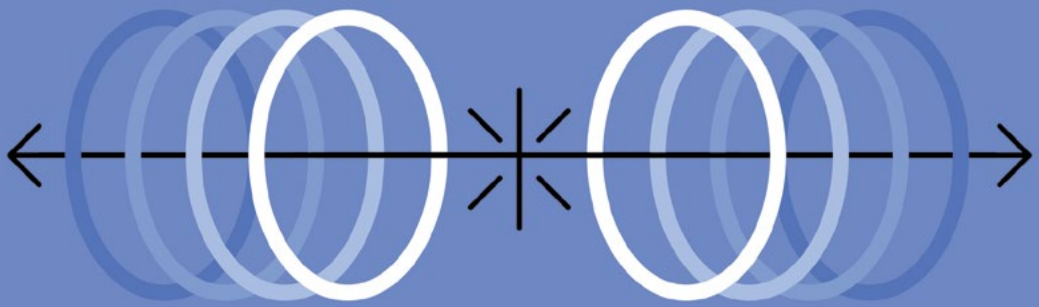
HARVARD BUSINESS REVIEW



Agentic AI

Orchestrating Intelligent Operations

HBR Analytic Services **WHITE PAPER**

Sponsored by

Deloitte.
Together makes progress

Orchestrating Agentic AI Operations

Agentic AI is rapidly redefining the very fabric of enterprise operations, moving beyond mere efficiency gains to unlock unprecedented opportunities for strategic advantage, market differentiation, and accelerated business growth. Yet despite the excitement, many organizations are struggling to strategically manage and deliver agentic AI at scale, often overlooking its potential as a catalyst for innovation and reinvention.

A core challenge lies in moving beyond isolated automation pilots to a holistic redesign of autonomous business systems for an era of abundant digital labor. Organizations are increasingly turning to third parties to help solve complex problems like managing the economics of digital arbitrage and variable artificial intelligence (AI) costs across IT service management, driving scalable governance and compliance across global finance systems, and connecting rapid AI innovation in customer contact operations with enterprise adoption. All this effort occurs while supporting workforce shifts and upholding human judgment where it matters most.

Next-generation managed service providers (MSPs)—including Deloitte through our Operate services—are emerging as powerful allies. They can bring a unique blend of business domain knowledge, technical acumen, and scaled capability. They serve as orchestrators—guiding organizations through the systematic deployment and ongoing optimization of agentic AI alongside human talent, crucially helping identify and exploit new and innovative business models to deliver business outcomes that matter.

Amid extensive discussions and writings on agentic AI, this Harvard Business Review Analytic Services paper, sponsored by Deloitte, offers a new and different perspective on how to succeed in the long-term management and refinement of this disruptive technology.

We encourage you to use the paper as a practical guide for considering and selecting the right next-generation MSP for your organization. Drawing on specialist insights and real-life accounts, it provides practical tips on building operational strategies that can deliver ongoing measurable results. The future of agentic AI success lies in the strategic orchestration of agents to transform your business. Read on for tangible insights on how to collaborate with MSPs to unlock enduring enterprise value.



Stuart Scotis
Global Agentic AI Leader
Deloitte



Doug Gish
Global Operate Leader
Deloitte

Agentic AI

Orchestrating Intelligent Operations

As organizations race to deploy agentic AI, many initiatives fail to deliver meaningful results, often due to a narrow focus on cost reduction rather than holistic strategic transformation. Gartner Research Inc. in Stamford, Conn., predicts that 40% of agentic AI projects will be canceled by the end of 2027. While technical challenges remain, the deeper obstacles to agentic AI implementation and ongoing operations are strategic and organizational, such as reimagining processes built for human execution and implementing governance structures for systems that act independently.

HOWEVER, FOR THOSE that navigate these complexities successfully, agentic AI promises not just operational excellence but also the power to reinvent business models, create new value propositions, and capture market share.

"The organizations that succeed will be those that treat agentic AI as a long-term operating model shift," says Kunal Bansal, chief digital officer at Mankind Pharma, a multinational pharmaceutical and health care products company headquartered in Delhi, India. "The shift requires patience, discipline, and partners who are able to evolve alongside them."

