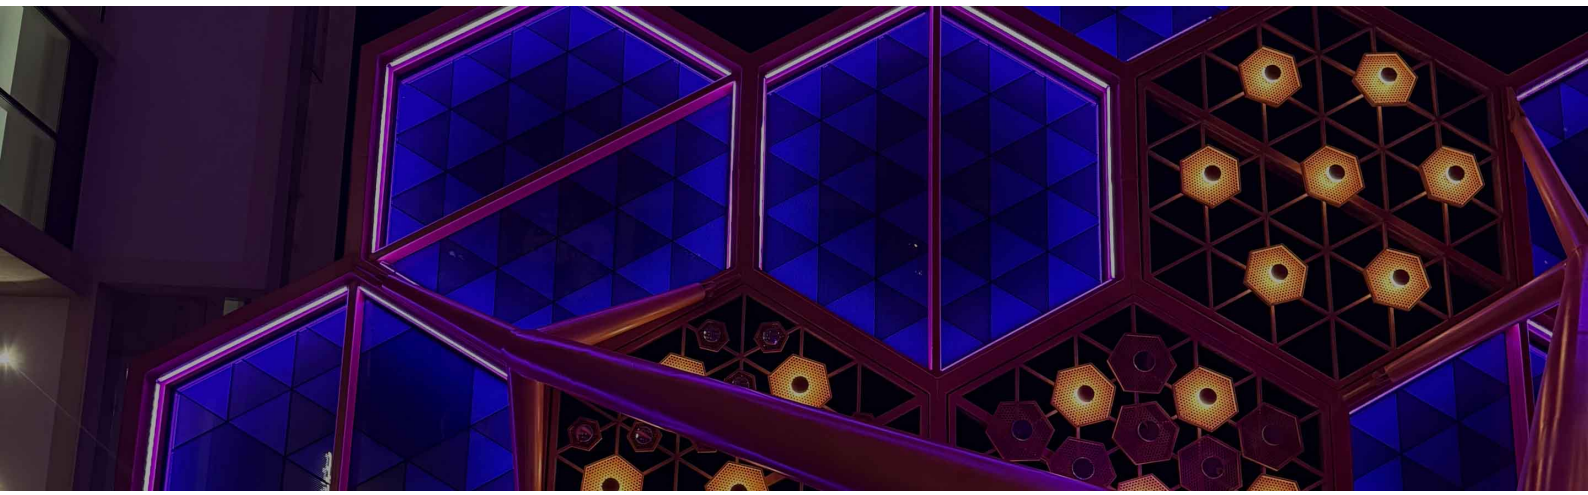

Kearney AI Trends Report 2026

Digital and Analytics / Article
February 27, 2026

Kearney works alongside leading technology innovators and three-quarters of the Fortune Global 500, which gives us a unique vantage point to see firsthand how AI is rewiring businesses. This report presents the AI trends alliance partners and consulting leaders believe will shape 2026.

Executive summary: the evolution from experimentation to enterprise intelligence

After years of pilots and productivity experiments, a decisive gap has emerged between AI ambition and execution. As per S&P Global Market Intelligence, 46 percent of AI proof-of-concepts were scrapped before deployment last year. The organizations breaking through this barrier share a common approach: building integrated decision fabrics that continuously sense, reason, and act across the entire value chain. Competitive advantage stems from how intelligently organizations orchestrate human expertise, proprietary data, and autonomous systems to deliver sustained business impact.

In 2026, a fundamental shift is under way in how enterprises approach AI. Three converging forces are driving this evolution:

Enterprise-grade agentic platforms have matured. AI systems can now orchestrate multi-step workflows, make contextual decisions, and operate semi-autonomously within governed boundaries. This moves AI from simple task automation to sophisticated process transformation.

The economics of AI have fundamentally changed. As access to models democratizes, differentiation shifts to data quality, domain expertise, and the ability to operationalize insights at scale.

Trust and governance have become table stakes. Regulatory scrutiny, operational risk, and ethical considerations demand transparent, auditable AI systems. Organizations must architect for observability, control, and human oversight from inception.

Organizations that successfully embed governed, transparent AI into core workflows will unlock unprecedented gains in growth, efficiency, and resilience. Defining leading metrics that demonstrate progress, clear baselining, and ongoing tracking of outcomes are crucial for every organization aspiring to successfully adopt AI.

