CLOUDERA

SURVEY REPORT

The Future of Enterprise AI Agents

Unlocking Autonomous Transformation in 2025

Table of Contents

Introduction 3

AI Agents: Ushering in the Future of Enterprise Technology 4

- Key Business Functions: Where AI Agents Add Value 5
- Approaches to Implementation 5
- Open Source vs. Closed Source 6

What's Holding Enterprises Back? 7

- Bias, Fairness, Ethics, and AI Governance 8

Industry-Specific Trends: AI Agent Adoption 9

- Finance and Insurance 9
- Manufacturing 10
- Retail and E-Commerce 10
- Healthcare 11
- Telecommunications 11

Conclusion: From Emerging Trend to Critical Enterprise Tech 12

- Recommendations for 2025 Agentic AI Adoption 12

About Cloudera 13

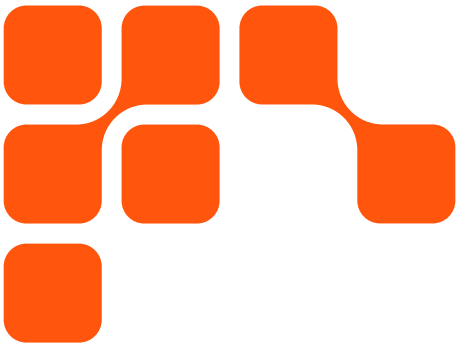

INTRODUCTION

Enterprise workflows are leveling up in 2025. Thanks to advances in generative AI (GenAI), large language models (LLMs), and natural language processing (NLP), agentic AI is transforming the way enterprises approach automation and decision-making, impacting everything from customer interactions to business operations.

AI agents are autonomous software systems that can reason, plan, and act on behalf of users. Many consider AI agents to be the next evolution of chatbots. Both use AI to receive input and take action based on that input; however, chatbots follow a predetermined workflow and are limited to handling finite scenarios and user inputs. AI agents—whether model-based, goal-based or multi-system—are more interactive and capable of handling complex tasks by applying reason and determining the best course of action autonomously.

When implemented correctly, agentic AI offers tremendous benefits such as increased efficiency, reduced costs, improved customer experience, and smarter (real-time, data-driven) decision-making. For CIOs and CTOs looking to drive innovation, agentic AI can accelerate their efforts.

To explore how organizations are investing in agentic AI, Cloudera surveyed 1,484 enterprise IT leaders across 14 countries. This report, fielded in February 2025, dives deep into adoption patterns, use cases, and sentiments around AI agents–including industry-specific insights across finance, retail, healthcare, manufacturing, and telecommunications—and looks at how enterprises are using AI agents to level-up their enterprise workflows in 2025.

AI Agents: Ushering in the Future of Enterprise Technology

The last two years have seen an adoption boom of agentic AI. A majority of respondents (57%) started implementing AI agents within the last two years, and 21% did so in the last year alone. This adoption rate reflects how quickly agentic AI has moved from concept to reality, likely catalyzed by advancements in AI technologies between 2023 and 2024.

This rapid rate of adoption shows that organizations view agentic AI as essential to their business competitiveness and hope to capitalize on their return on investment (ROI) sooner rather than later.

Consequently, expansion plans for agentic AI are nearly universal. An overwhelming 96% of respondents plan to expand their use of AI agents in the next 12 months, with half aiming for significant, organization-wide expansion.

As enterprises move agentic AI from pilot stage to broad deployment, they're exploring specific applications, with respondents most interested in performance optimization bots (66%), security monitoring agents (63%), and development assistants (62%).

83%
of organizations believe it's important to invest in agents to maintain a competitive edge within their industry.

2025
is a pivotal window for AI agent adoption as many companies move from experimentation to execution.

Performance optimization bots may take the form of an IT infrastructure bot that can dynamically adjust cloud resource allocation, database configurations, and server loads to optimize performance in real time. Security monitor agents can continuously analyze network activity, detect anomalies and autonomously respond to potential cyber threats. Development assistants can take the form of bots that generate, execute and refine test cases based on real-time code changes.

