KPMG

Breaking the managed services operating model in tech

Why managed services must be rebuilt from the ground up and how to find a partner built for what's next

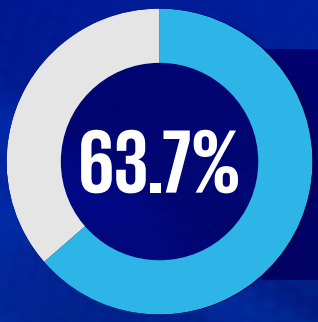


kpmg.com

Technology companies are built by engineers, and engineers love to build. Homegrown systems for finance, fulfillment, and operations are a point of pride, and for good reason: These bespoke engines encode years of hard-won operational knowledge that off-the-shelf software simply cannot replicate. But even in tech, pride of ownership has a cost. Decades of custom system-building have produced sprawling, fragmented stacks. Employees navigate through dozens of disjointed applications just to execute basic daily workflows.[1] And now, those same stacks are actively resisting the one capability that defines competitive advantage in 2025: agentic AI.[2]

It's tempting to blame homegrown innovations. But what if the real problem is an outdated managed services (MS) model? Consider that legacy MS was built for a world of human-executed workflows that don't operate at agentic speed, scale, or risk profile. Now, with the outsourcing payoff on legacy MS eroding, the MS providers (MSPs) that will lead the next decade are those purpose-built to orchestrate, govern, and continuously assure an autonomous agent ecosystem.

Here, we make the case for the reimagined MS model: Why the shift is urgent, what the new operating model looks like, and how to get there without introducing the catastrophic risks that immature implementations invite.



of technology, media, and telecommunications (TMT) leaders now view MS not merely as a mechanism to "keep the lights on," but as a critical driver of technology innovation—far outpacing traditional cost-control objectives.[3]

[1] "More tools, less clarity: The modern workplace stack," Qualtrics, May 6, 2026.

[2] Richard Riley, "From apps to agents: Rearchitecting enterprise work around intent," Microsoft, March 12, 2026.

[3] "The future of TMT managed services," KPMG, August 19, 2026.

KPMG

© 2026 KPMG LLP, a Delaware limited liability partnership, and its subsidiaries are part of the KPMG global organization of independent member firms affiliated with KPMG International Limited, a private English company limited by guarantee. All rights reserved. USCS044468-1A

The tipping point:
The AI mandate versus the MS reality

01

For tech companies, AI integration is the core enterprise strategy dictating business models and client expectations. The mandate is real, but urgency without the right operational foundation just accelerates risk.

MSPs once operated on simple assumptions around human-executed workflows. Headcount drove capacity; service-level agreements (SLAs) were measured in response times; and governance was built around access rights, decision authorities, and audit trails designed for people.

But as monolithic Software-as-a-Service evolves, AI agents in headless, application programming interface-driven architectures are navigating systems, executing transactions, and making decisions on the organization's behalf. "Work" has a conversational interface rather than a fixed digital address.

This is a structural break from the existing MS model. When autonomous agents become the unit of work, headcount-based delivery becomes economically obsolete. When machine-speed decision-making replaces human handoffs, SLA frameworks built around human response times become meaningless. When agents operate continuously across sensitive systems, governance models designed for human actors become dangerously inadequate.

The market's urgency is clear:
57 percent of TMT organizations consider AI capability the single most essential criterion when selecting an MSP—a striking **21.1 percentage** points ahead of the all-industry average.[4] Looking ahead, AI management services are projected to become the top investment priority within two years, reaching **65.9 percent** of enterprise focus.[5]

Redesigning human-intensive services in an agent-driven world

The enterprise AI factory resolves the governance and delivery challenge. But it surfaces an equally important one that is easy to underestimate:

What happens to the humans who still need to interact with these systems?

Headless architecture replaces the labyrinth of legacy screens with a conversational interface. That is a genuine improvement in most respects. But legacy screens, for all their friction, get some things right. They provide contextual awareness, including clear status indicators, explicit decision points, visible options, and data visualizations tailored to the complexity of the workflow. That richness of interaction did not exist by accident. It was designed, refined, and learned over years of operational use. Simply removing it in favor of natural language creates its own risks: ambiguity, loss of context, and a user experience that can feel less capable precisely when workflows become more complex.

