

Empowered Intelligence: The Impact of AI Agents

Publication date:

February 2025

Author:

Eden Zoller, Chief Analyst, Applied AI

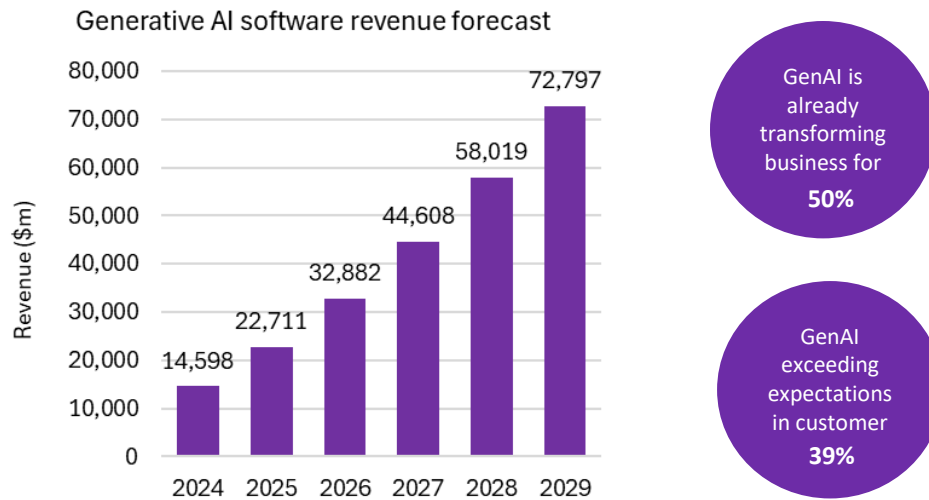
The rise of generative AI

Demand for generative AI (GenAI) solutions is high, and Omdia forecasts that GenAI software revenues will grow from \$6.8billion in 2023 to \$73bn by 2029, a CAGR of 38% (Figure 1).

In terms of market maturity, GenAI today has surpassed where the AI market was five years ago with predictive AI. According to Omdia's 2024 IT Enterprise Insights Survey, 36% of respondents have already deployed GenAI or are in the process of doing so, while 27 % are piloting the technology.

Overall, enterprises appear satisfied with the contribution of GenAI to the business: in Omdia's 2024 GenAI Enterprise Early Adopters Survey, half the respondent say that GenAI has already had a positive impact on their business. GenAI is also exceeding expectations across a range of business goals, notably improving customer support (39%), and enhancing business intelligence (26%).

Figure 1. GenAI is experiencing rapid growth, with positive impacts for users



Source: (Left) Omdia AI Software Market Forecasts –2H24
(Right) Omdia Gen AI Enterprise Early Adopters Survey, n=356

From AI assistants to AI agents: an innovation continuum

A key strength of GenAI is its capacity to adapt to a wide range of applications across diverse use cases with high impact results, with one of the most significant being the way that GenAI has transformed the pre-GenAI wave of AI assistants and chatbots. First wave AI assistants were rule-based and rely on predefined workflows and preprogrammed responses and narrow natural language understanding (NLP). Their capabilities were typically limited to specific tasks like answering FAQs, executing simple commands, or retrieving predefined information. GenAI assistants, powered by sophisticated large language models (LLMs), have improved traditional AI assistants by enabling them to better understand context and generate nuanced, human-like responses, often in near-real time. This makes GenAI assistants more versatile, and allows for more engaging, personalized multi-turn conversations.

GenAI assistants are now moving into the next phase of their evolution, driven by agentic AI technology. While agentic AI leverages LLM capabilities, this is augmented by several supportive technologies and methodologies such as a sense of memory, self-reflection, and tool use, which together enable agentic AI systems to seemingly operate with complete autonomy, or close to it.

Humans set goals for AI agents, but the agent determines how to achieve the goal, proactively performing supportive tasks with little or no human intervention. AI agents can react to changes in their environment (more ably if environments are well defined), responding dynamically and adapting in ways that enable them to effectively complete often complex tasks more effectively than their human counterparts. Unlike traditional AI assistants, AI agents can interact with and make use of external datasets, applications, even other agents and things in the physical world. Compared to a basic query

and response chatbot, AI agents are often built to retain data from interactions and tasks, which provides improved contextual and experiential understanding.

