Emerging Tech: Tech Innovators in Agentic AI

16 September 2025 - ID G00828290 - 21 min read

By: Danielle Casey, Alfredo Ramirez IV, Anushree Verma, Aakanksha Bansal, Akhil Singh

Initiatives:Emerging Technologies and Trends Impact on Products and Services; Increase Product Traction

The tech innovators in agentic AI are addressing emerging challenges in workflow design, agent collaboration, reliable performance and high autonomy. Product leaders must design differentiated agents by building innovative functionality that directly solves critical industry challenges.

Overview

Key Findings

- Chief product officers must plan for continuous innovation as a competitive requirement to achieve agentic AI adoption at scale by 2028. The innovations highlighted in this research will help drive this.
- Tools that automate agent workflow creation and facilitate collaboration between first- and third-party agents are tackling challenges in agent design and agent sprawl within organizations.
- Agent ops tools that allow enterprises to monitor, improve and enforce agent reliability are helping to solve the adopter trust challenges currently inhibiting buying decisions.
- Composite AI architectures augment agent accuracy, reliability and performance, enabling support for more complex, specialized and autonomous use cases than GenAI-only approaches.

Recommendations

- Focus R&D investments on solving the key technology challenges — agent communication, reliability and trusted autonomy — to drive scalable agentic adoption within the enterprise.
- Tackle the growing agent sprawl challenge within organizations by building support for first- and third-party agent collaboration. Failure to do so will introduce significant risk due to the inability to support adoption at scale and the associated revenue opportunities.
- Ensure reliable agentic AI performance by building agent ops tools that help monitor, evaluate and enhance performance. Reliable AI agents will drive near-term adoption of agentic AI.
- Build agent offerings capable of supporting advanced and complex use cases (such as prediction, highly autonomous applications and mission-critical implementations) by investing in composite AI architectures. Agents that go beyond LLMs and design a composite AI solution — for example, using machine learning models, vector databases or symbolic AI — can unlock new opportunities.

Analysis

Technology Description

Agentic AI refers to a class of system developed using various architectures, design patterns and frameworks, encompassing both single-agent and multiagent designs. These systems are capable of performing unsupervised tasks, making decisions and executing processes. They range from semiautonomous to fully autonomous and are software entities that utilize AI techniques to perceive their environments, make decisions, take actions, adapt and achieve specific goals in both digital and physical settings. Agentic AI is distinguished by the following characteristics: goal orientation and orchestrated collaboration; autonomy; perception and environmental interaction; reasoning and planning; tool use; learning and behavior adaptation; and memory (see Emerging Tech: The Key Defining Characteristics of Agentic AI).

The agentic AI market is rapidly growing, with significant future revenue at stake for the winners of this AI race. The high-risk, high-reward nature of the market is spurring innovation that addresses key challenges and advances agentic applications. The tech innovators featured in this report have pushed forward developments in workflow and communication tools, advanced AI agency and reliability features in agentic AI (see Figure 1).