Agentic AI rewires the enterprise operating model

The agentic AI market is experiencing explosive growth, projected to reach $10.41 billion in 2025. By 2030, the market is expected to reach $52.6 billion with 45 percent annual growth. This momentum reflects urgent enterprise needs: organizations face mounting pressure to automate complex workflows, reduce operational costs, and scale operations without proportional headcount increases.

Yet despite widespread interest, full deployment remains stagnant at only 11 percent as organizations grapple with integration complexity, security requirements, and infrastructure readiness. The real barrier isn't technology maturity, it's architectural thinking. Leading enterprises are moving beyond isolated pilots toward integrated decision fabrics, a mesh of governed data, fit-for-purpose models, and agentic workflows that continuously sense, reason, and act across revenue, cost, risk, and experience.

Value at scale requires five foundational pillars working in concert:

1. Data plane: high-quality, lineage-tracked data products spanning structured and unstructured content

2. Model plane: a governed portfolio of large language models (LLMs) and traditional machine learning (ML) with retrieval (RAG) over enterprise knowledge and human-in-the-loop for critical decisions

3. Orchestration: agentic patterns (plan # act # critique # learn) that automate multi-step work while remaining auditable

4. LLM ops: evaluation, telemetry, versioning, and cost governance to make GenAI repeatable and economical

5. AI office: an operating model that standardizes intake, prioritization, delivery, and measurement across business units

In 2026, AI becomes the enterprise's decision fabric—standardized, governed, and value#tracked—turning everyday work into continuously optimized, auditable workflows that improve growth, margins, and customer trust.

Brent Smolinski, Kearney partner, observes: "As enterprise-grade agentic AI platforms mature, we're seeing work itself being rewired. Core processes are shifting from linear, role-based workflows to dynamic, event-driven systems where AI agents handle routine decisions and humans step in for judgment, orchestration, and cross-functional problem solving. The real opportunity isn't just efficiency gains, but flatter organizations, more 'liquid' talent models, and AI that continuously upskills every employee as they work. Agentic platforms will also act as continuous learning companions, raising the pace at which employees acquire new skills. The net effect is a fundamental rewiring of how work gets done and how organizations are designed."

Guardrails first: trust through transparency and governance

As AI agents assume greater autonomy in enterprise decision-making, a critical question emerges: how do organizations maintain visibility into what these systems are doing and why? Enterprise trust in AI will increasingly be earned through evidence and control. As agents take on real decisions, organizations will require proof-of-work artifacts: human- and machine-readable records of reasoning, data, and actions to keep AI traceable and governable at scale.

At the same time, governance is becoming an execution layer that dynamically prioritizes and retires AI initiatives based on live ROI, risk, and adoption signals.

Proof of work becomes the foundation for trustworthy AI

For most of computing history, the division of labor between humans and machines was clean.

Machines executed structured systems: explicit, deterministic logic expressed through code, workflows, and rules. Humans supplied the unstructured capabilities: judgment, creativity, context, and meaning-making. The boundary was bright and stable.

Equally important: historically, when a human produced an unstructured artifact (for example, a narrative memo, a PowerPoint deck, or a strategy brief) that was the end of the road for computation. The machine

could store or format the file, but it could not interpret it. By contrast, when humans created structured artifacts like spreadsheets or databases, both people and machines could continue to read, manipulate, and verify them.

The emergence of large language models capable of interpreting and generating semi-structured content collapses this old separation. LLMs can now parse and produce artifacts that sit between structure and prose. But more importantly, tool calls let them invoke operations on those artifacts (for instance, querying data, executing analyses, triggering actions). This allows AI to engage with work products that were previously legible only to humans. Semi-structured systems create a shared representational layer that both sides can work with and act upon.

Software development is the first domain where these lines are blurred. The reason Copilot, Cursor, and other coding agents gained traction is simple: the output is code, both computers and programmers know how to interpret it. When an AI writes code, humans and machines can both evaluate what happened and why.

This shift illustrates what is often called proof of work. Coined by Fraser Kelton at Spark Capital, the term refers to the human- and agent-interpretable, verifiable representation of what an AI system did and why.

This pattern is expanding quickly. As AI participates in research, planning, decision support, and strategic exploration, organizations will require proof-of-work artifacts that keep reasoning legible across teams. At Growth Signals, innovation teams develop new concepts confidently because the artifact itself becomes a shared workspace. Teams can challenge the AI's analysis, add new sources, refine the logic and the AI can interpret those changes and build on them. A canvas that both humans and machines can read, evaluate, and modify together becomes something neither could produce alone. Proof of work turns opaque generation into collaborative cognition between human judgment and machine capability.

