Gartner

Innovation Insights for the AI Agent Platform Landscape

26 March 2025

By Leinar Ramos, Gabriele Rigon, et all.

The emergence of AI agents marks the next significant wave in AI innovation. Platforms for creating, deploying and managing AI agents introduce transformative changes to IT and business workflows. This research will help AI leaders navigate the complex and emerging AI agent platforms landscape.

Overview

Key Findings

- AI agent platforms help create and manage AI agents. They offer a diverse spectrum of options ranging from prebuilt agents to no-code builders, comprehensive development platforms and AI agent-training solutions.

- There is a rapid surge of interest in AI agents. Many vendors are contributing to the hype by engaging in “agent washing,” rebranding existing products, such as AI assistants, RPA tools and chatbots, to capture buyers’ attention without substantial agentic capabilities. This creates a fragmented market landscape that is hard to navigate.

- LLM-based agents are the most visible part of the market, but they are not the only way to build AI agents. For example, more sophisticated AI agents can be built into simulation platforms using techniques like reinforcement learning for specific use cases.

- AI agent platforms have broad applicability, with potential use cases across many business domains. However, their current use is mostly experimental, and only a few AI agents are in production that provide significant business value. Additionally, these platforms pose challenges related to reliability, security, regulatory compliance and other risks.

Recommendations

- Begin your AI agent journey by experimenting with prebuilt agents and no-code builders to navigate the extensive range of buy-to-build options without feeling overwhelmed. These options offer ease of development and integration, along with a lower skill barrier.

- Tackle the challenge of a fragmented market by identifying and focusing on the use cases where AI agents could be most valuable and feasible. Scrutinize vendor claims closely and request referenced deployments that are similar to your target use case.

- Broaden your strategy beyond LLM-based agents by evaluating AI agent training platforms that can help build AI agents for more specialized use cases, such as in robotics, energy management and industrial automation.

- Minimize uncertainty in proving business value by initially focusing on low-risk pilot use cases that deliver tangible business outcomes before committing to significant investments and organization wide rollouts. Involve business subject-matter experts in your pilot teams to maintain a strong emphasis on value, and ensure seamless integration of AI agents into existing workflows.

Introduction

AI agents have seen a rapid surge of interest. Gartner saw a 750% increase in AI-agent-related inquiries between the second and fourth quarters of 2024,[1] with Agentic AI becoming one of the biggest trends of the year (see Top Strategic Technology Trends for 2025: Agentic AI).

It is possible to build AI agents from scratch, manually integrating different components like AI models and code libraries. However, this requires a high level of expertise and is both challenging and time-consuming for most organizations. For this reason, AI agent platforms are emerging to provide a cohesive solution to develop and deploy AI agents.

Startups and existing vendors from adjacent markets such as automation, RPA, integration, conversational AI, low-code application platforms, analytics and decision intelligence platform vendors are jumping into this opportunity, creating a fragmented market that is hard to navigate for AI leaders.

Description

- **AI agents** are autonomous or semiautonomous software entities that use AI techniques to perceive, make decisions, take actions and achieve goals in their digital or physical environments.

- An **AI agent platform** is a cohesive set of integrated technologies designed to facilitate the creation, integration, deployment, optimization and management of AI agents. These platforms provide a range of creation or development options — from prebuilt agents with limited configuration capabilities, through no-code and code-first development, to training of the AI agents. These capabilities enable users with varying technical expertise and business requirements to implement and manage AI agents tailored to specific use cases.

The AI agent platform landscape is highly fragmented today. There is a spectrum of buy-to-build choices for organizations: from prebuilt agents to no-code agent builders, code-driven platforms, all the way to platforms that allow for AI agent training. These require different skillsets, as seen in Figure 1.

