

Technology, Media & Telecommunications Practice

Reimagining the value proposition of tech services for agentic AI

As enterprise customers seek support with agentic AI, tech services providers have a chance to turn a potential threat into an even bigger opportunity.

This article is a collaborative effort by Anuj Kadyan, Pallav Jain, Piyush Sharma, and Vikash Daga, with Anuja Rege and Radhika Sriram, representing views from McKinsey's Technology, Media & Telecommunications Practice.



After more than two years of navigating the transformative landscape of gen AI, technology services providers are now facing the emergence of a newer, more disruptive force to their business. Enterprises that have traditionally relied on these providers to manage their IT initiatives are now making significant investments in agentic AI, the next evolutionary stage of artificial intelligence. These organizations are cautiously optimistic that [agentic AI will deliver the top- and bottom-line growth](#) that gen AI has, to date, struggled to achieve. In response, most tech service players have started exploring use cases internally, such as agent-assisted software development, delivery management, and operations, as well as externally, including customer service, IT ticket resolution, and financial planning and analysis (FP&A) use cases.

Among the 200 C-suite executives across a wide range of enterprises that we recently surveyed, more than 80 percent of respondents report that they are already running pilots on agentic AI, with some progressing to scaled deployments.¹ This momentum is driving a notable increase in overall IT spending. Like gen AI before it, the potential downside is clear: As organizations leverage agentic AI to achieve significant productivity gains and bring more technology management in-house, the core business of technology services providers could face a 20 to 30 percent contraction.

Yet, this disruption also presents an opportunity. Enterprises are seeking support to navigate the complexities of adopting a joint human–agent operating model, and technology services providers are well positioned to play a critical role. If approached strategically, agentic AI could become a net positive for the sector, potentially driving an additional three percentage points of annual growth and unlocking \$100 billion to \$400 billion in incremental spending by the end of the decade.

Not that tech services players will have that new value pool all to themselves. The rise of agentic AI is blurring traditional boundaries among technology services providers, hyperscalers, software vendors, and a new, emerging class of disruptive, pure-play AI companies. To thrive in this dynamic environment, technology services players will need to fully embrace agentic AI, developing new capabilities and reimagining their ways of working.

Based on our survey and experience in the market, this article examines what that transformation will look like and require. It examines how the market is likely to evolve, identifies new roles and growth opportunities for providers, and outlines the steps to adapt core businesses and succeed in this new reality.

Agentic AI's impact on the tech services market

There is little doubt that leading enterprises are accelerating their push into agentic AI. According to our survey, roughly 12 percent of enterprises have already scaled deployments of the technology across multiple functions, and half are planning significant investments in scaled deployments over the next six months (Exhibit 1). This is translating into a sizeable increase in AI budgets for the current fiscal year. More than a third of respondents expect the annual outlays

¹The McKinsey Enterprise CxO Agentic AI Survey of 200 C-suite executives was conducted in July 2025 across Asia, Europe, and North America, spanning six industries and companies with annual revenues above \$500 million. Insights are complemented by McKinsey market-sizing models and NASSCOM data (2025 baseline, 2025–30 forecast).

to soar by more than 25 percent, and close to three-quarters say the jump will exceed 10 percent. This shift is fueling a broader increase in overall tech spending. Around three-quarters of all enterprises expect their total IT expenditures to grow, with the vast majority anticipating a 2 to 10 percent spike in the next two years (Exhibit 2).

Exhibit 1


More than half of enterprises are scaling or successfully piloting agentic AI and plan to invest significantly over the next six months.

Agentic AI adoption and investment

Adoption archetypes, % of enterprises



Share of respondents reporting their organization has invested or is planning substantial investments in agentic AI over the next 6 months, by archetype, %



Note: Figures may not sum to 100%, because of rounding.
Source: McKinsey Enterprise CxO Agentic AI Survey, July 2025 (n = 200)


McKinsey & Company

What the growth of agentic AI and its myriad capabilities means for tech services providers is not straightforward, as it ushers in two countervailing shifts. On one hand, traditional IT budgets that the sector relies on are expected to shrink over the coming years. On the other, emerging agentic AI workflow service lines and an expansion into new business function value pools have the potential to more than offset that legacy decline with an expanded total addressable market. With the emergence of agentic AI, the total tech services market is expected to grow by 4 to 7 percent annually over the next five years, reaching \$1.6 trillion to \$1.9 trillion by 2030, outpacing the previous, pre-gen/agentic AI estimates of 4 to 5 percent CAGR (Exhibit 3).

