



The Agentic AI Advantage: Finance agents that move the numbers

What to build, how to scale and how to measure value
KPMG Thought Leadership

February 2026



At KPMG, the future is agent-centric

Artificial Intelligence (AI) is moving from tools that only answer questions to AI agents (Agents) that can take action that is ordinarily done by Humans. In finance, this means software workers that read data, apply policy and accounting rules, act in enterprise systems, and learn within clear controls. These agents are not a trend but improve the work that finance teams do every day.

This paper, The Agentic AI Advantage, is written for finance leaders in Ghana. It explains what a finance agent is, where it fits in the Finance Delivery Model, and how leaders can begin the AI adoption journey in a safe and structured manner. The focus is on outcomes that matter, delivering higher quality financial data, achieving faster and more efficient financial close processes, strengthening internal controls, and providing sharper insights to support informed decision making.

Finance teams in Ghana are managing tight liquidity, rising expectations on compliance, and the shift to digital tax. Agents can help with real work. They can collect and validate electronic VAT invoices. They can run reconciliations and flag issues for action. They can support three-way match, supplier verification and payment readiness. They can apply credit rules, match cash, and prioritise collections.

They can prepare close checklists, draft simple narratives, and route tasks to the right people. They can build first draft forecasts, compare scenarios, and highlight the drivers that move results. They can gather data for sustainability reporting and track evidence for audit.

This is more than a technology choice - it is a leadership responsibility. Finance, technology, risk and internal audit teams must align on how agents will operate, what access they will have, and how outcomes will be monitored. Clear role definitions, comprehensive training, and steady guidance will be essential to help teams adapt to these changes.

Trust is non negotiable. Agents must work inside well-designed controls. Data quality, policy libraries, approvals, and segregation of duties must be clear. A human must stay in control. Monitoring, testing and model governance must be routine. Security and privacy must match the standards that boards and regulators expect.

It is important to keep focus. Agents are part of a wider finance strategy and they work best when the foundation is sound. Ensure leadership is aligned on the value you seek to achieve, and modernise your data and core processes to support this.

Start with a small set of use cases that have clear measures. Prove value, then scale in an orderly way through the service delivery model and the Target Operating Model. Report progress in plain language that builds confidence.

The journey will require effort, but the path forward is clear. Finance agents can help Ghanaian organisations accelerate processes, enhance trust in financial data, and enable professionals to focus on analysis and support executive decision making. I look forward to the dialogue this paper will support and to the ideas it will spark across our finance community.



Reindolf Annor
Partner, Accounting Advisory Services

Contents

04	Do you really know what a Finance AI Agent is	10	The spectrum of agents is evolving to meet critical needs
05	The era of agentic finance has already begun	14	Setting the foundation for your finance agent journey
07	The value at stake	23	Practical next steps to accelerate agent adoption
08	Four ways Finance AI Agents can unlock value	26	How KPMG can help



Do you really know what a Finance AI Agent is

The term "agent" is now widely used, creating the impression that everything qualifies as an agent, which is not accurate. To realise the potential of this rapidly evolving field, it is essential to first understand what an agent truly is.


A Finance AI Agent is a digital worker that fulfils finance objectives by taking meaningful and independent action.

It does this by combining reasoning from large language models with planning, orchestration, curated knowledge, data retrieval, specialist tools, and strong governance.

A finance agent can make decisions in real time, adapt to new situations, and improve through feedback and interaction.

What a finance AI agent does

- Reads both structured data, which is highly organised and easily searchable, like ledger entries, bank statements, inventory list, and payroll records typically residing in databases or accounting software and unstructured content, such as contracts, scanned receipts and emails which requires advanced tools like optical character recognition (OCR) and natural language processing (NLP) for interpretation.
- Takes action in systems. Examples include creating a draft journal, proposing a payment run inside limits, updating a vendor record, raising a ticket, or routing an approval.
- Learns from outcomes and user feedback but never steps outside defined controls or approval paths.
- Applies accounting policies, tax rules and process standards to make a decision.



© 2026 KPMG, a partnership established under Ghanaian law and a member firm of the KPMG network of independent member firms affiliated with KPMG International Limited, a private English company limited by guarantee. All rights reserved.





KPMG Quarterly Pulse Survey

Companies piloting AI agents



up from
37%
last quarter

Companies planning to put agents into production



only
11%
have implemented agents in production

Just as executives are beginning to adopt the game-changing capabilities of generative AI (GenAI) across their organisations, 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 organisations 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 make better decisions and act faster. 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.

