Innovation Spectrum: Agentic AI Leadership Requires Adaptive, Explainable & Collaborative Capabilities

5 November 2025 - ID G00840927 - 20 min read

By: George Brocklehurst

Initiatives:Emerging Technologies and Trends Impact on Products and Services

As agentic AI offerings evolve and enterprise buyers' expectations shift, yesterday's product differentiators are becoming table stakes. Platforms that don't adapt quickly risk irrelevance. To stay in the race, tech leaders must move beyond static automation to embrace adaptive, explainable and collaborative capabilities.

Overview

Key Findings

- Most agentic AI solutions still depend on periodic, manual retraining, creating hidden friction ("adaptation tax") that slows enterprise response and erodes automation ROI.
- Multiagent orchestration is now standard, but most platforms only coordinate tasks — not context or reasoning — so true collaboration across complex workflows remains limited.
- Lagging providers rely on siloed audit logs and access controls, lacking real-time oversight. This reactive approach fails to proactively support safe, scalable, trusted agentic operations.
- Providers with deep industry logic face obsolescence if they don't pair expertise with adaptive, learning-driven features; static orchestration will be commoditized without continuous improvement.

Recommendations

Product leaders working to differentiate their offerings with agentic AI capabilities must:

- Preempt agentic behavioral adaptation tax by investing in real-time and persistent learning mechanisms paired with memory that can capture and store relevant context and experiences.
- Improve agent coordination by building a unified context layer that enables agents to share and update information automatically, eliminating manual data handoffs.
- Maximize safety and efficiency by deploying real-time, context-aware oversight with risk-aligned thresholds that proactively monitor agent operations and trigger intervention only when necessary.
- Safeguard industry relevance by embedding domain expertise into orchestration through configurable agent libraries, expert-driven workflow design and deep enterprise integration.

Analysis

Overview

The very leading edge of agentic AI capabilities is entering a phase where the state of the art is not just about adding more autonomy or automation, but about managing the complex interdependencies between reasoning, perception, learning and memory. As enterprises push for greater value from agentic AI, the friction of behavioral adaptation — efforts and costs to keep systems relevant and effective — becomes a central challenge to sustaining cost-effective value.

Only those agentic solutions that enable persistent, scalable adaptation without prohibitive overhead will deliver sustainable and scalable ROI.

The future belongs to agentic systems that perceive across modalities, remember contextually, and collaborate transparently. These capabilities are essential for scaling automation, as enterprises must trust system decisions, understand their reasoning and rely on seamless agent collaboration. Explainability and transparent interaction are critical for continuous examination and tuning of system effectiveness, but they are more than technical features — they are prerequisites for organizational adoption, reduced operational friction and deeper user engagement. The next generation of enterprise agentic AI will be defined by how well platforms balance these forces, providing organizations with the control, visibility and interoperability needed to implement, scale, adapt and innovate with confidence. See Figure 1 for a visual of agentic AI capabilities across a spectrum from leading edge to lagging.