Figure 1: Tech Innovators in Agentic AI

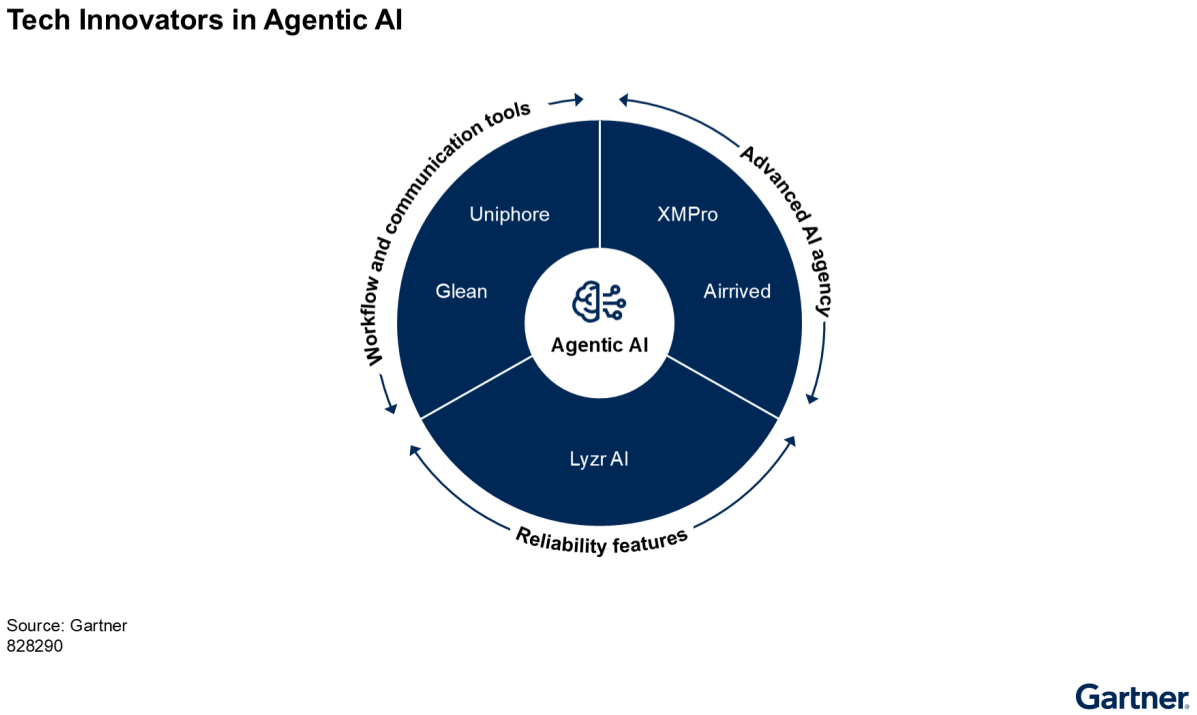

Source: Gartner
828290

Gartner

Table 1 lists the tech innovators profiled in this research and indicates which of their capabilities merit inclusion on this list. They may have strengths in other areas as well, but in Table 1 we highlight what we consider to be their strongest one or two points.

Table 1: Tech Innovators in Agentic AI

(Enlarged table in Appendix)

Vendor	Innovation	Market impact
Glean	Provides a third-party agent orchestrator that allows agent-to-agent workflow collaboration and permissioned third-party agent use of Glean-managed enterprise data.	Helps solve the agent sprawl issue that is emerging in the market and will facilitate the scalable adoption of agentic AI.
Uniphore	Offers a process discovery agent that intakes multimodal workflow information and autonomously designs custom agentic workflows.	Uses agents to discover process inefficiencies within the organization and automatically generate well-aligned agents for workflow automation.
Lyzr AI	Provides embedded agent reliability by autonomously reviewing, critiquing, and correcting an agent before it takes action.	Agent ops tools and techniques that improve agent reliability will grow customer trust and agent adoption.
Airrived	Offers expert cybersecurity agents with advanced capabilities that enable automated risk remediation.	Domain-specialized agents are reinventing software solutions across industries, unlocking new automation for productization.
XMPro	Composite AI (combining LLMs with robust data and ML capabilities) enables highly accurate and autonomous agents with capabilities such as prediction to inform agent actions.	Unlocks more autonomous and complex agentic applications.

LLM = large language model; ML = machine learning.

Source: Gartner, 2025 Agentic AI Case-Based Field Research

Glean AI's Third-Party Agent Orchestrator Solves Enterprise Agent Sprawl

Nature of the Innovation

Glean AI offers a no-code agent development platform, but the innovative aspect of its offering is its support for third-party AI agents. Using Model Context Protocol (MCP) server calls, Glean AI agents can connect to and invoke third-party agents through secured MCP tool calls to complete multistep tasks. More specifically, these third-party agents are registered in Glean and then considered by Glean's agentic reasoning engine in runtime as tools for planning and task execution.

