

EXECUTIVE INSIGHTS

What Every Executive Should Know About Agentic AI

The biggest threat facing your business is failing to see disruption coming. Senior executives navigate intricate decisions daily, confronting rapid market shifts, regulatory hurdles and unforeseen operational challenges that conventional technologies simply can't handle. And while over half (57%) of executives surveyed by L.E.K. Consulting¹ recognize artificial intelligence's (AI) potential to reshape their industries, many still rely on outdated tools ill-suited to managing today's realities.

AI has evolved significantly, from simple automation to sophisticated systems capable of understanding context, making autonomous decisions and continuously improving through experience. Agentic AI represents this latest stage of AI evolution, serving as a strategic digital ally that proactively anticipates disruptions, dynamically formulates responses and learns in real time.

Supported by robust governance and transparency, agentic AI enables your business to turn change into a sustainable competitive advantage. But what exactly makes agentic AI different? And how can your organization tap into its unique capabilities?

Understanding why agentic AI outperforms traditional automation

Most executives are familiar with automation: repetitive tasks handed off to software, simple decision trees executed by bots and straightforward business processes managed without human intervention. However, the limitations of traditional automation become increasingly evident when a company is faced with nuanced problems or sudden market disruptions. This is precisely where agentic AI differentiates itself — and where its true value emerges.

At its core, agentic AI operates through a clearly defined four-step cycle that traditional automation can't replicate (see Figure 1).

Perceive

Traditional automation solutions respond only to narrowly defined scenarios. Agentic AI, however, actively gathers and interprets data using multiple advanced AI techniques — including sophisticated natural language processing, computer vision and pattern recognition — to understand not just isolated data points but also the broader context surrounding each decision. This rich contextual awareness is guided by governance protocols that ensure data accuracy and ethical usage.

For example, within supply chain management, agentic AI monitors real-time inventory data across multiple warehouses, instantly detecting shortages or disruptions. Simultaneously, it analyzes contextual factors such as demand trends, regional weather patterns and logistical constraints, enabling truly proactive responses.

Reason

Going beyond basic logic, agentic AI leverages advanced analytical frameworks such as retrieval-augmented generation. For instance, in financial services, it evaluates complex financial data, integrates regulatory updates and formulates compliant, optimized trading strategies — much like a seasoned financial analyst, with built-in transparency to ensure auditability.

However, it's important to recognize that while these systems appear to reason, they do not truly understand in a human sense. Their outputs rely on learned statistical patterns rather than genuine semantic comprehension or self-reflective cognition. As a result, rigorous guardrails, validation processes and human oversight remain essential.

Act

Unlike rigid traditional automation, agentic AI dynamically executes interactions with external systems, always adhering to predefined safety and compliance guardrails. In operations, it autonomously manages high-volume invoice processing and accounts payable tasks, significantly enhancing operational efficiency without manual intervention unless exceptions are required.

