

AI Agents Reshape Process Automation

Idea In Short

Executives evaluating automation investments should treat AI agents as a distinct category from chatbots or robotic process automation, not an upgrade to either. The technology already delivers measurable results: a consumer goods company cut a week-long, six-analyst marketing review to under an hour with one employee, and a global bank reduced customer service costs tenfold. The decision that matters now is not whether to adopt AI agents but where to start. Leaders should pick one measurable process, build clean data pipelines and governance before scaling, and treat supervising agents as a core management skill rather than a technical afterthought. Organizations that build this foundation now, before multi-agent systems become standard, will compound their advantage. Those that wait will retrofit infrastructure under a system that competitors are already running.

Most businesses have already automated something. Scheduled emails, rule-based chatbots and workflow triggers that fire when a form is submitted have been standard for years. These tools saved time but never carried real intelligence: they followed rigid scripts that broke the moment something unexpected happened, and they required constant maintenance to stay useful. AI agents are a distinct category of technology, and the distinction matters for any executive weighing where to invest next. Unlike conventional automation, which executes a fixed sequence, an AI agent observes its environment, reasons through a goal and decides what action to take next. It handles multi-step tasks, adapts when conditions shift, and gets more accurate the longer it runs. The market for AI agents is projected to grow at a compound annual growth rate of 45% over the next five years, a figure that reflects adoption already underway rather than speculation about what might eventually be possible.¹

What an AI agent actually does

An AI agent is a software system that perceives information, reasons about a goal and takes action across connected systems to complete complex tasks without step-by-step human

direction. The term gets used loosely across vendor marketing, so it helps to ground it in what these systems do in practice rather than in how they are described. At the core, an AI agent runs a continuous cycle of observing, planning and acting. It gathers information from its environment, including user inputs, system data, application programming interface (API) responses and sensor readings, then uses a large language model (LLM) or a comparable planning component to decide what to do next. It then executes that decision through connected systems: updating a customer relationship management (CRM) platform, drafting a document, triggering a workflow, querying a database or handing a task to another agent.

What separates this from a workflow automation tool is memory and judgment. AI agents retain context across conversations and tasks rather than starting fresh each time. They distinguish between a routine situation and an edge case, escalate to a person when the situation calls for it, and refine their approach based on what worked before. IBM defines an AI agent as a system that autonomously performs tasks by designing its own workflow with the tools available to it, extending well beyond natural language processing into decision-making, problem-solving and direct action.² A well-built agent completes assigned tasks in the right order, with the right amount of human checking built in.

A practical illustration makes the shift concrete. A consumer goods company wanted to run its global marketing campaigns more efficiently. A process that previously took six analysts a full week was rebuilt around an AI agent, and the same work now takes a single employee less than an hour. The agent gathers marketing data on its own, analyzes campaign performance, drafts a recommendations report and updates media buying platforms once a person signs off on the changes. The analyst's job changed from doing the analysis to reviewing and approving it, which is a different kind of work entirely.

Where AI agents create measurable value

Growth from AI does not come from deploying agents everywhere at once. It comes from identifying the functions where intelligent automation produces outcomes that matter, such as faster decisions, lower costs, fewer errors and better customer experiences, and building outward from there. Decision support across functions is one of the more underappreciated applications, largely because it connects data that otherwise lives in silos. A procurement agent that pulls together supplier data, market pricing and inventory levels can flag material shortages before they disrupt production. A finance agent that monitors budget utilization in real time can adjust allocations automatically instead of waiting for a monthly review cycle.

Neither replaces the judgment of an experienced professional, but both close the lag between information becoming available and a decision getting made.

End-to-end workflow automation extends this further into processes that require interpretation rather than fixed rules. Robotic process automation improved speed on narrow, rules-based tasks, but it could not handle ambiguity. An HR agent that monitors engagement signals can identify employees at attrition risk and trigger a retention sequence before a resignation letter arrives. An operations agent can manage an order from receipt through dispatch, handling exceptions as they come up instead of queuing them for a person to review later.

