratory trends, and synthesizing multidisciplinary notes. But summarization alone does not change throughput. Health care work is procedural and distributed across platforms. It requires retrieval, reconciliation, routing, and escalation. Agentic capability provides that orchestration layer, he notes.

These systems are intentionally designed to support—not replace—clinical judgment. Within defined boundaries, agents can retrieve documents, administrate referrals, surface exceptions, and escalate tasks to human decision makers. They operate inside governed workflows rather than alongside them. “The system can suggest, highlight, and prioritize,” Sridharan notes. “But it does not make autonomous clinical decisions. Human expertise remains the arbiter.”

This bounded-autonomy model reflects GOSH’s broader strategy: embedding AI within governed clinical workflows while preserving human oversight. Rather than pursuing full clinical autonomy, GOSH’s strategy focuses first on operational and clinically adjacent workflows—what Sridharan

calls “the connective tissue of care delivery.” By reducing search time, minimizing duplication, and accelerating coordination, AI increases the time available for patient care.

Enforcing Trust and Transparency

Technology alone does not determine whether agentic AI succeeds. Organizational norms, professional trust, and day-to-day working habits likely shape whether people embrace or resist new systems. To build user confidence, GOSH is embedding role-based access controls, escalation logic, kill-switch capability, and full auditability into the platform architecture. Every AI-supported interaction is traceable, including what data was accessed, what output was generated, and how it entered a workflow. Clear accountability structures define who approves use cases and who intervenes when an exception is triggered.

“Agentic systems expand capability and risk simultaneously,” Sridharan says. “They increase the potential cyber surface area and introduce new governance complexity because they can act, not just analyze. If you do not design for bounded autonomy and accountability from the outset, you will struggle to scale safely.”

MIT’s Malone believes trust emerges when transparency, oversight, and performance align—and that alignment must be deliberately designed into the structure of work itself. “If an agent takes an action, someone in the organization should be accountable for that action,” he says. He compares the situation to owning a dangerous animal, where the owner is legally responsible for what the animal does. Similarly, organizations are responsible for the actions—and consequences—of their agents.

Thanks in part to this insistence on usability and trust, GOSH reports early signs of measurable impact, especially reductions in the time people spend searching for information. According to Sridharan, in initial agentic AI pilots, employees have been able to access critical data far more quickly and with less manual compilation. Still, Sridharan is

“If an agent takes an action, someone in the organization should be accountable for that action.”

Thomas W. Malone, Patrick J. McGovern Professor of Management, MIT Sloan School of Management, and founding director, MIT Center for Collective Intelligence

careful to frame these results in qualitative terms. “Productivity in health care isn’t just about speed,” he says. “It’s also about better decisions, fewer missed signals, and more time spent with patients.”

Stanford’s Brynjolfsson adds a longer-term perspective. He emphasizes that returns materialize only after organizations rethink how work is actually done. That rethinking requires organizations not only to redesign workflows but also to clarify who validates models, who monitors performance drift, and who intervenes when agents act outside expected parameters.

The takeaway is clear: Agentic AI scales safely only when governance scales with it. These guardrails make expectations explicit, clarify accountability, and preserve human judgment and organizational trust.

Understanding the Enterprise State

While GOSH illustrates how agentic AI can operate within a health care system, the same structural challenges emerge across industries wherever documents, cases, and governed workflows intersect. Tinglong Dai, the Bernard T. Ferrari Professor of Business at Johns Hopkins Carey Business School in Baltimore, argues that many AI systems struggle to account for where a case stands in a workflow and what constraints shape the next action. The deeper challenge is understanding the “enterprise state”—the evolving condition of documents, cases, workflows, and decisions at any given moment.

Drawing on his research into workflow design and decision sequencing, he emphasizes that AI must understand process context, not just document content, to function reliably. “Documents are not just static records,” Dai explains. “They exist within a workflow and reflect where you are in a case and what constraints apply.”

For example, when processing a health insurance claim, an AI system must know whether documentation has been validated, whether a prior denial occurred, whether the

clinical codes have been accurately recorded, and whether the case is under appeal. Claims documents do not exist in isolation. They signal progression, constraint, authority, and accountability within a multi-step approval process.


Brynjolfsson draws a distinction between content generation and process execution. “Agentic AI goes further [than generative AI] by autonomously executing complex multi-step workflows rather than just generating text,” he explains. This capability, he adds, “shifts the paradigm from simple task automation to comprehensive process augmentation,” enabling systems “to act on information rather than just retrieving it.”

In content-heavy enterprises, that evolution moves AI beyond reading a contract or claim toward participating in a carefully defined business process. When systems lack awareness of stage, permission, and constraint, Johns Hopkins’ Dai cautions, they may trigger actions out of sequence or bypass required approvals. These weaknesses often surface as organizations attempt to scale. When something goes wrong, he adds, the organization—not the technology provider—will be held responsible.