Glean also enables enterprises to scale the enterprise context used by Glean agents so it can be leveraged by third-party agents. This means that the specialized agent knowledge and context an organization builds within the Glean platform can be used to improve the performance of third-party agents. This effectively removes agent data silos and scales contextualized optimization.

Glean also innovates in security and data protection. Third-party agents are authenticated in the Glean platform, ensuring that third-party agents are permissioned to access relevant enterprise context. Glean Protect features for agents provide clear permission understanding through both UI and API, defining who can own, edit, view and execute an agent. Glean agents also follow user permissions, meaning the agents can only access and action data that users already have access to. These identity and access management features help support enterprise-scale agent usage.

Adopter Case Study

Miro, a collaborative online platform, used Glean enterprise search, Glean assistant and Glean agents to improve knowledge management and productivity across various teams. Miro integrated Glean's enterprise knowledge capabilities into its Innovation Workspace, allowing employees to access enterprise knowledge, documents and resources directly within Miro. This streamlined workflows and reduced the need to switch between applications. Miro has also utilized Glean's APIs to build deep integrations with its core product, enhancing the overall user experience.

The sales team uses Glean to generate sales outreach materials. Using a prospect's LinkedIn profile, the tool generates a custom email, reducing the time required to draft an outreach email by 80% and allowing the team to increase its daily outreach efforts. Miro has developed an AI agent for sales that significantly enhances the account research process. The AI agent analyzes internal reports and external web information and compiles relevant account information (such as industry, M&A activity and cost savings). This reduced the time it took to analyze one account from about two hours to 10 minutes, and helped the sales team identify customer problems in order to support and prepare for customer calls.

Miro's customer education team manages 500 help center articles, many of which become outdated due to UI changes and product releases. To address this, they leveraged Glean to identify articles affected by these changes, find relevant information, and propose rewritten drafts for new articles. This process saved one to two hours per article, creating up to 1,000 hours in time savings to update all the articles. Additionally, AI made this task feasible for their five-person team.

In total, Glean's platform is enabling Miro employees to enhance their efficiency and productivity. This project is in full deployment, with Glean agents used by all 2,000 employees.

Market Adoption and Impact

Agent sprawl is emerging as a significant challenge for adopter organizations due to the use of multiple agent-based solutions. Organizations' ability to manage and orchestrate multiple agents from various providers is key to scalable AI agent adoption. Vendors' ability to support third-party agent collaboration will quickly emerge as a key differentiator. This capability will be well-received by customers and will help AI agent providers stand out as a preferred platform.

Of course, agent-to-agent collaboration across platforms has varying degrees of complexity. Some agents simply pass along information or instructions (that is, one-way dialogue), few support multistep task processing (that is, two-way dialogue), and still fewer support complex levels of collaboration and permissioned data access. The more tools and features vendors build to support orchestrated and managed multiagent solutions, the better positioned vendors will be to capitalize on emerging agent opportunities.

Security is also a key component of agent collaboration and functionality that much of the market has yet to build. The issue of securing agents will quickly emerge as a key buyer consideration. Vendors that can secure and permission their agents will achieve scalable adoption ahead of those lacking this functionality.

Implications for Product Leaders

Agent sprawl is a rapidly emerging challenge that product leaders must be prepared to address soon. Failure to effectively orchestrate with third-party agents will prove severely limiting, if not competitively debilitating, within a year. Product leaders must start building this functionality now to ensure their agents can collaborate across platforms. This interoperability is critical to enabling scalable adoption of agentic AI within customer organizations, and will soon become a competitive necessity.

Enterprises will not adopt agents that cannot be permissioned and secured. However, most agent products lack this functionality. As the agent market matures, so will enterprise interest in adopting and scaling agents. Product leaders must start embedding security features and agent management and governance tools now or risk missing out on upcoming revenue opportunities.