To accommodate this shift, a new generation of managed service providers (MSPs) is emerging as critical orchestrators and innovation catalysts in successful agentic AI implementations. They are guiding systematic deployment, managing new and ongoing economics of digital labor, and helping govern autonomous systems at scale.

For business executives, however, the emergence of this breed of next-generation MSPs brings direct consequences to the criteria for selecting a service provider. Companies that evaluate MSPs the way they always have—on cost reduction, labor arbitrage, and service-level agreement-led agreements—risk choosing the wrong kind of provider at a moment when strategic alignment is imperative. Hiring a next-generation MSP challenges that paradigm because harnessing

HIGHLIGHTS

A new generation of managed service providers (MSPs) is emerging **as critical orchestrators and innovation catalysts** in successful agentic AI implementations.

The strongest next-generation MSPs combine three capabilities: **the ability to help define what should be built, the capacity to build it, and the operational strength** to run and improve it over time.

As organizations contend with the disruptive nature of agentic AI, the **critical need for artificial intelligence-proficient people and practical experience** to embed and run operations effectively signals a growing necessity for third-party service provider support.

Agentic AI is advancing so rapidly that no single organization can accumulate the experience needed to navigate it confidently.

agentic AI requires fundamentally different capabilities and measurement approaches.

This paper explores the role of next-generation MSPs in helping accelerate agentic AI ambitions. It covers how these providers can help reshape operations, embrace new economic models, and establish ongoing oversight. The paper also guides forward-looking executives on how to identify which service providers are genuinely equipped for the age of agentic AI and how to work with them most effectively.

Closing the Agentic AI Experience Gap

Agentic AI is advancing so rapidly that no single organization can accumulate the experience needed to navigate it confidently. "There's just not enough market evidence of guaranteed successful ideas," says Alex Bakker, an analyst and research leader at ISG Research, a global technology research and advisory firm based in Stamford, Conn. "Every enterprise has to choose where to pilot, match that to their own complicated environment, and then evaluate over a year or more to see if their approach is actually working."

That uncertainty is compounded by the scale of organizations' agentic AI strategies. Bakker's research found that most large organizations have between 1,000 and 2,000 applications in their environments and want to apply agentic thinking across their entire portfolio. But companies' internal teams are nowhere near equipped to deliver at that scale and, once deployed, handle the ongoing oversight of agentic solutions across the portfolio.

The capacity constraint is compounded by a knowledge problem that many organizations are only just coming to understand.

"Companies no longer have employees who truly understand how their own systems work," says Ronny Hendrych, a program manager at Siemens AG, a global technology and engineering company headquartered in Munich. "The highest risk is losing the institutional knowledge needed to maintain and build on what you have."

Next-generation MSPs with the power of agentic AI can fill that gap by providing the capacity to move faster and at greater scale than any internal team could manage alone. "Having somebody who has managed this kind of complexity before—that is the value, especially in a rapidly changing space like agentic AI," he says.

That cross-client and cross-industry experience becomes highly valuable in the early stages of an agentic deployment, before the organizational obstacles that derail most initiatives have a chance to take hold.

Next-generation MSPs can help enterprises with four categories of organizational debt that must be cleared before agentic AI can operate reliably at scale. The first is process debt, when workflows built for humans cannot be executed by artificial intelligence (AI) without redesign. The second is data debt, including fragmented or inconsistent information that prevents reliable decision making. The third is technical debt, including legacy systems that do not integrate smoothly with AI orchestration layers (the software that coordinates multiple agents working together). And the last is cultural resistance, which is the friction that surfaces whenever human roles shift.

Sequencing those early decisions correctly is what separates deployments and ongoing operations that build momentum from those that stall. "You should start with automation of tasks with a low cost of error and high volume," says Anil Vijayan, a partner at Everest Group, a Dallas-based global research and consulting firm. "Then get to more end-to-end multisystem processes with more meaningful business KPIs to measure, and ultimately to function-by-function transformation."