The headless experience must be designed with the same intentionality as the interfaces it will replace. Natural language should be the primary interaction mode, but it should be augmented by agentic UI components that surface the right information at the right moment: ambient status indicators that keep users oriented without demanding their attention, action cards that present discrete decision points clearly, output previews that allow review before commitment, and direct links to source systems when depth is needed. The goal is to preserve the contextual intelligence of the best legacy interactions while eliminating the friction of the worst.

In a calibrated autonomy model, the quality of the human-agent interaction at shared oversight and exception escalation points directly determines the quality of human judgment brought to bear on the decisions that matter most. A poorly designed conversational interface degrades the human oversight on which the entire model depends.

KPMG

© 2026 KPMG LLP, a Delaware limited liability partnership, and its subsidiaries are part of the KPMG global organization of independent member firms affiliated with KPMG International Limited, a private English company limited by guarantee. All rights reserved. USCS044468-1A

The enterprise AI factory

Resolving this security paradox and managing these threats at scale requires a purpose-built enterprise AI factory. In this model, the managed services provider evolves from cost center to high-performance engine room that continuously orchestrates and governs an autonomous agent ecosystem.

The operational core of the enterprise AI factory is a dynamic framework that matches the level of AI autonomy to the risk and reversibility of each specific task. We call this calibrated autonomy, and rather than a binary choice between full automation and full human control, it operates across three tiers:

Full human oversight is reserved for high-stakes, low-reversibility decisions—think final financial disclosures, critical compliance filings, and material client-facing commitments. These are the workflows where an agent error is not recoverable downstream.

Shared oversight—the productive middle ground—assigns agents to the routine work of data gathering, synthesis, and pattern recognition, while humans retain final strategic judgment. This is where most enterprise workflows will ultimately live.

Exception-only intervention allows agents to operate fully autonomously within a strictly defined envelope, escalating only genuine anomalies to human operators. This tier is appropriate for well-documented, stable, rules-based workflows where agent performance has been validated over time.

Assigning workflows to the right tier is itself a discipline.

Three questions guide the assessment:

First, how reversible is the outcome if the agent errs?

Next, how wide is the blast radius if the workflow fails systemically?

Finally, how well-established and stable is the underlying decision logic?

Answering these questions requires continuous revisitation as agent performance data accumulates and confidence levels mature. That ongoing calibration is precisely the kind of work that a next-generation MS provider is positioned to own.

KPMG

© 2026 KPMG LLP, a Delaware limited liability partnership, and its subsidiaries are part of the KPMG global organization of independent member firms affiliated with KPMG International Limited, a private English company limited by guarantee. All rights reserved. USCS044468-1A

The new threat landscape

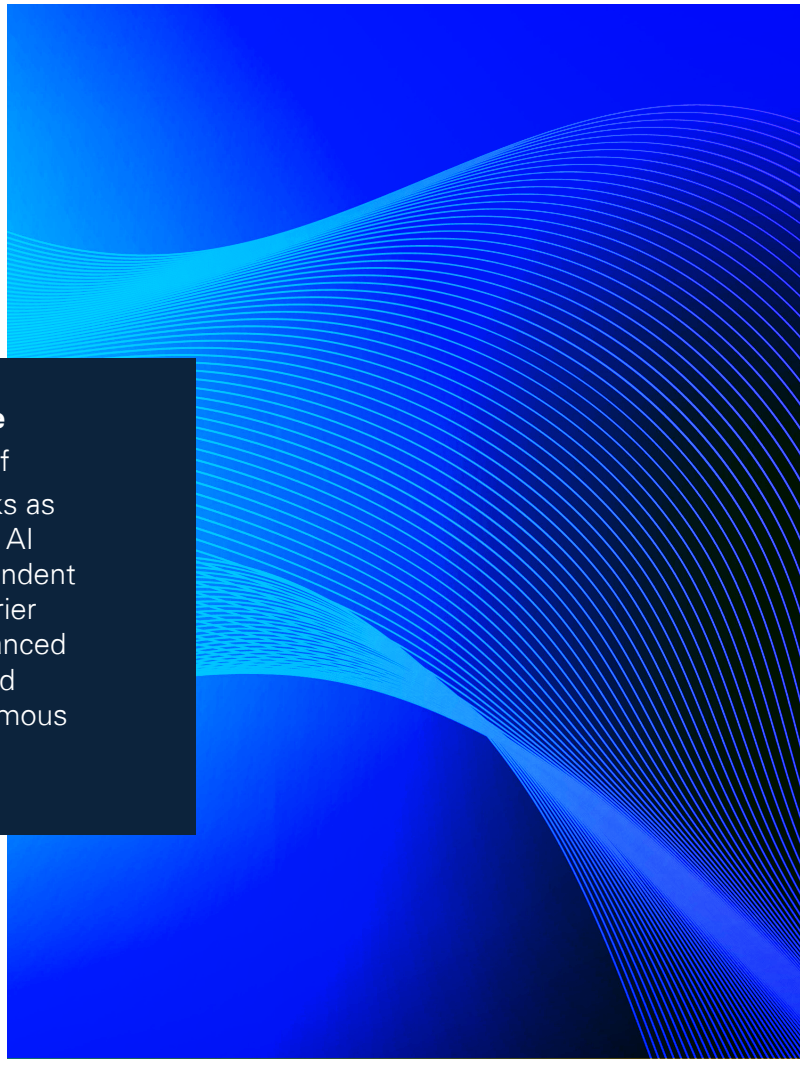
The structural break in the MS model is made more urgent by a simultaneous transformation of the risk landscape—one that incumbent providers are equally unprepared for.