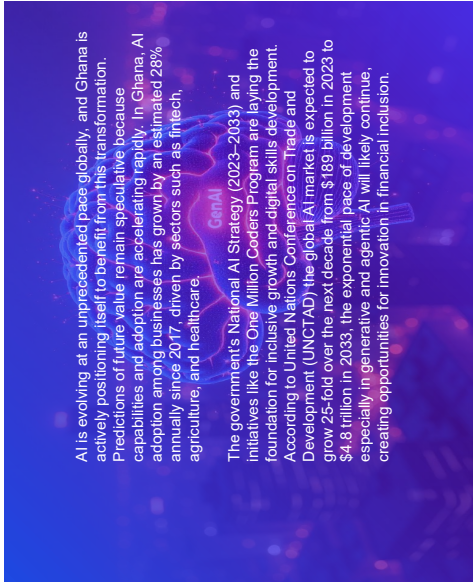
While still emerging and evolving, AI agents can already take on human tasks that were previously thought too complex to automate, and they may be a better choice for some activities previously automated with older AI tools. Hence, most companies in our quarterly pulse survey are already piloting AI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

¹KPMG AI Pulse Survey Q1, April 2025.
²KPMG AI Pulse Survey Q1, April 2025.

© 2026 KPMG, a partnership established under Ghanaian law and a member firm of the KPMG network of independent member firms affiliated with KPMG Network, a private English company limited by guarantee. All rights reserved.



The value at stake



The KPMG GenAI Value Assessment

In 2024, KPMG launched a 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 organisational digital twin for more than 17 million companies using more than 3 billion data points.



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.



[Reference Link](#) : National AI Strategy to Drive Innovation Launch, Ghana
[Reference Link](#) : Ministry of Communications, Ghana



© 2026 KPMG, a partnership established under Ghanaian law and a member firm of the KPMG network of independent member firms affiliated with KPMG International Limited, a private English company limited by guarantee. All rights reserved.

Four ways Finance AI agents can unlock value

AI agents can unlock greater enterprise value in at least four significant ways. The focus is on work that moves numbers every day across record to report, procure to pay, order to cash, plan to perform, and treasury.



Agents widen the aperture for automation

As agents grow in reasoning and collaboration, they can take on tasks, roles, and whole processes that once needed people. In finance this includes reading and checking invoices, running three-way match, preparing journals, reconciling subledgers to the general ledger, validating tax on transactions including electronic VAT, and applying credit rules before an order is accepted. After automation, agents can be directed to keep improving how they work. They can tune their interactions toward a stated goal such as faster close, higher match rates, lower days sales outstanding, or fewer control exceptions.




Agents do not sleep

Once agents automate a set of tasks, they can run them continuously with human supervision. This alone multiplies daily output. A finance team can keep reconciliations moving through the night, validate new invoices against policy as they arrive, match bank statements as files land, monitor working capital positions, and raise alerts on changes in cash, exposure, or counterparty risk. When several agents work in parallel they complete a period end checklist faster than a daytime team would, and the function can respond to off hour events, such as late batch postings or payment failures, without waiting for the next morning.



© 2020 KPMG LLP, a partnership established under Ghanaian law 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.

Document Classification: KPMG Confidential



Four ways Finance AI agents can unlock value - Cont'd

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, organisational structure, or lists of customers and suppliers change, the agentic system automatically adjusts the end-to-end process.

Agents convert knowledge into action

Generative tools can surface information. Finance agents go further by turning knowledge into coordinated action. This depends on capturing and structuring the finance expertise that often sits only in people's heads. Once that guidance is in policies, checklists, and playbooks, agents use it to act. They can draft standard narratives for management reports, prepare proposed adjustments with evidence, route items to the right approver, and prioritise collections based on payment behaviour and credit risk. In planning, agents can build first pass forecasts from drivers and assumptions, compare scenarios, and flag the inputs that move results the most. In compliance, agents can gather evidence and maintain links for audit.