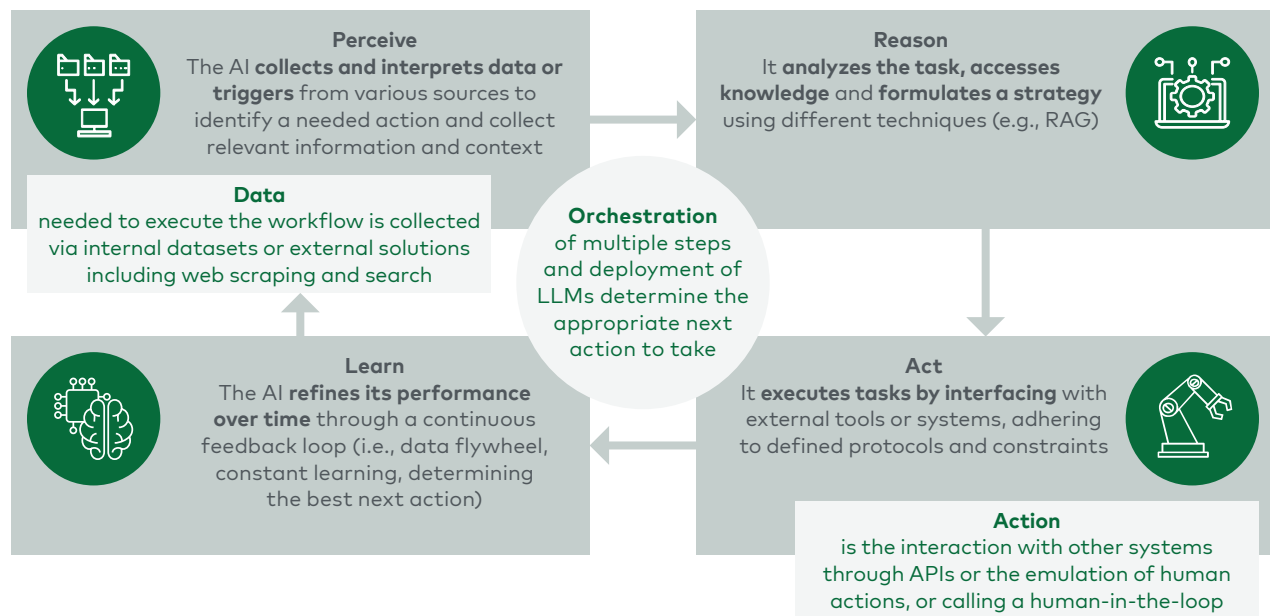
Learn

The most transformative aspect of agentic AI is its continuous learning through feedback loops. For example, in personalized marketing, it evaluates campaign responses to identify effective strategies, then iteratively enhances future campaigns — all while ensuring user privacy and

adherence to ethical marketing standards. Over time, this adaptive learning capability steadily improves marketing accuracy, efficiency and customer engagement. Stand-alone solutions like commerce platforms or finance systems often lack sufficient context to enable effective learning.

As a result, orchestrating data and actions across multiple systems becomes essential. Moreover, most agentic AI deployments today require semimanual retraining and human-curated oversight loops, underscoring the ongoing need for human judgment to manage operational risks.

Figure 1
The four-step cycle of agentic AI



Note: AI=artificial intelligence; LLMs=large language models; RAG=retrieval-augmented generation; APIs=application programming interfaces
Source: OpenAI; NVIDIA; IBM; L.E.K. research and analysis

Furthermore, these three foundational components enable and amplify agentic AI's unique capabilities:

- 1. Data** — Agentic AI harnesses internal datasets and external sources and gathers external data through web scraping, application programming interfaces (APIs) or other ways, always within stringent privacy standards, to produce rich, actionable insights and maintain real-time awareness. The effectiveness of agentic AI, however, directly depends on the quality of these underlying datasets. Ensuring robust data integration, accurate processing and thorough transformation within the data ecosystem — well before data reaches analytics systems — is therefore essential for maximizing AI's performance and reliability.

- 2. Action** — Agentic AI seamlessly executes tasks through secure interactions with external systems via APIs or human emulation, ensuring compliance and data security. Beyond mere automation, its autonomous capabilities extend to dynamic decision-making made capable by generative AI capabilities, scenario-based action planning, and adaptive responsiveness to unexpected changes or disruptions. This ensures operational agility and flexibility, allowing organizations to quickly capitalize on opportunities and minimize risk across diverse business processes.
- 3. Orchestration** — Agentic AI strategically coordinates complex workflows, determines optimal next actions and effectively manages multistep processes, guaranteeing cohesive, goal-driven execution at every stage. Effective orchestration involves managing specialized subagents that handle discrete tasks (e.g., document summarization, web search, knowledge base synthesis) and coordinate seamlessly through context-aware delegation. By continuously monitoring progress and outcomes across multiple workflows, agentic AI ensures alignment with broader strategic objectives, enabling sustained operational consistency, scalability and improved resource efficiency.

Building AI literacy across your organization

Successful implementation of agentic AI isn't solely a technological shift — it demands organizational readiness through AI literacy and a clearly defined, top-down AI strategy.

