




Agentic AI untangled: Navigating the build, buy, or borrow decision

January 2026





**Swami
Chandrasekaran**

Global Head of AI and Data Labs
KPMG US

The spectrum of agentic AI options can paralyze even the boldest leaders.

Every enterprise now faces the same question—should I build, buy, or borrow AI agents? Getting value from agentic AI isn't about chasing models for technology's sake; it's about redesigning how work gets done through a federated mix of agents anchored by a unified control system for trust and scale. This approach only works with a strong data and context foundation to govern, orchestrate, and continuously improve this new AI workforce. This paper provides practical strategies and a clear framework to help you lead by design and turn agents into advantage.

Agentic artificial intelligence (AI) is powering a \$3 trillion productivity revolution and nearly every enterprise wants in.¹ In fact, the gap between proof of concept and enterprise-scale transformation is narrowing: 42 percent of companies surveyed by KPMG have successfully integrated agents into their workflows, up from only 11 percent in early 2025.²

Agentic AI is growing at such breakneck speed that even well-resourced companies struggle to decide how to scale their use. Homegrown, bespoke agents are appealing for obvious reasons, but not every company can afford the in-house AI talent necessary to develop them. Buying an agent offers a direct route, although vendor complexities can cause tangles of their own, such as thorny questions around data ownership. And, although borrowing accelerates time-to-market and adds flexibility, commoditization becomes inevitable as agent ecosystems and marketplaces mature.

Success in agentification demands a clear strategy and focus on value. Without it, companies risk either a chaotic, high-risk ecosystem or an overly rigid, one-size-fits-all solution. In this paper, we help technology decision makers evaluate the build, buy, or borrow decision and explore ways to turn agentification opportunities into return on investment.

¹ Quantifying the GenAI opportunity, KPMG, June 2025.

² KPMG Quarterly Pulse Survey, Q3 2025.


The build, buy, or borrow question

Fear of being left behind has led many companies to buy off-the-shelf solutions, but we expect more companies to begin building their own solutions in-house or in partnership with other companies in the months to come. According to the Q3 2025 AI Pulse Survey, 57 percent of organizations favor a blend of building and buying AI agents, up from 51 percent in the second quarter of 2025.³

Further, as AI agents morph from breakthrough tools to true collaborators and orchestrators, it is likely that even more organizations will take a blended approach to acquiring them. In a blended approach, an enterprise might construct an internal agent and then augment or extend its capabilities with specialized modules or third-party components—adding advanced reasoning, industry or domain-specific expertise, and contextual intelligence to meet unique business needs.

“ According to our latest AI Pulse Survey, 57 percent of organizations favor a blend of building and buying AI agents, up from 51 percent in the second quarter of 2025.

³ KPMG Quarterly Pulse Survey, Q3 2025.





Deliver distinctive user experiences by **building** AI agents in house

Organizations that develop their own AI agents can boast product differentiation and user experiences that increase competitive advantage. This approach embeds security and compliance protocols into solutions, safeguarding sensitive data and helping ensure regulatory compliance. And, with a bespoke approach, even when organizations build an agent on a framework like LangGraph or Microsoft Agent Framework, they retain full ownership. That helps them grow the solutions organically as business and market needs evolve.

Who should pursue this path?

Building AI agents is generally best suited to organizations with advanced technical expertise, a robust and scalable infrastructure, well-defined governance, and a clear strategic vision for agent deployment. The process often takes longer than buying or borrowing, so be mindful about pulling key subject matter experts and information technology resources away from other business priorities. Further, this approach requires a substantial investment in software development, data infrastructure, and product management. That may make building more viable for larger companies versus smaller ones.

Ideal candidates also have continuous access to subject matter experts, cultivate an always-on evaluation mindset, and maintain a strong focus on data integrity and context engineering. Sectors like healthcare, life sciences, financial services, and Big Tech are more likely to use the build model as they process large amounts of sensitive data and have the financial resources to pursue this avenue.



Building AI agents in-house

To get started:

- 01** Conduct a thorough assessment of your organization's workflows, available data and data sources, and strategic objectives.
- 02** Collaborate with key stakeholders to define the specific business problems the agent should address, ensuring alignment with your long-term goals.
- 03** Assemble a cross-functional team that includes product managers, data engineers, software developers, and subject matter experts to map out the agent's architecture and required capabilities.
- 04** Tailor the solutions based on your own unique workflows, proprietary data, custom tools or model context protocol connectivity, and your unique context and strategic objectives.
- 05** Use modular AI frameworks and leading practices for context engineering to incrementally build, test, and iterate your agents—ensuring scalability, resilience, and ongoing alignment with evolving market demands.
- 06** Ensure that any agentic solution you build embeds your proprietary domain expertise, workflows, and compliance requirements.