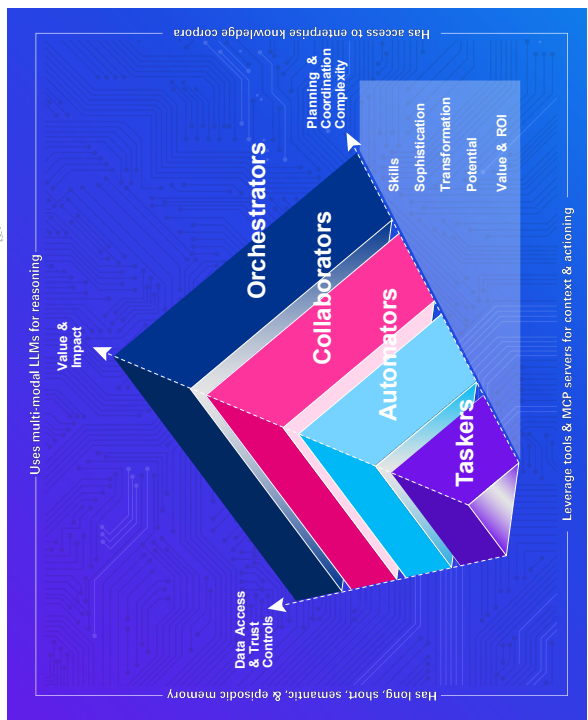
The spectrum of agents is evolving to meet critical needs

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 organisations 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 organisations 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, organisations 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 organisations that want to build their own customized agents instead of buying off-the-shelf products.





Tasks

Tasks execute well defined single goals with repetitive steps and low complexity. A human sets the objective and provides clear instructions. Tasks need to know what to achieve rather than every detail of how to achieve it. They rely on trusted data that can support many finance processes.



Automators

Automators break down more complex goals and run tasks that span many systems and steps. They use the tacit knowledge that sits with experienced staff and turn it into actions and rules. Automators can run end to end processes, manage dependencies, and call tools, models, or services as needed.

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 Ena to create an automated procure-to-pay system. Ena’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 Ena created a streamlined accrual solution for their joint client in which Ena pulled Coupa data that included invoices, 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.

Example Tasker applications

- Invoice validation that extracts fields, checks tax rules including electronic VAT, and flags exceptions for review
- Vendor due diligence that screens names against sanctions and adverse media and prepares a summary for approval
- Expense review that checks claims against policy and supporting documents before routing to a manager
- Bank statement classification that tags transactions for faster reconciliation
- Review employee expense reports for policy compliance and supporting documentation.

Example Automator applications

- Month end close that schedules and tracks accruals, prepares journals with evidence, runs reconciliations, and routes breaks to owners
- Procure to pay that captures invoices, performs three-way match, validates tax, posts to the ledger, and prepares payment proposals
- Order to cash that issues invoices, reconciles payments, applies cash, and sends reminders based on risk and behaviour
- Customer or supplier onboarding that adapts to missing information, performs checks, and updates master data
- Tax compliance that prepares draft returns, supports e VAT reporting, and packages evidence for internal review



Collaborators

Collaborators act as adaptive teammates. They work with people on multi dimensional goals and respond to feedback and context. They can combine several AI capabilities, keep context over long interactions, and help experts move faster while staying within policy.



Orchestrators

Orchestrators are advanced control towers. They co ordinate many agents, tools, and linked workflows at scale. They work in real time, allocate tasks, manage inter agent communication, resolve dependencies, and optimise resources. They rely on clear and explicit knowledge and strong Trusted AI protocols. Orchestrators can enable new finance service models across many entities and markets.

Example Collaborator applications

- Financial planning and analysis that co builds driver based models, tests scenarios, and explains movements in plain language
- Management reporting that drafts first pass commentary, prepares variance explanations with links to evidence, and supports board packs
- Audit readiness that manages prepared by client requests, tracks evidence, and summarises status for the controller
- Treasury support that watches positions, runs what if checks on rates and limits, and proposes actions for approval

Example Orchestrator applications


- Finance agents that manages intercompany reconciliation across subsidiaries and currencies by directing agents for currency conversion, account matching, and exception handling
- Regulatory and tax compliance that co-ordinates agents to track changes, prepare returns for each jurisdiction, check adherence, and escalate issues for review
- Multi-entity procurement optimisation that co ordinates sourcing, negotiation, and compliance agents across a group to achieve savings and reduce risk

