Emerging Tech: Agentic AI Maturity Roadmap

8 August 2025 - ID G00828291 - 19 min read

By: Danielle Casey, Eric Goodness, George Brocklehurst, Alfredo Ramirez IV, Akhil Singh, Anushree Verma, Aakanksha Bansal, Annette Jump

Initiatives:Emerging Technologies and Trends Impact on Products and Services

> Agentic AI will become a critical enterprise capability over the coming decade, as today's agents evolve and improve. Product leaders must plan for continued disruption by designing R&D roadmaps that prioritize the development of key capabilities for upcoming advanced and expert AI agents.

Overview

Key Findings

- Chief product officers are challenged with rapid decisions on when and how much to invest in agentic AI capabilities in the face of continuous disruption and the emergence of advanced and expert AI agents that heat up competition. The risk of obsolescence is concrete.
- Future AI agents will unlock new use-case opportunities by enabling higher levels of autonomy, complex workflow support, domain expertise and multimodal understanding, and advanced tooling for execution and collaboration.
- As AI agents mature, the current premium placed on agentic reliability will fade and require leading providers to invest in advanced reasoning and perception capabilities, such as navigating uncertainty and proactive decision making.
- The far future agency stage — agentic ecosystems — will accelerate the maturity of AI agents toward displaying full autonomy, emergent behaviors, embedded security and seamless orchestration across first- and third-party systems.

Recommendations

- Stop reacting to change, and plan for continued agentic AI disruption by ideating product features and company evolution against the agency maturity levels (Levels 0 through 5) and timelines.
- Innovate toward future AI agent autonomy and complexity by investing in advanced reasoning models and planning capabilities, multimodal understanding and environmental perception, domain-specialized tool use, and long-term memory.
- Target current market opportunity by focusing on agent reliability, accuracy and performance. Anticipate future demand for differentiating capabilities — such as proactivity, adaptivity and dynamic learning — that will unlock the next generation of more dynamic, intuitive and customized AI agents.
- Future-proof your AI agent strategy by assessing the capabilities of far future agents — such as full autonomy, distributed actioning across ecosystems and collective behaviors — and investing in enabling technologies in the mid to long term due to low market risk tolerance and tech immaturity.

Strategic Planning Assumptions

Through 2027, a majority of agentic AI deployments will use conditional autonomy with human-in-the-loop (HITL) oversight and review.

By 2028, only 15% of agentic AI deployments will be highly autonomous, expert agents, up from less than 5% in 2025.

By 2030, 30% of organizations will be using AI agents in three or more business units.

By 2033, 15% of agentic AI deployments will be fully autonomous, compared with less than 1% in 2025.

Analysis

The Agency Maturity Roadmap

This research presents Gartner's Agentic AI Maturity Roadmap (see Figure 1). It explores agentic AI's capability evolution from "novice" to "superior" agency over the coming decade. The roadmap outlines the technical evolution of "agency," which is the ability of AI to autonomously act, make decisions and pursue goals. Agentic AI is a high-risk, high-reward market. It is critical that product leaders understand how AI agents will mature and evolve to prioritize investments, develop differentiated capabilities and future-proof for upcoming change.

This document is one in a two-part series. Whereas this document maps agentic AI's capability evolution, the second document identifies the enabling tech at each agency stage and the impact on use cases (Emerging Tech: Agentic AI Innovation Will Foster Autonomous Business). The insights enumerated in these documents were informed by Gartner's 2025 Agentic AI Case-Based Research (CBR). [1]

Technology description: The agency roadmap builds on the foundational concepts in the Emerging Tech: The Key Defining Characteristics of Agentic AI. Agentic AI refers to a class of systems 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, 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.

This research will evaluate the following:

- How agentic AI will mature and evolve over the coming decade (see Figure 1)
- The defining features and functionality of this evolution (see Table 1)

Product leaders can use the agentic AI roadmap to evaluate current agentic capabilities and plan for emerging disruption.