Building an internal audit agent for rapid innovation

Problem: KPMG wanted to automate and streamline routine audit processes, such as data analysis and document review, as well as more substantive procedures including expense vouching, searching for unrecorded liabilities, and reviewing accrued expenses. The complexity and scale of developing an advanced AI-powered internal audit solution presented substantial investment and risk, while efficiency and compliance remained critical priorities.

Solution: KPMG partnered with Microsoft to cocreate the Clara global smart audit platform. This collaboration leveraged the extensive domain experience of KPMG and the advanced AI capabilities and cloud infrastructure. Agentic AI was integrated into Clara to standardize and automate repetitive audit tasks, enabling rapid deployment of innovative functionalities.

Benefit: This agentic development approach allowed KPMG to quickly implement leading AI features, minimize up-front investment, reduce development risks, and enhance the overall scalability of the platform. The result was improved operational efficiency, strengthened compliance, and increased client value through more consistent and streamlined audit procedures.



Increase speed-to-market by buying prebuilt agents

When urgent business needs demand proven tools, many organizations turn to ready-made AI agent solutions from trusted vendors. In fact, 58% of enterprises plan to customize existing AI models.⁴

These solutions typically offer vendor support, compliance assurances, and integration options that reduce the need for extensive in-house development. With hundreds of vendors currently providing agentic platforms—such as UiPath for business process automation, Moveworks for automated employee support, and DataRobot for predictive analytics—companies can achieve rapid deployment with immediate access to advanced capabilities. Vendor-supported solutions often include scalable cloud infrastructure, simplified maintenance, and prevalidated compliance with evolving industry regulations.

⁴ Aminu Abdullahi, AWS Study: Generative AI Tops Corporate Budget Priorities, Surpassing Cybersecurity, eSecurity Planet, May 7, 2025.

Who should pursue this path?

Prebuilt agents are especially attractive to organizations that lack sufficient in-house talent or AI expertise, as well as those looking to minimize up-front capital expenditures and ongoing research and development (R&D) investments. Buying AI agents is most appropriate for use cases where prebuilt agents closely match business needs; can be implemented with minimal adaptation; and don't elevate concerns about data confidentiality, residency, and compliance. Organizations need to be prepared for ongoing dependency on external vendors for updates, support, and long-term sustainability, particularly when it comes to the underlying model's decision logic. And since there are limitations to how extensively off-the-shelf solutions can be tailored to a company's unique domain, organizations following this path may need to explore alternative strategies to differentiate themselves in a crowded market.

To get started:

01 Assess the sufficiency of out-of-the-box features, the reputation and track record of the vendor, and the level of customization or integration required.

02 Identify core business requirements and use cases where AI can deliver the greatest impact. This involves mapping out workflows, compliance needs, data privacy considerations, and integration points with existing systems.

03 Conduct a thorough market assessment by reviewing multiple vendors and platforms, focusing on those with proven expertise, validated performance, and precertified compliance with industry standards.

04 Conduct due diligence of your short list—reviewing contract terms, service-level agreements, and ongoing support commitments.

Accelerating finance automation with agentic AI on Oracle Cloud

Problem: A global gaming company using Oracle Cloud Infrastructure (OCI) sought to modernize its finance and procurement operations. The client faced high manual workloads, fragmented workflows, and limited real-time insights. They needed a scalable solution that could automate key processes while ensuring compliance and audit readiness.

Solution: KPMG deployed a multiagent system built on Oracle Fusion applications and OCI's generative AI services. The solution included RAG agents for contextual invoice validation, Oracle Digital Assistant for conversational interfaces, and a data intelligence layer integrated with Oracle SaaS adapters. The KPMG Agent Factory model enabled rapid development and deployment of finance-specific agents, tailored to the client's operational needs.

Benefit: The implementation led to a 40 percent reduction in manual workflows, a 45 percent improvement in procurement cycle time, and enhanced transparency across finance operations. The client gained real-time insights, improved audit readiness, and benefited from the KPMG Trusted AI framework, which ensured ethical, secure, and scalable deployment.




