



The Agentic AI Advantage: Unlocking the next level of AI value

June 2025



At KPMG, the future is agent-centric

AI has reached another inflection point. Just as organizations begin to scale generative AI, a more transformative shift is already underway: the rise of AI agents.

Agents aren't just the next wave—they're a step-change. These systems will act independently, pursue goals, and operate across complex workflows. That shift changes everything: how work gets done, how decisions are made, and how value is delivered.

This paper breaks down what agentic AI really means: how to define it, what it enables, and where to begin. The perspectives are grounded in real-world experience—drawn from the work we're doing with clients and across our own business.

The time to prepare is now. Unlike cloud, which took nearly two decades to reshape enterprise infrastructure, this shift is playing out in months. Leading organizations aren't waiting—they're already investing, scaling, and positioning themselves to capture the benefits of AI agents at speed.

But this isn't just a tech conversation—it's a leadership one. Agents will demand tighter alignment between business and technology leaders than anything that's come before. Success will require a shared vision, coordinated execution, and a clear plan for embedding intelligence at scale.

That includes preparing your people. The organizations that lead will be those that help their workforce adapt—training teams to work with agents, clarifying evolving roles, and leading the cultural shift this transformation demands.

And as agents take on more autonomy and influence, trust becomes even more critical. A scalable strategy isn't just about what agents can do—it's about ensuring they operate reliably, transparently, and within guardrails that build confidence.

With all the hype around agents, it's easy to lose focus. Agents will be a critical part of AI strategy, but they're not the strategy. The leaders pulling ahead are moving fast, but with clarity—staying anchored on the “no regrets” moves: aligning leadership,

modernizing tech and data, upskilling their people, and embedding trust at every layer.

This shift won't be easy—but it will be defining. The decisions leaders make today will reshape industries and redefine competitive advantage. As we continue our journey, we'll keep sharing what we learn—and welcome your ideas and perspective as we collectively navigate what's next.

Peter Van den Spiegel
Partner
KPMG in Belgium

E: pvandenspiegel@kpmg.com
M: +32 (0)477 58 94 74



Nicolas Dubois
Partner
KPMG in Belgium

E: ndubois@kpmg.com
M: +32 (0)473 61 06 03



Contents

04 Do you really know what an AI agent is?

05 The era of agentic AI has already begun

07 The value at stake

08 Four ways AI agents can unlock value

10 The spectrum of agents is evolving to meet critical needs

14 Setting the foundation for your agentic AI journey

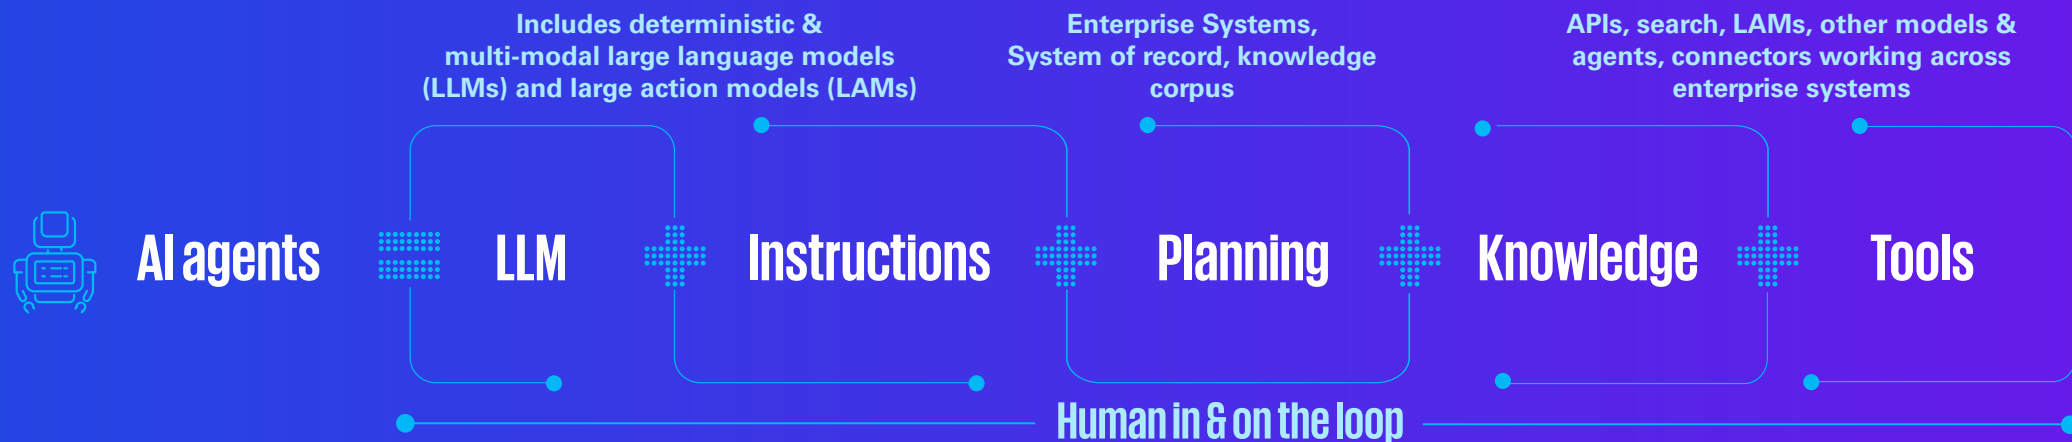
23 Practical next steps to accelerate agent adoption

26 How KPMG can help

Do you really know what an AI agent is?

Use of the term AI agent has exploded recently, and literally everyone is talking about them. It has gotten to the point that almost everything seems to be an agent of some kind. And that's just not true. Before you can begin to realize the incredible potential inherent in this quickly evolving new technology, you really need to understand what an agent actually is.

- AI agents are digital tools that fulfill organizational goals by taking meaningful, independent action. They accomplish this by blending advanced reasoning from LLMs with planning, orchestration, knowledge, data mining, curated tools, and careful governance.
- Agents make real-time decisions, adapt to new situations, and learn from their interactions and feedback.
- AI agents can ingest both structured and unstructured data.
- Potential real-world applications are mind-boggling. They range from basic task-oriented agents that can extract and compare data against standards, to end-to-end automation agents that could disrupt entire organizational value chains.



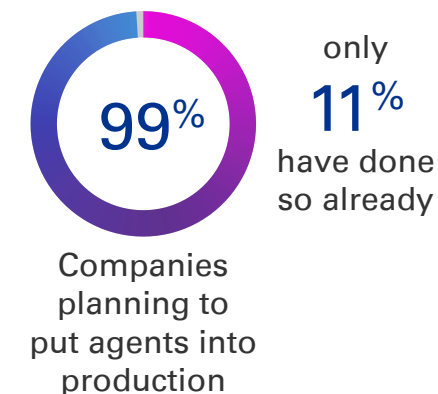
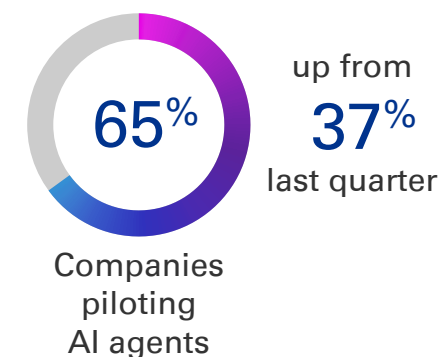
The era of agentic AI has already begun

Just as executives are beginning to scale the game-changing capabilities of generative AI (GenAI) across their organizations, a new and even more powerful AI-technology has emerged. While still early days, agentic AI is already showing such great promise, and evolving so rapidly, that no corporate executive, or board, can ignore its potential to transform their organizations and gain competitive advantage.