Figure 1: Innovation Spectrum — Agentic Solution Capabilities

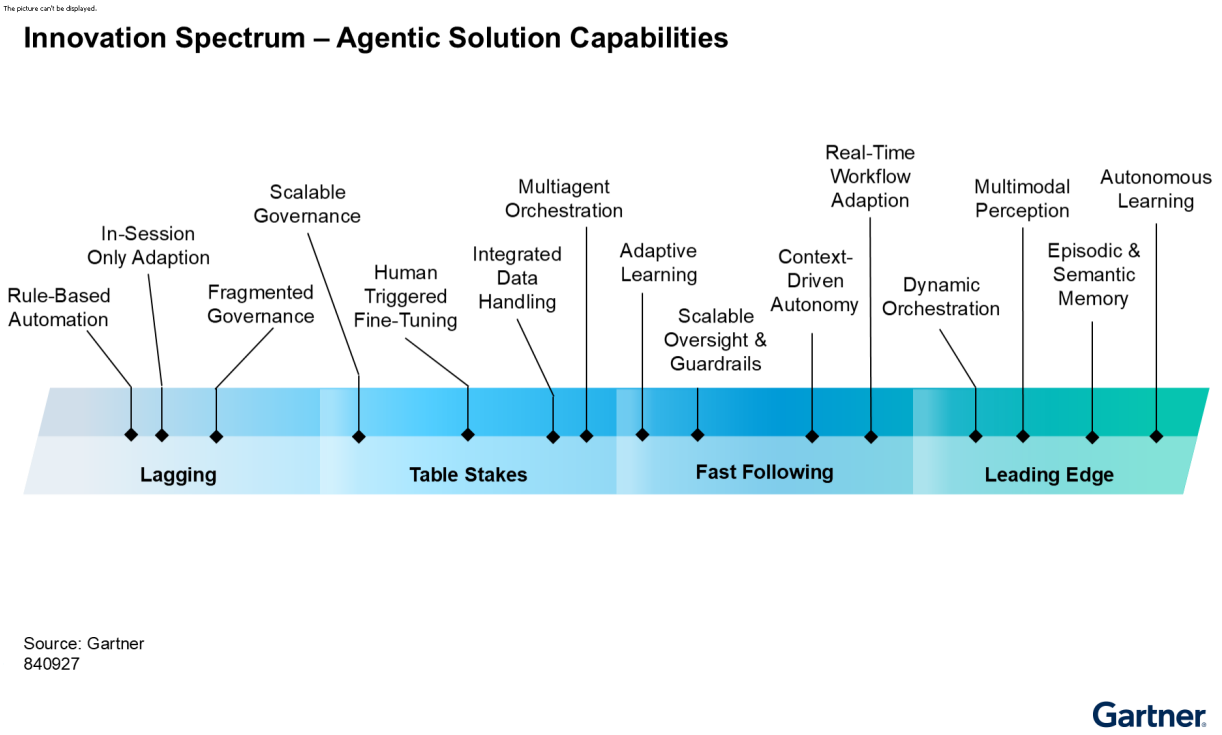

Source: Gartner
840927

Gartner

Product leaders who are intent on shaping the future of agentic AI must focus on pioneering these advanced capabilities because only those at the leading edge will be positioned to unlock the most complex, high-value opportunities as the market matures. While foundational agentic technologies remain relevant and valuable across the broader spectrum of deployments, this research is focused on the strategies and innovations required to advance the leading edge of agentic solutions into more sophisticated and impactful enterprise workflows.

The Future Belongs to Agents That Perceive, Remember and Collaborate

Across much of the market, advancing more complex agentic solutions are constrained by limited perception, fragmented memory and siloed execution — factors that not only hinder automation but also erode enterprise trust. When agents can't integrate and interpret diverse information, recall prior actions or collaborate across workflows, organizations face repeated errors, inconsistent decisions and a lack of transparency. This undermines confidence in agentic systems and stalls broader adoption, as users are reluctant to delegate critical processes to agents they cannot fully understand or rely on (see analysis on leading-edge capabilities addressing this need).

The inability to perceive broadly and remember contextually also means missed opportunities for proactive support, seamless knowledge transfer, and adaptive collaboration — capabilities essential for building trust at scale. Leading-edge capabilities are beginning to address these gaps, but most of the market is still catching up.

This research is based on analyzing 100 providers, [1] with solutions that would qualify as agentic per Emerging Tech: The Key Defining Characteristics of Agentic AI. These providers and the data used are skewed toward the leading edge of the market and the successes in agentic deployments. Also, these providers were stack-ranked based on their agentic capabilities. The top 10% were taken as leading edge, 70-90% as fast followers, 40-70% as table stakes, and <40% as lagging.

Leading-Edge Hallmarks Include Adaptive, Perceptive and Persistent Agents

Leading-edge agentic platforms enable agents with deep situational awareness and real-time adaptability. The result is greater operational resilience, accelerated innovation, and the foundation for collective intelligence that compounds over time.

Autonomous Learning and "Actionable Recall" Enable a Compounding Market Advantage

Persistent, organizationwide learning and memory are emerging in leading agentic platforms. Autonomous learning enables agents to refine actions, using techniques such as reinforcement learning, process mining and continuous adaptation from operational outcomes. Memory is synergistic, it's not just storage — it's structured recall. Agents retain episodic memory (what happened, when and why) and semantic memory (facts, relationships reasoning), allowing them to explain decisions, share knowledge across workflows and avoid repeating mistakes. For example, platforms like XMPro use dynamic knowledge graphs to make learning cumulative and accessible, so agents can recall best practices and context from previous events, not just raw data. AIRRIVED supports persistent adaptation by enabling agents to learn from ongoing data flows and user interactions, while Ema leverages feedback loops to improve agent responses over time. However, implementing persistent learning and memory requires careful attention to data quality, privacy and governance, as well as managing the risks of model forgetting or compounding errors if poor decisions are propagated.