Thinking Beyond Task Automation

As companies embrace agentic AI, they should be looking beyond the automation of individual tasks to designing autonomous workflows, thus reimagining entire systems and business functions. Consider a typical onboarding process.

According to Vijayan, if this process is approached as task automation, an agent might read a new employee's offer letter and extract basic information, such as role, location, and start date. Function redesign goes further. It creates the HR system record and triggers workflows for IT provisioning, payroll setup, and learning and development enrollment, all handled by separate agents handing off to each other.

"This approach redesigns workflows around key stakeholders, as opposed to being constrained by departments," Vijayan explains. "This enables end-to-end automation from the perspective of a moment that matters."

That kind of transition rarely happens all at once, experts warn. Most organizations need a period in which agent-based transformation runs alongside existing operations, with humans playing a critical human-in-the-loop role, ensuring oversight and collaboration in parallel with agents before a progressive handover occurs to incorporate many parts of an operation. Next-generation MSPs are increasingly being asked to manage that parallel run, thus maintaining continuity of existing operations while the new model is built alongside them and then helping run the new operations on an ongoing basis.

This end-to-end redesign is only possible when the MSP understands the function deeply enough to know where the handoffs should happen and where the exceptions will surface. "Agentic AI is very much tied to process," Vijayan says. "To train agents to execute a process, you need to get to the ground truth. And often that ground truth is in the head of an expert, not written down anywhere. The MSP needs to thoroughly understand the domain they're working in."

An end-to-end redesign—from technology executor to domain-informed orchestrator—has practical consequences for how MSPs structure their engagements.

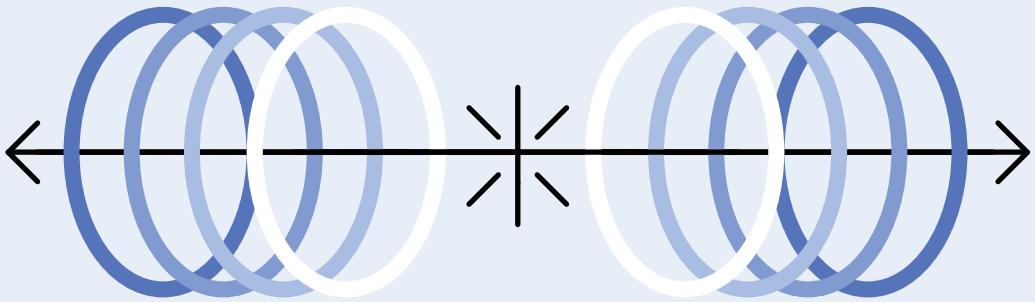
"Agentic AI is very much tied to process. To train agents to execute a process, you need to get to the ground truth. And often that ground truth is in the head of an expert, not written down anywhere. The MSP needs to thoroughly understand the domain they're working in," says Anil Vijayan, a partner at Everest Group.

The New Economics of Agentic AI

The underlying cost structure of agentic AI behaves differently from traditional software and technology budgets, which are often governed by licenses or virtual machines. A new unit of measure is the token, the unit of data an AI model processes when generating text, analyzing images, or reasoning across tasks.

The cost and scale of consumption of tokens vary significantly. Every agent interaction consumes tokens, each carrying a trackable cost that links technology choices directly to business value. This process makes AI spend often nonlinear, challenging traditional total-cost-of-ownership frameworks, which can be hard to manage internally and at scale.

The cost structure is only part of the economic challenge. The future state is new modeling for automation risk but also a move toward outcome-based and transformation-enabled commercial models. Conventional managed services contracts were often built and, in many cases, have remained around head count and fixed deliverables. "The 'automatability' of the process is not known ahead of time," ISG's Bakker says. "This means that every automation attempt is risky for the provider and the client, and right now the standard pricing models don't reward that risk correctly. The goal is to agree up front what the value is of attempting the automation so the provider can make a judgment call on where to focus and what to work on."