Scale quickly by borrowing AI agents from partners


For organizations in early adoption or scaling AI without major investment, codeveloping AI agents with specialized partners offers a practical, flexible path. In this model, companies cocreate agents with external consultants, technology vendors, or platform providers, blending their own domain expertise with their partners' technical capabilities. This approach enables businesses to rapidly access advanced agent functionality, benefit from industry-leading innovation, and minimize the up-front costs and risks associated with in-house development.


Who should pursue this path?


Partner-driven deployments can be especially advantageous in sectors where AI talent is scarce or where internal resources or infrastructure are stretched thin, such as in healthcare. Organizations should have access to resources that can help integrate borrowed agents with internal systems to ensure contextual fit and robust performance.

To get started:

- 

Carefully assess whether your strategic priorities, available resources, and compliance requirements align with a partner-driven approach. Critical decision points—such as the sufficiency of internal technical expertise, capital constraints, and the need for rapid deployment—should guide whether borrowing is the most practical strategy for scaling AI capabilities.
- 

Identify the business functions and use cases where external collaboration will close technical gaps. Shortlist potential partners with a proven track record in AI enablement, evaluating based on track records, industry certifications, and long-term support commitments.
- 

Prioritize data and risk sharing, intellectual property rights, and long-term management during contract negotiations.
- 

Address integration requirements, compliance needs, and security protocols to facilitate smooth deployment.

AI-powered demand forecasting for a global retailer

Challenge: A multinational retail company faced issues with inaccurate demand forecasts that resulted in frequent stockouts and excess inventory throughout its extensive supply chain. The company's traditional forecasting models struggled to keep pace with rapidly shifting consumer behavior, seasonal trends, and regional market differences.

Solution: KPMG addressed these challenges by implementing an AI-driven demand forecasting system powered by machine learning algorithms. This advanced solution analyzed a wide variety of data inputs, including historical sales data, weather patterns, local events, social media sentiment, and promotional calendars. The AI system continuously learned and adapted as new data became available, which enabled it to improve its forecasting accuracy over time.

Benefits: The implementation of this AI-powered solution led to a 15 percent reduction in inventory costs and a 30 percent improvement in forecast accuracy. The retailer significantly enhanced supply chain agility, allowing for faster responses to market changes and improved product availability—ultimately boosting customer satisfaction. An additional advantage was the AI model's ability to identify underperforming SKUs and recommend optimal pricing strategies based on regional demand elasticity.

Creating clarity from chaos

Agentic AI comes in many forms, each designed to fit specific business needs. Agents we call “Taskers” handle well-defined repetitive duties, making everyday tasks effortless. “Automators” power through complex workflows, transforming routine processes into seamless efficiency. “Collaborators” act as dynamic digital teammates, learning and adapting to drive innovation alongside human colleagues. “Orchestrators” serve as intelligent control towers, coordinating agents and resources to tackle large, multilayered objectives.

Choosing the right agent type isn’t just about matching features—it’s a strategic move that should align with the build, buy, or borrow decision tree (see Figure 1). By tracing each branch of the tree, you can methodically assess questions around expertise, capital investment, compliance needs, technical integration, and long-term innovation goals. The decision points help clarify whether internal development is feasible, whether off-the-shelf solutions offer sufficient validation and efficiency, if partnership with external specialists can accelerate deployment, or if a hybrid model best balances customization and scalability.