AI agents can complement each other within a multi-agent framework, provided there is a robust orchestration layer supporting it. For example, in the customer support context, an AI assistant could be optimized for more straight forward, level one support needs (e.g. product queries), while another AI agent might take over complex level two tasks (e.g., dispute resolution).

AI agents unlock wide-ranging benefits

Agentic AI has the potential to drive significant advantages across enterprise operations that are aligned with key investment drivers for GenAI (see Figure 2). This means that agentic AI will have a strong strategic fit within an enterprise's overarching goals for GenAI, enhancing capabilities and desired business outcomes. This should drive faster time-to-value and improve overall return on investment (ROI) by enabling companies to leverage existing GenAI expertise to tackle more complex use cases.

Figure 2. Key drivers for GenAI investments



Source: Omdia GenAI Enterprise Early Adopters Survey, n =356

Enhanced operational efficiency and cost savings

AI agents can automate tasks and optimize processes in highly adaptive, focused ways, resulting in faster and more efficient operations. This saves resources and frees human workers to focus on more creative, and/or higher value, strategic activity. Taken together, this can translate into cost savings for enterprises, amplified by the ability of agentic systems to adapt to unforeseen situations within the defined boundaries of their environment.

For example, AI agents could be used to autonomously manage and optimize inventory levels, predict demand, and reroute logistics based on real-time conditions. HR departments could use AI agents to completely automate routine tasks for HR professionals while empowering employees with personalized self-service support (e.g., real-time task guidance, answering complex questions, tailored training).

Improved performance and personalisation for a better user experience

AI agents are capable of superior, more sophisticated performance than traditional, one-dimensional AI assistants, particularly their ability to manage and execute complex tasks. This is because AI agents can engage with multiple assets, retain knowledge and experience. Notably, the ability of AI agents to collaborate is a unique feature that contributes to performance gains over one-dimensional chat assistants, for example by collaborating to synthesize data and perform interdependent tasks. Moreover, an AI agent can apply skills it has used in one context and apply them to other processes. Taken together these capabilities mean that AI agents can provide responses and handle tasks in ways that are more adaptable, accurate, contextually aware and personalised, which means a better user experience. This is particularly important in sales and marketing, and customer support use cases where customer expectations are high. A pleasing user experience is also crucial for enterprise employees using AI agents, as this can increase work satisfaction.

Sharper business intelligence and decision making

Agents can process and analyse vast amounts of data, often in near-real time, providing actionable insights on the fly. Real-time, data-driven insights help enterprises make informed decisions quickly, empowering them to be more agile. Enterprises in this position are better placed to capitalise on opportunities or mitigate risks as they arise. For example, in the financial services context AI agents could analyse real-time market data, predict trends, and make investment recommendations more efficiently than non-agentic systems, allowing enterprises to react to market changes faster than human analysts. In the procurement domain, AI agents could provide an efficient, fast means of vetting potential suppliers, collating and analysing data relating to a large number of suppliers.