Figure 1: Critical Insights on the AI Agency Roadmap

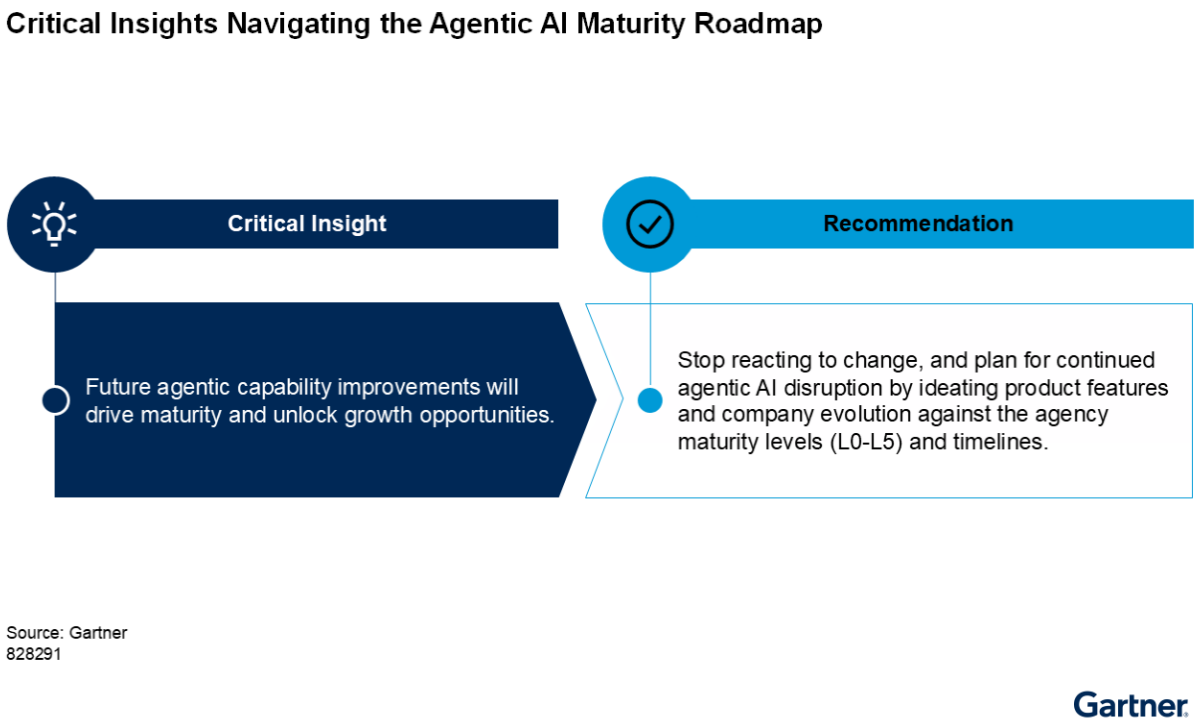

Figure 2 illustrates the agency progression from Level 0 (L0) chatbots to Level 5 (L5) agent ecosystems over the next decade. A majority of the market is currently at L2 — intermediate agency, as demonstrated by the proliferation of AI agent offerings. There are also emerging innovators with advanced AI agents and the beginnings of expert AI agent capabilities. Each agency rating considers four components: autonomy, complexity, communication and organization.

Autonomy — Considers the level of autonomous decision making utilized to include the degree of HITL involvement and whether the automation is driven by predefined processes or goal-based reasoning.

Complexity — Considers the level of complexity in the task, process or workflow being automated to include multimodal understanding, multistep processing and domain-specialized understanding.

Communication — Considers the evolution in AI agent collaboration across first- and third-party AI agent providers and, ultimately, across systems and undefined environments.

Organization — Considers the relationship between humans and agents within the enterprise.