In simple terms, GenAI builds “digital assistants” that help employees free up time to be more productive, effective, and creative, i.e., it augments humans by providing information and guidance to help them take better decisions and faster action. In contrast, agentic AI can fully automate what humans do, by taking actions on their behalf, in pursuit of prescribed business outcomes. Agents may loop in GenAI tools and other sources of data and insights, with the potential to deliver significantly more business value.

While still emerging and evolving, AI agents can already take on human tasks that were previously thought too complex to automate, and they may be a better choice for some activities previously automated with older AI tools. Hence, most companies in our latest quarterly pulse survey are already piloting AI agents (65 percent), and the number has almost doubled in three months (up from 37 percent last quarter).¹ So far, however, only 11 percent have put agents into production, although 99 percent plan to do so.²

KPMG Quarterly Pulse Survey



¹ KPMG AI Pulse Survey Q1, April 2025.

² KPMG AI Pulse Survey Q1, April 2025.

Agents' ability to automate complex tasks is rapidly growing, currently doubling every three to seven months. In coding, their capabilities have already grown from handling tasks that take humans seconds to complete to automating hour-long tasks. If this exponential growth continues, agents could be four times more capable in six months, and sixteen times more powerful in a year. This will enable them to move from automating human tasks to automating human roles, creating an entirely new talent pool of "digital co-workers."

Agents are also increasingly capable of communicating and collaborating with each other, further amplifying their business utility. By taking on a series of complex tasks as a group, "swarms"³ of agents with different skills may eventually automate entire business processes end to end. Once automated, agents can autonomously re-engineer, and continuously improve, these processes. This multi-layered approach can unlock new levels of operational performance and open the aperture for reimagining entire business models. At the same time, risks are higher and must be carefully managed to ensure value is unlocked safely and securely.

While these scenarios may sound futuristic today, we predict that most companies will eventually transform into hybrid organizations where both humans and AI agents collaborate seamlessly. The number of digital

\$3 trillion

KPMG estimates that agentic AI will be key to unlocking a staggering \$3T in corporate productivity annually

workers may initially be small but can grow rapidly and flex up and down with demand almost instantly. This will allow human workers to focus their time on the most productive, complex, and stimulating tasks, while learning to work seamlessly with agents in joint pursuit of unprecedented levels of performance.

The size of the prize for those getting it right could be astounding. Projections from IDC indicate that, over the next 10 years, generative AI alone will add nearly \$10 trillion to the global GDP.⁴ The value potential for agentic AI promises to be exponentially higher: KPMG estimates that agentic AI will be key to unlocking a staggering \$3 trillion in corporate productivity improvements annually, based on our proprietary research of the impact of AI on more than 17 million companies.⁵ We also anticipate that agentic AI can help unlock at least five percent of

EBITDA annually in labor productivity alone for the average Fortune 1000 company. The latter finding is echoed by many industry analysts. For example, according to a recent study from IDC, organizations adopting agentic AI can achieve an 18 percent improvement in employee productivity and satisfaction.⁶

This paper focuses on how to navigate the next phase of the AI journey, from a business leader's perspective, starting with the value at stake for those organizations that get it right. We also introduce various types of emerging agents and the KPMG TACO Framework™ for optimizing the use of agentic AI. Finally, we address the challenges ahead, including four key barriers to overcome, concluding with practical actions you can take today to improve your chances of success in the world of agentic AI.

³A "swarm" is a self-synchronized, efficient, effective team of agents. (Abbass Hussein and Mostaghim Sanaz, 2025, The road forward with swarm systems, *Phil. Trans. R. Soc. A.* 38320240135 <http://doi.org/10.1098/rsta.2024.0135>).

⁴Ritu Jyoti and David Schubmehl, The Business Opportunity of AI, IDC, November 2023.

⁵Quantifying the GenAI Opportunity, KPMG, 2025.

⁶Ritu Jyoti and David Schubmehl, The Business Opportunity of AI, IDC, November 2023.

The value at stake

AI is evolving so rapidly in terms of both advanced capabilities and lower costs that any predictions of future value must still be considered speculative. For example, in the last 18 months alone, usage costs for leading AI models with equivalent performance dropped from \$20.00 to \$0.07 per million tokens, a 240-fold improvement.⁷ The exponential pace of development for both generative and agentic AI will likely continue for the foreseeable future, with performance promising to improve at a faster rate than silicon chips did according to Moore's Law.

In terms of market spend on agentic AI, estimates from leading analysts are bullish and growing fast. In addition to the \$50B in market value many analysts predict, returns on agentic AI investments are also promising, according to an IDC study. Their research finds that for every \$1 a company invests in AI, it can realize an average of \$3.50 in return, and 5 percent of organizations worldwide are realizing an average of \$8 in return.⁸

Using our patent-pending GenAI Value Assessment model, KPMG has made a first attempt to quantify the addressable opportunity for agentic AI by industry, company, function, and roles. While it is too early to make reliable predictions, we took a conservative approach that only looks at labor productivity impact. Assuming that agentic AI will be the key to automating the high-complexity tasks in our model, this equates to \$3T in corporate productivity improvements and a 5.4 percent EBITDA improvement for the average company annually.⁹

⁷ AI Index 2025: State of AI in 10 Charts, Stanford Human-Centered Artificial Intelligence (HAI), April 7, 2025.

⁸ Ritu Jyoti and David Schubmehl, The Business Opportunity of AI, IDC, November 2023.

⁹ [Quantifying the GenAI Opportunity](#), KPMG, 2025.

The KPMG GenAI Value Assessment



In 2024, KPMG launched a massive study with the aim of developing the world's largest database to help quantify the potential impact of GenAI on labor productivity for enterprises across the world.



Leveraging our patent pending GenAI Value Assessment model, we built the equivalent of an organizational digital twin for more than 17 million companies using more than 3 billion data points.



Due to the availability of financial information, we zoomed in on the 7,000 public companies in the total sample, which collectively employ more than 72 million workers in 2,000 different roles, with average revenue of \$7.5B.



We graded the complexity of tasks exposed to AI into low, medium or high, using our internal AI models, based, e.g., on the degree of reasoning and interactions required.



Of these, we believe that agentic AI is best positioned to unlock value in high-complexity tasks.

Four ways AI agents can unlock value

AI agents can unlock greater enterprise value in at least four significant ways:



Agents widen the aperture for automation.