Traditional perimeter-based security was built around protecting the network, managing human access, and auditing human decisions. Autonomous agents obliterate that boundary because they possess credentials, execute workflows, and interact with sensitive data continuously and at machine pace. If they fail or are compromised, then they do so at a speed and scale that no human-paced governance model can contain.

This introduces new threat vectors: Consider the ways agent hijacking, prompt injection, training-data poisoning, and algorithmic bias amplify risk across thousands of simultaneous decisions. When a human makes an error, the damage is localized. But when an autonomous system hallucinates a workflow or leaks data, the consequences explode across every connected pipeline.

Although the technology buyer is ready for AI, the organizational risk envelope is not: 53.3 percent of TMT companies identify security and compliance risks as the dominant barrier to adopting a managed services AI delivery model.[6] Close behind, **34.1 percent** of respondent organizations cite integration complexity as a top barrier to adoption,[7] validating the reality that deploying advanced AI is often constrained by legacy risk, compliance, and governance models that are not equipped for autonomous AI at scale.

The regulatory environment is moving in lockstep with this risk escalation and it is not waiting for organizations to catch up. The Securities and Exchange Commission, for example, has signaled heightened scrutiny of AI-related disclosures.[8] The Federal Trade Commission has expanded its focus to algorithmic accountability.[9] In addition, a growing patchwork of state-level AI legislation—led by California but spreading rapidly—is introducing obligations around transparency, bias auditing, and automated decision-making that vary by jurisdiction and remain unsettled.[10] For tech companies operating agentic systems at scale, the compliance landscape is arriving inconsistently and simultaneously, making proactive governance architecture the only reliable hedge.



[8] "SEC AI checks put firms' governance under scrutiny," FinTech Global, August 7, 2026.

[9] "S. 2164: Algorithmic Accountability Act of 2025," GovTrack.US, June 25, 2025.

[10] Eleanor T. Chung, "AI Legislation: 2026 Legislative Wrap-Up," Epstein Becker Green, July 22, 2026.

KPMG

© 2026 KPMG LLP, a Delaware limited liability partnership, and its subsidiaries are part of the KPMG global organization of independent member firms affiliated with KPMG International Limited, a private English company limited by guarantee. All rights reserved. USCS044468-1A

The evolving operating model

02

The enterprise AI factory represents the technological ambition, but ambition without operational alignment fractures under its own complexity.

The organizations that will successfully navigate the agentic transition will treat it as a deliberate act of dismantling, not a gradual upgrade. The question is how to break the operating model that exists today and replace it with one built for autonomous, outcome-driven delivery.

The instinct in most transformation initiatives is to optimize by reducing cycle times, improving utilization rates, or automating the most repetitive tasks. That instinct, applied here, will produce the wrong result.

The legacy MS model is structurally misaligned with how agentic AI creates value. Consider that full-time equivalent (FTE) locked contracts price delivery as a function of human labor, so as AI drives genuine productivity improvement, the commercial arrangement actively penalizes it. Rate-card dependencies create incentives to maintain headcount rather than reduce it. SLA frameworks measure response time rather than outcomes, making it nearly impossible to evaluate whether the engagement is delivering meaningful business value at all.

