

# Emerging Tech: Customer Trust Is a Critical Barrier to Agentic AI Adoption

2 June 2025 - ID G00829218 - 12 min read

By: Danielle Casey, Alfredo Ramirez IV, Anushree Verma, Akhil Singh, Aakanksha Bansal  
Initiatives: [Emerging Technologies and Trends Impact on Products and Services](#)

Gartner research reveals that a top inhibitor to agentic AI adoption is a lack of customer trust. Vendors that offer observability tools, reliability controls and explainability features will emerge as near-term winners in the agentic AI market.

## Overview

### Key Findings

- C-level executives overseeing agentic AI products and services must recognize customer trust as a primary inhibitor of agentic AI adoption, or they risk losing agentic AI market share.
- Customer trust hinges on enabling explainable, guardrailed automation and human-in-the-loop controls that build user confidence. Over time, these features will facilitate the transition to trusted AI actions without human oversight.
- Low-agency use cases represent the near-term agent revenue opportunity. The market is currently favoring semisupervised, simple task automation over more autonomous, complex agentic task automation.
- Observability tools, reliability controls and explainability features are lagging behind agentic AI offerings. These key innovations will drive scalable agentic AI adoption and business success.

### Recommendations

Product leaders developing agentic AI offerings should:

- Embed explainability and reliability capabilities into the solution architecture by combining language models (LMs) for reasoning with classical AI techniques for task execution, such as machine learning (ML) models, neurosymbolic AI or rule-based logic.

- Build flexibility into agentic workflows by allowing users to choose between semiautonomous and fully autonomous agents that can be easily customized based on user risk tolerance. This capability is already a fundamental requirement for agentic adoption and building customer trust.
- Develop agentic AI use cases that align with customer demand by focusing on low-agency and simple use cases. This will build trust and credibility, paving the way for scaling to more advanced and high-agency applications in the future.
- Equip users with the ability to observe and monitor agent behavior by building tracing, reviewing, reporting and alerting mechanisms into their offerings. Utilize customized performance metrics and dashboards, LLM as a judge, and orchestration and reliability agents to enhance transparency and oversight.

## Strategic Planning Assumption

By 2028, less than 10% of agentic AI deployments will operate unsupervised, up from less than 1% in 2025.

## Analysis

### Technology Description

Agentic AI refers to various architectures, design patterns and frameworks for creating both single-agent and collaborative multiagent systems capable of unsupervised task execution. AI agents are autonomous or semiautonomous software entities that use AI techniques to perceive, make decisions, take actions and achieve goals within digital or physical environments.

Agentic AI differs from earlier agent systems (and robotic process automation [RPA]), which were traditionally guided by scripted workflows. Agentic AI systems may combine language models with classical AI technologies, such as machine learning or neurosymbolic AI. Agentic AI is characterized by agents imbued with the ability to:

- Access historized data views of context-based decisions, executed goals and workflows.
- Adjust goals, decisions, plans and actions dynamically in response to changing conditions.
- Exercise variable levels of autonomy for decision-making capabilities, ranging from semisupervised to fully autonomous.

- Adapt execution, including multiagent collaboration (such as multiagent generative systems), to learn to complete multistep processes within digital and physical environments (actuation).
- Use multimodal, multimodel capabilities for domain-specific context (such as natural language processing [NLP], time-series forecasting, accounting and image recognition) to effectively interpret and execute instructions.
- Utilize “reasoning” abilities for making context-based decisions and judgments.

---

*Agentic AI represents a significant, material advancement in intelligence and automation potential – impacting virtually all software.*

---

**Sample Providers:** Amazon; AIRRIVED; CrewAI; Glean; LlamaIndex; Leena AI; Lyzr; Orby.AI; Relevance AI; XMPPro.

## Primary Findings

According to Gartner’s 2025 Agentic AI Case-Based Research (CBR), customer trust is a top inhibitor of agentic AI adoption. The study interviewed 20 agentic AI providers, of which a majority (more than 50%) cited customer trust as a top challenge to driving customer adoption. (For more information on the CBR methodology, see the Evidence section.)