Every agent interaction becomes an opportunity for improvement, with lessons learned compounding over time. Organizations can progressively automate increasingly complex and variable tasks, adapt to subtle yet important shifts in workforce behavior, and build institutional knowledge that would otherwise be lost to turnover or silos. Persistent learning and memory will be a key factor in maintaining operational resilience and building an automation layer that is difficult for competitors to replicate.

Multimodal "Perception" Reduces Operational Blind Spots and Advances Decision Quality

While many platforms now support multimodal ingestion — from IoT sensors, documents, images, video and speech — advanced solutions are moving to perception. This means agentic solutions (and their data fabric) not only ingest, but also fuse and reason over multiple sources and modes of data in real time, correlating sensor alerts with maintenance manuals, extracting insights from visual inspections and combining real-time and historical data to spot emerging issues. For example, C3.ai agents fuse sensor data, drone images and business documents to detect equipment problems and recommend actions. AIRRIVED applies vision and language models to process documents, images and streaming alerts concurrently to detect risks such as data loss prevention violations. Accenture uses semantic layers and knowledge graphs to unify perception across data types improving shared agentic context awareness. This depth and breadth of perception helps agents connect events across systems, recognize patterns missed by siloed tools and deliver more reliable, adaptive decisions. However, integrating and synchronizing diverse data types at scale remains technically complex, and poor data alignment or gaps can lead to misleading signals, false positives or costly errors.

To address more complex, volatile and data rich environments, multimodal perception — and the ability to extend into new modalities — becomes essential for higher ROI automation. Agents equipped with these capabilities can automate tasks that require understanding both physical and digital context, respond faster to operational risks, and close information gaps that cause delays or errors. Organizations that invest in multimodal, context-aware agentic platforms will be able to automate more sophisticated processes and maintain better situational awareness, reducing blind spots and improving decision quality at scale.

Dynamic Cross-Domain Orchestration Streamlines Processes and Accelerates Business Outcomes

Dynamic orchestration in agentic platforms empowers agents to adapt tasks across workflows, systems and environments, integrating diverse resources in real time. In advanced solutions, this orchestration is supported by structured agent-to-agent communication sharing relevant context improving collaboration. Unlike rigid, predefined workflows, these agents can select optimal resources for each step, modify execution sequences as conditions change, and run parallel, specialized processes. Leading capabilities support integration with a wide range of tools and enable agents to collaborate across multiple digital, and in some cases physical, domains. For example, Accenture's AI refinery agents decompose complex business goals into actionable subtasks, which are then executed by specialized utility agents, while platforms like Lindy and Automation Anywhere allow agents to dynamically select, sequence and delegate tool use. This dynamic tool orchestration is emerging as a key differentiator for leading platforms, enabling automation of more complex and variable workflows. However, it also introduces new challenges around integration complexity and exception handling, which must be carefully managed to avoid operational risks.

Dynamic orchestration is core to enterprise agility, which will enable agents to coordinate across platforms and external networks as interoperability standards advance. As environments grow more complex, orchestration must evolve toward federated context sharing — facilitating selective, secure information exchange under strong governance. Providers like Accenture, Lindy, Automation Anywhere and XMPro now offer cross-platform orchestration via API integrations. However, significant deployment is yet to emerge. Providers investing in these capabilities will be more adaptable, but must also prioritize robust monitoring, exception management and fault tolerance to manage rising complexity.

Near-Term Implications and Actions

Recommendations here are critical for product leaders looking to lead in the AI race, but also for the roadmaps of those looking to fast follow:

- **Integration Complexity Requires Disciplined Oversight:** As organizations coordinate automation across increasingly diverse domains and data sources, complexity increases. To avoid new points of failure and maintain trust, providers must adopt disciplined practices in explainability, governance, and exception handling. Multiagent orchestration is already table stakes, therefore these capabilities now become an AI race priority for scale.

- **Continuous Learning and Context Engineering:** Organizations should establish explicit, actionable learning loops and implement federated context layers. These practices enable agentic systems to recall, explain, and share relevant context seamlessly — supporting auditability, collaboration, and the prevention of repeated mistakes. Investing in these capabilities ensures ongoing operational enhancement and institutional knowledge retention, which leads to the confidence enterprise customers need to reshape their organizations for transformational ROI.
- **Compounding Advantage Through Adaptable Automation:** Early adopters leveraging persistent learning, multimodal perception, and dynamic orchestration are building a compounding advantage in operational resilience. These capabilities are currently most prevalent among the top 10% of providers analyzed, but will rapidly diffuse to fast followers and then become table stakes within 18 months. As environments diversify, static or siloed automation is increasingly a liability.