Customer experience improves at a different scale entirely once agents anticipate needs rather than just responding to them. A service agent that resolves only the issue directly in front of it is useful on its own. One that anticipates what a customer will need next, based on account history, product usage patterns and past interactions, creates a noticeably different experience.

A leading global bank used AI virtual agents for customer interactions and reduced costs by 10 times, while improving response consistency

Research and development benefits follow a similar pattern in specialized fields, where agents shrink the gap between identifying a problem and producing a workable answer. A biopharma company used AI agents for lead generation, cutting cycle time by 25% and gaining 35% efficiency in drafting clinical study reports. In software development, one IT department used AI agents to modernize legacy technologies and lifted productivity by up to 40%. These gains come from removing the time-consuming scaffolding around specialized work, not from replacing the specialists doing it.

The infrastructure that determines whether agents work

Deploying an AI agent without the right foundation produces a system that performs well in a demonstration and fails once it meets production conditions. The infrastructure underneath the agent matters as much as the agent's own design. Clean, accessible data is

the starting point, since an agent that cannot reach reliable, current information cannot make a reliable decision. Unified data platforms and event-driven pipelines let agents sense and respond as conditions change instead of operating on a stale snapshot from hours or days earlier.

Interoperability determines whether an agent can function across departments or stays trapped in one. Open APIs and modular integrations mean an agent built for HR can communicate with a finance system without a custom integration project every time a new connection is needed. Without that interoperability, agents remain isolated tools rather than genuine collaborators across the organization. Deloitte's research on enterprise AI adoption found that agentic AI usage is set to rise sharply over the next two years, but oversight is lagging behind it, with only one in five companies reporting a mature governance model for autonomous agents.³

Governance and access control grow more important as an agent's autonomy increases. Clear rules about what an agent can do on its own, and where it must escalate to a person, are not constraints that slow the system down. They are what makes it safe to trust the system with more responsibility over time. Scalability should be planned for before it is needed, since moving from one agent to dozens requires computing environments and management tools built for that scale from the outset. Organizations that plan for scale from the beginning expand their AI capabilities faster and with fewer disruptions than those that retrofit infrastructure after a pilot succeeds.

How different industries are putting agents to work

The most instructive evidence for what AI agents can do comes from sector-specific deployments, where the constraints are well understood and the outcomes are measurable. In financial services, agents manage reconciliation, forecasting and reporting tasks that previously consumed significant analyst time around month-end close. Compliance monitoring agents track transactions continuously, detect anomalies and maintain audit-ready records, a task that would otherwise require a much larger team to perform manually with the same consistency.

In healthcare, scheduling agents coordinate clinician workloads while maintaining compliance with safety staffing standards, and diagnostic support agents assist physicians with data analysis, improving accuracy and reducing the time between data availability and a

clinical decision. In retail and consumer goods, marketing agents personalize offers based on real-time customer behavior, reducing churn among at-risk segments. One consumer packaged goods company used AI agents to produce blog content, cutting costs by 95% and reducing production time from four weeks to a single day.

In manufacturing, supply chain agents track shipments, predict delays from historical and real-time signals, and adjust production schedules in response. The value here goes beyond cost reduction, since these systems absorb disruption without requiring manual coordination across multiple systems. In legal and professional services, specialized agents such as Harvey handle contract review, document generation and case file analysis, trained on legal data and operating within compliance requirements to produce drafts and summaries that free attorneys to focus on judgment-intensive work.

Managing the organizational side of adoption

Digital transformation with AI is as much an organizational challenge as a technical one, and the gap between the two creates friction that undermines adoption. The technology can be ready well before the team is. The role of employees shifts rather than disappears under this kind of automation. New functions are emerging, including agent trainers, AI workflow designers and integration strategists, all of which require understanding agents well enough to configure, monitor and improve them. Teams brought into this process early develop confidence in the technology, while those who encounter it as a finished product with unexplained behavior tend to work around it instead of with it.