As agents become exponentially more capable in their reasoning and ability to collaborate, they can automate tasks, roles, and processes that could previously only be performed by humans. This expands the share of work that can be freed up for human talent, but also the effectiveness of the work over time. Specifically, after automation, agents can be instructed to continuously strive to optimize their interactions toward a specified goal.



Agents don't sleep.

Once agents have automated a set of tasks, they can be instructed to run these around the clock every day, as long as there is energy, data, and human supervision available. Just by doing the same work for 24 hours, instead of 8, an organization can improve daily productivity by three times. When multiple agents are working simultaneously, they can also accomplish tasks faster than a human team could. If the agentic team can do the same work, say in 1/3 of the time, the organization can deliver nine times more work per 24-hour day. This can also make the organization more agile, by responding more rapidly to trends that occur during off hours. For example, in manufacturing, agents can manage production lines by optimizing schedules and workflows and perform predictive maintenance to reduce downtime. When supply disruptions occur, procurement agents find alternatives, while production agents reconfigure schedules automatically.



Agents are wired for change.

Once a process has been automated by a swarm of agents, it also becomes easier and faster to adopt this workflow to ever-changing circumstances over time. The agents are trained and motivated to deliver specific outcomes but don't need detailed instructions on how to proceed or to take time away from delivery to get trained. In addition, they are not threatened by change, nor are they reluctant to collaborate, hence they require less change management. The agentic workflow can iterate and accommodate experimentation faster to get to an optimal outcome within the guardrails of their instructions. This can be a critical advantage in, for example, financial services, where month-end close agents identify anomalies and prepare journal entries, reducing close cycles. At the same time, vouching agents automatically cross-reference invoices against contracts, purchase orders, and receipts, flagging discrepancies without manual review. If reporting policies, organizational structure, or lists of customers and suppliers change, the agentic system automatically adjusts the end-to-end process.



Agents convert knowledge into action.

Although GenAI can significantly enhance productivity, the potential value from agentic AI can be much greater due to its ability to convert knowledge into coordinated action. As organizations deploy multiple agents, they need to feed them the right knowledge to drive the most effective outcomes. This requires systematically capturing and structuring the tacit expertise that typically resides only in employees' minds, making it available for the agents' use. For example, in retail store performance audits, audit agents continuously monitor store operations data, customer feedback, and store sales, flagging compliance issues and performance anomalies across locations. This occurs while agents automatically generate corrective action recommendations for regional managers.

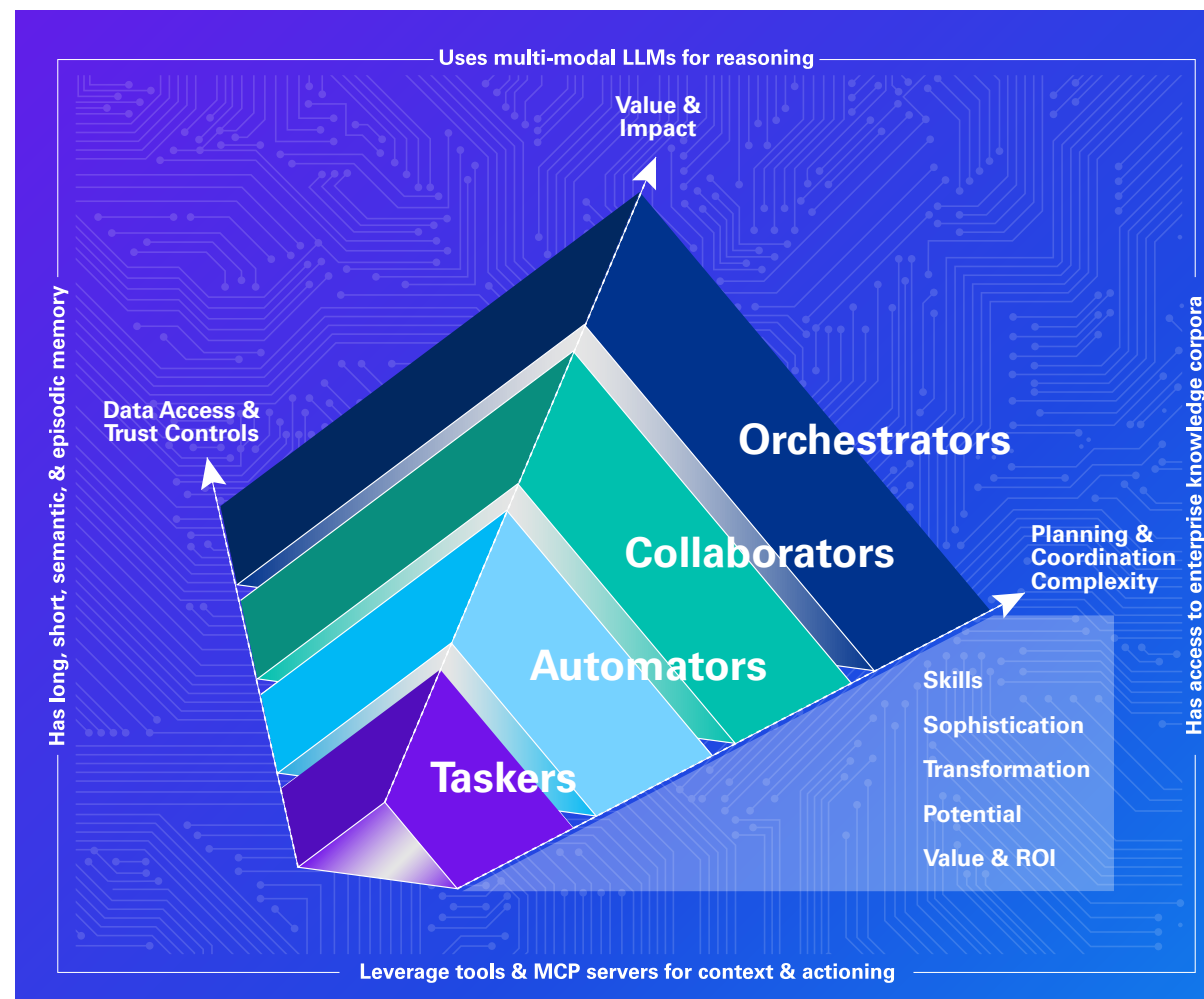
The spectrum of agents is evolving to meet critical needs

Determining where to use specific agent types can be complex. KPMG has developed a structured classification system called the KPMG TACO Framework™—Taskers, Automators, Collaborators, and Orchestrators—that helps organizations make sense of this evolving landscape and determine how best to leverage agentic AI.

Across categories, there is consistency in core computational components, e.g., foundation and reasoning models including LLMs, knowledge representation systems, contextual processing, memory architectures, and tool integration frameworks including model context protocol (MCP). However, key distinctions emerge in terms of the increasing complexity of goals these

agents can fulfill and the corresponding planning/orchestration capabilities they can execute.