Figure 1. A comparative guide to building, buying, or borrowing AI agents

Considerations and requirements	BUY: <i>When to acquire prebuilt AI agents</i>		BUILD: <i>When to develop AI agents in-house</i>		BORROW: <i>When to partner or codevelop AI agents</i>	
	You need rapid deployment and vendor support, and your requirements align well with available solutions		You need maximum control, differentiation, and security, and have the resources to support ongoing development and governance		You need flexibility and rapid access to advanced capabilities, and want to share development and risk with external partners	
Must-haves	Observability	Action-level audit trails	Error handling and rollback	AuthN/Auth	Least privilege access	Secure MCP with valid schemas
	HITL/HOTL escalation	Red-teaming testing	AI system cards	Choice of LLM/model	Measurable KPIs	Upkeep and monitoring
Strategic fit	<ul style="list-style-type: none"> • Low-to-moderate differentiation is acceptable* • Vendor solution meets more than 80 percent of needs 		<ul style="list-style-type: none"> • Core differentiation required* • High-value and long-term strategic asset 		<ul style="list-style-type: none"> • Medium differentiation where ownership is not essential 	
Knowledge and data sensitivity	<ul style="list-style-type: none"> • Minimal IP or data sensitivity is acceptable • Vendor environment is sufficient 		<ul style="list-style-type: none"> • High sensitivity/proprietary IP required • Full control, including data sovereignty, is retained 		<ul style="list-style-type: none"> • Moderate sensitivity is acceptable if partner manages risk 	
Customization and orchestration	<ul style="list-style-type: none"> • Minimal customization required • Vendor orchestration is sufficient 		<ul style="list-style-type: none"> • Deep customization and orchestration control required* • Full orchestration control (multi-agent, memory, tools) 		<ul style="list-style-type: none"> • Significant customization needed, but partner provides technical lift 	
Governance, safety, and evaluation	<ul style="list-style-type: none"> • Acceptable baseline governance and compliance • Adequate observability 		<ul style="list-style-type: none"> • Full governance control required • Internal evaluation pipelines are available (HITL/HOTL, audits, sandboxing) 		<ul style="list-style-type: none"> • Partner supplies and manages needed governance frameworks and safety tooling • Desire to offload risk to trusted partner* 	
Integration and interoperability	<ul style="list-style-type: none"> • Modest integration needs • Vendor supports enterprise connectors • Secure tool calling (MCP or equivalent) 		<ul style="list-style-type: none"> • Deep enterprise integration required • Custom MCP tools and A2A interoperability required 		<ul style="list-style-type: none"> • Integration complexity exists; partner handles tool-building and enterprise connectivity 	
Cost, risk, and lock-in	<ul style="list-style-type: none"> • Predictable total cost of ownership • Vendor lock-in is acceptable • BYO Model/LLM flexibility is available 		<ul style="list-style-type: none"> • Acceptable cost for sovereignty and for ownership • Desire to avoid long-term dependency 		<ul style="list-style-type: none"> • Lower up-front cost • Gainshare or outcome-based pricing reduces risk* 	
Skills, talent, and AgentOps maturity	<ul style="list-style-type: none"> • Minimal to low internal expertise required* 		<ul style="list-style-type: none"> • Strong engineering capabilities are available* • Mature AgentOps capabilities are in place* 		<ul style="list-style-type: none"> • Internal capability is insufficient* • Partner provides engineering and AgentOps support 	
Is best suited for ...	<ul style="list-style-type: none"> • Early or mid AI maturity orgs 		<ul style="list-style-type: none"> • Companies in high maturity, regulated sectors • Companies with continuous SME availability 		<ul style="list-style-type: none"> • Early maturity or resource-constrained orgs • Desire to de-risk before committing to build 	

*Key requirement

The path out of the maze

Agentic AI is an opportunity to redefine what's possible for your organization, but the real power lies in navigating the build-buy-borrow decision through a framework that transforms complexity into clarity. By rigorously evaluating whether to craft bespoke agents, harness tested external solutions, partner with industry specialists, or embrace a hybrid approach for maximum adaptability, leaders can help ensure that agentic investments deliver substantial and sustainable impact.

Start by honestly assessing organizational readiness for agentic AI:

Workforce capabilities:

Assess AI skills and readiness across the workforce, including the need for upskilling/reskilling and hiring specialized talent.

Technical infrastructure:

Ensure technological readiness, including the ability to integrate agentic AI with legacy systems.

Context engineering:

Focus on designing and implementing strategies to curate information in an agent's context window.

Security protocols:

Ensure end-to-end encryption, masked personally identifiable information, and deterministic system interactions are in place to protect organizational integrity and ensure policy adherence.

Regulatory compliance:

Ensure you have robust and compliant guardrails for AI deployment.

Ability to scale:

Assess the ease of expanding or scaling AI solutions across business units.

Finally, when you are ready to craft an agentic AI plan, map your agent selection directly to a structured decision tree. In doing so, you will take a major step toward solidifying the business case for using agentic AI to future-proof your business, transform operations, and deliver a leading customer experience.

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:

AI Strategy

Envision and develop your agentic AI strategy and business case with an actionable roadmap:

- Shape your agentic AI strategy and execution plan tailored to your unique starting point and business needs.
- Build a credible business case with quantified returns/metrics to help drive investments, executive buy-in, and project funding.