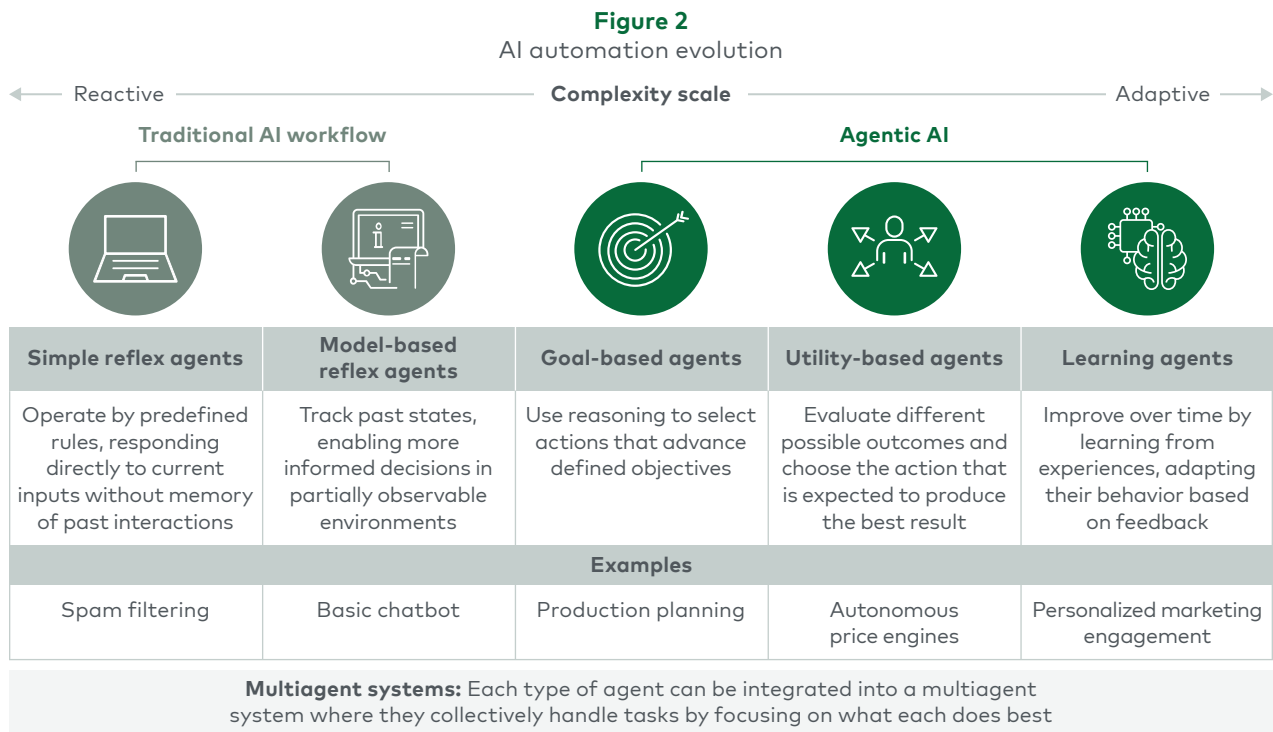
AI literacy involves understanding core concepts, appreciating ethical implications, recognizing AI's capabilities and limitations and identifying opportunities for effective human-AI collaboration.

But even robust AI literacy and strong technological foundations aren't enough without executive buy-in and strategic alignment. A cohesive enterprisewide AI strategy provides a critical framework, aligning departments around a shared vision, prioritizing high-impact initiatives and balancing targeted use cases with broader organizational transformation.

By investing in structured AI education programs, workshops, continuous learning opportunities and healthy cross-team collaboration, organizations empower employees to confidently engage with, manage and innovate alongside AI systems — positioning themselves for sustainable success in an increasingly AI-driven business environment.

Tracing the evolution from basic automation to agentic AI

To fully appreciate the significance of agentic AI, it's helpful to examine how AI has evolved from simple reflex agents to learning agents over time (see Figure 2).



Note: AI=artificial intelligence
Source: OpenAI; IBM; L.E.K. research and analysis

Early AI automation implementations were primarily simple reflex agents (e.g., basic automation systems such as spam filters), which react based strictly on predefined rules. Soon, model-based reflex agents (e.g., basic chatbots) — capable of limited contextual awareness and simple interactions — emerged.

Progress continued with goal-based agents, systems like production scheduling tools that proactively pursue defined objectives. Utility-based agents further advanced the field, optimizing outcomes based on evaluating potential scenarios — exemplified by pricing optimization engines.