Fast Followers Are Redefining Autonomy Through Context and Control

Fast follower platforms enable agents with the freedom to act, escalate or seek approval shaped by operational context, risk and user feedback. These agentic solutions can better balance autonomy, adaptability, and governance, setting the stage for even more advanced agentic innovation.

Context-Driven Autonomy Concurrently Increases Safety and Efficiency in Agent Operations

Agents in fast follower platforms are increasingly able to adjust their level of autonomy in real time, based on operational risk, task complexity, and user input. Rather than a one-size-fits-all model, these systems use dynamic permissions, escalation protocols, and confidence thresholds to balance independence with oversight. For instance, BusinessNext AgentNext and Datarobot platforms allow agents to handle routine tasks independently but escalate exceptions or seek approval for high-impact actions. This ensures automation aligns with enterprise risk tolerance and compliance requirements, reducing the likelihood of errors or unintended consequences. However, calibrating autonomy is complex (predicated on the quality of context) — too much independence can introduce risk, while too little stifles efficiency.

The challenge lies in designing agents that can accurately assess when to act independently and when to defer to humans. Poor calibration can lead to either excessive interruptions or risky decisions. As agentic automation expands, organizations will expect agents to self-assess their confidence, recognize their limits, and adapt to changing business priorities. Platforms that enable this context-driven flexibility will be better equipped to scale automation safely, maintain operational control, and build trust with users and regulators. This capability is set to become a standard requirement as enterprises move toward more autonomous operations.

Real-Time Workflow Adaptation Minimizes Disruption and Deployment Friction

Agents in fast follower platforms now adapt workflows dynamically as conditions change, using real-time information to adjust actions, escalate issues, or re-plan tasks. This adaptation is primarily focused on responsiveness within a single workflow or process. For example, IBM's planner agent can revise execution based on feedback, while Fairmarkit's agents autonomously manage sourcing events end-to-end, adapting workflows on the fly and handling exceptions or ambiguities by escalating only when necessary. This capability moves automation beyond rigid rule-following, allowing workflows to actively respond to operational changes, exceptions, and evolving business needs. The result is reduced downtime, fewer manual interventions, and improved process continuity. However, implementing real-time adaptation requires robust monitoring and exception handling to avoid process loops or error amplification. Deployments today are typically hybrid: rule-based with levels of agent initiative adaption.

Organizations that get workflow adaptation right will be better equipped to handle uncertainty, minimize disruption, and capture new opportunities as they arise. As agentic automation matures, the expectation will shift to systems that can sense, reason and adjust on the fly — making real-time adaptation a baseline for enterprise automation strategies. Fast followers are showing how adaptive workflows can support business agility, but success depends on clear escalation paths and well-defined boundaries for agent action.

Oversight to Collaboration: Organization-Specific Learning Through Human-in-the-Loop

Fast follower platforms systematically embed human expertise into agent workflows, treating human input as a targeted source of operational learning. Agents autonomously handle routine tasks but are designed to escalate exceptions or ambiguous cases to human reviewers. The feedback and corrections provided are captured and applied to similar future scenarios, enabling agents to improve within specific workflows. For example, Orby's invoice reconciliation agents route low-confidence matches to human auditors, who validate and correct outcomes; these interventions are then used to refine agent logic for subsequent cases. Datarobot's agents prompt users for feedback at key steps, supporting real-time course correction. This approach ensures automation remains flexible and safe, with human expertise driving incremental improvement where it matters most.

Looking ahead, platforms that make human-in-the-loop collaboration explicit and actionable will be well-positioned for enterprises seeking resilient automation. The ability to capture, structure and reuse human feedback supports higher decision quality and operational trust, particularly in complex or regulated environments. However, the scope of learning remains bounded by the specific process or domain, and generalization across workflows typically requires additional human oversight or retraining. Providers who invest in structured feedback loops will deliver automation that adapts to operational nuance but may lag in broad, autonomous improvement.