IT leaders are looking for these improvements in the AI agents they use:

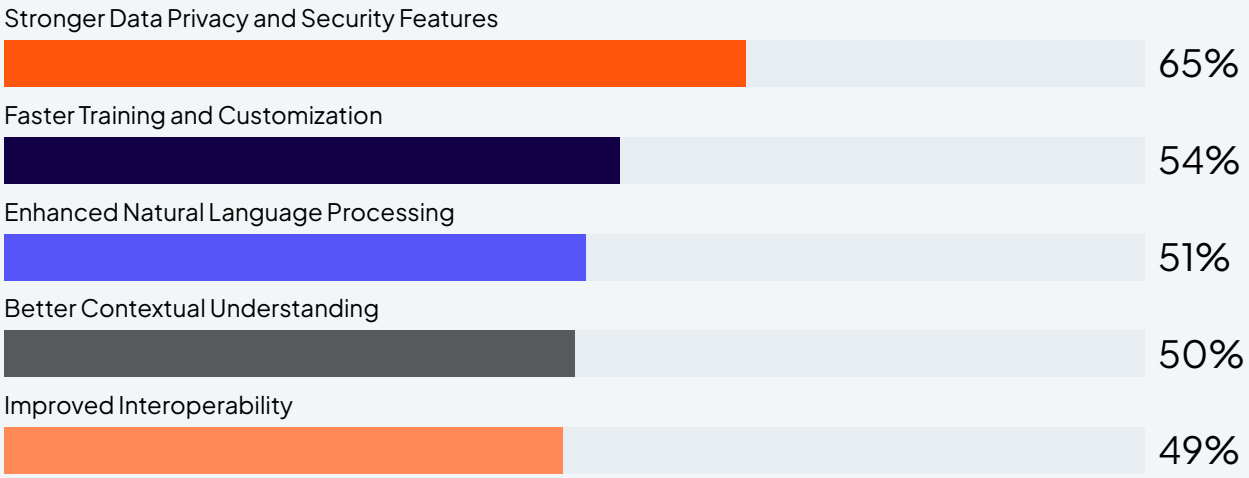

Key Business Functions: Where AI Agents Add Value

The synergy between AI agents and GenAI paves the way for investments in agentic AI to deliver immediate ROI. 98% of surveyed organizations are either already using agentic AI to orchestrate GenAI use cases or plan to do so in the near future. In fact, 85% say their prior investments in GenAI have prepared them well to implement AI agents. Enterprises should view agentic AI as a natural next step to capitalize on their GenAI investments, especially as agentic AI vendors increase the availability of out-of-box API integrations.

Early deployments of AI agents tend to focus on IT and customer-facing operations, according to respondents. When it comes to business functions, AI agents are most embedded in IT operations (61%), but customer support (18%) and marketing (6%) are also key areas of adoption. In terms of concrete applications, AI agents are most used for customer support (78%), process automation (71%), and predictive analytics (57%). For enterprise companies that have embedded agents into IT operations, they are most likely to also use agents for customer support, operations, and marketing, respectively.

These top use cases show that many companies start adoption in well-defined, ROI-driven domains, and internal functions, such as IT helpdesk agents and DevOps assistants. Both areas provide ample opportunity for automation to quickly augment human effort, driving tangible results. For instance, agentic AI can help enterprises stay ahead of cyber attackers by leveraging predictive analytics to foresee potential attack vectors and suggest mitigation strategies before a threat fully manifests.

Approaches to Implementation

How are enterprises building agentic AI? 66% of respondents are using enterprise AI infrastructure platforms to develop and deploy their AI agents, and 60% are taking advantage of agent capabilities embedded within their existing core applications. These preferences underscore enterprises' need for scalable, reliable infrastructure, as they opt to establish their AI agents as close to private data as possible.

81%

of enterprises are leveraging agentic AI to enhance their existing GenAI models.

AI agents are most commonly used for:

78%
Customer Support

71%
Process Automation

57%
Predictive Analytics

What's next?

65% of enterprises want to see stronger data privacy and security features in their AI agents.