Uniphore Automates Workflow Design for Accurate and Easy Agent Development

Nature of the Innovation

Uniphore's recent acquisition of Orby AI (on 28 August 2025) adds to Uniphore's Business AI Cloud unit a no-code agent development platform that innovates with its "process discovery" automation tool. This tool is exposed to the user as an agent that runs in the background and observes how tasks are regularly performed, including the necessary steps, applications and document interdependencies. The agent then analyzes the frequency of user tasks and identifies the best opportunities for efficiency improvements. Next, the agent generates executable automation in the form of an agent that the user can modify for deployment.

More specifically, Uniphore's process discovery agent is powered by a multimodal model that ingests text (standard operating procedure documents), video (screen recordings) and digital monitoring data to understand processes and digital tasks. Using Uniphore's Large Action Model (LAM), an interactive representation of the observed processes is created. This allows for deep AI-led analysis of the workflow to identify patterns and automation opportunities. This understanding is translated into "digital actions," which are the steps required to automate a task. Uniphore's LAM then creates a digital workflow and writes the code needed to automate the process itself. This code is what the agent uses to execute tasks.

Adopter Case Study

A real estate company used Uniphore's agents to help automate invoice and purchase order reconciliation by matching invoice data against receipts. The organization processes 36,000 invoices per year (or 3,000 per month), a highly manual, time-consuming process that is prone to errors and requires significant resources. This is because the invoice reconciliation process requires people to match and verify invoice data against other financial records such as purchase orders, receipts and bank statements across multiple ERP and invoice management systems with varying vendor templates. These complexities created overhead costs and accuracy issues, leading to delayed vendor payments and compliance concerns.

Uniphore's agents were used to reconcile data changes in these systems and fill out invoice processing forms; at this point, a human in the loop (HITL) reviewed the information and approved submission. This organization was able to reduce invoice processing and reconciliation time from 750 hours per month to 150 hours per month by using Uniphore AI agents to fully automate 64% of invoices and partially automate 20% of invoices. Additionally, they reported a 15% improvement in accuracy, more timely vendor payments and improved compliance.

Market Adoption and Impact

Many AI agent development platforms focus on offering easy-to-use, no-code interfaces. Most of these platforms currently use conversational interfaces for agent design, which lowers the skills barrier to agent adoption. However, there may remain a delta between how a process should be automated and how a user designs the agent. This mismatch can contribute to underperforming agents and poor accuracy. Automation tools that help uncover hidden inefficiencies and optimize agent workflows will help drive enterprise adoption, accelerate time to value and improve agent performance. Automation tools that help with pre- or post-agent deployment will help vendors differentiate themselves from competing no-code agent development platforms.

Implications for Product Leaders

Automation tools that help users more easily build better agents will create differentiation in a crowded market. Systems that can effectively uncover operational inefficiencies and design agentic solutions to remediate them offer clear business differentiation. Using agents as tools to identify and remove operational inefficiencies and reduce costs is a strong value proposition. Product leaders should invest in multimodal functionality and advanced reasoning models to improve agent design and performance.

Lyzr AI Innovates for Agent Reliability and Performance, Facilitating Trusted Agentic Adoption

Nature of the Innovation

Lyzr AI provides a full-stack enterprise agent infrastructure platform and innovates with responsible AI and safe AI (input and output) management features. For safe AI, user prompts and LLM outputs are evaluated to perform four checks: personally identifiable information (PII) redaction, bias and fairness detection, toxicity check and prompt injection. The safe AI modules — powered by named entity recognition and zero-shot classification language models — either reject or redact the data. Next, the LLM response is inspected by the Responsible AI module that sits within its core agent architecture. This module is powered by deep learning algorithms that are highly deterministic and performs the following functions:

- Reflection: Instruction adherence
- Groundedness: Fact adherence
- Context relevance: Retrieval accuracy
- LLM as a Judge: One model as a reviewer agent
- LLM as a Panel: Multiple models as reviewer agents

The module checks and adjusts the response for accuracy before surfacing the information to the customer. If any of these checks fail, the failure is noted and the response is regenerated before any generated content is presented to the user. In effect, Lyzr has embedded responsible and safe AI into its agent architecture.