Exhibit 2

A majority of enterprises expect to increase both total and outsourced IT spend in the coming years because of agentic AI.

Expected change in IT spending over the next 2 years, as a result of agentic AI, % of respondents



Note: Figures may not sum to 100%, because of rounding.
Source: McKinsey Enterprise CxO Agentic AI Survey, July 2025 (n = 200)


McKinsey & Company

However, that positive outcome is far from guaranteed. The erosion of the core business is a significant threat that could well become existential for those tech services players that don't make the necessary adjustments and repositioning in response to agentic AI. Agentic AI is expected to drive a 20 to 30 percent compression in the core tech services market across existing service lines, primarily led by two forces: employee productivity gains from agent deployment and the world's largest enterprises increasingly turning to global capability centers, which service providers typically only play an initial role in standing up, to build and manage agentic capabilities in-house. More structurally, the compression will reshape the nature of

Exhibit 3

While agentic AI could drive a more than \$300 billion compression in legacy tech services, it could also enable an expansion in newer value pools.

Global tech services market, \$ trillion



Source: NASSCOM; McKinsey analysis

McKinsey & Company

services work itself, shifting the work performed by employees from execution tasks (such as monitoring and processing) to enablement-focused activities (such as strategy design and stakeholder engagement). To simply maintain revenue (or minimize revenue loss) from existing service lines, traditional tech services players would be forced to move into adjacent digital service lines, including cloud, digital product engineering, and data and AI services, intensifying competition in the already crowded segment.

The intricacies of successfully implementing AI mean that tech services players may still have a critical role to play for their clients. While enterprise focus on in-house agentic AI capabilities and deployments is rising, with more than three-quarters of surveyed organizations reporting the use of gen AI in at least one business function, a similar share have reported seeing [no material impact on earnings](#). Our survey highlights three core challenges hindering agentic AI

adoption at scale: the complexity of integrating agentic AI with existing systems, limited internal technical expertise or talent, and growing concerns about [agent security vulnerabilities](#). As a result, agentic AI offers tech services providers two primary avenues for new growth:

- *Agentic AI workflow services*: Our survey revealed that enterprises expanding their usage of agentic AI expect tech services players to perform multiple roles across the AI stack, including infrastructure support, multiagent architecture design, talent and change management, and rapid prototyping. This could spark a new pool of enterprise spending, worth nearly \$200 billion, on an array of digital services revolving around such critical work as orchestration, agent engineering, agent security, and governance (Exhibit 4). To seize this new opportunity, services players will have to leapfrog into AI and develop end-to-end capabilities, including expertise in LLMs and multiagent systems.
- *Business function transformation*: The other significant new value pool, potentially worth \$100 billion to \$400 billion in annual spending, lies in transforming core business functions through joint human and agent operating models, particularly in knowledge roles.

According to our survey, enterprises on average believe as much as 15 to 30 percent of their current roles' work could be taken on by agents over the next three years, and they expect to make their agentic investments in four key areas: technology and engineering function agents, for jobs like testing and migrating software code and handling root cause analysis or DevOps; customer-facing agents, for content creation, sales pitch assistance, and client onboarding; back-end function agents, for such activities as call center coverage, legal services, and ticket routing; and vertical-specific process agents, including patient care management, fleet routing, claims authorization, and credit risk assessment. These investments have a direct impact on traditional spend pools across direct labor, R&D, and sales, general, and administrative expenses. Notably, more than 70 percent of this opportunity is projected to be driven by five major industries that are particularly well suited to automation: financial services, retail, high tech (including software and hardware), healthcare, and manufacturing (including automotive).

To capture this expanded opportunity, tech services players will need to evolve their traditionally fragmented model and bring together their core capabilities across functions such as operations, process reengineering, data and AI services, and software engineering. One of the leading providers, for instance, has already merged previously separate service lines into a single AI-centric business unit.

These overarching shifts are also leading to a fundamental redrawing of the competitive landscape for tech services providers and the other companies that play in that space. Based on the survey and our experience, we expect three main structural changes to occur as part of this transition.

Blurring of traditional archetype boundaries: The clear lines that once separated software vendors from service providers are becoming fuzzier as tech services players provide “service as software” through agentic operations and custom agent development. In parallel, independent

Exhibit 4

Tech services players may tap into a new pool of enterprise spending across three stages of AI workflow transformation.