This progressive framework provides organizations with a methodical approach to matching their business requirements with the appropriate level of agent sophistication. As types and complexity of agents are evolving rapidly, organizations should also prepare for how agents will collaborate and communicate with each other, an eventuality around which industry leaders are already aligning. Finally, having a centralized “AI factory” where new agentic capabilities can be conceived and developed will be critical, particularly for organizations that want to build their own customized agents instead of buying off-the-shelf products.





Taskers

Taskers execute well-defined singular goals that feature repetitive tasks and minimal complexity, such as reviewing invoices or screening for compliance risks. They rely on foundational data that can be used across multiple applications and business processes. The “human in the loop” supplies detailed instructions, outlining the steps and the actions the agent must execute. Taskers require guidance on what needs to be achieved, rather than how to achieve it.

Example Tasker applications

- **Customer service:** Categorize incoming support tickets based on keywords and assign to the appropriate team.
- **Vendor due diligence:** Screen vendor names against global sanction lists and flag potential issues.
- **Invoice validation:** Extract and verify data against compliance rules.
- **Legal affairs:** Convert past legal rulings into a standardized format.
- **Administrative:** Review employee expense reports for policy compliance and supporting documentation.



Automators

Automators deconstruct more complex goals and manage tasks that span multi-system workflows that can comprise applications, systems, and functional areas. Such agents draw from “tacit knowledge,” which is knowledge contained in workers’ minds based on their past experience with the organization. Automators execute end-to-end processes, orchestrate related tasks, manage dependencies between tasks, and dynamically engage tools, APIs, or MCP servers as needed.

Example Automator applications

- **Supply chain:** Streamline workflows by integrating order and transportation management systems, as well as supplier APIs.
- **Healthcare:** Automate claim adjudication by integrating hospital billing systems, payer systems, and regulatory databases.
- **Finance:** Accelerate month-end close by managing accruals, reconciliations, and journal entries across multiple systems.
- **E2E order-to-cash:** Automate invoice generation, payment reconciliation, and reminders.
- **Customer onboarding:** Adapt dynamically to missing info or compliance checks.

Case study: Online trading company streamlines procure-to-pay process with multi-agent “automator”

An online trading company turned to KPMG and agentic AI platform company Ema to create an automated procure-to-pay system. Ema’s product is a multi-agent system that can connect to enterprise tools to fulfill complex, multi-step tasks using both structured and unstructured data. In this case, KPMG and Ema created a streamlined accrual solution for their joint client in which Ema pulled Coupa data that included invoice purchase orders and goods received, identifying anomalies that can then be escalated to a human for reconciliation. The result was the company will now be able to complete an accrual process in a few hours, instead of the 10 days of effort by three full-time employees that were historically required.



Collaborators

Collaborators function as adaptive AI teammates, working interactively with humans to achieve multi-dimensional goals while dynamically responding to user feedback and evolving contexts. These agents integrate multiple AI capabilities (machine learning, NLP, speech, computer vision) and can work with both structured and unstructured data sources to provide expertise-driven support and augmentation. Unlike simpler agents, collaborators maintain contextual awareness and long-term memory across extended complex interactions, enabling them to function as true digital colleagues rather than passive tools.

Example Collaborator applications

- **Product development:** Synthesize user feedback, suggesting feature prioritization and generating release plans with timelines and resource estimates.
- **Marketing:** Brainstorm campaign ideas with marketing teams, refine ad copy based on style preferences, and automatically schedule social media posts.
- **Financial planning and analysis:** Work alongside analysts to develop scenario models, evaluate assumptions, and provide insights based on historical data patterns, while adapting to the specific analytical approach preferred by each team member.



Orchestrators

Orchestrators are an advanced class of agentic system designed to coordinate multiple agents, tools, and interdependent workflows at scale. These agents operate as intelligent control towers, dynamically selecting from a broad ecosystem of specialized agents and services to achieve complex, cross-functional objectives. Such agents rely on explicit knowledge, which is more objective, quantifiable, and often technical. Operating in real time, orchestrators enable dynamic resource optimization, multi-agent choreography, task delegation, inter-agent communication (leveraging protocols such as the agent-to-agent [A2A] protocol), and the reconciliation of dependencies across systems. Preparing for their adoption requires addressing the technology stack and codifying robust Trusted AI protocols. As the most sophisticated of current AI agents, orchestrators have the potential to truly unlock new business models, foster cross-industry ecosystems, and stimulate economic transformation at scale.