When asked which technologies they currently use or plan to use to build agentic AI, respondents identified the following:

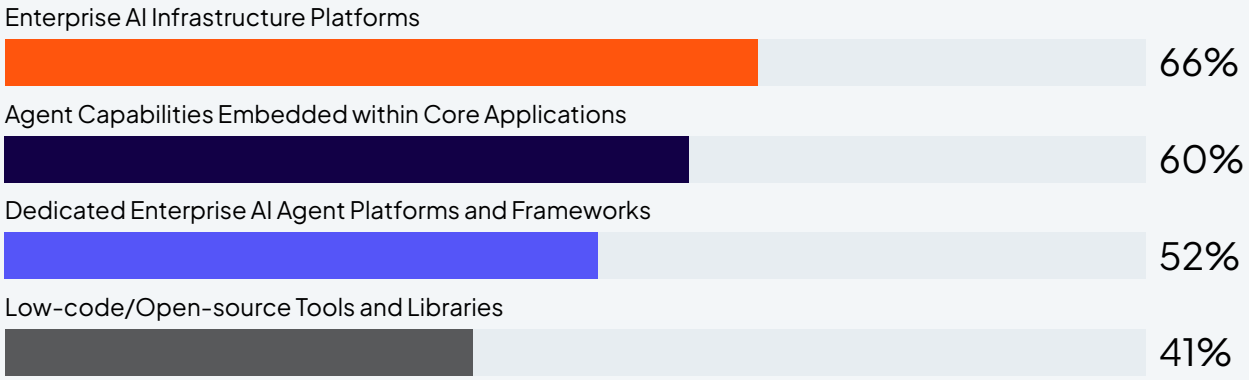

Open Source vs. Closed Source

A key trend in the evolution of agentic AI is the rapid rise of open-source LLMs. Once considered to be lagging behind their closed-source counterparts, today's open models are now rivaling—and in some cases surpassing—the performance of proprietary systems in crucial enterprise use cases.

Recent breakthroughs from communities working on models like Llama, Mistral, and DeepSeek have shown that open-source LLMs are a more cost-effective route for enterprises to take. For enterprises operating on a large scale, the economic implications are significant: if a proprietary model provides only marginal gains (for example, a 1% improvement in benchmark tasks) but comes at 10 times the cost, the business case simply does not hold up.

Additionally, open-source models provide enterprises with advantages that proprietary vendors often cannot match: deployment flexibility. Most closed models are tied to specific public clouds or APIs, which can create challenges related to data sovereignty, vendor lock-in, and integration with existing infrastructure. In contrast, open models can be self-hosted or deployed in private or hybrid cloud environments, making it easier to align with enterprise security and compliance requirements.

From cost efficiency to greater control over data and deployment, the strategic advantages of open-source models are clear.

What's Holding Enterprises Back?

Despite enthusiasm, enterprises face notable hurdles in adopting agent AI — ranging from implementation difficulties to ethical and organizational concerns.

When asked to rank their top concerns in adopting AI, respondents pointed to data privacy concerns (53%), followed by integration with existing systems (40%) and high implementation costs (39%). These findings illustrate that trust and compatibility issues are primary roadblocks, as enterprises worry about safeguarding sensitive data and transforming legacy environments.

To be fair, 37% of surveyed enterprises report that integrating AI agents into current systems and workflows has been very or extremely challenging. This finding points to integration as a pain point for large organizations with complex IT ecosystems. In other words, deploying AI agents is not a plug-and-play endeavor.

For this reason, deploying and managing agentic AI demands skilled professionals and proper infrastructure. The most effective way for organizations to begin leveraging agentic AI is by evaluating their existing infrastructure to ensure it meets the necessary requirements, focusing on data management, security, and compliance standards. Equally important is training teams to effectively manage and deploy AI agents, starting with small-scale implementations to assess their impact before expanding on a larger scale.

> **My highest priority for expanding AI in the organization is to help with the integration of data security such as privacy, security, and staff adaptation training."**
>
> SURVEY RESPONDENT

Key challenges in adopting AI agents:

53%
Data Privacy Concerns

40%
Integration with Existing Systems

39%
High Costs

34%
Lack of Expertise

32%
Ethical or Regulatory Concerns

30%
Governance Concerns

Bias, Fairness, Ethics, and AI Governance

Agentic AI has the potential to transform user experiences and business operations, but it also carries the risk of undermining consumer and employee trust.