Figure 2: The Agency Maturity Roadmap — Where We Are Now and Where We Are Going

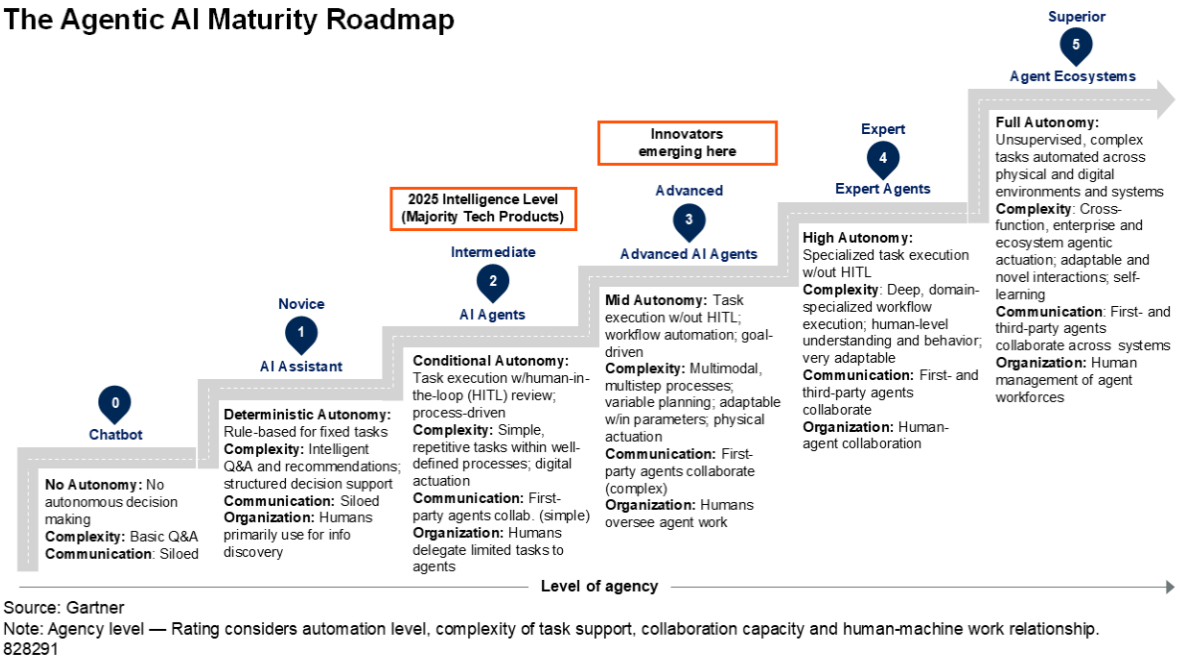

Source: Gartner
Note: Agency level — Rating considers automation level, complexity of task support, collaboration capacity and human-machine work relationship.
828291

Gartner

Critical Insight: Future Agentic Capability Improvements Will Drive Maturity and Unlock Growth Opportunities

As Figure 2 illustrates, a majority of current product development is focused on Level 2 agency — AI agents. These agents are characterized as having conditional autonomy and are focused on more simple and repetitive tasks and predefined processes. Gartner has observed these agents being used to automate expense reports, identify security vulnerabilities, automate insurance claims review and processing, and offload customer Q&A and basic support tasks. There is also significant venture capital investment in this space (see Emerging Tech: Techscape for Startups in Agentic AI).

Level 3 agency — Advanced AI agents represent emerging innovation by advancing agents' automation capabilities to execute without HITL controls and across multistep processes, up to entire workflows. These agents are goal-oriented, which equips them with more flexibility in reasoning through both defined and undefined tasks, enabling a level of adaptability not present in Level 2 AI agents. Another notable capability step in Level 3 agency is the advancement from agent-to-agent communication to collaboration, where agent negotiation, delegation and reconciliation occur.

Level 4 agency — Expert agents have yet to emerge, as underlying technologies are not mature enough to support the level of automation and workflow complexity characteristic of this agency stage.

Level 5 agency — Ecosystem agents are a far-future vision of AI agents characterized by full autonomy, high adaptivity, and the ability to negotiate, execute and otherwise transact across ever-changing environments.

As outlined in Emerging Tech: The Key Defining Characteristics of Agentic AI, agentic AI has seven primary dimensions:

- **Goal orientation and orchestrated collaboration:** Driven by predefined or dynamically evolving objectives and works to meet goals effectively and efficiently.
- **Autonomy:** Has the technical capabilities and functions to operate independently to make decisions and take actions without persistent human oversight for every step. Autonomy is often bounded by a set of structured protocols, checks (such as workflow contingencies), constraints (such as legal and regulatory regimes) and HITL controls that ensure that agentic AI systems operate within safe, ethical and legally compliant boundaries.
- **Reasoning and planning:** Possesses the ability to reason about situations, formulate multistep plans and adapt plans based on new information or obstacles.
- **Perception and environmental interaction:** Senses and interprets its environment (digital or physical) and interacts with it through actions or communication.
- **Tool use:** Can leverage external tools, APIs, software or information sources to accomplish its goals.
- **Learning and** behavior adaptation: Agentic AI is stateful and can access and use memory for this purpose. It improves knowledge and performance over time through experience, feedback or new data, adjusting agentic behavior and strategies.
- **Memory:** The agentic AI system's ability to store and recall past scenarios and actions to improve decision making, perception and performance.