Stages and categories of enterprise AI transformation, by potential tech services offerings

Advisory	Development and integration			Managed services		
						
AI strategy and opportunity mapping	Architecture and model orchestration	Data for AI services	Agent development	Deployment, ops, and monitoring	AI security	Responsible AI and governance
• Use case identification • Agent workflow design • Process engineering	• Retrieval-augmented generation (RAG) agent core • Multiagent orchestration systems	• Data architecture for agents • Data connectivity and workflow enablement • RAG and semantic retrieval services • Unstructured data activation	• Persona-driven agent build • Fine-tuning	• Performance monitoring • Learning and feedback services	• Zero-trust agent access controls • Prompt injection defense • Agent behavior constraints	• Trust and compliance • Human in the loop • Escalation management

Source: NASSCOM; McKinsey analysis

McKinsey & Company

software vendors, tapping their software-as-a-service (SaaS) workflow and business function expertise, and hyperscalers, leveraging their scale and infrastructure stack, are offering managed agentic solutions, bespoke agent deployment services, and agentic marketplaces. For example, while service providers are already launching agentic platforms focused on technology and engineering functions, software providers are moving toward managed agentic services that combine software with operations.

Rise of new disruptors: Agentic AI's capabilities are creating room for a new class of disruptors that have often been limited to a specific function or vertical. This includes tech services players, with their deep domain specialization, engineering strength, and a platform-led offering. However, it is also a valuable new opportunity for LLM providers and other pure-play AI platform providers that can, for instance, build domain-specific agents that threaten traditional services value pools. For instance, a popular AI start-up claims its customer service agents can capably

resolve up to 90 percent of inquiries and interactions without any human intervention, dramatically raising an enterprise's resolution rate and reducing the customer effort.

Emergence of new partnerships: As a general rule, services providers used to pair up with other enterprise vendors once they had reached a certain scale in the market. That is less of an option in the rapidly evolving agentic era, where small, specialized software companies, including agentic orchestration providers and vertical-focused vendors, already play key roles in the end-to-end ecosystem. To deliver comprehensive, best-of-breed solutions, expanding the breadth of partnerships from the get-go has become paramount for all players. This includes not only tech services companies partnering with hyperscalers to develop comprehensive, customized offerings but also strategic SaaS alliances designed to build cross-functional solutions (like seamless integrations between enterprise platforms). For instance, two leading industry players are building the foundation for interoperable enterprise agents that collaborate seamlessly across different SaaS ecosystems, a model in which a finance agent could pull data from enterprise resource planning, analyze it in spreadsheets through an AI assistant, and trigger HR actions in a workspace collaboration tool.

Reengineering tech services for the agentic AI era

Tech services players are responding to this changing landscape in a number of ways. Some are pursuing a more cautious transition, maintaining their focus on automation through generative AI use cases while starting to explore the emergence of agentic capabilities. Others are engaging in a more wholesale embrace of platform-led delivery, embedding orchestration layers to manage agents (and agentified business processes) at scale. A more advanced segment is already building modular, domain-specific agents that can be quickly customized and deployed. And yet another cohort is racing to partner up with the likes of hyperscalers and LLM providers to augment their capabilities and offer joint solutions.

Figuring out the most suitable, long-term positioning in this new agentic environment will take some time. The right value proposition isn't likely to be one-size-fits-all; even the same provider may play a distinct role across different verticals and domains. The right fit will depend on several factors, including domain expertise, access to enterprise data, strength of client relationships, maturity of agentic services, and platform capabilities.

Based on our survey and experience, we think that four primary value propositions will emerge for tech services providers in the agentic era:

- *'Agentic AI enabler' (cloud, data, agent services):* Provide foundational infrastructure and agent services for other parties (in-house or external) to build, deploy, and scale agentic AI solutions. Success in this role requires deep technical and architectural expertise, strong partnerships with hyperscalers (given their scale and infrastructure position) and LLM providers, and the ability to deliver custom AI/agentic services.

- *'Packaged agent implementer'*: Leverage a strong partner ecosystem to deploy third-party agent solutions that address common business challenges with minimal customization, underpinned by a prebuilt agent library, rapid deployment frameworks, and seamless integration solutions.
- *'Custom agent developer'*: Create bespoke agentic AI solutions tailored to an enterprise's unique and often complex business context, which calls for domain expertise and access, agentic-stack development, and design capabilities.
- *'End-to-end workflow disruptor'*: Act as a transformational partner by going beyond agent development and deployment to redesigning and implementing entire customized workflows, requiring deep domain and enterprise depth to drive large-scale change.