When AI is trained on historical data, it can reinforce societal biases unintentionally and influence outcomes. AI bias has already negatively impacted real-world workplaces:

Healthcare: A Cautionary Case Study in Lifecycle Bias

A recent Yale study underscores the depth and pervasiveness of bias in medical AI. Researchers found that algorithmic bias can enter at every phase of the AI lifecycle—from data curation and model design to implementation and post-deployment use. One critical insight: diagnostic systems trained on non-diverse datasets underperformed on underrepresented populations, leading to misdiagnosis and compromised care. In high-stakes fields like healthcare, even minor skew in model training can create systemic inequity at scale. This is a clear mandate for organizations to prioritize data diversity, embed lifecycle auditing, and enforce transparency protocols in every model pipeline.

Defense: When Algorithmic Bias Meets Strategic Decision-Making

The risks compound in the defense sector. As detailed by the International Committee of the Red Cross, AI-powered military decision support systems can encode and amplify bias in ways that directly influence tactical and strategic operations. Bias in scenario simulation or system design can shape lethal outcomes, especially when human oversight defers to machine-driven inference. Critically, the ICRC identifies four distinct points of failure: bias in training data, model architecture, operational use, and post-use evaluation. These insights are directly applicable to enterprise AI: when agents are deployed in high-impact environments, clarity around bias, accountability, and auditability is not optional—it's imperative.

Four Principles of Trustworthy AI

Transparency

Fairness

Accountability

Security

Over half of respondents (51%) have significant concerns about AI bias and fairness. For enterprises scaling AI agents, the lesson is clear: responsible AI is not just a governance function—it's a design principle. Bias is not confined to flawed data; it can manifest in how workflows are structured, how intent is interpreted, and how outcomes are evaluated.

To establish trust, organizations must prioritize data quality, ensure model robustness, and adopt ethical decision-making practices. This involves thoroughly testing AI models to eliminate biases, implementing strong data governance and security measures, and conducting regular audits to maintain trust throughout the AI system's lifecycle.

Garbage In, Garbage Out

Data quality and availability issues are significant technical challenges in AI implementations, more so than algorithm accuracy or infrastructure limitations.

My biggest challenge is to make sure the algorithms and systems don't have biases that could harm the company."

SURVEY RESPONDENT

To combat bias, enterprises are taking additional steps to govern AI agents responsibly. Many respondents (38%) are instituting multiple processes, such as human review, diverse training data, and formal fairness audits. Another 36% have introduced some bias-check measures, such as periodic human reviews or bias-detection tools. That said, a notable minority (14%) admit they've only taken minimal or ad hoc steps so far to combat bias. This shows that while many organizations are formalizing AI ethics, others are just beginning.

Encouragingly, though, 80% of organizations are extremely or very confident in the transparency and explainability of their AI agents' decisions, suggesting this is a priority for enterprises, their vendors, and their teams.

Industry-Specific Trends: AI Agent Adoption

While the full survey includes over 1,400 enterprise respondents, the industry-specific insights highlighted here reflect a focused subset of participants from key sectors: manufacturing, finance and insurance, retail, healthcare, and telecommunications.

While some trends in agentic AI adoption are universal, there are nuances by industry. For example, when looking at barriers to adoption, data privacy is a top worry across sectors. But when breaking it down by industry, highly regulated industries such as finance and healthcare show a greater concern for compliance and transparency, whereas the technology sector is more concerned with integration complexity and talent gaps. Each industry must address a unique mix of obstacles—technical, organizational, and ethical—when rolling out AI agents.

Finance and Insurance

Organizations in the finance and insurance industries are most focused on using agentic AI for risk mitigation, compliance, and client service automation. Survey respondents selected the top use cases as fraud detection (56%), risk assessment (44%), and investment advisory (38%).