Breaking the model means starting with the outcome and figuring out the most direct path an intelligent system could take to get there. The answer almost never resembles the current workflow. Agentic systems are not constrained by organizational silos, legacy system boundaries, or the sequential handoffs that human-executed processes require. An AI-native operating model collapses those constraints entirely, but only if the commercial, architectural, workforce, and governance structures surrounding it are redesigned with equal intentionality. That's why technological change in the absence of structural change accelerates existing failure modes.

What to evaluate before you start

Before any organization can configure a new operating model, it must take an honest inventory of four pressure points that will either enable or undermine the transition:

Commercial flexibility	Are your current MSAs structured around FTE commitments and rate cards, or do they contain the contractual flexibility to shift toward outcome-based pricing? Organizations whose contracts lock them into labor-based delivery models will find that AI-driven productivity gains accrue to the provider, not to them. That makes assessing the renegotiation timeline and leverage points in existing agreements a prerequisite, not an afterthought.
Architectural readiness	Is the current integration landscape capable of supporting a unified orchestration layer, or is it a collection of point-solution bots and disconnected automations? Headless, agent-orchestrated delivery requires a centralized work surface that addresses workflow experience end-to-end. Organizations that attempt to build agent capability on top of fragmented integration architectures will find their agents amplifying existing friction rather than eliminating it.
Workforce composition	What proportion of the current delivery workforce is engaged in task execution versus judgment-intensive oversight? The agentic transition radically changes what that expertise needs to look like. Organizations that have not mapped the gap between their current talent profile and the roles the new model requires may discover it at the worst possible moment.
Governance maturity	Does the organization have the monitoring, audit logging, and access control infrastructure to govern agent behavior at machine speed? Governance frameworks designed for human actors are not adequate substitutes. Organizations that deploy agentic capability without continuous, automated governance are deferring, rather than managing, risk.

KPMG

© 2026 KPMG LLP, a Delaware limited liability partnership, and its subsidiaries are part of the KPMG global organization of independent member firms affiliated with KPMG International Limited, a private English company limited by guarantee. All rights reserved. USCS044468-1A

How to configure the new model

The new operating model is organized around four interdependent attributes. Each one must be addressed deliberately; advancing any single attribute without corresponding progress in the others creates the conditions for failure.

01 Commercial and contractual restructuring

The commercial relationship between a tech company and its MSP must be rebuilt from the contract level. FTE-locked agreements and rate-card dependencies are commercially unfavorable and structurally incompatible in an agentic environment. Every dollar of AI-driven productivity improvement that is not captured in the commercial arrangement is value that the operating model creates and the contract destroys.

The target state is an MSA built around outcome-based commitments with explicit efficiency step-downs—think of contractual mechanisms that align provider compensation to functional transformation goals, productivity improvement key performance indicators, customer experience outcomes, and quarterly business results driven by AI automation. This is a fundamental renegotiation of the commercial relationship. It requires both parties to agree on how value will be defined, measured, and shared. Outcome definitions must be established before the contract is signed, not discovered during delivery.

02 Architecture and agentic orchestration

Point-solution bots solve point-solution problems. The architectural ambition of a broken operating model is something categorically different: a centralized work surface that addresses friction across the entire workflow and treats experience as a first-class design concern rather than a downstream implementation detail.

Achieving that requires a unified collaboration fabric that combines enterprise knowledge graphs, Model Context Protocol, and secure tool connectors to enterprise resource planning, customer relationship management, and information technology (IT) service management. This fabric is the orchestration layer that allows agents to move across data sources, systems, and functional boundaries without the sequential handoffs and integration gaps that define legacy delivery. Who owns, governs, and assures the orchestration layer is a distinct strategic question that must be answered explicitly rather than inherited from existing vendor relationships.

03 Workforce re-skilling

The new model creates a significant and urgent demand for a different kind of worker. Offshore and nearshore delivery talent must transition from task execution to intelligence-assisted supervision. Clear boundaries must be defined for where intelligence will broker value across complex workflows and where human decision-making retains primacy based on trust, confidence, and consequence. This is a workforce redesign, and it belongs in the operating model from the start.

04 Trust, governance, and guardrails

Trust is the precondition for adoption, and adoption is the precondition for value. An agentic operating model that cannot demonstrate trustworthy, auditable behavior across its workflows is destined to stall.