Example Orchestrator applications

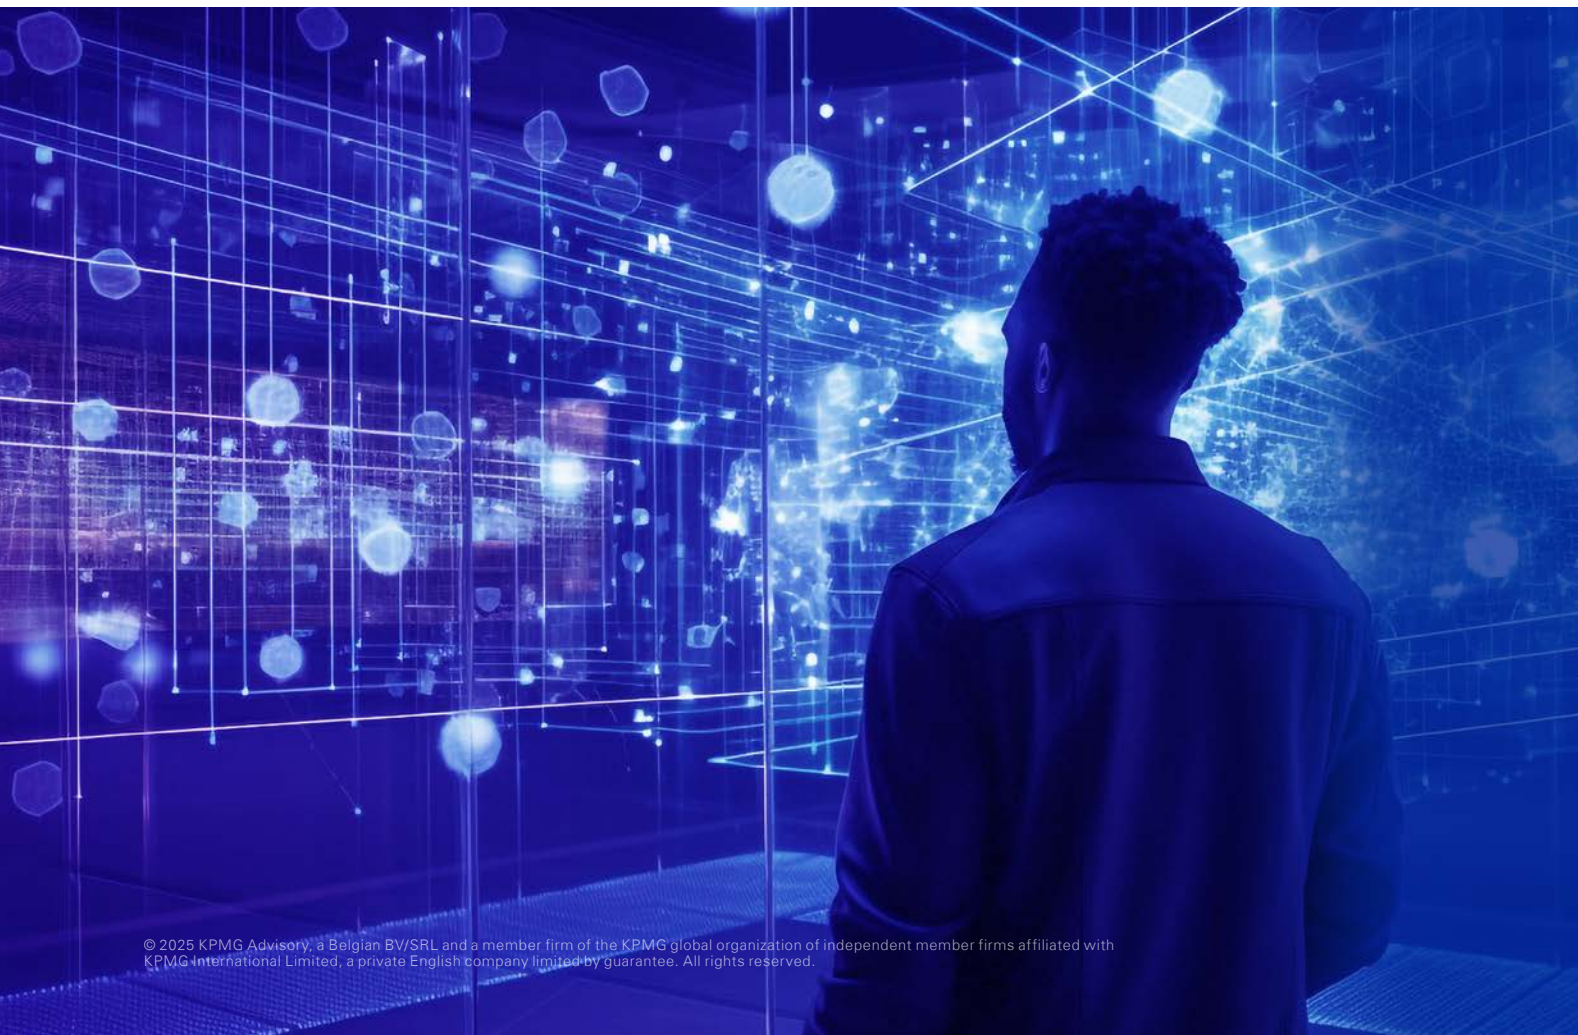
- **Finance:** Manage inter-company reconciliation processes by orchestrating agents for currency conversion, account matching, and exception reporting.
- **Regulatory compliance:** Orchestrate multi-jurisdictional compliance monitoring by assigning agents to review local regulations, ensure adherence, and escalate violations for manual review.
- **Multi-entity procurement optimization:** Orchestrate sourcing, negotiation, and compliance agents across global subsidiaries—both within and outside of the enterprise.

Each TACO agent—whether a Tasker, Automator, Collaborator, or Orchestrator—is inherently a multi-agent system. Each is composed of several coordinated agents and sub-agents working together to achieve a stated goal, rather than functioning as a single, monolithic unit. All TACO agents involve some level of orchestration or coordination, though the complexity of that coordination may vary.

KPMG TACO Framework™ overview

	Taskers	Automators	Collaborators	Orchestrators
Overall complexity	Low	Low-to-medium	Medium-to-high	High
Primary use	Execute singular goals with one or many tasks.	Automate goals that span multiple systems and process areas.	Human/AI partnerships, with AI acting as a teammate for problem solving.	Multi-agent coordination across ecosystems in pursuit of complex goals.
Planning capabilities	Simple: Planning with prompt-based execution including chaining and gated logic.	Basic: Several sub-goals, deterministic logic with decision points, and adaptive prompts.	Intermediate: Adaptive and contextual planning with human collaboration.	Sophisticated: Multi-agent coordination with contingency planning and resource optimization.
Value propositions	Free up workforce for higher-value tasks. Efficiency, innovation, and new ways of working.	Coordinate multiple goals across systems. Reimagine workflows and processes. Cycle time reduction. Improved customer experience.	Enhance workforce creativity and innovation. Create unlimited human-multiplier effect. Improved employee experience.	Drive new revenue streams and models. Unleash economic transformation.
Required knowledge and tools, including MCP needs	<ul style="list-style-type: none">• Explicit goal instructions• Standardized API connectors for specific actions• MCP: Minimal complexity with direct tool invocation	<ul style="list-style-type: none">• Complex, cross-application and procedural goal instructions w/ deep knowledge across systems• Enterprise API connectors w/ authentication handling• MCP: Moderate complexity w/ sequential tool operations	<ul style="list-style-type: none">• Availability of curated data and knowledge corpus with domain expertise• Goal expression w/ constraints and preferences• Several advanced connectors w/ error handling• MCP: Sophisticated/high complexity for handling tool interdependence and state management	<ul style="list-style-type: none">• Comprehensive knowledge and goal expression both outside of and within the enterprise• Knowledge of complex policies; sophisticated connector ecosystem w/ dynamic discovery• MCP: Advanced orchestration w/ parallel tool execution and complex state reconciliation

Setting the foundation for your agentic AI journey



Agentic AI appears to be moving faster and promises to be even more disruptive than generative AI. At the same time, it's still very early. The corporate path from possibility to operational reality is paved with uncertainty and still lacks bridges over many obstacles, some known and others yet to be discovered. But as Einstein said, "In the middle of difficulty lies opportunity." Companies that succeed first may gain advantages while others struggle to keep pace. The time between hype and hyper-performance is often underestimated but may also be shorter than we imagine. Winners and losers could be determined in 12 to 36 months. To improve their odds, every company can start the journey now, look ahead to imagine the future, and prepare to start their agentic AI journeys by addressing four key foundational issues now.



Setting the foundation for your agentic AI journey

Strategy

In the words of renowned physicist Niels Bohr, “Prediction is very difficult, especially if it’s about the future.” But we are willing to bet that agentic AI will both force, and enable, organizations to be more strategically agile. It will be necessary to review and challenge corporate strategies—often formed years ago and reviewed only annually—more frequently, deeply, and rapidly than ever before. In many industries, barriers to entry are falling as labor becomes digital and knowledge becomes free.¹⁰ At the same time, new digital moats¹¹ are appearing, allowing, e.g., access to large, high-quality datasets and affordable power. To compete in a market reshaped by AI, many will need to rethink their entire business from the ground up—including the customers they serve, the value they provide, and how they compete.

Review your enterprise strategy.

Instead of the annual off-site strategy meeting, we suggest a more frequent review cycle to cope with the level of uncertainty and speed of change around agentic AI. Use scenario building to inform your choices and include at least three components in the approach: quantify the value opportunity at stake (for the current state business), assess the risk of others disrupting you, and consider opportunities to disrupt others. Companies may start with an analysis of opportunities at their own organization but must also consider the impact of all their customers, suppliers, and competitors who may also be adopting agentic AI.