Financial institutions, in particular, are leveraging agentic AI to strengthen security and offer smarter advisory services. For instance, a critical component of AI trust is authorization, meaning that AI systems can only access the data they are permitted to use. The challenge becomes even greater when AI agents pull data from multiple sources to generate more complex responses. For example, if a high-net-worth individual asks an AI-powered advisor to structure an asset class that balances medium risk with high returns, the AI must not only analyze vast amounts of data but also ensure that it is authorized to retrieve and process each piece of information.

> **"The biggest challenge is implementing an AI agent in a way that enables that touch of personalization and care you typically expect from human agents. That's the challenge facing many customer support teams today."**

SURVEY RESPONDENT

Regional Trend

Financial institutions in North America and Europe are leveraging AI fraud detection more than other parts of the world.

Manufacturing

The manufacturing sector is embracing AI agents to support operational efficiency and supply chain optimization. In fact, roughly half of manufacturing organizations are exploring AI agents for supply chain optimization (49%), process automation (48%), quality control (47%). AI agents have the capacity to monitor production lines for defects or to intelligently reroute supply chain logistics when disruptions occur. While nearly half of responses focused on supply chain optimization, the potential for agentic AI to support autonomous decision-making in production and operations is an emerging area to watch. Agentic AI has the potential to fuel Industry 4.0 objectives, directly impacting enterprises' bottom lines.

One of the most important challenges to solve in oil and gas operations is how to properly monitor safety risks at oil rigs. Personnel on offshore oil and gas sites are far more likely than the average American worker to die on the job. Traditionally, health and safety teams rely on manual inspections, sending contractors on-site to assess risks and determine if conditions require intervention. This process is not only time-consuming but also error prone, as incidents continue to occur despite rigorous protocols. AI presents a transformative opportunity here. By leveraging AI-driven agent workflows for Health, Safety, and Environment (HSE) management, organizations can analyze historical data, detect patterns, and predict potential hazards before they escalate. Instead of relying solely on human inspections, AI can automate risk assessments, reducing the need for physical site visits and ultimately saving lives.

Regional Trend

AI-driven supply chain optimization is a top priority in the UAE and Indonesia especially, according to respondents.

In manufacturing, respondents are exploring the following AI agent applications:

Retail and E-Commerce

In retail and e-commerce, AI agents are primarily used for customer service (50%), price optimization (49%), and demand forecasting (48%). Combined, there is an enormous opportunity for agentic AI to fuel personalization and richer customer satisfaction across shopping experiences.

For instance, AI agents can analyze a customer's browsing history, preferences, and past purchases to recommend products tailored to their tastes and needs and increase the likelihood of purchase. AI can also curate personalized content, such as special

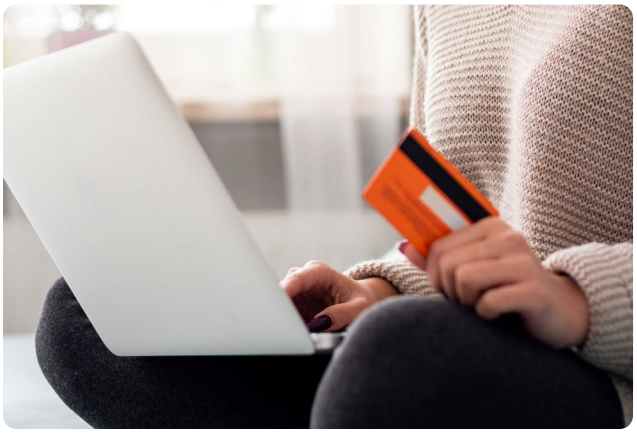

offers, emails or ads, based on a customer's browsing behavior and past interactions. These customized approaches not only nurture consumers along the sales funnel, but also free up human workers for more strategic initiatives.

Regional Trend

Brazil, the US, and the UK lead the way for AI-driven customer engagement strategies in retail.