By the end of 2026, enterprises that embrace these shifts will dramatically increase the productivity of knowledge work through more transparent, collaborative, and adaptive AI systems.

Responsible AI becomes the governance fabric

servicenow

As organizations accelerate their adoption of AI, governance has become the cornerstone of enterprise trust and sustainable ROI. Without clear oversight, transparent operations, and continuous monitoring, AI deployments can introduce operational, compliance, and ethical risks that erode confidence and limit scalability. The imperative is clear: strong governance, observability, and life cycle control are essential for transforming AI from scattered experimentation into reliable, high-value capability.

Responsible AI requires unified visibility across all AI assets—agents, models, prompts, datasets, and workflows—regardless of where they operate (such as ServiceNow's AI Control Tower). Organizations need end-to-end observability into where AI is running, how it behaves, and its impact on productivity, revenue, customer experience, and risk. This requires a single source of truth that tracks and manages not only internal AI deployments but also distributed solutions across platforms, cloud environments, and vendor ecosystems.

By integrating AI governance with existing business service management and embedding risk and compliance workflows, leading organizations are creating authoritative systems of record for compliant, transparent AI operations. Policies, accountability, and business alignment extend consistently across both internal and external AI environments.

For an AI center of excellence (CoE), this control layer is indispensable. It consolidates strategy, governance, and execution so that policies, risk controls, life cycle processes, and guardrails can be defined once and applied universally—across IT, HR, customer service, and other business domains, as well as across external AI providers. This empowers the CoE to evolve from an advisory function into an operational intelligence hub that continuously tracks ROI, monitors model drift and agent performance, and quickly mitigates emerging compliance or ethical risks regardless of where AI workloads reside.

Effective governance also enables organizations to scale AI responsibly. By orchestrating agentic operations and ensuring alignment with enterprise priorities, teams can eliminate redundant or unsanctioned "shadow AI," prioritize high-value initiatives, and coordinate complex interactions across multiple agents. With comprehensive visibility, governance, life cycle management, and orchestration spanning enterprise-wide AI systems, organizations can operate mature, data-driven AI operating models that deliver trustworthy, measurable, and scalable outcomes.

Open standards power the agentic ecosystem

The standardization of agentic frameworks is fundamentally reshaping how enterprises approach AI adoption in 2026. The convergence on open, interoperable protocols represents a critical inflection point, driven by urgent market demands for interoperability, scalability, and reduced integration complexity.

Without interoperable protocols, each agent-to-tool or agent-to-agent interaction requires bespoke integration, creating scalability bottlenecks that delay maturation and increase opportunity costs. Open standards such as Model Context Protocol (MCP) eliminate this integration overhead, enabling organizations to focus on business value rather than technical plumbing. Early adopters are already experiencing measurable benefits:

- Sixty to 80 percent reduction in manual exception handling work for enterprise resource planning (ERP) systems
- Twenty to 60 percent productivity gains already seen with AI agents across industries
- Thirty percent reduction in operational costs projected by 2029 as agentic AI autonomously resolves 80 percent of common customer service issues
- Fifteen percent of work decisions expected to be automated by agentic AI by 2028, up from 0 percent in 2024

In 2026, architecting agents is less about a single smart model and more about leveraging an open, interoperable ecosystem. Standards like MCP ensure consistent communication. Agent-to-agent protocols enable collaboration between autonomous systems. Orchestration frameworks coordinate workflows. Enterprise-grade platforms provide secure, scalable infrastructure.

Business outcomes validate this approach: Lyft achieved an 87 percent reduction in average resolution time for customer and driver support requests. Thomson Reuters modernized 1.5 million lines of code monthly—a 4x velocity increase—while cutting costs by 30 percent.

Open standards create the foundation for an interoperable ecosystem where hyperscalers such as AWS provide secure infrastructure and foundational services; ISVs embed agentic capabilities into vertical solutions; and system integrators deliver implementation helping enterprises navigate organizational complexity while ensuring governance, compliance, and trust. Together and empowered by open standards, these partners create a flywheel effect with faster innovation, scaled adoption, and more sophisticated multi-agent systems that tackle complex business challenges.

As we progress through 2026, open standards are removing the primary barrier to agentic AI scaling: integration complexity. Organizations can now confidently invest in agent development, knowing their solutions will interoperate across vendors and evolve with emerging capabilities.