Table 1 outlines the capability progression of these seven agency dimensions over future agency stages. More specifically, it provides a detailed view of agentic AI's evolution at each maturity phase — L1 through L5 — in the Agency Maturity Roadmap.

Goal orientation and orchestrated collaboration — The goals of agents will significantly evolve over the coming decade. The goal-focused nature of each agent stage is reflected in the autonomy, reasoning, perception, tool use and learning characteristics in Table 1. Notably, agent orchestrators (or workflow agents) play a key role in agent performance, workflow automation and multiagent architectures. In fact, almost half of the agentic AI CBR participants highlighted their approach to orchestration when either explaining their agentic technology or how it differentiates. Importantly, as agents advance, so will their ability to collaborate. Currently, AI agents have limited collaboration functionality. Evidence of truly collaborative agents (such as the ability to collaborate with third-party agents across multiple steps in a workflow) will appear at L3 — advanced AI agents — due to advancements in communication protocols. However, L5 — ecosystem agents — will evolve the definition of "collaboration" to require first- and third-party agent transactions, negotiation and multistep task completion across multiple systems.

Autonomy — As agentic AI progresses from L2/L3 to L4/L5, there will be a seismic change in agent autonomy and environment. Primarily, AI agents will progress from low autonomy with conditional actioning to high or full autonomy with no HITL requirement. AI agents will also evolve from being utilized across front- and back-office systems (such as IT, HR, ERP and CRM) to being deployed in regulated and domain-specialized systems (such as clinical decision support systems in healthcare, trading platforms in finance and product management systems in manufacturing), as well as physical infrastructure, such as factory floors, energy grids and traffic systems. Future agent autonomy will significantly change in both depth (the complexity of what can be automated) and breadth (the scale of usage across systems and infrastructure).

Reasoning and planning — Agents analyze information, draw conclusions and formulate plans to achieve goals. Advancements in reasoning models and planning capabilities will improve agents' ability to act independently. A key difference between today's agents and advanced AI agents is the extensibility of automation to workflows, undefined tasks and multimodal environments. However, it's not until expert and ecosystem agents that key capabilities are unlocked that allow agents to reason through uncertainty and exhibit spontaneity.

Perception and environmental Interactions — Advancements in agent perception and environmental understanding will change agents from rigid, reactive software entities to more adaptive and proactive. Today, agents are limited in their ability to perceive and interact with unknown systems, various data types, environmental context and unforeseen variables. Comparatively, future agents will be able to identify, evaluate, interpret and combine multiple types of information (image, voice, video and touch) from various sources (cameras, microphones, industrial assets, iPads and CRM systems) to understand environmental conditions, asset states and events. Advancements in an agent's ability to perceive will allow them to be more adaptive, anticipating change, navigating unforeseen opportunities and challenges, and taking preemptive actions. These far-future capabilities will introduce emergent behavior, such as creative problem solving and some of the "intellectual" qualities characteristic of artificial general intelligence.

Tool use — The main upcoming change in tool use will be the shift from flexible plug-and-play APIs to more adaptive and scalable tool use. This could include multitool orchestration and adaptive tool selection and chaining based on real-time agent reasoning. Also, API integration will become increasingly complex and domain-specialized as agents become more expert. For example, multistep, multimodal tool use will require agents to call particular tools, depending on the data modality and process. Agents will also use memory to optimize tool calls based on previous interactions. In the final agency stage, agents will autonomously discover tools, reconfigure tool chains or even build their own tools.

Learning and behavior adaptation — The methods agents use to learn and adapt their behavior will significantly change over the next three years. The current use of HITL feedback and manual, offline fine-tuning will change to be more dynamic and adaptive. Agents will continually learn and adapt from new data and interactions, as well as apply knowledge gained in one domain to other tasks. This means that agent training will evolve from being mostly static and offline to a dynamic, online process that renders agents continuously evolving and improving based on context or application. This fundamental shift in how agents learn will improve their ability to navigate new tasks and information, as well as support more probabilistic behaviors.

Memory — There are two main upcoming changes in memory. In the near term, agents will go from supporting short-term memory and session-only context retention to long-term memory and recollection of past interactions. This retention of information is linked to future agents to learn from past experiences, better support complex or ongoing tasks, as well as prove user-based personalization. In the longer term, agents will demonstrate ecosystem memory or memory sharing between agents, enabling collective learning and distributed problem solving.