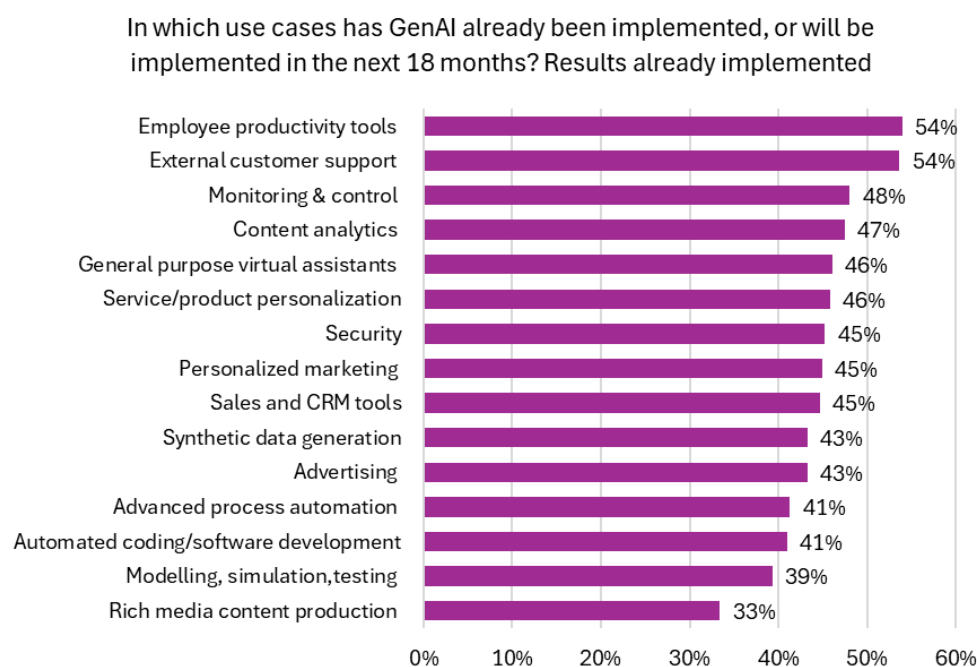
Fresh opportunities for innovation

Agentic AI has the potential to foster fresh approaches to product/service development and innovation, which has the potential to unlock new revenue streams and strengthen competitive advantage. AI agents can improve upon existing techniques used in the research and analysis phase of new product/service development, such as retrieval-augmented generation (RAG). In the case of agentic AI, two or more agents could interrogate multiple data sources to identify, and analyze parameters like customer feedback, and competitor offerings. Agents can use these insights to recommend new features or improvements aligned with goals set by the human product/service lead. This accelerates the innovation cycle. There is also potential for AI agents to both improve and streamline A/B testing across marketing, product development, and customer services. This allows enterprises to experiment, fine tune and iterate at speed.

AI agents can power multiple use cases

GenAI assistants have already been welcomed by enterprises with high levels of implementation (see Figure X) that span a variety of use cases, from customer care (e.g. customer care chatbots), employee productivity tools that help workers in their daily tasks, through to general purpose assistants capable of handling a variety of broad-based functions. AI agents will be able to enhance all the use cases currently supported by AI assistants, thanks to the ability of agentic AI to automate even complex tasks while handling those processes with little or no supervision. The use cases where AI agents are expected to have the biggest initial impact in terms performance gains and ROI are customer care, sales support, procurement, human resources and cybersecurity.

Figure 3. Key drivers for GenAI investments



Source: Omdia GenAI Enterprise Early Adopters Survey, n =356

IBM

HR departments can struggle to keep pace with employee needs, and with managing the recruitment and onboarding of new staff, challenges that can be acute for large, international organisations and young, inexperienced companies. IBM has a decade of experience developing AI automation for HR functions and uses solutions systems internally, with good outcomes. For example, IBM's AskHR AI assistant enables its own employees to complete tasks 75% faster and resolves 90% of inquiries without escalation. More recently, IBM developed HR agents (available in private preview) that are built on watsonx Orchestrate and are designed to perform complex, multi-step tasks (e.g., employee onboarding, support, learning). IBM HR agents integrate with existing HR platforms and can be customised without coding expertise. HR agents feature pre-built skills such as automated email drafting, employee leave management.

Source: IBM, Omdia

Avid Solutions

Avid solutions develops advanced technology for sustainable food production systems in space and on earth. The company provides consulting services to clients, which involves customer onboarding, project management, document processing, and expense reporting. Avid Solutions uses IBM watsonx Orchestrate to automate these process, with quantitative benefits including a 25% reduction in the time it takes to onboard new customers, and a 10% reduction in errors caused by manual project management processes. Going forward, the company expects watsonx Orchestrate to reduce customer onboarding times by 50%, and manual project management errors by 20%. Avid Solutions has also experienced qualitative benefits such as improved employee satisfaction through the automation of repetitive tasks, and greater customer satisfaction because of faster response times. Source: IBM, Avid Solutions, Omdia