Shape your agent strategy.

In conjunction with the broader enterprise strategy, and parallel with the GenAI adoption journey, companies need to determine the level of urgency to start building an agentic workforce. Decide whether you want to be an early adopter, fast follower, or active monitor. Then, pinpoint the highest-value areas for initial deployments within the organization, e.g., function, workflow, or types of tasks. These analyses require quantification of the incremental value that agents can unlock, assessment of the incremental risks, and understanding of agent readiness from both a technology and a people perspective.

Evolve your partner ecosystem.

As agents move AI into the core operations of the enterprise, reliance on technology providers will increase. Your choice of partners, and your importance to them, could materially impact your speed to market.¹² Partnerships with leading AI-centric companies across both Big Tech and startups will be critical and are already in high demand. Don’t overlook new and smaller potential partners that can bring a competitive edge within a specific function, industry, or area of use.

¹⁰Understanding AI Agents and Their Business Impact, Wall Street Journal, April 18, 2025.

¹¹Moats are AI technologies to protect market share and profitability from competitors. (Source: How Big is Your Digital Moat? S&P Global Market Intelligence, August 16, 2019.)

¹²Accelerate growth and innovation with partner ecosystems, KPMG, 2025.



Workforce

Setting the foundation for your agentic AI journey

By taking the time now to envision the impact of agentic AI on the organization, companies can better manage the necessary shifts in roles, skills, and workforce structure. This foresight will help organizations cope with the magnitude of change and increase the chances of gaining a competitive edge.

Codify the ways you work.

Human activities can be divided into structured work, which follows standardized processes with clear rules, and tacit work, which involves applying judgment, creativity, and experience. Structured work offers near-term potential for agents and can be aided by existing software tools for process mining and documentation. Tacit work is harder to automate, but agentic AI may have a larger impact over time, with several start-ups emerging to capture tacit work. The more work that can be codified now, the faster and greater the impact of agentic AI will be. In a recent KPMG AI Pulse Survey, 78 percent of respondents said they would use agentic AI to analyze complex data sets, which requires tacit work, while 66 percent plan to employ agents for routine administrative tasks.¹³ Establishing systematic methods to capture and codify subject matter experts' tacit knowledge to enable agents' effectiveness will become an operational necessity and will likely lead to competitive advantage.

¹³ KPMG AI Pulse Survey Q1, April 2025.

What about the workforce? Driving adoption and engagement

The adoption of GenAI has been a gradual process, often hindered by workforce resistance, skepticism, and a "swivel chair" approach, since workers have to pivot from their day-to-day workflow to use GenAI. The introduction of agentic AI is changing this dynamic. By embedding AI agents into the unique workflows and tasks of different roles, organizations can realize much greater adoption, acceptance, and utilization.

Even given agentic AI's relative ease of use, some employees may be reticent to work side by side with AI agents. In a recent KPMG AI Pulse Survey, 45% of respondents said their employees were resistant to change.* Therefore, it is necessary to establish a robust change management program that considers the impact of AI agents on team dynamics, morale, and a culture of collaboration and innovation. Further, it is important to communicate that integration of AI agents into the workforce is not about replacing human jobs; it is designed to enhance the capabilities of the existing workforce.

Since AI agents will be embedded into the daily workflows of professionals and automatically combined with action taking, organizations will be positioned to realize greater workforce specialization and retention, higher trust in outcomes, and enhanced innovation. Agents promise to drive a 30 percent increase in workforce efficiency and a 25 percent reduction in operational costs by 2027, according to a Harvard Business Review report.**

Organizations that adopt both GenAI and AI agents while concurrently dissecting and reconstructing workforce roles will be more likely to realize significant productivity and capacity gains. A recent study surveying more than 100,000 workers from 11 GenAI-connected occupations found that AI tools can reduce working times by 50 percent for a third of employees' job tasks.** According to a recent KPMG AI Pulse Survey, the functions organizational leadership thinks would benefit most from agentic AI include information technology (76%), operations (74%), risk and compliance (56%), finance (39%) and marketing (35%).*

*KPMG AI Pulse Survey Q1, April 2025.

**Mark Purdy, What is agentic AI, and how will it change work? Harvard Business Review, December 12, 2024.



Workforce cont.

Setting the foundation for your agentic AI journey

Change the way you change.

Adoption of AI in the flow of work is proving more challenging than anticipated for many companies. While agentic AI relies less on humans changing their day-to-day behavior, similar challenges will eventually apply to the initial transfer of work and interactions between digital and human co-workers. Traditional top-down change-management approaches, relying on upskilling, training, performance metrics, and accountability, have not yet proven effective for AI. Companies should instead explore behavioral change approaches that feature role-modelling, peer-to-peer learning, and psychological safety as key ingredients in a cocktail of simultaneous interventions designed for their unique cultures.

Imagine a fluid hybrid organization.

As agents join the digital workforce, they will become part of the organizational structure alongside humans. This will require new ways of reporting, goal setting, and performance management for agents and for humans. Adding to the challenge is that, in the future, the number of digital workers could conceivably change by hundreds or thousands within seconds. For humans, their roles, reporting lines, and performance metrics may also need to change significantly and more frequently; an apt analogy would be conducting a major merger integration every year. Companies that rely on offshoring may find it possible to reshore work at a fraction of the cost, if fully automated with agents. Even starting a fully agentic business unit, or new start-up company, may prove possible with few or no humans from the start.

Following are some examples of how humans and AI agents can work together:



In finance, teams no longer need to consciously engage with a GenAI system during month-end close processes. Instead, agents can automatically identify unusual transactions, prepare journal entries, and alert controllers when human judgment is needed.



In the consumer and retail sector, AI agents can serve as virtual shopping assistants, providing hyper-personalized customer interactions. This allows human employees to focus on more strategic tasks, such as developing new products and services.



In the life sciences sector, AI agents can collaborate with human researchers, leading to faster breakthrough treatments, particularly for rare diseases and hard-to-treat cancers, that might not be possible with human effort alone.

An enhanced workforce, empowered by agentic AI, can drive significant advancements and innovation, positioning organizations at the forefront of their industries.



Governance/trust

Setting the foundation for your agentic AI journey

AI agents inherently require more stringent controls and an even greater commitment to Trusted AI principles since they operate with relative autonomy. Errors that do occur can also be amplified at a greater scale with AI agents that may propagate and repeat unfavorable actions before human intervention. Hence, without a more robust governance program, organizations will be unable to realize the full value from agentic AI or move as quickly as competitors that do.

Elevate security and privacy.

In addition to standard security considerations for AI, such as transparency, explainability, data privacy, and compliance, organizations adopting agentic AI should conduct regular stress testing, bias detection, and fail-safe mechanisms to prevent unintended consequences. Imagine the planner within an agent coming up with inconsistent plans every time it executes, or, even worse, the wrong tools being used for action taking. Unlike human errors, which are typically contained incidents, a single error in agentic AI programming or decision-making logic can lead to widespread consequences, potentially affecting numerous outcomes simultaneously before issues are even detected.