Table 1: Agentic AI Capability Dimensions by Agency Stage

(Enlarged table in Appendix)

Capability Dimensions/Agency Stage	L1, AI Assistant	L2, AI Agent	L3, Advanced AI Agents	L4, Expert Agents	L5, Agent Ecosystems
Goal orientation and orchestration	Enterprise knowledge Q&A Basic enterprise application handoff	Simple task automation Limited, well-defined processes Simple collaboration among first-party agents	Complex, multistep task automation Broad processes Complex collaboration between first-party agents Emerging collaboration with third-party agents Scaled, multiagent architectures (in-platform)	Domain-specialized knowledge (industry, function and use case) and process automation Seamless collaboration between first- and third-party agents	Business-level, spontaneous swarm objectives Seamless, cross-system process automation Ecosystemwide system interaction First- and third-party agents negotiate and transact across systems
Autonomy	No autonomy Limited, rule-based actioning	Low autonomy Conditional actioning Simple processing (limited front- and back-office systems)	Mid autonomy Multistep digital processing across a diversity of front- and back-office systems (e.g., IT, HR, ERP, CRM) Emerging physical actuation (sensors, Internet of Things devices, assets)	High autonomy Actioning in domain-specific and regulated systems On-device applications (e.g., laptops and phones) (that is, runs on edge versus cloud)	High to full autonomy Unprompted/self-defined action Physical actuation in critical infrastructure
Reasoning and planning	Recommendation-focused	Follow rules prescribing predefined tasks Limited reasoning capabilities Simple, multistep planning	Reason through undefined tasks with multimodal input Novel responses Reactive adaptation of plans to new info	Reason through complex, undefined tasks Readily navigate uncertainty Domain-specific planning and judgment Complex, long-term planning	Reason through end-to-end, cross-system processes Spontaneous and self-initiated reasoning Dynamic adaptation and replanning across agent ecosystems
Perception and environmental interaction	Limited environment sensing	Highly reactive Retrieval-augmented generation (RAG) contextual awareness	Emerging adaptive perception Multimodal sensing and combined data understanding	Highly adaptive Nuanced understanding of complex environments	Anticipates and proactively navigates uncertainties Holistic view of large or complex environments
Tool Use	Single-step and linear tool use Tool use is hard-coded	Plug-and-play APIs Tool use is flexible, but limited	Multitool orchestration (sequence or parallelize calls for complex tasks) Tool utilization memory for call optimization Adaptive tool chaining based on intermediate results	Domain-specialized integrations expertly used Multistep, multimodal tool use	Autonomous tool discovery Autonomously reconfigure toolchains Agents build their own tools
Learning and behavior adaptation	Models trained on labeled data	Reinforcement learning with human feedback (RLHF) Outcome explainability Model fine-tuning Learning is mostly "offline" and static Behavioral change is limited due to scripted flows	Self-reflection Contextual adaptation Online and real-time learning Active learning (proactively seek feedback)	Transfer learning accelerates adaptation to new tasks Multimodal learning More probabilistic behavior	Self-adaptation Emergent behavior Collective intelligence for effective problem solving Enforce operational standards and adapt as the ecosystem evolves
Memory	Stateless, single-session context retention	Short-term memory Session-only content retention Structured and unstructured data	Long-term memory Contextual and semantic memory (recall context, intent and user preferences) Self-updating memory	Agents choose storage strategies (e.g., what to cache, archive or forget) Shared and synchronized across multiple agents and modalities (text, audio, video and sensor data)	Distributed memory (access and synchronize memory across devices, platforms or agents) Self-reconstructing and generalizable memory architecture Ecosystem memory for collective access

Source: Gartner (August 2025)

Near-Term Implications for Product Leaders

Goal orientation and orchestrated collaboration — The goals of agents will change significantly over the coming decade, and future R&D investments must proactively reflect this vision. Today, most agentic offerings support simple task automation or L2 agency. Emerging innovators are working to support more complex and domain-specialized tasks, multimodal workflows, and end-to-end processing (that is, primarily L3 agent goals). Examples of such vendors uncovered in Gartner's 2025 Agentic AI CBR project include XMPro, AIrrived, Lyzr and Relevance.

Most of today's agentic offerings were reactively built and are poorly defined, fueling market confusion around agentic AI. Vendors need to consider the purpose and future objectives of their AI agents and build an R&D roadmap that delivers on this vision and differentiated capabilities.