The governance architecture required here is an embedded capability designed to match the speed of the systems it oversees. That means continuous model monitoring, automated audit logging, and input/output content filtering operating as standard components of every agent pipeline. It means strict role-based access control governing what agents can access, execute, and modify. And it means ensuring that every agent interaction leaves a deterministic audit trail capable of satisfying Sarbanes-Oxley (SOX), the General Data Protection Regulation, and applicable industry regulations. Common failure modes, such as hallucination, data leakage, or compliance breach, are machine-speed incidents that propagate before human review can contain them. Real governance intercepts them before they become incidents at all.

KPMG

© 2026 KPMG LLP, a Delaware limited liability partnership, and its subsidiaries are part of the KPMG global organization of independent member firms affiliated with KPMG International Limited, a private English company limited by guarantee. All rights reserved. USCS044468-1A

What to test and measure

Validating that the new model is delivering against the desired outcomes requires a disciplined measurement architecture built before deployment.

In the immediate term, the right questions are operational:

- Are agent pipelines performing within defined behavioral boundaries?
- Are exception escalation rates declining as agent confidence matures?
- Are consumption costs tracking within planned thresholds, or are variable costs spiking in ways that threaten the economic sustainability of the model?
- Are the humans operating in shared oversight and exception management roles making better decisions, faster and with greater confidence, than the model they replaced?

In the mid-term, the questions shift to outcomes:

- Is the commercial arrangement delivering the efficiency step-downs that were contractually committed?
- Is the workforce transition producing the domain validation and exception management capability the model depends on?
- Is the governance architecture surfacing the right signals to senior leaders without requiring them to navigate the operational complexity beneath?
- Is the orchestration layer genuinely collapsing the workflow friction it was designed to eliminate (or has it introduced new integration dependencies that replicate the problem in a different form)?

The organizations that will lead the next decade are those that treat these measurements as a continuous feedback loop. This loop should drive ongoing calibration of where agents operate autonomously, where humans retain final judgment, and where the model itself needs to be redesigned. That ongoing calibration is precisely the kind of work that a next-generation MS partner is positioned to own. It is also the work that separates a managed services relationship built for this moment from one that is merely dressed to look like it.

KPMG

© 2026 KPMG LLP, a Delaware limited liability partnership, and its subsidiaries are part of the KPMG global organization of independent member firms affiliated with KPMG International Limited, a private English company limited by guarantee. All rights reserved. USCS044468-1A

A phased roadmap to reimagining the MSP role

03

For tech companies, operating model pressure is arriving at a pace that rewards early movers and penalizes those who wait for the landscape to settle. The phases that follow are a sequenced, time-bound roadmap that organizations can act on immediately. Each phase builds the foundation for the next, and the milestones within each are designed to be testable, not aspirational.

Phase 1
Diagnostic and contract decoupling (months 0–3)

The first move is conducting an honest if perhaps uncomfortable audit of what the current managed services arrangement is delivering and what it is contractually preventing.

Begin by mapping existing functional MS contracts against a single question: Are these agreements structured to capture the value that AI-driven productivity improvement will create, or are they structured to protect FTE commitments and rate-card revenues? The answer, in most cases, will be the latter. Identifying the contracts, clauses, and renewal timelines is the diagnostic output that makes everything that follows possible.

In parallel, identify the high-volume, high-SOP workflows within each function. These are the processes with the clearest decision logic, the most repetitive execution patterns, and the highest potential for agentic displacement. Ranking them by volume, reversibility, and blast radius gives the organization its pilot candidate list for Phase 2. It's also the negotiating rationale for Phase 1 contract conversations.

The deliverable from this phase is a set of outcome metrics and AI-infused contractual clauses, ready for vendor renewal discussions, that reframe the commercial relationship around business results rather than labor inputs. Organizations that complete this phase have changed their negotiating position permanently. Those that skip it will be locked into an operating model they will eventually outgrow.

Phase 2
Agentic lighthouse pilots (months 3–6)

The second move is testing the highest-friction areas where the evidence of value will be undeniable.

Lighthouse pilots are production deployments of multiagent systems running alongside legacy teams, in real workflows that generate real performance data. The target areas are those where experience has been most consistently poor and the workflow logic is well-documented: Think IT incident autoremediation, journal entry verification, or Human Resources tier-1 triage. These are environments where agent accuracy can be validated against known ground truth, where human-in-the-loop guardrails can be calibrated in practice rather than in theory, and where the gap between the legacy and agentic delivery models becomes visible.