Nathan Bell, Kearney partner: "As AI agents take on high-stakes decisions, oversight is shifting from step-by-step intervention to exception-based supervision. Governance is no longer a gate—it's becoming an active execution layer that continuously evaluates which AI initiatives to scale, pause, or retire based on real-time performance signals. Retaining the 'human in the loop' of AI governance will remain crucial to ensure the 'predefined' guardrails of AI initiatives are maintained and adhered to."

Human–AI collaboration: augmentation over automation

From rethinking the nature of agentic AI to reimagining supply chain planning, the trend is clear: augmentation, not replacement, drives superior outcomes. Organizations that design AI systems as collaborative partners, not autonomous replacements, unlock superior outcomes while maintaining the judgment and contextual understanding that humans provide.

From deterministic automation to probabilistic intelligence

mindcorp.ai

Prior to GenAI, AI systems were largely deterministic: given the same input, they would produce the same output every time. Expert systems excelled at predefined tasks, rules-based workflows, and predictable outcomes. Agentic AI, by contrast, is based on GenAI and is more non-deterministic. It operates probabilistically, along with using rules, meaning it evaluates multiple possible actions and selects among them based on probabilities, context, rules, and goals.

This probabilistic approach enables agentic AI to be more adaptive, creative, and context-aware. Even when presented with the same prompt, it may generate different responses depending on situational factors, prior steps taken, or evolving objectives. This flexibility allows it to handle ambiguity and complexity in ways traditional automation cannot, but it also introduces new challenges and requires new approaches to building systems that actually behave reliably.

Another defining characteristic of agentic AI is its ability to move beyond passive response generation. These systems can reason through problems, plan sequences of actions, and execute tasks across tools or environments. Instead of simply answering a question, an agent can decide what needs to be done, in what order, and how to do it, and can even use tools or work with other agents in the process. This is great for automation of low-level tasks, but for higher-level knowledge work, agentic AI does not replace human decision-making. Its strength lies in augmenting human capability—accelerating analysis, exploring options, and handling intermediate steps that would otherwise consume time and attention.

A useful mental model for agentic AI is that of an "intelligent intern." Like a strong intern, an AI agent can add meaningful value: it can research, draft, analyze, test ideas, and take initiative within defined boundaries. It brings reasoning and action together in a way that feels proactive rather than reactive. Yet it still requires guidance, oversight, and clear expectations—direction on goals, constraints, quality standards, and when to escalate decisions.

As agentic AI becomes more prevalent, organizations will need to rethink how work is designed. The focus will shift from automating individual tasks to orchestrating human–AI collaboration. Success will depend less on full autonomy and more on effective supervision, clear instructions, and well-defined roles between humans and agents.

To realize full value from AI capabilities, organizations must design a progressive journey that moves agents toward becoming trusted team members.

The adaptive supply chain: AI agents redefine supply chain planning and empower, not replace, planners

KINAXIS®

AI in the supply chain is shifting from automation to augmentation, replacing chaos, not planners. Kearney partner P.S. Subramaniam notes that AI agents are redefining planning not by replacing planners, but by removing the noise and complexity that hold them back. Embedded within a concurrent planning environment, AI agents continuously sense change, explore scenarios, and surface trade-offs, while operating as transparent, governed systems that elevate human judgment rather than obscure it.

The future belongs to adaptive supply chains built on true human–AI partnership. Deep planning expertise—such as that of a strong MPS planner—is increasingly scarce, arguably rarer than advanced AI skills. When that expertise is amplified with AI intelligence and automation, planners can move beyond firefighting to orchestrating higher-value decisions across service, cost, and resilience, fundamentally changing the game.

The implications impact supply chains and beyond. Hagen Goetz Hastenteufel, Kearney partner, observes: *"Many companies are quietly restructuring 20 to 40 percent of support and management roles. Not to cut costs, but to increase performance. As AI starts to manage coordination and reporting, the question becomes: what must humans still do, and how do you design around it?"*

The value chain reimagined: domain-specific AI applications

As AI matures beyond horizontal productivity tools, competitive advantage increasingly comes from embedding intelligence into specific value chain functions. Organizations are deploying specialized AI agents that combine proprietary data, domain expertise, and contextual understanding to transform how work gets done in procurement, manufacturing, and logistics.

Judgment and data: the next frontier of continuous competitiveness

As AI is adopted universally, the benefits accrued to each enterprise could be uniformly distributed, leading to a new normal and a competitive stalemate—unless enterprises deliberately determine how they can set themselves apart through resetting their ambitions, the quality of the data they use, and the expertise and operating models they apply to interpret it.

