



News Release

New Research: Why Enterprise Agentic AI Stalls Before It Scales

2026-07-07

Study of 1,000 global senior technology and data leaders uncovers what's blocking enterprises from making the leap from personal AI to organizational AI

SAN DIEGO, July 7, 2026 /PRNewswire/ -- **Teradata** (NYSE: **TDC**) today released findings from a commissioned Wakefield Research study of 1,000 senior technology and data leaders across six global markets. The report, "Arrested Automation: Why Agentic AI Stalls at the Enterprise Level," finds that while enthusiasm to deploy agentic AI is near-universal, foundational data systems were not built for agents and need rethinking to deliver the ROI organizations expect.

Many of the hurdles outlined in the report (and summarized below) are easier to understand by recognizing the need to move from personal AI — tools like chatbots and writing assistants that help individuals work faster — to organizational AI, which works on behalf of the whole company using shared knowledge, appropriate access levels, and well-designed governance. The returns enterprises are chasing don't happen until AI operates at the organizational level.

The report introduces the **Agentic AI Maturity Index** to track where organizations stand on that journey and charts a path forward through what it calls Autonomous Knowledge — enterprise data with enough context, lineage, and governance for AI agents to act on it reliably at scale.

The Agentic AI Maturity Index: Where Enterprises Actually Stand

This four-stage framework maps where organizations stand: Experimenting, Developing, Building, and Operationalizing, which is where AI is executing multi-step workflows with measurable business impact. Currently, only **7%** of the global enterprises have reached the final stage where tangible outcomes occur. The majority (**68%**) remain in Experimenting or Developing, where context fragmentation — when data exists but carries no usable

meaning for agents — is a major limiting factor.

Notably, **69%** of C-suite executives say their organization is already operating with agentic AI, while only **57%** of VPs say the same.

Report Breadth: Industry and Country Comparisons

The report breaks down findings across industries including healthcare, financial services, IT, manufacturing, and retail, and across six markets: the United States, United Kingdom, France, Germany, Japan, and Saudi Arabia. The agentic AI challenge is a global phenomenon, but not a uniform one. The research points to several barriers.

The ROI Gap

Nine in ten (90%) senior technology leaders expect to increase their agentic AI investments over the next 12 months; yet **nearly two-thirds (63%)** report they have seen no more than a small or emerging positive return on those investments to date. The gap between investment and returns is not a lack of ambition, but a data foundation that was built for human users, not autonomous AI agents.

"Individual productivity gains — faster code, better drafts, quicker research — are real benefits, but they don't show up on the P&L in a way that justifies significant infrastructure investment. The ROI executives expect requires agents operating at the organizational level: automating decisions, executing workflows, driving measurable business outcomes. Most organizations are measuring enterprise AI ROI against personal AI infrastructure — and wondering why the numbers don't add up."

-Louis Landry, Chief Technology Officer at Teradata

Context Fragmentation

At the core of the agentic AI stall is context fragmentation: enterprise data that lacks the meaning, lineage, and governance AI agents need to act reliably across an organization. According to the report, **77%** of executives report that **20%** or less of their enterprise data is sufficiently described and contextualized for agents to use. And **78%** find it challenging to unify data and knowledge across business functions so agents can reason across the full enterprise.

The top two barriers leaders cite — data lacking the necessary metadata, context, and relationships (**43%**) and data fragmented across systems that cannot be connected in real time (**42%**) — point to the same root problem. The challenge isn't how much data organizations have, but whether that data carries enough meaning to be trusted when agents use it. When it cannot, the pilot does not make it to production. **40%** of tech leaders report that more than **40%** of their AI pilot projects fail to reach production because infrastructure systems were never built for

autonomous use. Only 15% of organizations are successfully getting 80% or more of their AI pilots into production.

"The goal of contextualizing your entire data estate is likely the wrong goalpost, and chasing it is part of why organizations stall. Instead, identify the highest-value portion of your data, structured and/or unstructured, and focus on getting that portion fully described, governed, and agent-ready. If most of the data is unusable, the answer isn't to fix all of it at once. It's to be ruthlessly selective about where you start."

- Josh Fecteau, Chief Data and AI Officer & Chief Information Officer at Teradata

The Action Bridge

Even when organizations make progress on context fragmentation, implementing autonomous action is still hard. 60% of leaders report decision paralysis on durable infrastructure decisions. The hesitation may not be about technology selection (though 30% are worried about vendor lock-in) but instead a lack of trust in what's being deployed. Until organizations trust the data their agents are operating on, they won't let those agents act autonomously. 51% of leaders cite accuracy and reliability of outputs as a significant deployment barrier.

There is also a location problem. AI output currently lives outside the systems where consequential work actually happens. When intelligence is surfaced inside a tool or app where someone is already working, action follows. When it lives in a separate dashboard, it usually does not. Both problems stem from the same deficit: data that lacks enough context, lineage, and meaning to be trusted.

The Path Forward: Autonomous Knowledge

The report identifies Autonomous Knowledge as what organizations need to move from personal AI to organizational AI. It outlines a phased approach: audit and contextualize the highest-value portions of the data estate, embed governance directly into the data layer, and build for architectural portability. Organizations that have done this are already seeing returns. Those that have not are still waiting for their pilots to reach production.

About the Research

Arrested Automation: Why Agentic AI Stalls at the Enterprise Level was conducted by Wakefield Research on behalf of Teradata. The study surveyed 1,000 senior technology and data leaders at the vice president level or above, at companies with a minimum of 500 employees, across the United States (500), United Kingdom (100), France (100), Germany (100), Japan (100), and Saudi Arabia (100). Fieldwork was conducted between March 23 and April 5, 2026.

To download the full report, visit: <https://www.teradata.com/insights/white-papers/why-agentic-ai-stalls-enterprise>

About Teradata

Teradata empowers enterprises to turn intelligence into autonomous action, grounding AI agents in deep business context and trusted data. As AI agents multiply, Teradata is the context foundation, governance layer, and performance backbone that companies need now. The Teradata Autonomous Knowledge Platform puts AI into production across cloud, on-premises, and hybrid environments.

The Teradata logo is a trademark, and Teradata is a registered trademark of Teradata Corporation and/or its affiliates in the U.S. and worldwide.

MEDIA CONTACT

January Machold

january.machold@teradata.com

View original content to download multimedia:**<https://www.prnewswire.com/news-releases/new-research-why-enterprise-agentic-ai-stalls-before-it-scales-302818887.html>**

SOURCE Teradata