---

*Lack of customer trust is inhibiting adoption of high agency offerings in favor of human-in-the-loop automation.*

---

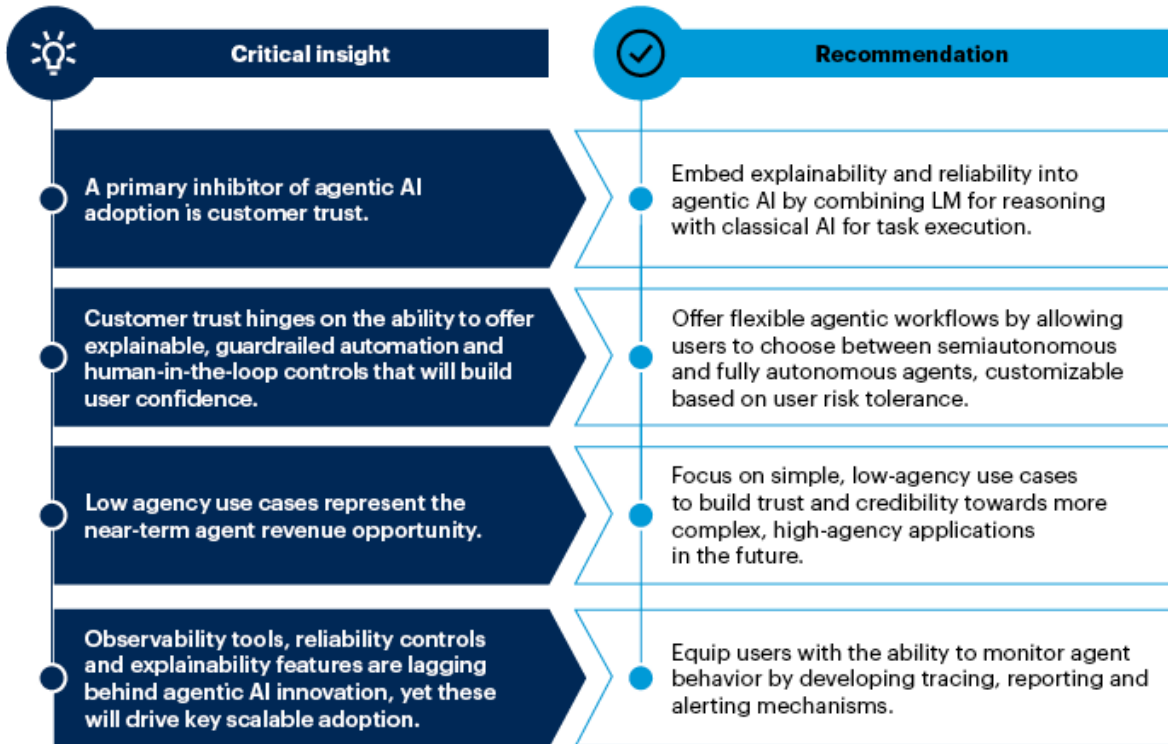
The near-term winners in the agentic AI market will offer “trusted agency,” which is a system that adopter organizations have confidence in to perform automation and deliver value.

This preliminary CBR finding has profound implications for the agentic AI market (see Figure 1). Providers must act accordingly to capitalize on both current and emerging agentic AI revenue opportunities or risk being left behind.

Figure 1: Critical Insights for Agentic AI Providers

**Critical Insights for Agentic AI Providers**

Agentic AI providers must tackle customer trust to drive adoption



Source: Gartner 2025 Agentic AI Case-Based Research

AI = artificial intelligence; LM = language model

829218\_C

**Critical Insight: Customer Trust Is Your Top Challenge to Overcome in the Next Year**

Considering the challenges with generative AI (GenAI) in 2024 around accuracy, hallucinations and value delivery, it is unsurprising that customer trust is reported as a major agentic AI adoption challenge. Agentic AI is the next evolution of GenAI that executes tasks. The reality is that agentic AI is an emerging technology that the market does not yet fully trust to perform accurately and deliver solid value. According to the [Emerging Tech Impact Radar: Generative AI](#), agentic AI is approximately six years from reaching early majority (~16%) customer adoption. The trust challenge that agentic AI presents is neither insignificant nor fleeting, and it will impact both technology capabilities and use cases.