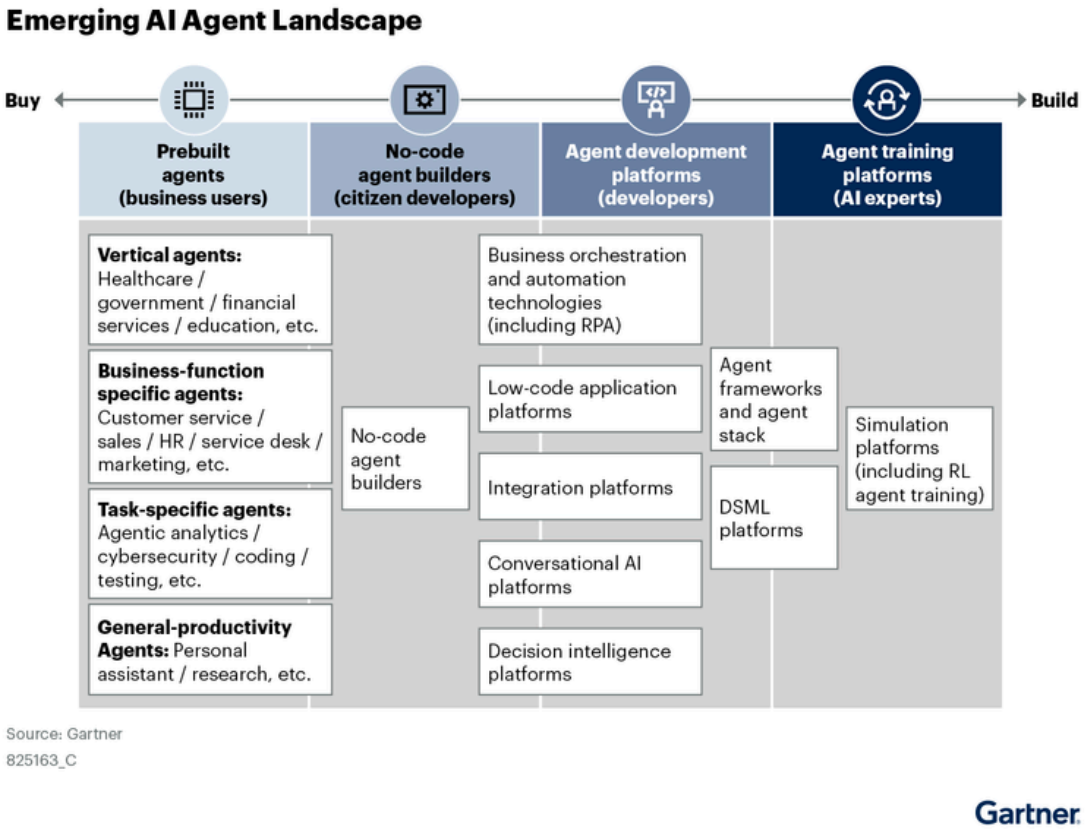

The Four Main Categories

- **Prebuilt agents**: There is a rapidly growing ecosystem of prebuilt agents that organizations can directly use, ranging from general-productivity agents, task-specific agents, business-function-specific agents to industry-specific vertical agents. These can often be configured to some degree, but they are positioned to be used directly by business users for the target use cases without requiring technical expertise.

- **No-code agent builders**: These are SaaS-delivered products that offer an integrated environment to build, publish and manage AI-powered agents without using any coding. They are used primarily by business technologists, such as citizen developers and citizen data scientists. They can build agents, in a no-code way, to improve individual productivity, standardize business workflows and drive more autonomous operations.

- **Agent development platforms**: These platforms enable developers to create and manage AI agents through code, configuration and metadata, offering flexibility and control for implementing complex logic and integrations. Many of the tools provide low-code, graphical or natural language development experiences. This category includes both low-code and pro-code offerings (see Note 1 for a definition). It encompasses both new AI agent frameworks and existing platforms like conversational AI platforms, business orchestration and automation technologies, RPA and others that now offer agentic AI capabilities.

- **Agent training platforms**: These platforms enable AI experts to train AI agents, including their underlying AI models, to perform specific tasks. They enable a higher degree of customization than what is possible in the building of AI agents by assembling pretrained AI components. A prominent example is simulation platforms, where agents can be trained and tested, often via reinforcement learning, in a virtual environment. These platforms enable development and refining of complex behaviors in agents — particularly in robotics and gaming — including their collaboration in multiagent approaches.

LLM-based AI agent platforms are the most visible part of the market, but they are not the only way to build AI agents. Alternatives include training AI agents via reinforcement learning in simulation environments, and building AI agents from more deterministic approaches — including the use of logic-based models and knowledge graphs.

Benefits & Uses

Benefit of AI Agent Platforms

- **Ease of use:** AI agent platforms simplify the development and deployment of AI agents, reducing friction and speeding up the transition from prototype to production. They offer built-in privacy, security, compliance and analytics, and often support multiagent scalability. With no-code, low-code and pro-code options, these platforms are accessible to users with varying technical skills, democratizing AI technology and enabling wider team participation in AI initiatives.