Avoid potential ethical violations.

Throughout any AI implementation, but particularly with agentic AI, organizations should keep ethical considerations top of mind. Be sure to consider ethics during every step of the development process, not as a bolt-on at the end. Establish ethical guidelines and governance protocols early to mitigate potential issues with bias and fairness. Expand on your existing ethics policies to include the specific risks around AI agents. Insist upon transparency into how all AI agents are built and monitored. Ensure the data you enter into agentic AI systems is free of bias to avoid perpetuating harmful stereotypes or exclusionary practices.

Put humans on-the-loop.

While generative AI benefits from having a human *in* the loop between the advice generated by AI and the action taken, this is not the case for the more autonomous agents. Instead, companies need to ensure that agentic designs incorporate human-*on*-the-loop principles, where humans supervise and monitor but do not directly intervene and approve every action. Instead, they oversee the process and have the authority to step in to override the system if necessary.



Technology, data, and security

Setting the foundation for your agentic AI journey

While the extensive use of data by AI agents is still to come, companies need to ensure that their proprietary data is accessible and of sufficient quality for agents to ingest. As their use of agents expands, organizations must also ensure that agents can interact with each other effectively and have digital identities that can be managed as co-workers within company HR systems.

Start building your agentic supply chain.

To source agents, organizations currently have three primary options: build a tailored custom solution, buy a pre-built agent, or partner with an external vendor. Each approach has its own benefits and considerations, and the choice should align with specific organizational needs, resources, and strategic goals, as well as the complexity of the agents required. Many will find it advantageous to create a multi-faceted agent supply chain with a mix of more than one of these options that can be applied across different parts of the organization over time.

How to acquire agents: Build, buy, or partner?

As C-suite executives, and particularly the CFO, consider the most economically viable way to acquire AI agents, they will want to consider the pros and cons of each of the options. Here's a detailed look at each one:

Build: Customization and control



Building AI agents in house allows organizations to tailor them to their specific needs and integrate them seamlessly with existing systems and processes. The organization will also own the intellectual property output by the agent. On the other hand, due to the complexity of this process, organizations will need an internal and/or external team with deep expertise in AI, data science, and software development to guide the process. For example, a large financial services firm decided to build its own AI Tasker to automate invoice processing. By developing the agent in house, the firm was able to customize the agent to handle specific financial regulations and internal processes relevant to their organization.



Technology, data, and security cont.

Setting the foundation for your agentic AI journey

Ensure a robust data foundation.

Organizations sit on a wealth of proprietary data that many are not able to access effectively today. To succeed in the agentic era, organizations will need a modern, cloud-enabled data platform and robust data foundation, with high-quality data that is accurate, accessible, unbiased, and enriched with context. Establish robust data governance protocols to ensure that data is collected, stored, and used ethically and securely and in compliance with regulatory requirements and data privacy standards. The MCP plays a critical role in enabling agents to dynamically access and inject relevant data into their reasoning processes. Finally, organizations should enforce enterprise-grade data governance—securing data integrity, enabling safe access, protecting sensitive information through DLP controls, mitigating bias, ensuring data quality, and maintaining compliance with privacy regulations such as GDPR, HIPAA, and CCPA.

Buy: Rapid execution and efficiency



Buying existing agentic AI capabilities is the quickest and most straightforward option and is the preferred method of acquiring agents, according to 67% of respondents to a recent KPMG AI Pulse Survey.* On the one hand, this option is the most expeditious, as it involves acquiring pre-built AI agents from established vendors, which can be deployed immediately to address specific business needs. It can also be more cost-effective than building from scratch and has the added benefit of offering access to well-tested and proven AI solutions. On the other hand, this option is not without its challenges: Off-the-shelf AI agents have limited customization options, so, over time, they may become obsolete for certain customers. There are also potential challenges integrating these tools with existing technology stacks, as well as the potential that the agent won't have data security capabilities that are on par with the organization's own standards.

Partner: Access to advanced expertise



Partnering with external vendors combines the benefits of both building and buying. This approach involves collaborating with AI experts to develop and implement AI agents, leveraging their expertise and infrastructure, while sharing risk and cost. Challenges include less control over the development and customization of AI agents compared to building in house, as well as the potential for higher costs. As with any technology partnership, third-party risks need to be top of mind and built into the service level agreement. Protections need to be in place to protect your company if a partner pivots on the development plan originally agreed upon, if a system fails, if incorrect data is fed into the system, etc. The latter is particularly critical as plugging highly sensitive information into an AI agent without guardrails in place could lead to a highly damaging data breach.

Finally, as you evaluate which avenue to take, key questions to consider include: Does the organization have the in-house expertise and resources to build AI agents? Are pre-built solutions sufficient, or is there a need for highly tailored AI agents? Do agents need to be deployed quickly? What types of agents align best with long-term strategic goals? Is there an operating model in place to upkeep and maintain these agents?

*KPMG AI Pulse Survey Q1, April 2025.



Technology, data, and security cont.

Setting the foundation for your agentic AI journey

Enable consistent agent-to-agent interoperability.

Agent-to-agent communication must move beyond custom integrations or brittle handoffs. Organizations can adopt A2A standards, such as the emerging protocol from Google, which provides a standardized framework for secure, explainable, and auditable agent interactions. A2A supports message passing, intent negotiation, and structured collaboration between agents—enabling use cases such as multi-agent task orchestration, agent chaining, and dynamic role assignment. Adopting standards for agent interoperability ensures that enterprise agents remain modular, composable, and future-proof, while avoiding vendor lock-in and duplicated effort.

Common standards for cross-agent communication and external data access

At present, what an AI model “knows” is fed into the system during an initial procedure called “pre-training,” which typically happens only once or during periodic updates. In these cases, datasets have a cutoff date, and functionality only comprises user prompting, a running conversation history, and some external information pulled into the conversation.

As agents develop at a rapid pace, a concurrent effort is underway to allow them to communicate with each other across organizations, while continuously connecting with and drawing from external information sources, tools, and services. While still a work in progress, Anthropic’s Model Context Protocol (MCP) has established standards and a platform that eliminate the need for custom integrations and enable more capable, context-aware AI applications.

Core components of MCP include:



MCP client

An AI app or tool that allows MCP servers to access external data and functionality via a standardized protocol



MCP server

Programs that use MCP protocols to leverage specific capabilities from data sources, systems of record, agents, etc.



MCP transport

A communication layer that handles the exchange of messages between MCP clients and MCP servers

Ultimately, this protocol will allow organization to implement smaller, more targeted and efficient AI systems that interact relatively seamlessly with other agents and with external data without the need for tailored fine-tuning.

Source: Beni Edwards, MCP: The new “USB-C for AI” that’s bringing fierce rivals together, Ars Technica, April 1, 2025



Technology, data, and security cont.