In retail, respondents are exploring the following AI agent applications:

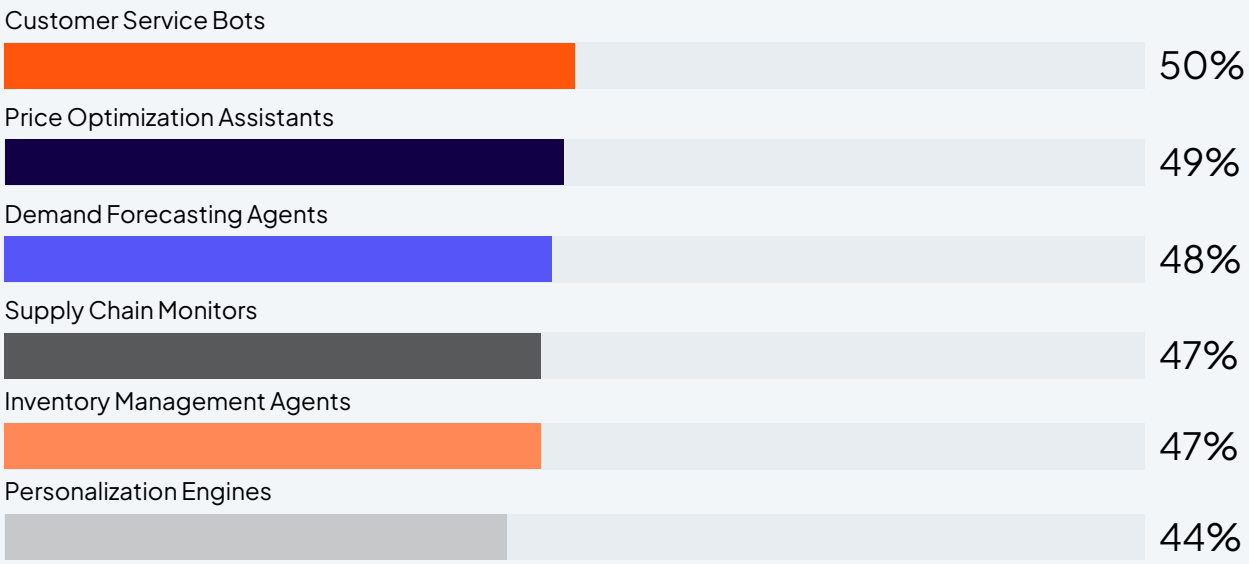

Healthcare

Today, appointment scheduling (51%), diagnostic assistance (50%), and medical records processing (47%) are the top use cases for agentic AI in healthcare. AI agents have immense potential to transform the healthcare industry by improving patient care, streamlining administrative tasks, and assisting providers in decision-making. This would ultimately drive efficiency and better outcomes for patient care.

For instance, AI agents can assist medical professionals in diagnosing and treating patients by providing evidence-based recommendations. A diagnostic assistance agent trained on thousands of X-ray images might detect early signs of pneumonia or lung cancer that are not immediately visible to the human eye and highlight where the radiologist should examine more closely. In this way, agentic AI can help physicians diagnose conditions more accurately and quickly.

Telecommunications

The teleco sector shows strong interest in AI agents for customer-facing and network operations use cases. Respondents are particularly exploring customer service and experience agents, with 49% looking into customer support bots and 44% looking into customer experience agents. Another 49% are exploring security monitoring bots, reflecting the importance of network security and fraud detection in this industry.

AI agents can leverage historical data—such as usage patterns, billing history, and customer support interactions—to predict which customers are at risk of leaving and why. For instance, a teleco company might use AI to identify customers who have reduced their monthly plan usage or who have contacted customer service multiple times within recent months with unresolved issues. The AI agents can flag customers as "at-risk" for churn, allowing the company to intervene early.

CONCLUSION

From Emerging Trend to Critical Enterprise Tech

2025 is the year of agentic AI, and the time to invest is now. Enterprise IT leaders are eager to harness the power of AI agents to drive efficiency, customer engagement, and innovation across their organizations.

As all industries race toward an even more autonomous age, early adoption is imperative to avoid falling behind competitors. Most companies are already deploying agents or poised to do so within the year.