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

Overall complexity		Taskers		Automators		Collaborators		Orchestrators	
		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 optimisation.	
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		• Explicit goal instructions • Standardized API connectors for specific actions • MCP: Minimal complexity with direct tool invocation		• Complex, cross-application and procedural goal instructions with deep knowledge across systems • Enterprise API connectors with authentication handling • MCP: Moderate complexity with sequential tool operations		• Availability of curated data and knowledge corpus with domain expertise • Goal expression with constraints and preferences • Several advanced connectors with error handling • MCP: Sophisticated high complexity for handling tool interdependence and state management		• Comprehensive knowledge and goal expression both outside of and within the enterprise • Knowledge of complex policies; sophisticated connector ecosystem with dynamic discovery • MCP: Advanced orchestration with parallel tool execution and complex state reconciliation	

Setting the foundation for your finance agent journey



Agentic AI is moving fast and may be more disruptive than generative AI. It is still early. The path from possibility to day-to-day use has uncertainty and many obstacles, some known and others yet to come. In the middle of difficulty lies opportunity. Early movers can gain an advantage while others struggle to keep pace.

The gap between excitement and high performance is often underestimated and may be shorter than we expect. Outcomes for leaders could be set within twelve to thirty six months. To improve the odds, every organisation can start now, look ahead, and prepare by addressing four foundations.

2

Setting the foundation for your agentic AI journey

Strategy

- Prediction is difficult, especially about the future. What is clear is that agentic AI will both force and enable greater strategic agility. Many organisations will need to review and challenge strategies that were formed years ago and refreshed only once a year. Barriers to entry are falling as labour becomes digital and knowledge becomes widely available.
- At the same time, new digital advantages are emerging, including access to quality data and affordable compute. To compete in a market reshaped by AI, leaders may need to revisit the customers they serve, the value they provide, and how they compete.

Review your enterprise strategy.

- Shorten the strategy cycle to match the speed of change. Use scenarios to test choices. Quantify the value at stake for today's business. Assess the risk that others disrupt you. Consider how you might disrupt your own market.
- Look beyond your organisation to your customers, suppliers, and competitors who may also adopt agents. Anchor these choices in your Finance Target Operating Model so that process, people, service delivery model, technology, performance insights, and governance move together.


Shape your agent strategy.

- Decide your posture. Early adopter, fast follower, or active watcher. Select a small set of high value areas to begin. Examples include record to report, procure to pay, order to cash, plan to perform, and treasury.
- For each area, size the value, identify the risks, and check readiness across data, policy, controls, and people. Use simple measures that a board will understand such as close days, match rates, aged debt, cash forecasting error, and exceptions cleared.

Evolve your partner ecosystem.

- Your speed will depend on partners. Most leaders plan to buy pre built agent solutions and combine them with selective in house build.
- Choose providers you trust and where you matter to them. Balance large platforms with specialist partners who know finance.

¹Understanding AI Agents and Their Business Impact, Wall Street Journal, April 16, 2025.
²What is the Impact of AI on the Financial Markets? Market Share and Profitability from Competitors, (Source: How Big is Your Digital Moat? S&P Global Market Intelligence, August 16, 2019).
³Accelerate 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 organisation, companies can better manage the necessary shifts in roles, skills, and workforce structure. This foresight will help organisations 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. 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, organisations can realize much greater adoption, acceptance, and utilisation.
- Even given agentic AI's relative ease of use, some employees may be reluctant 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, organisations will be positioned to realize greater workforce specialisation 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.
- Organisations 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-corrected occupations found that AI tools can reduce working times by 50 percent for a third of employees' job tasks. According to KPMG AI Pulse Survey, the functions organisational leadership thinks would benefit most from agentic AI include information technology (76%), operations (74%), risk and compliance (56%), finance (59%) and marketing (53%).*

*KPMG AI Pulse Survey Q1, April 2025.
**Mark Paddy, What is agentic AI, and how will it change work? Harvard Business Review, December 12, 2024.



Workforce

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 behaviour, 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 organisation.

- As agents join the digital workforce, they will become part of the organisational 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 order to cash, agents apply credit rules, prioritise collections using payment behaviour, and prepare dispute packs.



In procure to pay, agents run three way match, validate tax including VAT, and route holds to the right approver.

An enhanced workforce, empowered by agentic AI, can drive significant advancements and innovation, positioning organisations 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, organisations 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, organisations 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, organisations 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 in-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.





Setting the foundation for your agentic AI journey

Technology, data, and security



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, organisations must also ensure that agents can interact with each other effectively and have digital identities that can be managed as co-workers within company finance systems.

Start building your agentic supply chain.