Near-Term Implications and Actions

Recommendations here are important for product leaders looking to differentiate on Agentic performance, which will be further enhanced by agentic domain specialisation:

- **Oversight Mechanisms Must Be Context-Sensitive, Not Uniform:** One-size-fits-all guardrails can create unnecessary friction or risk. Prioritize oversight where the business impact or uncertainty is highest, and allow routine, low-risk actions to proceed with minimal intervention.
- **Agent learning will capture subtleties in process drift:** As agents adapt in real time, changes in decision logic may go unnoticed until they accumulate, leading to silent shifts in business outcomes or compliance posture. Use targeted scenario testing — focused on stretch and edge cases — to catch drift or bias that may not show up in routine audits.
- **Workflow Exceptions Will Expose Process Gaps:** As agents adapt workflows in real time, organizations will uncover previously hidden exceptions and edge cases. Use these as opportunities to refine processes and clarify escalation responsibilities, not just as errors to fix.

- **Start with High-Variance, Low-Regret Workflows:** Pilot adaptive autonomy and real-time workflow adaptation in processes where mistakes are recoverable and the value of flexibility is highest. Use lessons learned to inform broader deployment in more sensitive areas.

Multiagent Systems, the New Table Stakes in Agentic AI

Table-stakes providers set multiagent orchestration, integration of both structured and unstructured data, and automation of cross-functional workflows as baseline expectations. As these capabilities become commoditized, the real differentiation will shift to platforms that can deliver more adaptive, real-time, and self-improving agentic systems — raising the bar for what enterprises will demand from the next wave.

Multiagent Systems Mark the New Baseline for Enterprise Automation

Multiagent orchestration is now a baseline capability for enterprise AI platforms, but most solutions remain hybrid — combining agent-driven automation with rule-based workflows and human oversight. While agent teams routinely collaborate, share context and decompose tasks into modular steps, real-world deployments show that adaptability is often bounded by configuration and governance. For example, platforms like OpenstreamAI and Appian enable agents to coordinate across domains and integrate with enterprise tools, but critical decisions and workflow changes typically require human input or follow preset rules. Features such as automated handoffs, shared memory and explainable reasoning have become expected, yet true end-to-end autonomous orchestration is rare. The market has moved past single-agent assistants, but fully adaptable multiagent collaboration is still an emerging frontier.

Supervised Learning Remains the Norm, Autonomous Learning Remains Elusive

Learning and adaptation in enterprise agentic platforms are rapidly evolving, but most offerings still fall short of persistent, autonomous improvement. Today's table-stakes capabilities focus on supervised, feedback-driven learning — where agents update prompts, instructions, or memory-driven context based on explicit human input or implicit agent needs. Some platforms support continuous rating and prompt updates, while others enable agents to adapt outputs in response to changing data or workflows. However, unsupervised or cross-session learning, where agents autonomously generalize from experience or transfer knowledge across domains, remains rare. Real-world deployments show adaptation is typically bounded by privacy controls, governance and the need for traceability, with persistent memory and learning limited to specific workflows or user-driven actions.

Scalable Governance and Transparency for Entry Level Trust

Real-time governance and transparency are now essential elements in agentic AI platforms, moving from afterthoughts to table-stakes requirements. Modern solutions increasingly embed auditability, traceability, and policy controls directly into agent workflows for real-time guidance. This includes persistent memory with versioning, explainable reasoning logs and granular access controls that allow organizations to monitor, review and constrain agent actions. Real-world deployments show that enterprises prioritize the ability to track what agents have done, why decisions were made, and how data was used — especially in regulated industries. Platforms that offer federated context layers and governance dashboards are addressing these needs, ensuring agents operate within defined boundaries and organizational policies.

Near-Term Implications and Actions

Recommendations here are critical for product leaders wanting to keep up with the AI race, however, differentiation will be dependent on agentic domain specialization:

- **Multiagent is now the baseline, not a differentiator:** Users will increasingly expect platforms that can orchestrate across multiple agents and integrate diverse data sources by default. Make seamless agent coordination and data integration the default experience — don't position "multiagent" as a unique value proposition, instead focus on how the level of agent collaboration improves outcomes.
- **Governance is a business enabler:** Users increasingly see fine-grained controls and real-time workflow visibility as essential for expanding automation, not just for compliance. Provide intuitive tools to monitor, audit and adjust agent behavior; treating governance and transparency as core.
- **Incremental learning is a hidden drag on customer outcomes:** Manual retraining or workflow tweaks limit adaptability, which will slow innovation. This will be difficult to recognize until it manifests as user friction. Identify where your product still relies on manual learning; invest in persistent, in-context learning to reduce customer effort and continual improvement, e.g., self-supervised model pipelines that fine-tune and optimize performance based on system-generated signals.