Amidst enthusiasm, there are still critical challenges around data privacy and ethical AI that organizations must address. IT leaders must prioritize trust, transparency, and seamless integration to realize the full benefits of agentic AI. Here is how to start making headway:

Recommendations for 2025 Agentic AI Adoption

Strengthen your data foundation and integration capabilities. The top barriers—data privacy, integration, and data quality—all point to the need for robust data management. Enterprises should ensure they have modern data architecture and unified platforms that can securely handle the volume and variety of data that AI agents require. A unified data platform can simplify the integration of AI agents by bringing data and AI together and mitigate privacy and security risks by providing consistent controls.

Start with high-impact projects to deliver immediate ROI, and grow from there. Begin with AI agent projects that have clear value and relatively contained scope (for example, an internal IT support agent). The high-value, low-complexity that are fast-to-value will help prove the ROI, learn from those implementations, and then scale up. Survey respondents focused on customer support and process automation as initial use cases, suggesting these areas are good launch pads because they address real pain points and have measurable outcomes. Don't attempt a high-risk use case first, as it could create irreparable damage. After an initial project, organizations can iterate and add more complexity.

Establish accountability first, and then build governance and ethics frameworks. As organizations deploy agentic AI, the first and most critical step is assigning accountability. AI agents don't just surface information—they make decisions, act on behalf of users, and influence outcomes. That's why enterprises must clarify: who is responsible for an agent's performance? Is it the developer who built it, the business owner who uses it, or the operations team that oversees it? For example, if a sales agent recommends leads, someone must be accountable for verifying whether those leads are qualified. Without this layer of ownership, agentic systems risk becoming black boxes that no one truly governs—and trust erodes quickly. Once accountability is established, enterprises must build robust governance and ethics frameworks to support safe, transparent, and equitable AI operations. These frameworks should include mechanisms to audit bias, ensure transparency in agent decision-making, and regularly review agent behavior against enterprise policies and user expectations. Governance should also reflect the full lifecycle of agent development—acknowledging that bias can enter not only through data, but also during design, deployment, and even post-use review.

Upskill teams and foster a culture of human-AI collaboration. The successful adoption of agentic AI depends on teams that can keep pace with a rapidly shifting landscape. This means going beyond basic training to cultivate hybrid skill sets—people who can not only build and integrate AI agents, but also understand their reasoning, limitations, and evolving capabilities. Enterprise teams must be equipped to implement change at speed. With new agentic frameworks, orchestration tools, and reasoning models emerging weekly, adaptability is critical. Organizations should prioritize hands-on, continuous learning, encouraging experimentation and knowledge-sharing across roles.

Learn More

In the not-so-distant future, agentic AI will be woven into the fabric of every data-driven organization. Enterprises that lead the charge will turn the promises of autonomous AI agents into practical, transformative outcomes.

To learn how Cloudera - the data and AI intelligence platform for today and tomorrow - can empower your enterprise with agentic AI, click here.

About Cloudera

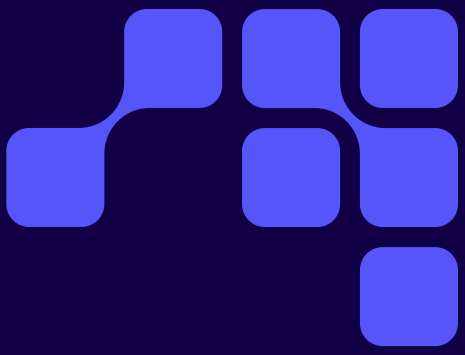

Cloudera is a hybrid platform for data, analytics, and AI. With 100× more data under management than other cloud-only vendors, Cloudera empowers global enterprises to transform data of all types, on any public or private cloud, into valuable, trusted insights. Our open data lakehouse delivers scalable and secure data management with portable cloud-native analytics, enabling customers to bring GenAI models to their data while maintaining privacy and ensuring responsible, reliable AI deployments. The world's largest brands in financial services, insurance, media, manufacturing, and government rely on Cloudera to use their data to solve what seemed impossible—today and in the future.

To learn more, visit Cloudera.com and follow us on LinkedIn and X. Cloudera and associated marks are trademarks or registered trademarks of Cloudera, Inc. All other company and product names may be trademarks of their respective owners.

CLOUDERA

Cloudera, Inc. | 5470 Great America Pkwy, Santa Clara, CA 95054 USA | cloudera.com