- **Integration:** By supporting connectivity with various tools, data sources or other AI agents, AI agent platforms facilitate the integration of AI agents with each other or into existing systems, such as bots, workflows, rules and decision-making processes. This integration is crucial for AI agents to significantly impact organizational workflows and drive real business value.

- **Governance and risk management**: AI agent platforms often offer features like agent authentication and authorization, cost and usage monitoring, and guardrails to ensure reliability and fairness. These capabilities allow organizations to maintain more control over their AI agent deployments and manage their risks in a more scalable way than they could do otherwise.

- Customization: AI agent platforms offer a range of development options, allowing users to tailor AI agents to specific use cases. This customization ensures that agents are aligned with organizational goals and can effectively address unique challenges.

These benefits vary depending on the type of AI agent platform. Figure 2 illustrates some of the trade-offs between these different types of solutions.

AI Agent Platform Trade-Offs

Magnitude: Low ○ ◔ ◑ ◕ ● High

		Approach			
		Prebuilt agents	No-code agent builders	Agent development platforms	Agent training platforms
Benefits	Ease of use	●	◕	◑	◔
	Integration capabilities	●	◕	◕	◔
	Governance capabilities	◑	◕	◕	◑
	Customizability	◔	◑	◕	●

Source: Gartner
825163_C

Gartner

As solutions move from prebuilt agents toward agent-training platforms, their ease of use decreases and integration capabilities typically reduce. However, organizations gain an increased ability to customize AI agents more deeply for their use cases as they move toward the right-side categories shown in Figure 2. Finally, governance capabilities, a mixture of built-in governance features and the ability to customize governance controls are currently stronger in no-code agent builders and agent development platforms.

LLM-Beware of "agent-washing." Many vendors have an AI agent story. Many of those stories are a rebranding of existing products, such as AI assistants and chatbots, to capture market attention.

Use of AI Agent Platforms

- **Business process automation**: AI agents play a pivotal role in business process orchestration and automation by autonomously executing tasks that range from routine data entry and invoice processing to more complex activities like dynamic case management and decision making. By integrating AI agents into workflows, organizations can significantly reduce human intervention, leading to increased efficiency and accuracy. These agents are capable of handling unstructured data, making informed decisions and managing interactive processes, thereby enhancing the overall agility and responsiveness of business operations.

- **Customer service and support (CSS):** AI agents in CSS are mostly interactive and semiautonomous, and are primarily aimed at enhancing back-office operations and augmenting CSS representatives. In those use cases, as components of agent assist products, they can offer real-time, proactive advice to human CSS representatives and improve ticket management. In customer-facing self-service use cases, they can be leveraged embedded in, or in combination with, virtual assistants, to provide 24/7 support across various channels and languages. They can manage conversational user journeys that demand flexibility, such as offering product recommendations based on dynamic catalogs and varying customer preferences. They can also manage more formalized processes such as password resets or order tracking when backed by prebuilt workflows.

- **Cybersecurity**: Cybersecurity AI agents automatically process events and alerts, updating existing policies with new detection or enforcement rules, and suggesting and potentially deploying vulnerability remediations following an automated scan.

- **Data and analytics**: Agentic analytics is a process of data analysis that applies AI agents across the data-to-insight workflow, orchestrating tasks semiautonomously or autonomously toward stated goals that support, augment and automate insights.

- **Decision intelligence**: Decision intelligence platforms (DIPs) are software used to create solutions that support, automate and augment decision making of humans or machines, powered by the composition of data, analytics, knowledge and AI techniques. AI agents for decision intelligence include roles like decision model designer, information gatherer, decision steward, decision trigger (event) alert, decision flow generator, decision executor/maker, decision quality/outcomes watcher, decision network organizer, decision-network optimizer and event adapter.

- **IT operations**: AI agents can execute tasks from routine development to complex incident handling with limited human intervention. An example would be an incident management agent, that while perhaps enriching the incident with other information, coordinates its processing with a configuration management agent to develop and instantiate a fix to a recently identified critical problem.

- **Marketing**: Personalize customer messaging agents automate campaigns, tailoring content and optimizing ad targeting which can be delivered autonomously. By automating routine tasks and providing insights, AI agents enhance marketing efficiency and effectiveness.

- **Sales**: AI sales prospecting agents can autonomously perform many important and repetitive sales tasks in multiple modalities (i.e., text, voice, video), including creating personalized outbound emails, responding and qualifying inbound leads, and scheduling meetings. Many prebuilt agents can autonomously or semiautonomously perform tasks typically performed by a sales development representative (SDR).