---

*Customer trust will remain a top barrier to agentic AI adoption in 2025. How providers manage this challenge will prove crucial to their market success.*

---

Meanwhile, Figure 2 depicts the current and future market demand for agentic AI.

## **Low Agency – Current Customer Demand**

Currently, most agentic AI use cases are focused on the “low agency” end of the spectrum. Low agency is characterized as being more static and reactive in nature, automating simple tasks, operating in simple environments, and including some level of human supervision. Low-agency use cases tend to be simple and low risk.

Characteristics of low-agency use cases include:

- **Static** – Controls placed on LMs for reliability limit adaptivity.
- **Reactive** – Primarily used to automate predefined tasks within structured workflows, allowing actions to be executed when certain conditions are satisfied.
- **Simple tasks** – Remedial, data-heavy or labor-intensive tasks are ideal candidates for agentic AI.
- **Simple environment** – Tasks are executed within well-defined and preestablished digital environments. The simplicity enables faster time-to-deployment due to lower technology and business barriers.
- **Supervised** – LMs are used to reason to a recommended action. Subsequently, a human will either perform the action, review the recommendation and authorize the LM to proceed, or evaluate and verify agentic automation after task execution, thereby creating a data feedback loop. Regardless of the approach, there is a level of human input and supervision that renders the application semiautonomous. Currently, human-in-the-loop controls predominate.

Current examples of low-agency use cases in the market include:

- Contract and account reconciliation
- Expense reimbursement processing
- Insurance claims review and processing

- Tasks in the HR recruitment process
  - Purchase order and know your customer (KYC) processing
  - Lead verification and sales prospecting
  - Due diligence reporting and assessment generation for mergers and acquisitions (M&A)
- 

*Lack of customer trust and technology immaturity has adopter organizations favoring “low-agency” use cases in the near term.*

---

## High Agency – Future Customer Demand

As market trust of agentic AI grows and the technology improves, “high-agency” use cases will increasingly appear. High agency is characterized as being adaptable and proactive, having the ability to automate complex goals and being deployed in complex environments for autonomous operation. High-agency use cases will be more complex and tackle riskier operations compared to their low-agency counterparts.

Characteristics of high-agency use cases include:

- **Adaptive** – Adjusts goals, plans and actions based on new and evolving conditions.
- **Proactive planning** – Anticipates future events (such as needs and opportunities) and initiates actions via goal-directed behavior.
- **Complex goals** – Handles complex tasks with competing requirements, goal negotiation and multilayered data inputs, typically facilitated via multiagent generative systems (MAGS).
- **Complex environment** – Both simple and complex tasks will be executed in digital and physical environments characterized by dynamic and versatile information, time-bound pressures, access controls and stakeholder considerations.
- **Autonomous** – Possesses the agency to execute tasks independently, without human input or supervision.

Current examples of use cases with medium to high agency in the market include:

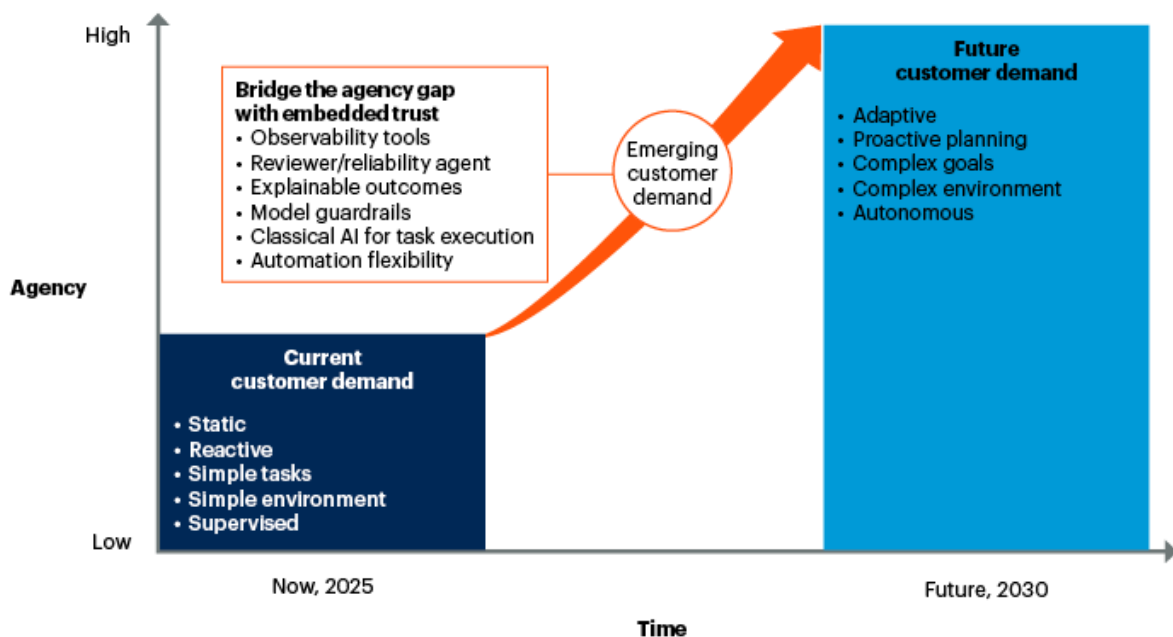
- IT agent for autonomous diagnostics, troubleshooting and solution implementation.