Adopter Case Study

A SaaS platform provider used Lyzr AI to create a white-labeled AI hiring agent that was then offered as a solution to its customers, who are primarily HR leaders. The agent was developed to help improve the SaaS provider's customers' experience by helping automate the HR recruitment and hiring process. This multiagent solution is composed of four agents:

- Job description agent: Using hiring requirements and the company format, this agent generates a job description.
- Resume shortlist agent: Candidate applications are then analyzed by this agent and filtered against shortlist criteria.
- Meeting scheduler agent: Job interviews are scheduled with candidates on the shortlist by identifying the availability of the hiring manager and sending timeslots to candidates.
- Feedback collection agent: The hiring manager provides feedback to the agent post-interview, and the agent then drafts and sends rejection emails.

This use case is currently in limited deployment.

Market Adoption and Impact

The market currently places a premium on agentic reliability. Agents that can deliver accurate and consistent outcomes are highly desirable across industries and use cases, particularly in regulated ones. Agents that are reliable, responsible, and can be observed and checked by adopters will earn customer trust.

Customer trust currently presents a significant go-to-market challenge for the emerging AI agent market. This is inhibiting both current adoption and future scaling of AI agents. For more details, see Emerging Tech: Customer Trust Is a Critical Barrier to Agentic AI Adoption.

Implications for Product Leaders

Embedded reliability and responsibility address one of the foremost challenges in AI agent adoption. Yet, most agent providers lack this functionality, making it a key differentiator.

There is significant market interest in AI agents, but adoption remains low due to lack of trust. Providers that build a “trustworthy by design” approach into their agent offerings will be well-positioned to win customer trust and revenue.

XMPro Uses Composite AI to Deliver Fully Autonomous Agents for Managing Physical Infrastructure

Nature of the Innovation

XMPro’s composite AI architecture positions its agent offerings to tackle highly autonomous and complex use cases. XMPro uses digital twins to provide real-time data observations of the environment. The AI agents use LLMs with robust retrieval-augmented generation (RAG) for contextual knowledge. This contextual knowledge is informed by data inputs from IoT devices, sensor data, contextual enterprise data and engineering systems. Vector and graph databases support memory retrieval and help with reflection and planning. Causal AI analysis and ML models are also used to provide predictive insights. Together, this information is utilized by AI agents to make decisions, which can either be fully autonomous or semiautonomous via HITL control. This robust, multifaceted agent architecture is key to XMPro’s focus on using AI agents to manage decisions.

XMPro’s use of digital twins and predictive AI supports applications in predictive maintenance, infrastructure management, forecasting and more, where agents can further extend the value chain by providing reasoned recommendations based on this extensive and detailed information, as well as autonomously implemented actions. Agents built on top of a composite AI architecture can operate at high or full autonomy and perform digital as well as physical actuation, as demonstrated by XMPro’s innovation.

Adopter Case Study

An Australian water utility is using XMPro’s fully autonomous, collaborative, multiagent offering to help run its facility. This is a production pilot of a multiagent solution for remote infrastructure management. The solution uses digital twins to provide a digital mirror of plant data and operations, and language models to interact with and action plant data to perform automation. There are six agents in total being used to help run this facility:

- Water quality agent: Monitors water quality by analyzing various indicators, such as pH levels and chemical levels, against health and safety standards.