Dai emphasizes that AI must understand where a document sits within a governed process, not just what it contains. In practical terms, that requires consistent metadata, standardized key fields such as dates and approval status, and interoperable systems that exchange information seamlessly. Without that structural clarity, AI systems cannot reliably determine a case’s status or initiate the appropriate next action.

Redesigning Organizations for the Agentic Era

Malone approaches agentic AI through the lens of organizational design. “Most major information technologies change the way we coordinate, not just the way we communicate,” he asserts. These current views build on ideas he has explored for more than a decade at the MIT Center



“Documents are not just static records. They exist within a workflow and reflect where you are in a case and what constraints apply.”

Tinglong Dai
Bernard T. Ferrari Professor of Business
Johns Hopkins Carey Business School

“The goal should always be to maximize the total value of the human-machine partnership rather than simply minimizing labor costs.”

Stanford Digital Economy Lab’s Brynjolfsson

for Collective Intelligence on how humans and machines collaborate. In his 2019 book, *Superminds*, he examines how decision authority, coordination mechanisms, and shared knowledge structures enable groups to solve complex problems. Those early insights foreshadowed agentic AI systems that do not merely generate outputs but also coordinate tasks. The core question, then and now, is not simply how well machines can follow instructions but how people and machines can be organized to think and act together more effectively.

In *Superminds*, Malone explores how humans and computers can think together to solve problems neither could address alone. Agentic AI can reduce coordination costs, he argues, such as deciding what should happen next, who should do it, and how tasks should be sequenced. In document- and workflow-intensive environments, AI agents can act as coordinators, identifying dependencies and initiating actions within defined boundaries.

Organizations should establish clear ownership and define operating boundaries as agents scale. That emphasis on structured coordination and responsibility did not emerge in response to today’s AI boom, Malone maintains. It builds on years of research into how humans and machines coordinate decisions and share responsibility.

Conclusion

Agentic AI marks a turning point in enterprise transformation. When grounded in a structured information architecture, governed by clear accountability, and embedded in operational workflows, AI agents can augment human capability rather than displace it. But as the survey findings suggest, many organizations are still in transition, convinced of AI’s collaborative potential yet early in embedding it into core business processes.

For GOSH, the agentic enterprise is not a technology experiment. It is an operational redesign. The hospital’s experience illustrates a broader lesson for regulated industries:

AI delivers value not at the periphery of workflows but when embedded inside them in a manner that is governed, bounded, and aligned with professional judgment.

That principle extends beyond health care. The same foundation that enables clinicians to trust AI-generated insights—structured information, defined workflows, and clear accountability—also determines whether agentic systems can function effectively in other document- and process-intensive environments. Many AI systems can read documents but fail to grasp where a case stands in a governed process or what constraints shape the next action. As Dai explains, documents “exist within a workflow” and reflect “where you are in a case and what constraints apply.” Without that context, AI remains limited to surface-level analysis rather than true operational participation.

The deeper transformation, however, is organizational. “The real impact is not just that AI helps people do the same work faster,” Malone observes. “It changes how work is organized.” Agentic AI reshapes coordination itself, altering who decides what happens next and how responsibility is structured.

Ultimately, the objective is not just cutting costs but creating value. Brynjolfsson argues that “the goal should always be to maximize the total value of the human-machine partnership rather than simply minimizing labor costs.”

Together, these observations point to a common conclusion: Realizing the promise of agentic AI requires more than rolling out new software tools. The agentic enterprise takes shape when organizations deliberately align people, content, systems, and intelligent agents within redesigned workflows. When organizations embed AI within governed document- and content-driven processes, they improve compliance, strengthen decision quality, and give professionals the capacity to focus on higher-value work.

In the end, the true embrace of agentic AI will depend on a uniquely human sentiment: courage. As Brynjolfsson notes, “Employees must move past the fear of replacement and foster a culture of augmentation where they feel empowered, not threatened, by these tools.”

Endnotes

- 1 Alex Singla, Alexander Sukharevsky, Bryce Hall, et al., "The State of AI in 2025: Agents, Innovation, and Transformation," McKinsey & Co., November 5, 2025. <https://www.mckinsey.com/capabilities/quantumblack/our-insights/the-state-of-ai>.



VISIT US ONLINE

hbr.org/hbr-analytic-services

Harvard Business Review Analytic Services is an independent commercial research unit within Harvard Business Review Group, conducting research and comparative analysis on important management challenges and emerging business opportunities. Seeking to provide business intelligence and peer-group insight, each report is published based on the findings of original quantitative and/or qualitative research and analysis. Quantitative surveys are conducted with the HBR Advisory Council, HBR's global research panel, and qualitative research is conducted with senior business executives and subject-matter experts from within and beyond the *Harvard Business Review* author community. Email us at hbranalyticservices@hbr.org.