- Multiagent solution for factory remote infrastructure management.
- Enterprise software development life cycle (SDLC) management – to include proactive code testing, code iteration and patch generation.
- Predictive maintenance scheduling and parts inventory management for manufacturing and utilities.
- Supply chain logistics optimization.

Product leaders must mind the agency gap depicted in Figure 2 by focusing on low-agency use cases in the near term, while planning for high-agency use cases in the long term. Customer trust and technological maturity are key factors that will help vendors to transition from low- to high-agency applications.

**Figure 2: Bridge the Agency Gap Between Current and Future Market Demand**

### Bridge the Agency Gap Between Current and Future Market Demand Illustrative



Source: Gartner  
829218\_C

### Near-Term Implications for Product Leaders: Early Agentic AI Winners Will Offer Trusted Agency

Providers that invest in technologies, features and capabilities that facilitate the trusted use of agentic AI will emerge as near-term winners. There is massive competition and emerging revenue opportunities in the agentic AI market. Providers capable of driving adoption of their agentic AI offerings now will benefit from scaling their revenue streams ahead of their peers. To achieve this, providers should evaluate how they can offer trusted agency. There are currently several observed approaches – offering observability and reliability controls, combining classical and GenAI techniques, and embedding flexibility in agentic solutions.

**Observability and reliability controls:** Some vendors report using neurosymbolic AI, guardrails or a reliability/reviewer agent to verify the work of other AI agents, to demonstrate designed control and embedded reliability. For example, Lyzr has a reliability agent, Beam AI uses neurosymbolic AI in its architecture and AIRRIVED provides advanced agentic controls. There is also the emerging use of LLM as judge and orchestrator agents. Less sophisticated approaches include using a data feedback loop for agent monitoring, improvement and ongoing customer personalization.

Dashboards play a key role in allowing users to observe and monitor AI agent performance and task automation. It is important that vendors offer tools for agent observability and performance improvement so that enterprises can ensure agents are behaving correctly. As users develop trust in AI performance, implementations will progress from human-in-the-loop reviewers to unsupervised automation. Don't underestimate the importance of observability and reliability tools in building customer trust and educating users about agentic systems. Currently, these tools lag behind agentic AI innovation.

---

*Plan for significant innovation around agentic AI tooling for observability, reliability controls and reporting. This tooling will play an increasingly important role in building customer trust and driving adoption of agentic automation.*

---

**Classical AI techniques embed reliability:** Many interviewed providers did not use GenAI for task execution; rather, they used classical AI technologies, such as ML models and rule-based logic. For example, using LM for reasoning but ML for task execution. This hybrid approach to agentic AI embeds transparency and reliability into task automation.

Hybrid AI not only unlocks immediate technology trust benefits but also better positions vendors for high-agency use cases. LMs are not designed to do everything and often perform best when used according to their strengths and combined with other AI technologies. A GenAI-only approach is not ideal for providing the complexity and controls needed for more advanced use cases (such as prediction) or for deployment in environments (such as factory floors). Examples of vendors pursuing combinatorial innovation include AIRRIVED, Lyzr and XMPPro.