The most effective organizations treat AI as a complement to human understanding, not a substitute for it. Agents can flag patterns and opportunities faster than any team could but knowing how to act on them still requires the experience and perspective of people who understand procurement, operations, and the business context. Further, it is clear that humans and human-led teams are still superior in seeing patterns and opportunities where there is no pre-existing template to learn from.

Vel Dhinagaravel, CEO of Beroe, explains: "Technology gives us the reach to see everything, but not the judgment to know what matters. The organizations that pull ahead are the ones that can filter the noise, recognize opportunity, and act before others do. Many companies still set thresholds for when to act—for example, only pursuing sourcing opportunities above $500,000 or cost savings above $50,000. With the right data and insight, that limit disappears. The ability to capture hundreds of small, overlooked opportunities every day is where continuous competitiveness becomes real."

Traditional planning horizons will be dramatically disrupted, and we should expect to see companies moving to dynamic budgets which provide them with the ability to seize opportunities and pivot on a dime.

Data quality will define the next wave of competitive advantage. Publicly available datasets are easy to access, which means they no longer offer meaningful differentiation. Companies that rely only on what everyone else can see will move at the same speed as their competitors. Curated, proprietary data that reflects the nuances of specific markets and supply networks will become a key advantage for organizations that want to outpace the market, because an AI agent can only be as good as the data it has access to.

Suketu Gandhi, co-leader of global strategic operations at Kearney, adds: "Continuous competitiveness is becoming the measure of how prepared an organization is for the future. The ability to stay ahead comes from constant awareness and the discipline to act on what matters most. Progressive procurement teams are moving beyond traditional KPIs to a new model of always-on optimization. The results are significant where function-level ROI across procurement, planning, and supply chain now often exceeds $100 million, a benchmark that shows what continuous improvement at scale can deliver."

Continuous competitiveness depends on the combination of human insight, domain expertise, contextualized data, and advanced technology working together. The future belongs to organizations that see AI not as a replacement for expertise but as a force multiplier for it and a way to unite intelligence, experience, and action into one continuous source of advantage.

Beroe's curated intelligence on markets and suppliers supports Kearney's procurement and supply chain engagements, helping teams make faster and more confident procurement decisions.

Physical AI moves from experiment to infrastructure

Physical AI marks a paradigm shift, moving beyond virtual cognition to embodied intelligence that perceives, reasons, and acts in the real world. Unlike traditional AI confined to digital boundaries, physical AI fuses advanced robotics, sensor networks, and generative intelligence to create adaptive systems capable of autonomous decision-making in dynamic physical environments.

This evolution is already transforming mobility, manufacturing, and industrial ecosystems. Autonomous mobility platforms, predictive factories, and GenAI-enabled production systems demonstrate how intelligence engineered at scale delivers tangible business impact beyond prototypes. By leveraging digital twins, edge intelligence, and multi-modal AI architectures, enterprises achieve real-time responsiveness, operational resilience, and sustainability at unprecedented levels.

The market validates this shift. Physical AI is projected to grow from $371.7 billion in 2025 to $2.4 trillion by 2032, driven by human–machine partnerships and edge computing frameworks. For technology service providers, this represents a $300 billion opportunity to orchestrate AI-native ecosystems through robotics-as-a-service (RaaS) and digital twin-as-a-service (DTaaS).

Yet capitalizing on this opportunity requires more than technology deployment—it demands reinvention of enterprise operating models. Three imperatives for leaders:

1. Industrialize AI at scale. Integrate physical AI into core workflows using modular platforms, simulation environments, and governed data fabrics. This ensures safety, compliance, and interoperability across global value chains.

2. Design for human–machine collaboration. Physical AI augments rather than replaces human expertise. Embedding AI into robotics and assistive systems unlocks productivity while maintaining ethical and transparent practices—a principle consistent across all successful AI deployments.

3. Build innovation ecosystems. Co-innovation hubs like TCS Paceports and Agile Innovation Centers enable rapid prototyping and collaborative ideation, accelerating time-to-market while de-risking complex implementations.

Roland Scharrer, Kearney partner, says: "Physical AI will shift from isolated robotics and automation pilots into core enterprise infrastructure, where intelligence is embedded directly into assets, factories, fleets, and supply networks. The real inflection point is not autonomy alone but closed-loop reinforced and accelerated learning, linking perception, simulation, and action through digital twins, edge intelligence,

and governed data flows that operate in real time. Leaders need to increasingly treat physical AI as an operating model change, not a technology layer, and prioritize redesign of processes around outcomes like uptime, safety, and sustainability. The advantage will accrue to organizations that pair embodied intelligence with strong governance and ecosystem partnerships, enabling human–machine collaboration at scale while turning physical operations into continuously learning systems."