AI Jumpstart

Jump-start your agentic AI journey and navigate AI disruption to your advantage:

- Capitalize on our repeatable and replicable approach to rapidly produce solutions that automate and augment operations to unleash the full value of AI.
- Accelerate the path to value from proof-of-concept to launch to adoption at scale.
- Safely experiment with AI agents and scale adoption of tested use cases across your organization.

AI Workforce

Transform your organization to thrive in an AI environment:

- Unlock the full potential of agentic AI for, and with, your people during the journey.
- Augment your workforce with AI agents, enabling more strategic work and accelerating efficiencies.
- Reshape your workforce and define agentic AI governance to help realize the full value of your investments.
- Offer personalized adoption and upskilling experiences to your workforce, so AI can be infused into everyday work.

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

AI Technology


Build sustainable AI and agentic AI solutions, along with the underlying data infrastructure:

- Use our replicable, proof-of-concept AI factory approach to pilot and scale agents using a breadth of resources that include our Advisory services, KPMG TACO Framework™, advanced tools, established and emerging technologies, and targeted training.
- Streamline the integration of agentic AI frameworks, platforms, and accelerators.
- Help ensure rapid rollout of AI tools and the data required to run them.
- Bolster and enhance your technology infrastructure to support the infusion of AI and AI agents at scale across your organization.

AI Trust

Safely introduce AI agents and scale across the enterprise:

- Manage risk, security, and compliance to enable safe AI and agentic AI rollouts.
- Help ensure your agentic AI solutions are ethical, secure, and compliant by following our Trusted AI framework.
- Deploy agentic AI boldly, transparently, and with confidence by adhering to our 10 ethical AI pillars.



Authors



Swami Chandrasekaran

**Global Head of AI & Data Labs,
KPMG US**

E: swamchan@kpmg.com

Swami is a technology executive with over 25 years of experience in enterprise-scale digital, AI, and cloud transformation. As global head of AI and Data Labs at KPMG, he 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, synthetic data, enterprise discovery and search, and hardware-optimized solutions, while also chairing the KPMG AI Technology Review Board to help ensure trusted and scalable AI adoption.



Prasad Jayaraman

**Partner, Advisory,
KPMG US**

E: prasadjayaraman@kpmg.com

Prasad is a Partner in KPMG Advisory focused on innovation, emerging technology, and enterprise transformation. He works with senior leaders to interpret market and technology signals, translate them into strategic choices, and design operating models for what comes next. His work sits at the intersection of AI, platforms, and organizational change, with an emphasis on helping large enterprises move from experimentation to durable capability. Prasad is known for connecting disparate ideas into clear narratives that inform investment, build conviction, and drive action.



Abhinav Raghunathan

**Manager, Solutions Architect,
KPMG US AI Center of Excellence**

E: abhinavraghunathan@kpmg.com

Abhinav is a lead at the KPMG US AI and Data Labs, where he drives the development of innovative agentic AI applications and services across areas like trust, data, and models. Specializing in trusted and responsible AI, he also runs the Ethical AI Database (EAIDB). Before joining KPMG, he played a pivotal role at a financial services firm, where he developed some of the company's first generative AI applications and published their first technical paper on the subject.

We would like to thank our contributors:

Lisa Bigelow, Donna Ceparano, Tammi Ryley, Joseph Schnering, Rich Schulte, and Lara Volpe

Irish Contacts



Colm O'Neill

**Partner,
Head of Consulting**

E: colm.oneill@kpmg.ie



Owen Lewis

**Partner,
Head of AI & Velocity,
Head of Banking & Capital
Markets, Global Lead
for Financial Services
Cost Transformation**

E: owen.lewis@kpmg.ie



Rory Timlin

**Partner,
AI & Data Practice Lead**

E: rory.timlin@kpmg.ie

Related thought leadership



[The agentic AI advantage:
Unlocking the next level of AI value](#)



[Shadow AI is already here:
Take control, reduce risk, unleash innovation](#)



[Data governance in the age of AI](#)

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

Please visit us:



[kpmg.com](https://www.kpmg.com)



Subscribe

© 2026 KPMG, an Irish partnership 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.

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 endeavour 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 registered trademarks of KPMG International Limited ("KPMG International"), a private English company limited by guarantee.

If you've received this communication directly from KPMG, it is because we hold your name and company details for the purpose of keeping you informed on a range of business issues and the services we provide. If you would like us to delete this information from our records and would prefer not to receive any further updates from us please contact unsubscribe@kpmg.ie.

Produced by: KPMG's Creative Services. Publication Date: January 2026. (JOB NO:12227)