Most recently, learning agents introduced a revolutionary capability: continuous improvement through experience, notably used in personalized marketing and recommendation systems. However, agentic AI represents the current pinnacle of this evolution. It combines advanced perception, nuanced strategic reasoning, autonomous action and ongoing self-improvement, delivering a level of adaptability, intelligence and responsiveness previously unattainable by earlier AI generations.

Crucially, this advanced autonomy demands strengthened ethical oversight and robust governance frameworks to responsibly manage the increasingly independent actions of AI systems. Understanding this evolution sets the stage for exploring precisely where agentic AI shines in practical, real-world scenarios.

Exploring practical use cases for agentic AI

Agentic AI's capabilities may sound impressive in theory, but its true value becomes evident through real-world applications. Businesses across diverse sectors are already leveraging agentic AI to achieve substantial operational gains, drive efficiency and enhance customer engagement.

In **customer service**, brands like Uber and Bell Canada employ agentic AI to autonomously manage written customer inquiries, providing accurate responses swiftly and efficiently — all while ensuring transparency in customer interactions and adherence to strict data privacy standards. Similarly, within sales and marketing, companies such as Adobe and Alaska Airlines leverage autonomous AI agents to independently manage personalized customer interactions and optimize outreach efforts — driving higher engagement and conversion rates, with embedded compliance measures to safeguard consumer trust.

Operations and supply chain management also significantly benefit. For instance, BMW and Cradle deploy agentic AI for forward-demand planning and automated supply chain execution, reducing operational costs and enhancing logistical efficiency, supported by clear human oversight protocols for critical decision points. In cybersecurity, firms like Palo Alto Networks and AppOmni use agentic AI to proactively detect threats and analyze malware, swiftly addressing security risks while maintaining strict data security standards and governance protocols.

The **software development** sector equally embraces the power of agentic AI, as seen with companies such as Workato that utilize autonomous agents to optimize processes from software building and testing through telemetry data analysis. This not only accelerates innovation and development cycles but also ensures greater accuracy and quality, guided by robust frameworks that ensure ethical AI deployment and data security.

Clearly, the practical applications of agentic AI span a wide range of industries and operational contexts. The examples above highlight its versatility and demonstrate its potential to profoundly transform organizational outcomes.

But realizing these transformative benefits requires a balanced understanding of both opportunities and risks. When strategically applied, agentic AI can actively mitigate many of today's most pressing business challenges, as illustrated in the following risk-reward framework:

Strategic risk	How agentic AI helps
Market volatility	Offers real-time perception and adaptive responses
Regulatory complexity	Has built-in transparency and compliance logic
Talent shortages	Automates repetitive workflows
Innovation bottlenecks	Enables rapid iteration and scaled experimentation
Siloed systems and insights	Orchestrates cross-functional agent collaboration

To fully capitalize on agentic AI's strategic potential, organizations must also be ready for common implementation challenges. Anticipating and proactively managing these challenges ensures smoother adoption and maximizes return on investment.

Navigating agentic AI implementation

Deploying agentic AI involves proactively managing orchestration complexity, addressing performance and accuracy concerns inherent in early-stage analytical models, ensuring rigorous oversight for safety and reliability, and establishing robust frameworks for data security, privacy and compliance.

Performance and accuracy remain significant considerations. The analytical engines behind agentic AI systems are still developing and are susceptible to errors, false information and poor planning in multistep workflows. Therefore, continuous model improvement is crucial for broader and deeper integration in high-stakes applications, necessitating ongoing transparency and regular performance audits.

Additionally, deploying agentic AI solely within isolated systems (such as a stand-alone customer relationship management system) typically provides limited results. Effective implementation requires integrating multiple data sources to capture comprehensive context and leveraging outcome-based feedback loops to enable continuous learning and optimization.

Orchestration adds another layer of complexity. Successful implementation requires specialized subagents quarterbacked by a central orchestrator agent, capable of delegating tasks effectively. But effective orchestration also depends on teams clearly dividing workflows into tasks that autonomous agents can successfully execute. Leading organizations address these challenges by leveraging robust orchestration platforms offering the context-aware management of multiple agents. Palo Alto Networks, for example, integrates human-in-the-loop validation processes to ensure accuracy and compliance, significantly reducing operational risks.