Dun & Bradstreet

Dun & Bradstreet (D&B), a leading provider of business data and analytics, wanted to provide customers with more comprehensive, instant insights into supplier risk evaluation at the procurement phase. Towards this, D&B worked with IBM to develop D&B Ask Procurement, using the capabilities of the Llama 3 and Mistral LLMs and IBM's watsonx Orchestrate, integrated on IBM Cloud. The D&B Ask Procurement assistant can query and analyse detailed data about suppliers, eliminating time-consuming, manual data entry, increasing efficiency and reducing error rates. D&B Ask Procurement provides a 10-20% reduction time for procurement tasks, leading to significant cost savings. D&B Ask Procurement also provides real-time monitoring capabilities to anticipate potential supplier issues, which means pre-emptive actions can be taken. Source: IBM, Dun & Bradstreet, Omdia

Transformation brings challenges

Complexity and costs

Getting started with AI agents can be challenging, particularly for firms that lack AI expertise or sufficient resources. Implementing AI agents requires access to deep troves of high-quality data that must be managed, a process that can be demanding and resource intensive. Integration poses another challenge, as AI agents must seamlessly connect with core enterprise systems, as well as with partners and broader ecosystem platforms. Beyond initial deployment, scaling AI agents and managing their ongoing operation present further challenges. Enterprises must ensure that infrastructure can handle increased inference and data management workloads while maintaining performance and reliability, which can mean increased costs.

The solution...

Enterprises that want to pilot and/or implement AI agents can call on pre-built AI agents that are optimized for specific use cases and designed to be integrated with existing systems for easy, faster deployment. Enterprises can also build customized AI agents with ease and speed by leveraging low-code solutions to design, deploy, and manage agents tailored to their unique needs.

Enterprises should invest in data infrastructure solutions capable of managing large volumes of data efficiently and cost-effectively across diverse deployment scenarios. Data observability is

another critical requirement of agentic AI systems, enabling enterprises to monitor and manage the quality, reliability, and performance of data across different systems and processes.

Managing AI agents is simplified with multi-agent orchestration solutions, which allow enterprises to oversee and scale all their agents from a unified platform. This streamlined experience ensures that multiple agents can be monitored, and optimized efficiently, supporting scalability and flexibility as business needs evolve.

Developers planning to build AI agents should also look for tools that supports the above capabilities, along with foundational elements such as reconfigured SDKs, APIs retrieval augmented generation (RAG). Environments that support optimized development workflows of this kind this will reduce complexity, development times, and operational costs.

Data privacy and security

Ensuring data security and privacy while allowing autonomy is a significant challenge. AI agents may require access to large datasets that are disparate and located across interconnected systems. These disparate data sources can contain sensitive data. Moreover, agentic systems may autonomously share data with other systems or make decisions about data usage. Interconnected systems can also heighten cybersecurity vulnerabilities, for example multiple end points can amplify attack vectors. While agentic AI systems need to be adaptable, there can be safety risks in certain scenarios when AI agents are able to access and modify system-level environments autonomously. In this situation, AI agents may take actions with the potential to cause harm or make decisions with unforeseen consequences.

The solution...

Data privacy governance and safeguards must be baked into the architecture of agentic systems from the outset – privacy by design takes on new urgency in the agentic domain. Strong data management tools are essential, particularly those supporting data observability (see under *Complexity and costs*). Although not a common approach yet, Multi-Party Computation (MPC) could be useful for multi-agent systems, allowing agents to collaborate without revealing underlying data.

Agentic AI should only be deployed with adequate guardrails in place, and this needs to be handled in a way that gets the balance right between allowing adaptability while maintaining control. Enterprises must carefully evaluate the level of autonomy granted to AI agents and what this could mean for potential security risks, which will vary by use case. Considerations include how much access AI agents have to external data, other agents and tools, and the degree to which human oversight may be needed.

Reliability

Agentic AI, while highly capable, can still produce inaccuracies, ‘hallucinations’, or unexpected outcomes. Such errors can create repeatability issues for developers and undermine trust and pose safety risks if users act on incorrect information.

The solution...

The accuracy and reliability of agentic systems can be improved by feedback mechanisms. This can be in the form of humans on hand, rather than constantly in the loop, that agents can call for further direction if encountering uncertainties. Humans can review any outcomes that may have the potential to be problematic, inaccurate and so on. Other agents can also provide direction and feedback mechanisms, which can be a means of reducing the need for humans to carry out this task.