As the landscape evolves, tech services players will need to move from their current siloed, service-line-centered engagement models and reposition themselves as a holistic provider within each of the four roles. Each play demands a different set of capabilities across dimensions, and understanding where to double down is critical (Exhibit 5).

Building new agentic services capabilities

Agentic AI is not only transforming how technology is built and delivered but also how enterprises buy it. As organizations shift from experimentation to scaled deployment, and a growing array of players compete to guide them through the journey, the criteria for selecting service partners will evolve rapidly. Enterprises we surveyed cite six core factors in choosing to work with a services provider: the ability to customize solutions, the strength of their partnership ecosystem and intellectual property (IP), a consultative sales engine, domain expertise, line of business-focused delivery, and outcome-based pricing and commercial models.

At the outset, tech services players seem relatively well positioned in many of those areas, especially given their deep enterprise context and experience and vertical know-how. But they will need to go beyond incremental adjustments to assert a strong position in the market. We have already presented an [overarching agentic-transformation framework](#); following that playbook, we see five foundational capabilities that will distinguish leading service providers in this AI/agentic era:

1. *Reimagine positioning and offerings around agentic-first opportunities*: Success in the agentic era starts with redefining the core value proposition and reinventing services to align with emerging value pools. This includes building a diverse portfolio of vertical business offerings (for example, agent-led claims management), horizontal AI-led solutions (such as a sales coaching agent or FP&A copilots), and foundational capabilities, including agent orchestration services. As part of this process, it is critical to clearly articulate the business outcomes and bankable productivity commitments, backed by domain-specific

Exhibit 5

Four primary agentic AI value propositions are likely to emerge for tech services players, each requiring different capabilities.

Tech services agentic AI value propositions, by key factors and capability strengths

Capability strength ■ Low ■ Medium ■ High

Key factor	Agentic AI value propositions			
	Agentic AI enabler	Packaged agent implementer	Custom agent developer	End-to-end workflow disruptor
Domain expertise and access to enterprise data	■ Broad domain understanding needed to enable foundational infrastructure and services	■ Minimal domain and data access needed due to limited customization	■ Deep domain expertise and privileged data access for agent tuning and contextualization	■ Critical requirement for designing agent-led workflow solutions
Engineering depth (eg, process engineering, data, and AI)	■ High depth needed across cloud, data engineering, and orchestration engineering	■ Integration of prebuilt agents across processes	■ Strong engineering depth required to build tailored agents and integrate across enterprise	■ Deep process and AI engineering to embed agents into reimagined workflows
Strength of agentic services (eg, agent orchestration, agent security, governance)	■ Core strength required for supporting enterprises on agent services	■ Limited requirement for deploying prebuilt agents on external platforms	■ Requires full-stack build and orchestration but can leverage agentic AI enablers	■ Focus on embedding and building multiagents and less on managing them
Proprietary platforms and agents	■ Can leverage partnerships with hyperscalers to scale orchestration and deployment	■ Requires reusable agent libraries and deployment tool kits	■ Needs modular agent development frameworks and tuning tool kits	■ Platforms to enable integrated workflow and agent transformation
Talent and change management	■ Needs scaled engineering talent (cloud, infrastructure, AI) but minimal org change focus	■ Light functional talent requires rollout teams and basic client enablement	■ Requires niche AI talent (large language model, prompt, orchestration) and cocreation pods	■ Demands cross-functional squads and strong change management for enterprise-wide adoption

McKinsey & Company

credentials and repeatable agentic use cases. Providers will also have to be careful to avoid overindexing on experimental copilots without scalable solutions or commercial models, which can limit enterprise adoption.