Three questions must be answered before a lighthouse pilot is considered successful. First: where does the experience need to reside in the new operating model and is it better than what it replaced? Next, who owns the change and will take responsibility for adoption at the workflow level? Finally, does the governance architecture hold under real operating conditions?

The output is a proven agent architecture, a calibrated autonomy model, and a governance framework that the organization can scale with confidence in Phase 3.

KPMG

© 2026 KPMG LLP, a Delaware limited liability partnership, and its subsidiaries are part of the KPMG global organization of independent member firms affiliated with KPMG International Limited, a private English company limited by guarantee. All rights reserved. USCS044468-1A

Phase 3
Delivery model pivot

(months 6–12)

Phase 3 is where the operating model structurally breaks. The legacy delivery tiers that agent pipelines have now demonstrably outperformed are decommissioned. L1 and L2 offshore tiers, built for human-executed transactional work, give way to scaled agent orchestration across enterprise platforms.

This is the phase that requires the most organizational courage, because it is the phase where the commitments made in Phases 1 and 2 become irreversible. The retained delivery workforce is formally transitioned into oversight and continuous workflow refinement roles.

The operational infrastructure for this phase centers on collaboration control towers: unified visibility layers across enterprise subsystems that surface agent performance, exception queues, and outcome metrics aligned to annual business goals. This human decision-making architecture allows senior leaders to govern a machine-speed operating environment without being consumed by it. Functional owners can see what their agent pipelines are doing, where they are degrading, and what decisions require their judgment, without navigating the complexity beneath.

The measure of success in this phase is whether the delivery model is now organized around outcomes rather than activities—and whether the governance infrastructure can validate that claim continuously.

Phase 4
Ecosystem at scale

(months 12–18+)

The final phase is a self-optimizing agentic delivery ecosystem that continuously improves against the outcomes it is accountable for, across the entire supplier and internal service portfolio.

At this stage, the shift is from managing agent pipelines to governing an ecosystem. That means enterprise-wide agent performance benchmarks that define what good looks like across the entire interconnected agent environment, including speed of change, productivity, and cost of service.

Organizations that reach this phase have replaced the broken MS model with something that compounds. Every agent interaction generates performance data. Every performance data point improves calibration. Every calibration improvement raises the baseline against which the next wave of automation is measured. The MS relationship becomes a strategic partnership whose value grows with the maturity of the ecosystem it is building.

A note on sequencing

These phases are designed to be sequential, but the boundaries between them are not rigid. Diagnostic work begun in Phase 1 will surface findings that reshape Phase 2 pilot selection. Governance frameworks validated in Phase 2 will inform Phase 3 contract activation. The roadmap is a discipline rather than a timeline.

KPMG

© 2026 KPMG LLP, a Delaware limited liability partnership, and its subsidiaries are part of the KPMG global organization of independent member firms affiliated with KPMG International Limited, a private English company limited by guarantee. All rights reserved. USCS044468-1A

The one thing worth building yourself

The innate desire to build has produced some of the most sophisticated proprietary systems in the technology world. It is ultimately the right instinct, but today, the question of what to build comes down to what versus whether.

For decades, the answer included everything: the finance systems, the fulfillment engines, the operational tooling, and the MS arrangements layered on top of them. Building it all was a source of pride, and in many cases, a source of genuine competitive advantage.

The agentic shift changes the calculus. The organizations that will lead the next decade are those that build the right things. The orchestration complexity, governance burden, and assurance infrastructure of the agentic enterprise are not problems that reward proprietary solutions.

They reward specialization, maturity, and the kind of purpose-built capability that only comes from an MS partner whose entire architecture was designed for this moment.

What is worth building, and protecting and investing in, is the capability that only your organization can create: the proprietary workflows, the bespoke business logic, the agent interaction models that reflect your specific competitive advantage. That is where the builder's instinct belongs.

Everything else? Leave it to your MSP.

KPMG

© 2026 KPMG LLP, a Delaware limited liability partnership, and its subsidiaries are part of the KPMG global organization of independent member firms affiliated with KPMG International Limited, a private English company limited by guarantee. All rights reserved. USCS044468-1A

How KPMG can help

About KPMG Managed Services

Business transformation is the path to sustained advantage. But transformation is not a fixed destination; it's an ongoing journey. How can you continually evolve your business functions to keep up with ever-changing targets?