Regarding collaboration, today's frameworks fall short of what is required for effective agent negotiation, particularly between first- and third-party agents. Much work is still required for the level of interoperability that L4 and L5 agency levels predict. In the near term, many more protocols and frameworks will come to market, where innovations will compete for market share. This will be followed by market consolidation as usage concentrates around the most performant offerings and agentic communications standardize.

Autonomy — There is a drastic difference in the autonomy capabilities of L2/L3 agency compared with L4/L5 agency. At the L2/L3 agency phase, vendors will be focused on driving the adoption and subsequent scaling of AI agents. Success at this stage will be linked with AI agent accuracy, reliability and explainability. Expert-level agency (L4) will unlock lots of new and innovative applications and unique value outcomes that go beyond productivity and align with solving business problems.

Over the next five years, most agents will be focused on internal front- and back-office operations and optimization. This is an easily reachable agentic opportunity, as the key pain points of adopter organizations are linked with IT, HR, ERP and CRM systems. However, multistep, multimodal and domain-specialized processing will become more critical to vendors' ability to compete as the market progresses and develops deep specialization at the advanced and expert agents' stages.

Vendors' near-term priority is driving the adoption and scale of current AI agents and emerging advanced agents (L2/L3). Future adoption of expert and ecosystem agents (L4/L5) will require both significant technical innovation as well as trust in agents to perform at high autonomy in deeply specialized, regulated and complex environments.

Reasoning and planning — Key advancements in reasoning and planning and associated features will not be realized until L3 — advanced agents and beyond. Today's language models are quite limited in their ability to reason through undefined tasks, navigate uncertainty, proactively engage with domain-specialized tasks, or render reliable judgments on static or dynamic scenarios. Emerging model innovation will help enable these capabilities.

Perception and environmental interaction — The current market places a premium on agent reliability versus adaptability. This is because accurate and reliable performance and explainable outcomes are key for trusted automation. Agentic AI is an emerging technology that has yet to gain the market's trust. Trusted automation will help vendors drive customer adoption and the scale of AI agents. See Emerging Tech: Enabling Agentic AI Adoption With a 'Trustworthy-by-Design' Strategy for more details.

As agents begin to scale and then evolve to expert-level agency, the premium placed on reliability will shift to proactivity. There will be demand for features and functionality related to agents' ability to anticipate and adapt to their environments. Though the true promise of agentic AI includes proactive actions, navigating uncertainty and providing automation in unforeseen events, the technology simply does not yet exist.

Tool use — Tool use will become more advanced and intuitive as AI agents mature. It will progress to support scalable agent adoption within the enterprise, as well as more complex and specialized tasks. Notably, advancements in tool use will be enabled by the wider adoption of open-standard protocols to simplify tool integration and use. Standardization will help facilitate both the scale and specialization of AI agents.

Learning and behavior adaptation — The current primary techniques for agent learning are offline and static — RLHF and model fine-tuning. This will soon change. Providers must start investing now in more dynamic learning techniques, such as self-reflection, contextual adaptation and transfer learning. Self-reflection is a learning technique where the agent evaluates and diagnoses its own performance for adjustment. In contextual adaptation, an agent adjusts its behavior based on the characteristics of a situation. Transfer learning is also promising, where agents will apply knowledge gained in completing a task to inform future actions. The learning approaches will unlock new levels of agent performance and customization unachievable by today's static, offline learning techniques. The future of learning is online and dynamic.

Memory — Do not underestimate the importance of memory, as progress here is linked to two other key capability dimensions: learning and behavior adaptation and automation. However, progress around memory retention and sharing will also introduce governance requirements around knowledge access, provenance, privacy and usage.

More importantly, future dynamic memory architectures will potentially change how data is managed and used by agents. Memory architectures that autonomously structure and optimize based on long-term learning goals will enable agents to more effectively and seamlessly use "dirty data," removing the data-ready approach for AI agent projects.

Recommended Actions for the Next Six to 18 Months

- Stop reacting to change, and plan for continued agentic AI disruption by ideating product features and company evolution against the agency maturity levels (L0 through L5) and timelines.
- Innovate toward future AI agent autonomy and complexity by investing in advanced reasoning models and planning capabilities, multimodal understanding and environmental perception, domain-specialized tool use, and long-term memory.
- Target current market opportunity by focusing on agent reliability, accuracy and performance. Anticipate future demand for differentiating capabilities, such as proactivity, adaptivity and dynamic learning, that will unlock the next generation of more dynamic, intuitive and customized AI agents.
- Future-proof your AI agent strategy by assessing the capabilities of far-future agents — such as full autonomy, distributed actioning across ecosystems and collective behaviors — and investing in enabling technologies in the midterm to long term due to low market risk tolerance and tech immaturity.