The cost and scale of consumption of tokens vary significantly. Every agent interaction consumes tokens, each carrying a trackable cost that links technology choices directly to business value.

> **“What we’re trying to solve for is an agreed-on mechanism that still gives some predictability around progressive automation.”**
>
> Alex Bakker, analyst and research leader, ISG Research

The challenge is sharpest around bespoke processes. Currently, when a client has a custom workflow built on proprietary tools and legacy data integrations, there is no standard mechanism that rewards a service provider for automating it proactively. ISG is working to address that through what Bakker calls autonomy-level pricing.

This approach uses predefined pricing tiers established at the start of a multiyear contract. It gives providers clear incentives to pursue automation without requiring a change order and governance review every time an improvement is proposed. “What we’re trying to solve for is an agreed-on mechanism that still gives some predictability around progressive automation,” he says.

Governing Autonomous Systems

Governance remains a critical challenge for organizations adopting agentic AI. A September 2025 Gartner survey of 360 IT application leaders worldwide found that 13% strongly agreed that they have the right governance structures in place to manage agentic AI. The core design decision for organizations comes down to determining where human judgment remains essential and where autonomous operations can be trusted.

The governance challenge is further complicated in regulated industries, where agents may be making or influencing decisions that require full explainability and auditability. “There may be places where you will need to explain why a particular decision was made this way,” Everest’s Vijayan says. “In those cases, you may not be able to use a neural network, which doesn’t provide the same level of explainability that other approaches do. MSPs can take those calls at the appropriate process and task level and help build out that structure to minimize risk.”

The stakes of getting those calls wrong are highest in environments where agentic decisions have physical consequences, such as a rail system where a malfunctioning operation can lead to a significant safety incident. “In operational technology, there is no undo function,” Siemens’ Hendrych says.

The human dimension of these deployments is equally consequential. When agents are introduced into frontline workflows, employee acceptance matters as much as the technical architecture.

At Mankind Pharma, field managers were initially cautious about the new agentic AI system for the company’s pharmaceutical sales organization. They became concerned that agent-generated recommendations would reduce their autonomy and replace their judgment rather than support it. “The solution was designed as decision support, not decision replacement,” Mankind’s Bansal explains. “Recommendations are transparent, contextual, and explainable rather than prescriptive.”

Over time, as managers experienced reduced preparation time and more consistent team visibility, the response shifted from cautious evaluation to practical appreciation.

Determining where human judgment remains essential also requires its own systematic approach for managing now and in the future. Vijayan describes a phased transition that starts with humans reviewing every decision an agent makes, then shifts to spot-checking a sample of decisions as confidence grows, and eventually reaches a point where humans step in only when something goes wrong.

The questions organizations then need to consider are: How is this human-in-the-loop process managed consistently over time as agentic operations evolve and grow? Can this process be managed in-house or would a next-generation MSP add more value in helping orchestrate it?

Choosing the Right Provider

The criteria for identifying the right service provider start with how the engagement is structured. “If the contract is not focused on business outcomes, you can end up with misaligned incentives that ultimately don’t hit the mark,” Vijayan says.

Agentic AI initiatives cut across business functions and technology teams. Partnerships work best when providers can coordinate across these boundaries rather than operate in silos.

Beyond the contract, he points to two criteria that separate genuine orchestrators from vendors with strong marketing: differentiated intellectual property and proof of scaled deployments—not just pilots but also evidence of implementations that reached production and are delivering ongoing measurable outcomes.

Bansal describes the evaluation process at Mankind Pharma. "We looked well beyond technical skills," he says. "We prioritized partners who demonstrated strong systems thinking, a clear point of view on governance and operating models, and the ability to cocreate alongside our teams rather than simply deliver against fixed requirements."

Effective partnerships require providers that can both implement and operate agentic systems. This prerequisite means experience implementing the technology—including such areas as integrating data, automating workflows, and managing models—as well as the operational capacity to keep systems performing as business needs evolve. Beyond this competency, the strongest service providers act as innovation catalysts, bringing foresight into emerging AI capabilities and proactively suggesting how agentic AI can be leveraged for new product development, market expansion, or competitive disruption.