---

*A GenAI-only approach to agentic automation may inhibit your ability to build customer trust and scale your offering to more complex use cases.*

---

**Agentic solutions flexible by user risk thresholds:** Every agentic AI platform and solution provider interviewed supports flexibility in the models used, tasks supported and automation allowed. Allowing the user to determine where they want agentic automation versus human-in-the-loop intervention allows customers to customize agentic AI systems to their specific risk thresholds. These configurations can be adjusted over time as trust in agentic AI increases. Many vendors are also working to offer prebuilt agent use cases, in addition to providing easy-to-use tools for build-your-own agents. Others offer services to assist with data cleaning and agent customization. Examples of vendors offering flexible agentic AI platforms include Lindy, Orby AI and Supervity.

---

*Prioritize value delivery over “agentic” labels. Trusted, performant solutions with human-in-the-loop flexibility will out-compete truly “agentic” (highly autonomous) use cases.*

---

As agentic systems become more autonomous, security of AI agents will become more important, particularly as it relates to trusted agent actions in collaborative environments.

### **Recommended Actions for the Next Six to 18 Months**

- Embed explainability and reliability into your agentic AI offerings by combining LMs for reasoning with classical AI for task execution, such as machine learning (ML) models, neurosymbolic AI, rule-based logic, or other classical AI approaches. In addition to driving customer trust, this approach will improve your ability to tackle more complex and high-risk use cases.

- Build flexibility into agentic workflow development by allowing users to choose between semiautonomous and fully autonomous agents that can be easily customized based on user risk tolerance. This capability is already table stakes for agentic adoption and will help build customer trust.
- Develop agentic AI use cases that align with customer demand by initially focusing on low-agency and simple use cases. This approach helps build trust and credibility, paving the way for scaling to more advanced and high-agency applications in the future.
- Drive immediate agentic AI adoption by identifying common AI use cases and using agentic AI to provide additional value.
- Equip users with the ability to observe and monitor agent behavior by building tracing, reviewing, reporting and alerting mechanisms into their offerings. Utilize customized performance metrics and dashboards, LLM as a judge, and orchestration and reliability agents to enhance transparency and oversight.

## Acronym Key and Glossary Terms

|       |                                    |
|-------|------------------------------------|
| AI    | artificial intelligence            |
| CBR   | case-based research                |
| GenAI | generative artificial intelligence |
| HR    | human resources                    |
| KYC   | know your customer                 |
| LM    | language model                     |
| M&A   | mergers and acquisitions           |
| MAGS  | multiagent generative systems      |
| ML    | machine learning                   |
| NLP   | natural language processing        |
| RPA   | robotic process automation         |
| SDLC  | software development life cycle    |

## Evidence

Gartner's agentic AI CBR project ran from December 2024 through March 2025. As part of this research, over 20 vendors participated in two interviews each. The first focused on the technology and the second on the case studies. This insight was derived from the first interview on the technology, where vendors were asked about the main customer barrier to adoption and primary go-to-market challenges.

---

## Recommended by the Authors

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

[Top Strategic Technology Trends for 2025: Agentic AI](#)

[Emerging Tech: 6 Ways Generative AI Will Disrupt and Transform Enterprise Applications](#)

[Emerging Tech: Tech Innovators in Generative AI Workflow Orchestration](#)

---

© 2025 Gartner, Inc. and/or its affiliates. All rights reserved. Gartner is a registered trademark of Gartner, Inc. and its affiliates. This publication may not be reproduced or distributed in any form without Gartner's prior written permission. It consists of the opinions of Gartner's research organization, which should not be construed as statements of fact. While the information contained in this publication has been obtained from sources believed to be reliable, Gartner disclaims all warranties as to the accuracy, completeness or adequacy of such information. Although Gartner research may address legal and financial issues, Gartner does not provide legal or investment advice and its research should not be construed or used as such. Your access and use of this publication are governed by [Gartner's Usage Policy](#). Gartner prides itself on its reputation for independence and objectivity. Its research is produced independently by its research organization without input or influence from any third party. For further information, see "[Guiding Principles on Independence and Objectivity](#)." Gartner research may not be used as input into or for the training or development of generative artificial intelligence, machine learning, algorithms, software, or related technologies.