Lagging Capabilities Position a Make-or-Break Moment for Future Relevance

At the lagging edge risk management and operational readiness have taken precedence over exploiting technology value. Future relevance will rely on moving beyond static automation, blending industry expertise with enhanced reasoning and adaptability.

Fragmented Governance Limits Agent Oversight

Lagging platforms typically provide only basic governance features — such as access controls, audit logs and manual approval steps — without unified or transparent oversight across agentic workflows. These controls are often fragmented, requiring administrators to monitor multiple systems or rely on manual interventions to prevent errors. For example, some systems log user actions and restrict access based on roles, but lack real-time visibility into agent decisions or automated compliance checks. Unlike advanced audit trails, which capture the full sequence and rationale of agent actions across workflows, basic audit logs are siloed and offer limited context, making it difficult to understand or trace agent behavior in real time. Governance in these environments is reactive, designed primarily to catch mistakes after they occur, rather than proactively enabling agents to operate safely and autonomously within defined boundaries.

Manual Learning Limits Agent Progress

Learning in lagging platforms remains manual, and is rarely persistent across sessions. Most systems rely on in-session human-driven feedback, prompt adjustments or configuration changes to improve agent performance. For example, agents may update their instructions or outputs in response to explicit user input, or developers may fine-tune models based on periodic reviews. However, these adaptations are typically isolated to specific sessions or workflows and do not accumulate into broader, transferable knowledge. Agents lack the ability to autonomously generalize from experience or share learning across tasks, resulting in repeated mistakes and missed opportunities for continuous improvement.

Automation Is Limited by Rigid Task Boundaries

Lagging platforms focus on discrete, narrowly defined tasks — such as patch management, code deployment or workflow execution — operating within rigid boundaries. These solutions typically rely on rule-based automation, where agents follow preset criteria and sequences, with limited ability to handle exceptions or adapt to changing requirements. For example, patching systems may automate update deployment based on success rates or severity, but require manual intervention for unusual cases or broader workflow changes. Similarly, code automation tools can execute scripts and manage deployments, but rarely orchestrate actions across multiple domains or respond dynamically to unforeseen issues. This approach delivers reliability for routine operations but restricts the scope and impact of automation.

Near-Term Implications and Actions

Recommendations here are a priority for product leaders that need to catch up and get in the AI race:

- **Static automation faces growing pressure:** Reliability and trust are valued, but platforms relying on rigid controls and narrowly scoped automation risk obsolescence as enterprise buyers shift toward solutions that learn and adapt. Identify where your automation is locked into static rules or manual oversight, and begin piloting adaptive features — such as feedback-driven updates — in low-risk workflows. Start by targeting areas with frequent manual interventions or exceptions.

- **Industry expertise can unlock future value:** Deep domain integration remains a unique strength; platforms that embed industry logic and compliance can leap ahead if they combine this with adaptive reasoning or AI-driven improvements. Use your domain-specific workflows as a foundation to experiment with new reasoning models or adaptive features, focusing on areas where regulatory or operational context is a barrier for more generic solutions.

- **Partnerships can accelerate progress:** Platforms that have not invested heavily in advanced AI can still move quickly by leveraging open-source models or partnering for pilots, rather than building everything in-house. Seek partnerships or pilot programs that allow you to integrate adaptive or reasoning-driven AI features, using these as a way to test, de-risk, and accelerate your capability development without sacrificing reliability.

Evidence

[1] Validated against interview data from 235 deployment stories across 100 providers. Deployment stories were extracted from two sources outlined below:

- Gartner's agentic AI case-based research (CBR) project which ran from December 2024 through March 2025. As part of this research, 24 vendors (capturing 101 deployment stories) 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 the case studies, to include the business problem addressed, the processes that were changed, the outcomes achieved and the unmet needs.
  - This document was primarily informed by insights from the first interview.

- Vendor briefings with a further 76 providers gathering a further 124 deployment stories were conducted between January 2025 to June 2025.

Recommended by the Author

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

Use Gartner's Agentic Compass for Effective Product Leadership

Emerging Tech: The Key Defining Characteristics of Agentic AI

Emerging Tech: The Key Defining Characteristics of a Domain-Specific Language Model

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

Agent-to-Agent Economy: Position Now for Autonomous Business Shift