Recommendations

Successful AI agent deployments require a mix of connected capabilities that require investment considerations and strategic planning, and now is the time to do it. During this process, the following points are worth keeping front of mind:

- **Orchestration is a critical success factor for multi-agent systems.** Orchestration can abstract complexity and ensures agents collaborate effectively and in alignment with overall goals they have been set.
- **Good agentic systems require a comprehensive and consistent view of corporate data.** Any investment in a unified semantic data layer (metadata about all available data sources) will ensure agentic systems can better understand business syntax and meaning from across both projects and business departments.
- **Lean into flexible platforms that support your specific needs and available resources.** This could be options for pre-built skills optimized for specific use cases or business functions, solutions that enable easy, low code customization, and/or opportunities for rapid scaling of agentic systems.
- **As with any AI investment, AI agents need strong governance.** This will include guidance and tools for supporting responsible AI, mechanisms for the continuous monitoring and evaluation of AI agent performance and outcomes in real-time.
- **Double down on security.** Seek partners that have embraced security-by-design practices and that have baked security functions into agentic AI solutions from the ground up.

Appendix

Introducing IBM Agents

IBM makes it easy for clients to embrace the era of agentic AI through a comprehensive and modular AI Agent stack that offers pre-built domain agents, custom-built agents, and native multi-agent and assistant orchestration.

IBM watsonx Orchestrate is the low-code studio for building, deploying and managing custom, pre-built and 3rd party AI assistants and agents to automate work across the enterprise. Clients can realize quick time to value by accessing a catalog of pre-built agents designed for high-ROI domain use cases. Pre-built agents are powered by 1000s of out-of-the-box skills integrating with 80+ leading enterprise applications for HR, procurement, sales and customer service. In addition, IBM offers both pro-code and low-code custom agent configuration frameworks so developers and domain experts can launch agents and assistants on top of your existing workflows, processes and tools.

Alongside watsonx Orchestrate, clients can leverage IBM watsonx.ai, an enterprise-grade AI developer studio, to equip developers with the foundational models and services required to build, deploy and manage custom agents optimised for your specific enterprise use-cases. These services support developers across the end-to-end development lifecycle, offering a collection of APIs and SDKs that integrate deeply with the existing open-source ecosystem. With these services, developers are empowered to build performant, reliable and trustworthy agents that integrate into a clients specific business systems, data and processes. Developers can also leverage our intuitive low-code tooling, powered by the same foundational APIs and services, to get started and accelerate time-to-market.

Through a seamlessly integrated, native multi-agent orchestration engine, builder teams can deliver a unified experience to their end users that incorporates agents, assistants, and digital automations built within watsonx Orchestrate and watsonx.ai.

Author

Eden Zoller, Chief Analyst, Applied AI
askananalyst@omdia.com



Omdia Commissioned Research

This piece of research was commissioned by IBM.

Citation policy

Request external citation and usage of Omdia research and data via citations@omdia.com.

Omdia consulting

The Omdia research, data and information referenced herein (the "Omdia Materials") are the copyrighted property of TechTarget, Inc. and its subsidiaries or affiliates (together "Informa TechTarget") or its third party data providers and represent data, research, opinions, or viewpoints published by Informa TechTarget, and are not representations of fact.

The Omdia Materials reflect information and opinions from the original publication date and not from the date of this document. The information and opinions expressed in the Omdia Materials are subject to change without notice and Informa TechTarget does not have any duty or responsibility to update the Omdia Materials or this publication as a result.

Omdia Materials are delivered on an "as-is" and "as-available" basis. No representation or warranty, express or implied, is made as to the fairness, accuracy, completeness, or correctness of the information, opinions, and conclusions contained in Omdia Materials.

To the maximum extent permitted by law, Informa TechTarget and its affiliates, officers, directors, employees, agents, and third party data providers disclaim any liability (including, without limitation, any liability arising from fault or negligence) as to the accuracy or completeness or use of the Omdia Materials. Informa TechTarget will not, under any circumstance whatsoever, be liable for any trading, investment, commercial, or other decisions based on or made in reliance of the Omdia Materials.

CONTACT US

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

customersuccess@omdia.com