Evidence

[1] 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, among other things. The second interview focused on case studies, to include the business problem addressed, the processes that were changed, the outcomes achieved and the unmet needs. This research 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: The Key Defining Characteristics of Agentic AI

Emerging Tech: Agentic AI Innovation Will Foster Autonomous Business

Emerging Tech: Enabling Agentic AI Adoption With a 'Trustworthy-by-Design' Strategy

When to Use or Not to Use AI Agents

Table 1: Agentic AI Capability Dimensions by Agency Stage

Capability Dimensions/Agency Stage	L1, AI Assistant	L2, AI Agent	L3, Advanced AI Agents	L4, Expert Agents	L5, Agent Ecosystems
Goal orientation and orchestration	Enterprise knowledge Q&A Basic enterprise application handoff	Simple task automation Limited, well-defined processes Simple collaboration among first-party agents	Complex, multistep task automation Broad processes Complex collaboration between first-party agents Emerging collaboration with third-party agents Scaled, multiagent architectures (in-platform)	Domain-specialized knowledge (industry, function and use case) and process automation Seamless collaboration between first- and third-party agents	Business-level, spontaneous swarm objectives Seamless, cross-system process automation Ecosystemwide system interaction First- and third-party agents negotiate and transact across systems
Autonomy	No autonomy Limited, rule-based actioning	Low autonomy Conditional actioning Simple processing (limited front- and back-office systems)	Mid autonomy Multistep digital processing across a diversity of front- and back-office systems (e.g., IT, HR, ERP, CRM) Emerging physical actuation (sensors, Internet of Things devices, assets)	High autonomy Actioning in domain-specific and regulated systems On-device applications (e.g., laptops and phones) (that is, runs on edge versus cloud)	High to full autonomy Unprompted/self-defined action Physical actuation in critical infrastructure

Reasoning and planning	Recommendation-focused	Follow rules prescribing predefined tasks Limited reasoning capabilities Simple, multistep planning	Reason through undefined tasks with multimodal input Novel responses Reactive adaptation of plans to new info	Reason through complex, undefined tasks Readily navigate uncertainty Domain-specific planning and judgment Complex, long-term planning	Reason through end-to-end, cross-system processes Spontaneous and self-initiated reasoning Dynamic adaptation and replanning across agent ecosystems
Perception and environmental interaction	Limited environment sensing	Highly reactive Retrieval-augmented generation (RAG) contextual awareness	Emerging adaptive perception Multimodal sensing and combined data understanding	Highly adaptive Nuanced understanding of complex environments	Anticipates and proactively navigates uncertainties Holistic view of large or complex environments
Tool Use	Single-step and linear tool use Tool use is hard-coded	Plug-and-play APIs Tool use is flexible, but limited	Multitool orchestration (sequence or parallelize calls for complex tasks) Tool utilization memory for call optimization Adaptive tool chaining based on intermediate results	Domain-specialized integrations expertly used Multistep, multimodal tool use	Autonomous tool discovery Autonomously reconfigure toolchains Agents build their own tools
Learning and behavior adaptation	Models trained on labeled data	Reinforcement learning with human feedback (RLHF) Outcome explainability	Self-reflection Contextual adaptation Online and real-time learning	Transfer learning accelerates adaptation to new tasks Multimodal learning	Self-adaptation Emergent behavior Collective intelligence for effective problem

		Model fine-tuning Learning is mostly “offline” and static Behavioral change is limited due to scripted flows	Active learning (proactively seek feedback)	More probabilistic behavior	solving Enforce operational standards and adapt as the ecosystem evolves
Memory	Stateless, single-session context retention	Short-term memory Session-only content retention Structured and unstructured data	Long-term memory Contextual and semantic memory (recall context, intent and user preferences) Self-updating memory	Agents choose storage strategies (e.g., what to cache, archive or forget) Shared and synchronized across multiple agents and modalities (text, audio, video and sensor data)	Distributed memory (access and synchronize memory across devices, platforms or agents) Self-reconstructing and generalizable memory architecture Ecosystem memory for collective access

Source: Gartner (August 2025)