KPMG Managed Services can help — by handling knowledge-intensive processes across your enterprise on a subscription, as-a-service basis. This outcome-based approach has the potential to reduce your total cost of operations by as much as 15 to 45 percent, in addition to driving priorities like resilience, customer and employee retention, and stakeholder trust.

Taking you beyond traditional managed services, KPMG firms combine advanced technology with functional, process and sector expertise — plus smart analytics, data governance, change management, and alliances with software providers — to operationalize your growth ambition. It's this strategic collaboration that makes the difference between transactional and transformational services, helping you gain competitive advantage, sustain it through ongoing transformation, and limit disruption and risk.

© 2026 KPMG LLP, a Delaware limited liability partnership, and its subsidiaries are part of the KPMG global organization of independent member firms affiliated with KPMG International Limited, a private English company limited by guarantee. All rights reserved. USCS044468-1A

Authors



Vijay Subramanyam
Principal
Advisory Technology, Media, and Telecommunications (TMT) Consulting Industry Leader
KPMG US

Vijay leads the firm's consulting strategy and delivery for clients across the TMT sector, helping them navigate disruption, accelerate digital transformation, and unlock sustainable growth. With more than two decades of experience in management consulting, Vijay brings extensive knowledge in enterprise transformation, business model transitions, technology strategy and modernization, risk and compliance, cyber security and managed services. He works closely with C-suite leaders to drive innovation through AI-powered solutions, data-driven decision-making, and agile operating models tailored to the unique challenges of the TMT industry.

Vijay is also the OpenAI Platform Leader for KPMG in the United States. He is responsible for helping clients realize the value of AI-Native enterprise transformations using OpenAI platforms.



Anuj Mathur
Principal
TMT-Managed Services Lead
KPMG US

Anuj leads the firms TMT-Managed Services offering and has created sustainable business impact for Global Fortune 500 clients by designing and executing finance transformation strategies leveraging Gen AI technologies, process re-engineering, operational excellence, and leading outsourcing/ shared services operations. With a global mindset and cross-cultural experience, Anuj has the ability to build senior client leadership connect and deliver creative financial solutions. Anuj brings deep finance functional knowledge, transformation and project management experience to any organizational change initiative. He is a certified Lean Six-Sigma Black Belt and brings a unique blend of quality, change management, and operations expertise to the transformation cycle.



Aditya Rath
Principal
TMT Line of Business
KPMG US

Adi leads our customer and operations offering focused on revenue and growth agenda for TMT clients. He has worked across multiple geographies bringing in different dimensions of growth, challenges which are local and constantly engages with the revenue and growth leaders to transform sales, service and marketing functions. In an ever-evolving AI world, he doubled downs as a forward deployed growth partner of his clients focused on identifying disproportionate growth lever which can be unlocked using AI, reduce channel friction while transforming customer journeys and drive returns on AI investments.

© 2026 KPMG LLP, a Delaware limited liability partnership, and its subsidiaries are part of the KPMG global organization of independent member firms affiliated with KPMG International Limited, a private English company limited by guarantee. All rights reserved. USCS044468-1A

Contributors

The authors thank Lisa Bigelow, Peter Centofante, Emah Madegwa, Sarah McEwen, Stacie Rich and Kerry Workman for their contributions to this paper.

Related TL



[The future of TMT managed services: From AI ambition to secure reality]



[KPMG Quarterly AI Pulse Survey: Technology]



[Clients rank KPMG #1 for overall quality of work in TMT]

Some or all of the services described herein may not be permissible for KPMG audit clients and their affiliates or related entities.

Learn about us: **kpmg.com**

The information contained herein is of a general nature and is not intended to address the circumstances of any particular individual or entity. Although we endeavor to provide accurate and timely information, there can be no guarantee that such information is accurate as of the date it is received or that it will continue to be accurate in the future. No one should act upon such information without appropriate professional advice after a thorough examination of the particular situation.

© 2026 KPMG LLP, a Delaware limited liability partnership, and its subsidiaries are part of the KPMG global organization of independent member firms affiliated with KPMG International Limited, a private English company limited by guarantee. All rights reserved. The KPMG name and logo are trademarks used under license by the independent member firms of the KPMG global organization. USCS044468-1A