- To source agents, organisations 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 organisational 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 organisation 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 : Customisation and control

- Building AI agents in house allows organisations to tailor them to their specific needs and integrate them seamlessly with existing systems and processes. The organisation will also own the intellectual property output by the agent.
- On the other hand, due to the complexity of this process, organisations 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 customise the agent to handle specific financial regulations and internal processes relevant to their organisation.





Technology, data, and security

Ensure a robust data foundation.

- Organisations sit on a wealth of proprietary data that many are not able to access effectively today. To succeed in the agentic era, organisations 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 model context protocol (MCP) plays a critical role in enabling agents to dynamically access and inject relevant data into their reasoning processes.
- Finally, organisations should enforce enterprise-grade data governance securing data integrity, enabling safe access, protecting sensitive information through data loss prevention (DLP) controls, mitigating bias, ensuring data quality, and maintaining compliance with privacy regulations such as Data Protection Act and guidance from the Data Protection Commission.

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 customisation 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 organisation'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 customisation 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 organisation 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

Strengthen agent identity and security.

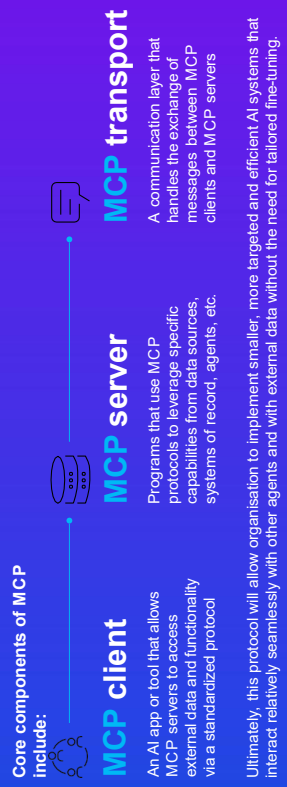
- Each agent must have a unique identity, with scoped permissions, credentials, and continuous oversight. Equally essential is runtime isolation, ensuring agents operate in segmented environments where data exposure, tool misuse, or anomalous behaviour can be detected and contained.
- Every agent action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

Source: vTop 10 for LLM Applications and Agentic Systems, The Open Web Application Security Project (OWASP) Generative AI LLM Security Guidance and Initiatives

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 organisations, 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.




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





Practical next steps to accelerate agent adoption

- Many global companies have taken the first steps on their agentic journeys by starting to explore or pilot agents. However, majority have not yet actively deployed any agents in their operations.
- 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 organisation for the arrival of digital co-workers, e.g., putting in place new performance metrics for both digital and human workers.



© 2026 KPMG, a partnership established under Ghanaian law and a member firm of the KPMG network of independent member firms affiliated with KPMG Network, a private English company limited by guarantee. All rights reserved.



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 organisation the technology can provide the most value and, therefore, how pilot projects should be approached. Organisations venturing into AI-enhanced operations can attempt one or more of the following three approaches:

- **Focus on “hot spots”:** Target areas where the organisation 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 close reconciliation and consolidation. Use this pilot as a scalable model to apply agentic AI to other functions across the organisation.
- **Broad utilisation:** 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 organisations 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 organisations 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 organisation, 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, organisations 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 organisation's AI principles.





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.

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, organisations will be well-positioned to supercharge not only their operations, but also their entire businesses.



How KPMG can help

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



We help turn agentic AI into your competitive edge, working side-by-side item strategy to execution, with a people-centric approach grounded in trust. By optimising 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.



Reindolf Annor

Partner, Accounting Advisory
Services
Mob: +233 (0) 501 324 312
Email: raannor@kpmg.com



Nathaniel Adjin-Tettey

Associate Director, Accounting
Advisory Services
Mob: +233 (0) 501 324 285
Email: nadjin-tettey@kpmg.com



Frank Osei Tutu

Assistant Manager, Accounting
Advisory Services
Mob: +233 (0) 543 710 059
Email: fosetutu@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 on such information without appropriate professional advice after a thorough examination of the particular situation.

The KPMG name and logo are trademarks used under license by the independent member firms of the KPMG global organisation.

Powered by KPMG International

© 2026 KPMG, a partnership established under Ghanaian law and a member firm of the KPMG global organisation of independent member firms affiliated with KPMG International Limited, a private English company limited by guarantee. All rights reserved.

Document Classification: KPMG Confidential