- Reservoir ops agent: Coordinates with the water quality agent and adjusts water treatments to ensure water quality.
- Maintenance agent: Monitors pump sensor data for predictive maintenance and coordinates with the reservoir ops agent to ensure equipment downtime does not inhibit operational needs.
- Safety and compliance agent: Monitors critical systems, documents key reporting information and ensures regulatory compliance.
- Data flow agent: Ensures data quality across systems by validating data, maintaining records and generating operational insights for use by both AI agents and human operators.
- Energy efficiency agent: Forecasts when it is most cost-effective to pump.

This multiagent architecture helps solve key challenges related to the plant's remote location, primarily:

- Performing water quality checks more frequently than is physically enabled in remote locations.
- Reducing disruption from asset downtime and associated costs through accurate maintenance scheduling that aligns with asset downtime windows and operational needs.
- Providing data insights, recommendations and next best actions that help ensure reliable plant operation despite limited physical site access.

This use case was implemented in 2024 and is currently still in pilot.

Market Adoption and Impact

Many agentic AI providers today take a generative AI first or only approach, which can undermine their ability to support more complex and autonomous agentic applications. Vendors that invest in a composite AI architecture that leverages different types of models based on their strengths (such as causal AI, ML models, small language models, advanced reasoning models, digital twins and symbolic AI) can tackle more robust use cases both now and in the future.

Composite AI architectures will help mitigate many of the accuracy, reliability and performance challenges posed by LLM solutions. Agent providers that employ such approaches may observe faster adoption and scaling of their agents, as well as an advantage in servicing more regulated or complex environments.

Composite AI architectures will also position agent providers to pursue future physical AI opportunities. The current focus on digital actuation is expected to shift to include physical actuation, in which AI agents make decisions that change the condition of physical assets — for example, autonomously turning off machines, deciding when to bring them back online, and redistributing asset workloads based on real-time information.

Implications for Product Leaders

The future of AI agents is characterized by high autonomy, complex decision management and physical actuation. Vendors' ability to compete in this rapidly emerging future will depend on today's R&D investment decisions. Robust tech architectures are key to future AI agent competition and value expression. Future-proof for more complex use cases, such as prediction, simulation and physical actuation, by investing in composite AI architectures.

Airrived Offers a Composable Multiagent Architecture for Cybersecurity Practitioners to Deploy Domain-Specific Security Agents

Nature of the Innovation

Airrived offers a no-code platform for the development and management of cybersecurity agents to cybersecurity practitioners. Airrived provides a library of prebuilt security agents as well as tools to create enterprise-specific agents. Airrived offers customization tools, such as RAG, reinforcement learning from human feedback (RLHF) and low-rank adaptation (LoRA) fine-tuning. Airrived's multiagent solution uses a composite AI architecture that includes open-source large language models, large vision models, domain- and task-specific LLMs, machine learning models, knowledge graphs and vector databases. This multifaceted model and data architecture is unique in its ability to support domain-specialized automation and complex use cases with high accuracy.

These cybersecurity task-focused agents are designed to support security analyst and network security roles, can converse and reason through dynamic cybersecurity data and environments, and can act and provide automation in response to cybersecurity events. Airrived agents can connect to virtually any application or existing cybersecurity software, such as Palo Alto Networks, CrowdStrike, Wiz, Zscaler, Snowflake, Splunk, AWS and Microsoft, to monitor, identify and grade security risks. They also provide recommended actions for remediation that security analysts can approve and that the agent autonomously implements. The domain-specialized nature of these agents supports complex applications in cybersecurity environments and enables fully autonomous or semiautonomous solutions for threat remediation.

Adopter Case Study

A U.S.-based insurance provider used Airrived's cybersecurity agents to build multiple solutions for their security operations, threat intelligence and security engineering teams. A mesh of agents was created to boost operational efficiency and save time. First, conversational AI and generative search agents were used to provide a unified conversational interface that interacts with security tools such as Splunk and CrowdStrike EDR and Recorded Future threat intelligence. Second, a reporting agent was used to help generate the daily creation of internal advisory documents. This reduced the time it took threat intel analysts to create reports by 30%. Lastly, ticket classification and policy agents were used to autonomously classify security tickets using historical ticketing data. This reduced ticketing overhead for the security engineering team by 50%.