- **Software engineering**: AI agents in software engineering step beyond AI code assistants to act as autonomous autopilots for software development and testing tasks, autonomously performing labor-intensive tasks across the software development life cycle (SDLC).

- **Robotics:** In robotics, AI agent platforms (particularly simulation platforms) are improving how machines perform tasks and interact with their surroundings. AI agents allow robots to carry out functions autonomously, from basic repetitive tasks in manufacturing and logistics to complex operations in autonomous vehicles and drones. By using AI agents, robots can adjust to changing environments, make quick decisions and work alongside humans.

Risks

- **Reliability issues**: Despite the tooling available in these platforms to test and monitor AI agents, their performance for many use cases will not be reliable enough, given the current state of the technology. This is especially relevant for AI agents that depend on a pure LLM-based architecture and are not at least partially backed by prebuilt workflows. These reliability issues could impact the expected ROI from investments in these platforms.

- **Security risks**: The integration of AI agent platforms into existing systems and data sources, and an increased dependency on LLM services provided by third parties, can create new security vulnerabilities, expanding the potential security threat surface and potential data exposure. The complexity of multiagent systems can amplify these security risks by making it more challenging to monitor, identify and remediate vulnerabilities.

- **Agent anarch**y: Many vendors are aiming these platforms to a very broad population of end users, raising the prospect of AI agent anarchy if not properly governed, with many poorly orchestrated agents that could collide and cause vicious loops of automated interactions. Without proper planning and supervision, agent anarchy is a near-term and serious threat to the stability of enterprise applications.

- **Regulatory and governance risk**s: By making AI agents easier to build in many parts of the organization, these platforms can lead to a proliferation of use cases, some of which might not comply with existing regulations. Some use cases may be subject to additional obligations in some geographies, such as several that may be conceived in HR as per the EU AI Act standards (see Getting Ready for the EU AI Act, Phase 2: Risk-Assess & Categorize). As many parts of the organization develop AI agents, maintaining governance and ensuring regulatory compliance will be a challenge.

- **Vendor lock-in**: Relying on a specific AI agent platform can make it difficult to migrate AI agents onto other platforms or run agents in other systems. The ease-of-use of these platforms comes with increased dependency on specific vendors.

- **Interoperability**: Current AI agent platforms are typically not compatible with each other due to a lack of commonly accepted data and protocol standards. This limits organizations' ability to have AI agents from separate platforms interact with each other in a workflow.

- **Cost overruns**: It can be challenging for customers of these platforms to accurately predict and have control over usage of their agents, leading to potentially higher costs than originally expected. Pricing mechanisms also vary widely by vendor, making it hard to compare the cost of similar use cases.

Recommendations

- To navigate the extensive range of buy-to-build options without feeling overwhelmed, begin your AI agent journey by experimenting with prebuilt agents and no-code builders. These options offer ease of development and integration, along with a lower skill barrier. Consider AI agent pro-code development platforms and training platforms when more customization is required for your use case.

- Tackle the challenge of a fragmented market by identifying and focusing on the use cases where AI agents could be most valuable and feasible. Scrutinize vendor claims closely and request referenced deployments that are similar to your target use case. Evaluate new agent capabilities added to existing platforms that have already been deployed in your organization like RPA, automation, conversational AI, analytics and integration platforms.

- Broaden your strategy beyond LLM-based agents by evaluating AI agent training platforms that can help build AI agents for more specialized use cases, such as in robotics, energy management and industrial automation. Alternatives include the training of AI agents via reinforcement learning in simulation platforms, and building AI agents from more deterministic approaches — including the use of logic-based models and knowledge graphs.

- Minimize uncertainty in proving business value by initially focusing on low-risk pilot use cases that deliver tangible business outcomes before committing to significant investments and organization wide rollouts. Involve business subject-matter experts in your pilot teams to maintain a strong emphasis on value, and ensure seamless integration of AI agents into existing workflows.

Evidence

[1] Analysis from Gartner client inquiries for the keywords “Agentic AI” or “AI Agents” between the dates of April 2024 and December 2024. This observed a 750% increase in inquiries from the second quarter (April to June 2024) to the fourth quarter (October to December 2024).

Note 1: Low-Code Tools

Low-code tools facilitate application development by providing dedicated graphical user interfaces (GUIs) that allow both citizen and professional developers to utilize visual programming abstractions, reducing the need for extensive coding. Pro-code tools allow professional developers to extensively employ programming languages, such as Java or Python, to enable the creation of highly customized functionalities and integrations.