Setting the foundation for
your agentic AI journey

Strengthen agent identity and security.

Each agent must have a unique identity, with scoped permissions, credentials, and continuous oversight. Equally essential is runtime isolation, ensuring agents operate in segmented environments where data exposure, tool misuse, or anomalous behavior can be detected and contained. Every agent action must be auditable, revocable, and attributable. Organizations must also proactively mitigate emerging security threats, such as memory poisoning, agent hijacking, and identity spoofing.¹⁴ Remedies include implementing execution logging, identity attestation, runtime containment, and real-time monitoring to ensure agent behavior remains secure, explainable, and compliant.

Source:

¹⁴Top 10 for LLM Applications and Agentic Systems, The Open Web Application Security Project (OWASP) Generative AI, LLM Security Guidance and Initiatives.

Practical next steps to accelerate agent adoption

At the time of this paper's writing, many companies in the US have taken the first steps on their agentic journeys by starting to explore or pilot agents. However, the majority (88 percent) have not yet actively deployed any agents in their operations.¹⁵ No matter where you are on your agentic AI journey, it may matter less where you begin than how rapidly, effectively, and securely you move forward.

To be successful, companies need a clear vision of how to incorporate agents into their overall AI transformation journeys, including a solid rationale for where and when to start an approach to scaling across key functions over time. Most companies would do well to accelerate their

data transformation efforts and their Trusted AI governance given higher agentic risks. Talent leadership will be busy preparing the organization for the arrival of digital co-workers, e.g., putting in place new performance metrics for both digital and human workers.

Six steps to success

¹⁵KPMG AI Pulse Survey Q1, April 2025.

To accelerate an agentic AI program, leaders should consider the following six next steps:



1. Articulate the vision

Start with a clear vision and strategy for agentic AI integration. The vision should articulate how AI agents will transform the business, desired outcomes, key business goals to pursue, and how to start supplying AI agents. Identify areas where agents can address pain points and add the most value, such as routine tasks, complex workflows, or customer interactions.



2. Start agentic pilots

Executives starting an agentic AI journey should consider where in the organization the technology can provide the most value and, therefore, how pilot projects should be approached. Organizations venturing into AI-enhanced operations can attempt one or more of the following three approaches:

Focus on “hot spots”: Target areas where the organization has previously conducted successful GenAI pilots, especially those with high user engagement. Those will likely have the highest potential for value creation, as well as the most mature AI infrastructures and skilled personnel.

Go deep into a function: Embed AI agents extensively within a specific, highly repetitive, low-risk function, such as finance. Use this pilot as a scalable model to apply agentic AI to other functions across the organization.

Broad utilization: Identify a critical cross-functional value stream that is central to business operations and implement AI agents throughout this end-to-end process.



3. Prepare to scale agents within key functions and/or workflows

Consider whether there are a few AI agent platforms that can be piloted quickly to achieve early wins that can build executive buy-in and momentum. Use tools like the KPMG TACO Framework™ to identify the types of AI agents that can best address your needs and to determine which departments and processes are most appropriate for agentic AI applications. Eventually, develop a plan to scale to an enterprise-wide strategy, involving multiple agent types working in tandem.



4. Evolve your trusted AI governance playbook for agents

It is critical that organizations expand their existing AI governance programs to account for agentic AI, with mechanisms in place to stay ahead of legal and regulatory considerations. Consider creating a living catalog of all AI agents to track their purposes, dependencies, and performance metrics for effective scaling and governance. Develop robust operational processes to manage and maintain AI agents, including monitoring systems, performance metrics, and continuous improvement methods.



5. Implement Trusted AI evaluations specific to AI agents

As organizations align their AI governance practices with increasing regulatory conventions and leading standards, a rigorous evaluation and reporting approach for AI agents will be essential. AI system cards document how technology is used in an organization, acting as single, authoritative sources for the results of AI system evaluations that detail intended use, data considerations, AI components, and limitations. Further, before launching an agentic AI system into production, organizations should conduct thorough testing, which increasingly includes “purple teams” that run rigorous offensive testing and defensive evaluation. Inputs from these processes ultimately inform a trust score, which reflects an AI system’s adherence to the organization’s AI principles.

While new types of AI agents are being introduced at a rapid pace, it is important to remember that this is only the beginning of the beginning. By laying the groundwork today for the advent of the age of agentic AI, organizations will be well-positioned to supercharge not only their operations, but also their entire businesses.



6. Establish agentic talent performance metrics

Put metrics in place to evaluate the impact of pilot projects, as well as mechanisms for gathering feedback and learnings from stakeholders during the pilot phase. Continuously measure and optimize the performance of AI agents. Develop clear performance metrics to evaluate their effectiveness, including error rates, processing times, and customer satisfaction. Use data and feedback to continuously improve performance, including refining algorithms, updating training data, and implementing new features. It should be noted that AI agents implement telemetry to monitor their own service level agreement compliance, system health, tool performance, and anomalies.

How KPMG can help

We help turn agentic AI into your competitive edge, working side-by-side from strategy to execution, with a people-centric approach grounded in trust. By optimizing the collaboration between what humans and agents can do today, we can unlock the full potential of what will be possible tomorrow and help create lasting value together.

We offer a portfolio of services tailored to help drive your agentic AI journey:



Authors



Swami Chandrasekaran
Global Head of AI & Data Labs
E: swamchan@kpmg.com

Swami leads and executes the firm’s AI strategy across Tax, Audit, Advisory, and other functions, serving 200,000 knowledge professionals worldwide. He directs and oversees various R&D efforts and initiatives covering AI architecture, advanced knowledge assistants, AI Agents, domain-tuned Small Language Models (SLMs), synthetic data, enterprise discovery and search, and hardware-optimized solutions. He also chairs the KPMG AI Technology Review Board to ensure trusted and scalable AI adoption.



Per Edin
Global AI Sales Acceleration Leader for Advisory and US AI Go-to-Market Leader for the Tech, Media and Telecom Sector
E: pedin@kpmg.com

Per holds dual global and US leadership roles for AI: He is responsible for the alignment of AI client offerings, messaging, and go-to-market approach across the global advisory network, and for driving AI-related growth in the US Technology, Media & Telecom sector. His work has included leading the firm’s efforts to help clients navigate the GenAI market disruption, developing a portfolio of new AI services, AI thought leadership, and key accelerators like the GenAI Value Assessment.



Todd Lohr
Head of Ecosystems and National Operations Leader for Advisory
E: tlohr@kpmg.com

Todd develops and executes the KPMG ecosystem strategy, which includes coordination of corporate joint ventures, technology alliances, suppliers, downstream partners, start-ups, academia, and capital investors. Todd collaborates with KPMG teams to create new market offerings, develop the most innovative solutions for clients, and extend our reach into new markets and sectors by leveraging the latest technologies.

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

kpmg.com/be



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.

© 2025 KPMG Advisory, a Belgian BV/SRL and a member firm of the KPMG global organization of independent member firms affiliated with KPMG International Limited, a private English company limited by guarantee. All rights reserved.