Industry-specific knowledge separates providers that understand a sector's realities from those that simply know terminology. "There are MSPs out there who are very good in life science and not so much in discrete manufacturing," Hendrych says. "You should think about whether you're new territory for them."

Agentic AI initiatives cut across business functions and technology teams. Partnerships work best when providers can coordinate across these boundaries rather than operate in silos. Successful partnerships also require providers with the ability to scale. "You need somebody who has a large amount of experience deploying those systems in the right scale of your enterprise," according to Hendrych.

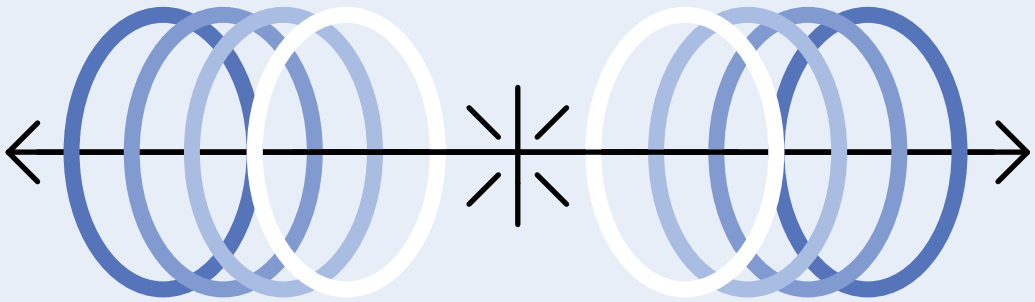
Crucially, this scalability extends to the ability to rapidly prototype, test, and deploy ongoing agentic AI solutions and functional agents across the enterprise, accelerating the pace of innovation and time to market for new offerings. In addition, this requirement means operating models that can accommodate current systems while evolving as needs change, and for organizations with global operations, it also indicates operational resources and capabilities that can maintain consistent performance across regions.

The strongest next-generation MSPs combine three capabilities: the ability to help define what should be built, the capacity to build it, and the operational strength to run and improve it over time.

Once the service provider is selected, the nature of the relationship matters as much as the selection process. "The partnership dynamic for agentic AI has been fundamentally different from traditional vendor relationships," Bansal says. "Our MSP relationships are structured around joint teams, frequent feedback loops, and shared responsibility for outcomes and targets—not just delivery milestones. This creates a partnership model grounded in shared ambition rather than contractual execution."

Conclusion

As organizations contend with the disruptive nature of agentic AI, the critical need for AI-proficient people and practical experience to embed and run operations



"Our MSP relationships are structured around joint teams, frequent feedback loops, and shared responsibility for outcomes and targets—not just delivery milestones."

Kunal Bansal
Chief digital officer
Mankind Pharma

> **“What surprised me is MSPs’ willingness to invest with their own money on developing agentic AI.”**
>
> Ronny Hendrych, program manager, Siemens AG

effectively signals a growing necessity for third-party service provider support.

Next-generation MSPs can bring a distinct blend of industry and domain knowledge, technical acumen, and scaled capability. Crucially, they can serve as orchestrators in the agentic AI ecosystem—guiding organizations through the systematic deployment and, importantly, the ongoing optimization of agentic AI to deliver tangible long-term results.

“What surprised me is MSPs’ willingness to invest with their own money on developing agentic AI,” Hendrych says. “Ten years ago, they would need to see a really good return on investment before committing to a new technology, but now they understand they have to invest heavily and try out agentic AI themselves. For me, this is the best indication that we all believe agentic AI is a real human industry disruption and MSPs will play a key part in orchestrating that disruption.”

Next-generation MSPs, with their deep skill base, scalability, and commitment to cocreation, are uniquely positioned to help guide enterprises on their agentic AI journey, helping ensure initiatives deliver not just efficiencies but also sustainable advantage and enduring business value.

HARVARD BUSINESS REVIEW

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.

CRE77061126

Copyright © 2026 Harvard Business School Publishing.