2. *Build proprietary solutions to orchestrate, adapt, and scale:* Players can build proprietary platforms or solutions to create vertical and domain-specific IP while enabling adaptable tech stacks. Developing modular agentic AI architecture with differentiable components that can be reused across clients is key to this effort. An integrated platform, for example, combines multiple LLMs, allowing for composable agent deployment across development, testing, and design workflows. Providers should not underestimate the importance of built-in agent observability, governance, and adaptability, without which firms risk deploying opaque systems that behave unpredictably across client environments. That can lead to compliance breaches, quality lapses, and costly remediation efforts that erode credibility and margin.
3. *Lead with consultative, domain-driven go-to-market models:* The shift from IT to business-led buying means providers need to rely on a combination of consultative selling and domain knowledge, but that is just the start. In this new era, the strict dividing line between buyer and seller will become fuzzier, as forward-thinking players proactively embrace rapid cocreation with their customers to help unlock adjacent business opportunities. Cross-functional squads including consultants, prompt engineers, and data specialists can build and iterate agentic prototypes directly with clients. As part of this approach, providers should avoid positioning AI as a “horizontal capability,” with domain specificity a critical success factor. It’s equally important to rely on agile development, as the more traditional sequential delivery model can slow adoption and weaken impact.
4. *Redesign operating model and talent:* The shift to an agentic paradigm necessitates a [reimagined operating model centered on human–agent collaboration](#). Talent strategies should focus on rapid AI reskilling and upskilling, fostering a culture of continuous learning, and optimizing human–agent collaboration to achieve maximum productivity and value creation. As part of this reorientation, providers can build a clear, AI-native delivery model with defined human–agent handoffs, supported by agent operation centers and new roles specifically designed for overseeing, (re)training, scaling, and collaborating with autonomous agents. Another AI platform, for example, is rolling out human–agent delivery pods with centralized governance through an “AI command center” that tracks agent performance, retraining cycles, and exception management—ensuring scalable, governed adoption across clients. Successful human–agent collaboration can be challenging to pull off. Without deliberate planning for adoption and aligned incentives, human behavior may not change (or change enough); teams may bypass agents, duplicate work, or revert to manual processes, preventing significant productivity gains.
5. *Reinvent commercial models to align with impact:* As clients increasingly expect clear productivity and margin benefits from service providers, the traditional time-and-material models risk rapid margin erosion. More than 70 percent of enterprises we surveyed expressed a preference for alternative pricing models, suggesting players may need to move toward new monetization strategies. These could include subscription-based models or fixed-price models. One offering, for instance, is a scalable, token-based subscription model that charges a flat monthly fee per pod, combining engineers and AI agents—turning

Find more content like this on the
McKinsey Insights App



Scan • Download • Personalize



IP into a billable asset and ensuring predictable, SaaS-like revenue. As gen AI/agentive AI use cases scale, gain-share models linking fees directly to measurable impact such as cost savings will likely become more prevalent. While overall productivity gains will also be an important variable, full-time-equivalent productivity may need to be deemphasized or entirely delinked from commercial models, given how much AI automates workflows and processes. Whatever precise commercial model enterprises land on, it will have to feature a transparent pricing framework and clear proof of realized value. If either of those is lacking, clients may well end up disputing gains, which can drive margin erosion and undermine trust in AI-led delivery.

For tech services players already trying to figure out how to handle the disruptive force of gen AI, agentive AI presents a potentially greater opportunity—or threat. While the technology's capabilities could help to depress the industry's once reliable core IT management business, enterprise customers' struggles with incorporating and profiting from agentive AI could fuel a lucrative new chance for value creation. A growing roster of both incumbent and upstart competitors has its eyes on the same prize, however, and service players can't assume they will walk away with their fair share.

Tapping these new avenues for growth will require service providers to undergo a challenging transformation. In the near term, tech services firms will need to rapidly build agentive capabilities and position themselves as trusted learning partners to enterprises, offering the technical expertise, infrastructure, and guidance needed to scale responsibly. The focus should be on demonstrating tangible value through initial deployments and iterating rapidly. Over the long term, sustained leadership will require a more structural response. Tech services firms will need to cultivate deep domain expertise, combine capabilities across service lines to deliver integrated solutions, secure unique access to enterprise data, and master hybrid human–agent operational models. Only by embarking on such a comprehensive overhaul can tech services players hope to succeed in becoming enterprises' favored guides along their own agentive journeys.

Anuj Kadyan is a senior partner in McKinsey's New Jersey office, where **Piyush Sharma** is a partner; **Pallav Jain** is a senior partner in the Atlanta office; **Vikash Daga** is a senior partner in the Gurugram office; and **Anuja Rege** is an associate partner in the Mumbai office, where **Radhika Sriram** is a consultant.

The authors wish to thank Budakausik Vedula, Dhruv Shah, Gaurav Sharma, Kanika Gupta, PVS Anurag, Raghavan Janardhanan, and Shruti Kumar for their contributions to this article.

This article was edited by Daniel Eisenberg, an executive editor in the New York office.

Copyright © 2025 McKinsey & Company. All rights reserved.