Safety, reliability and data security are critical, particularly as agentic AI undertakes real-world responsibilities such as executing purchases or managing fraud prevention. Errors in these

systems carry tangible consequences, necessitating robust human oversight and stringent data security measures, especially when sensitive data is involved.

Proactively managing these challenges not only streamlines AI implementation but also directly supports achieving strategic early-adopter advantages. Early adopters like Uber and BMW demonstrate the benefits of integrating mature foundational AI models and have achieved substantial operational agility and scalability underpinned by comprehensive ethical governance frameworks. Conversely, delays in adoption can lead to increased operational inefficiencies, higher costs, severe talent shortages and stifled innovation — as evidenced by historical patterns in sectors like retail and manufacturing.

Ultimately, effectively navigating implementation challenges is about more than mitigating risks; it's about capturing strategic opportunities. Organizations that address these challenges decisively, through transparent, ethical practices and human-centered oversight, position themselves to lead rather than follow, shaping their industries through innovation and proactive adaptation.

Realizing the full promise of agentic AI

Turning strategic ambition into real-world results requires clarity about your organization's best path forward. Organizations must strategically choose between deploying purpose-built agents — ideal for rapid deployment in focused, specialized areas — and more-flexible platform-based solutions, which offer broad, scalable applicability across diverse business functions.

Translating your strategic vision into real-world impact and tangible results requires a structured approach to implementation, anchored in a well-defined AI strategy. This strategy should prioritize the right initiatives and thoughtfully balance targeted use cases with broader organizational transformation and growth. Consider the following steps to guide your organization:

- **Identify high-impact work outcomes for agentic AI:** Identifying the most obvious opportunities to deploy agentic AI requires an understanding of the work outcomes that teams execute. Breaking down roles into discrete tasks and aligning agentic deployment against those outcomes allows for orchestration and extensibility. Successful prioritization of agentic deployment must balance top-down strategic guidance from leadership with bottom-up insights from empowered teams. Together, these perspectives help pinpoint workflows where autonomy, adaptability or scale deliver measurable, strategically relevant business value, ensuring alignment across the entire organization.
- **Choose your entry path:** Selecting the best starting point for agentic AI adoption depends on your organization's specific context, immediate objectives and long-term vision. Clearly

evaluate whether purpose-built agents, designed for targeted impact in specific workflows, or broader platform-based solutions, suited for scalable implementation across multiple functions, align more closely with your strategic goals and internal capabilities.

- **Assess organizational readiness:** Implementing agentic AI effectively requires an honest evaluation of your technical and organizational foundations. Carefully consider your data ecosystem, including integration and transformation capabilities, the presence of a unified data platform, API connectivity, specialized AI expertise and the maturity of your governance structures. Beyond the technical and data readiness, ensure that your teams are ready to become AI managers and think about work and tasks in a way that can be abstracted to AI agent co-workers.
- **Prioritize data quality and management:** Successful agentic AI depends fundamentally on robust, well-managed data. Executives should ensure the organization follows structured data management practices outlined in industry standards such as DAMA-DMBOK, the authoritative guide for data management and governance. Adhering to its recommended domains — including data quality, metadata management, data integration and governance — helps establish the strong data foundations required to fully unlock agentic AI's potential.
- **Pilot strategically:** Rather than deploying broadly all at once, begin with well-defined pilot programs. Clearly establish pilot objectives, integrate robust measurement frameworks to monitor effectiveness and include proactive human oversight and transparency measures. By iteratively refining your approach based on pilot insights, you can confidently scale agentic AI initiatives with significantly reduced risk.
- **Scale intelligently:** Transition from isolated pilot projects into broader coordinated ecosystems using robust orchestration frameworks. Successfully scaling agentic AI demands thoughtful integration across diverse business functions, continuous ethical review and compliance monitoring, and active development of internal capabilities to manage multiple autonomous agents. This deliberate approach ensures long-term strategic success and sustainable growth.

To move from strategic intent to impactful action, the next step is choosing the best entry point for your organization's specific context.