If the ticket was classified as low or medium risk, the policy agent generated a firewall policy that was then applied to the Palo Alto Networks NG Firewall. If the ticket was classified as high risk, a change management request was created in ServiceNow. This reduced firewall change request times for low- and medium-risk rules from seven days to two minutes. It also automated the creation of change management requests, resulting in a 30% overall time saving for security engineers. This allowed them to focus on other tasks, such as threat modeling and retroactive threat hunting. Based on these initial results, the insurance provider expects a return on investment (ROI) from the $400,000 contract within 12 months of deployment via operational efficiency gains and measurable productivity improvements. Notably, this adopter also aims to reduce security ops headcount by 50% in the next three to five years.

Market Adoption and Impact

Agentic AI is and will continue disrupting existing enterprise software. Multiagent architectures are being used in cybersecurity offerings, software development products, conversational AI platforms, enterprise productivity software and more. Product leaders can use agentic AI to develop superior products and gain a competitive edge by augmenting existing software or by entering markets where agents can unlock new value currently missing from legacy software. Domain-specialized agents will erode the market share and profit margins of existing software markets.

Domain-specialized agents are also key to unlocking high-value use cases that are complex and highly autonomous. This is because domain-specialized agents can reason, perceive information and environments, and actuate outcomes that are reliably accurate. Specialization is essential to driving deeper automation that is trusted, useful and adopted by the market. As product leaders work on building more specialized and autonomous agents, measuring customer time to trust will be important for driving adoption and achieving profitability (see Emerging Tech: 'Time to Trust' Is the New Vital Agentic AI Metric).

Implications for Product Leaders

The AI agent market will become increasingly specialized over the next five years. Product leaders must plan on building or adapting their existing solutions using agentic design patterns, as well as plan an R&D and use-case strategy focused on specialized automation. Future revenue opportunities will be concentrated around role, industry and use-case expertise.

Evidence

Gartner's agentic AI CBR project ran from December 2024 through March 2025. As part of this research, 24 vendors participated in two interviews each. The first interview focused on the agentic products, technology capabilities, R&D roadmap and adoption challenges, and more. The second interview focused on the case studies, including the business problem addressed, the processes that were changed, the outcomes achieved and the unmet needs. This document was informed by insights from both interviews.

Recommended by the Authors

Some documents may not be available as part of your current Gartner subscription.

Emerging Tech: Top-Funded Startups in Agentic AI

Emerging Tech: Top-Funded Startups for Domain-Specialized Agentic AI

Table 1: Tech Innovators in Agentic AI

Vendor	Innovation	Market impact
Glean	Provides a third-party agent orchestrator that allows agent-to-agent workflow collaboration and permissioned third-party agent use of Glean-managed enterprise data.	Helps solve the agent sprawl issue that is emerging in the market and will facilitate the scalable adoption of agentic AI.
Uniphore	Offers a process discovery agent that intakes multimodal workflow information and autonomously designs custom agentic workflows.	Uses agents to discover process inefficiencies within the organization and automatically generate well-aligned agents for workflow automation.
Lyzr AI	Provides embedded agent reliability by autonomously reviewing, critiquing, and correcting an agent before it takes action.	Agent ops tools and techniques that improve agent reliability will grow customer trust and agent adoption.
Airrived	Offers expert cybersecurity agents with advanced capabilities that enable automated risk remediation.	Domain-specialized agents are reinventing software solutions across industries, unlocking new automation for productization.
XMPro	Composite AI (combining LLMs with robust data and ML capabilities) enables highly accurate and autonomous agents with capabilities such as prediction to inform agent actions.	Unlocks more autonomous and complex agentic applications.

LLM = large language model; ML = machine learning.

Source: Gartner, 2025 Agentic AI Case-Based Field Research