Leadership alignment is what allows scaling to happen at all. When the chief information officer, chief financial officer and department heads share a coherent view of what AI agents are supposed to accomplish, investment decisions move faster and rollout efforts keep their momentum. When they do not share that view, individual deployments succeed in isolation without producing the cross-functional benefits that make AI genuinely valuable at the enterprise level. Gartner projects that task-specific AI agents will appear in 40% of enterprise applications by the end of 2026, up from less than 5% in 2025, a pace of change that outstrips most organizations' governance planning.⁴

Supervising AI agents is becoming a standard management skill rather than a specialized technical one. As agents take on more responsibility, the people overseeing them need to understand what an agent is optimizing for, recognize when it operates outside the scope it

was designed for, and know when to override it. Harvard Business Review's research on scaling agentic systems argues that organizations should treat agents like team members, with a defined identity, limited authority, trusted information sources and audit trails, rather than as standard software.⁵

Companies that adopt this mindset and introduce autonomy gradually will be far more likely to capture the benefits of agentic AI without exposing themselves to costly mistakes

Moving from pilot to scale without stalling

The most common failure mode in AI agent adoption is a successful pilot that never expands beyond its original team. The pilot works because it is small, well-supported and closely watched. Scaling fails because the organization treats it as a technical problem rather than an operational one that requires the same discipline applied to any other capital investment.

A practical approach to scaling looks like this:

1. Start with a use case where success can be measured clearly, such as reducing a specific process time or improving a defined service metric
2. Build a minimum viable agent and test it against real conditions before expanding its scope
3. Integrate the agent into core systems early rather than running it in parallel with existing processes
4. Track performance using concrete indicators, including cost savings, turnaround time, error rates and customer satisfaction
5. Expand to new domains only after the first deployment is stable and the governance model is established

This sequence keeps risk contained and ensures each stage of scaling builds on evidence rather than assumption. Organizations that follow it tend to reach meaningful scale faster than those that attempt a comprehensive rollout from the start.

Where AI agent development is headed

The near-term direction for AI agents is multi-agent collaboration, where individual agents stop operating independently and start coordinating within networks toward shared goals. A finance agent negotiates payment terms with a procurement agent, while a compliance agent validates the transaction against regulatory requirements, all within roughly the time it currently takes a person to draft one email in that same chain. MIT Sloan Management Review's research on managing agentic systems found that a clear majority of experts, 69%, believe this shift requires genuinely new management approaches rather than adapted versions of existing ones, while a smaller group cautions against treating AI as exceptional in ways that shift accountability away from the people deploying it.⁶

This distributed model of AI in business will reshape how decisions get made and how organizations structure the work of coordination itself. Companies investing now in the data infrastructure, integration capability and governance practices needed to support individual agents are positioning themselves to take advantage of multi-agent architectures as those mature. Companies that wait will find themselves building foundations under a system that is already running elsewhere in their industry.

Earlier automation efforts focused on efficiency alone. AI agents extend that focus into responsiveness and organizational intelligence, representing the first wave of technology that improves both at once rather than trading one for the other. As agents take on a greater share of execution, the differentiator between competing organizations increasingly becomes the quality of the strategic thinking directing them. Building that capability is not a one-time initiative but an ongoing, structured effort across teams and leadership levels, and it is why many enterprises are now pairing agent deployment with formal strategy development rather than treating the two as separate workstreams.

In an environment where intelligent systems are becoming standard infrastructure, the advantage belongs to organizations that combine them with consistently applied, high-quality strategy.

- 1Seizing the agentic AI advantage
- 2What are AI agents
- 3The state of AI in the enterprise
- 4Gartner predicts 40% of enterprise apps will feature task-specific AI agents by 2026
- 5To scale AI agents successfully, think of them like team members

- 6Agentic AI at scale: Redefining management for a superhuman workforce

Summary

AI agents differ from earlier automation in kind, not degree. They observe, reason and act across connected systems, adapting to conditions that would break a scripted workflow. The evidence spans functions and industries: procurement teams flagging shortages before they disrupt production, biopharma teams cutting lead-generation cycles by a quarter, and legal teams offloading contract review to specialized agents. None of this happens without clean data, interoperable systems and clear governance over what an agent can decide alone. The organizational shift matters as much as the technical one, since employees need to become agent trainers and supervisors rather than bystanders. As agents take over execution, the differentiator between competitors becomes the quality of the strategic thinking guiding them, not the technology itself.