Strategic entry paths for agentic AI adoption

Organizations looking to adopt agentic AI strategically have distinct paths, each aligned to different operational needs and strategic priorities:

- **Purpose-built agents** — These specialized agents are optimal for organizations seeking immediate, targeted impacts in specific areas such as customer service automation,

cybersecurity or supply chain management. They offer rapid, out-of-the-box deployment with straightforward integration, enabling a quick return on investment and minimal disruption to existing systems. However, their specificity can limit their scalability in other use cases, potentially leading to fragmented AI strategies over time.

- **Agentic platforms** — Designed for enterprises pursuing flexibility and extensive scalability across diverse business functions, agentic platforms provide comprehensive frameworks to build, customize and manage multiple AI agents cohesively. While the implementation of these platforms involves more complexity and requires robust technical and organizational readiness, the resulting infrastructure significantly enhances the organization's agility and innovation capabilities. Supported by rigorous ethical governance and transparent operational oversight, these platforms ensure responsible and sustainable agentic AI integration.

Determining the right entry path demands careful consideration of the organization's immediate goals, long-term AI vision, internal capabilities and readiness for comprehensive governance and oversight.

Your next steps with agentic AI

The shift from traditional automation to fully autonomous agentic AI isn't merely technological — it's strategically essential. Agentic AI is increasingly enterprise-ready and is supported by mature technologies, a growing vendor ecosystem and accelerating enterprise demand. Embracing agentic AI today unlocks tangible competitive advantages such as greater agility, reduced costs and accelerated innovation. Delaying, however, could amplify operational inefficiencies and erode market position. In innovation-driven sectors, delayed adoption may significantly heighten competitive risks, while legacy organizations risk becoming further entrenched in outdated processes and frustrated by diminishing returns.

The choice is clear: Take strategic control of your future by investing decisively in agentic AI, building robust governance and scaling ethically and confidently. Don't just adapt to change — define it.

We would like to acknowledge Shivam Sharma, Hunter Reynolds and Tess Wrigley for their contributions to this piece.

For more information, please [contact us](#).

Endnote

¹LEK.com, "AI Has Moved Beyond Hype — It Is Driving Performance and Deployment Is Widespread." <https://www.lek.com/insights/dig/us/ar/ai-has-moved-beyond-hype-it-driving-performance-and-deployment-widespread>

About the Authors

**Chuck Reynolds**

Chuck Reynolds is a Managing Director and Partner in L.E.K. Consulting's Boston office and a member of the Digital practice. Chuck has extensive experience in digital strategy across various areas, including digital commerce, customer engagement, agile, direct-to-consumer sales, data, and the application of artificial intelligence/machine learning. He has particular expertise in digital strategies that enhance growth and profitability through customer engagement.

**Pawel Jablonski**

Pawel Jablonski is a Partner in L.E.K. Consulting's Dubai office. Pawel brings over 17 years of strategy consulting experience, serving as a bridge between business and technical executives to ensure that technology solutions are seamlessly aligned with business objectives. He specializes in delivering highly innovative and transformative projects for top-tier clients, particularly in financial services (retail and corporate banking, regulatory), as well as in government, military and defense.

**Darren Perry**

Darren Perry is a Managing Director based in L.E.K. Consulting's Boston office and is Head of the Digital practice. He has more than 20 years of technology and strategy consulting experience, is a founding member of L.E.K.'s Technology practice and currently serves on the firm's Global Board of Directors. Darren advises clients across all sectors, with particular expertise in areas such as AI transformation, disruptive new product strategy and digital ways of working.

About L.E.K. Consulting

We're L.E.K. Consulting, a global strategy consultancy working with business leaders to seize competitive advantage and amplify growth. Our insights are catalysts that reshape the trajectory of our clients' businesses, uncovering opportunities and empowering them to master their moments of truth. Since 1983, our worldwide practice — spanning the Americas, Asia-Pacific and Europe — has guided leaders across all industries, from global corporations to emerging entrepreneurial businesses and private equity investors. Looking for more? Visit www.lek.com.

L.E.K. Consulting is a registered trademark of L.E.K. Consulting LLC. All other products and brands mentioned in this document are properties of their respective owners. © 2025 L.E.K. Consulting LLC