Autonomous agent for logistics spend: from collaboration to decisioning

Logistics and supply chain leaders have largely invested in visibility, workflow automation, and analytics, but they remain hampered by siloed systems, manual workflows, cumbersome third-party onboarding, and decision latency. With global complexity rising (multi-modal carriers, geopolitical shocks, sustainability mandates), these firms are still operating at the "co-pilot" layer: humans interpret dashboards, build connections, then act.

The future agenda demands embedding business-specific, autonomous agents into the logistics stack—decision-making systems that operate as employees, not just tools. To quickly solve these challenges, businesses need to avoid investing in large-scale tech implementations, instead deploying quick, simple agents using a framework of existing agentic workflows. A self-serve agent studio platform allows enterprises to spin up domain-specific agents (for example, control towers, supplier collaboration, global trade exception management) with enterprise data, governance, and monitoring.

These agents ingest multi-structured data (invoices, contracts, rate-cards, shipment flows, PO/STO data), interface across carrier and supplier networks without the need for heavy EDI/API integrations, benchmark in real time (contracted, empaneled, spot/data-provider, network carriers), and execute decisions: reroutes, carrier swaps, procurement pivots, invoice recoveries. In short: connectivity # visibility # action # autonomy. The key implications for leaders:

- Investments must shift from building dashboards to deploying agent workflows that reduce human intervention, shorten decision cycles, and adapt dynamically.
- Providers that offer only a "TMS plus analytics" play will fall behind; firms need systems that act, learn, and continuously optimize.
- ROI metrics evolve, beyond cost-to-serve and on-time delivery to include decision-cycle time, autonomous-execution rate, and learning rate of the agent.
- Organizations must reorient talent and governance: logistics management becomes oversight of agents, carriers become the extended network of the AI agent, and data architecture becomes mission-critical.

Korhan Acar, Kearney partner, said, "Visibility alone is no longer enough. In logistics, strategic thinking is the real unlock, because value rarely resides in the transactional layer, but in how short- to long-term strategy is built and executed. By moving from dashboards to domain-specific agents, Freehand significantly reduces time to value and frees teams to focus on the work that actually creates competitive advantage by integrating quickly and learning continuously."

Conclusion

The trends outlined in this report converge on a singular insight: AI is no longer a technology initiative—it is a fundamental business transformation. Organizations that architect AI as core infrastructure, governed with rigor and deployed with strategic intent, will establish decisive competitive advantages. Those that continue to treat AI as a collection of point solutions will find themselves increasingly outpaced.

Three principles distinguish leaders from laggards:

1. Integration over isolation. The most impactful AI implementations connect across functions, creating decision fabrics that orchestrate insight and action enterprise-wide rather than optimizing narrow processes in silos.

2. Augmentation over automation. Superior outcomes emerge when AI amplifies human judgment—handling complexity and routine while preserving the creativity, context, and strategic thinking that only people provide.

3. Governance as enabler. Trust and transparency aren't constraints on AI adoption; they are prerequisites for scale. Organizations that build observability, control, and ethical frameworks from the start move faster and further than those that bolt on governance as an afterthought.

Maximizing the impact of AI requires business transformation, not just technology adoption. Organizations need to pursue integration strategies where AI becomes the enterprise operating system, rather than a bolt-on solution. The question is no longer whether to invest in AI, but how to boldly reimagine your business around it.

About this report

This report was developed by Kearney in collaboration with our alliance ecosystem that includes leading AI and technology partners. For more information about Kearney's AI capabilities and alliance ecosystem, please contact the authors.

The authors would like to thank our alliance partners for their valuable contributions to this report.

Authors

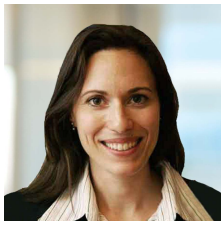

Alanna Klassen Jamjoum
Partner

Shivam Kumar
AI Alliances Lead

About Kearney

For 100 years, Kearney has been a leading management consulting firm and trusted partner to three-quarters of the Fortune Global 500 and governments around the world. With a presence across more than 40 countries, our people make us who we are. We work impact first, tackling your toughest challenges with original thinking and a commitment to making change happen together. By your side, we deliver—value, results, impact.