

The next big thing in AI is agents, but is your data ready?

Story by Derek Slager

[Artificial intelligence](#) continues to fundamentally change how we do business, and in the past year, a new innovation has entered the spotlight. AI agents are being adopted at record speed across organizations, from [marketing](#) to data management to customer service, with the promise to streamline decisions, engage customers, and boost productivity for companies to drive business value.

We've seen AI agent launches from companies across all sizes and industries. In May, Google announced it would incorporate AI agents in its searches, while Microsoft also announced a plan to use AI agents to help its users search the web. The use of AI agents is surging across industries, from finance and healthcare to car dealerships.

In fact, Boston Consulting Group predicts that the market for AI agents will grow at a 45% CAGR over the next five years. Gartner has also estimated that 80% of common customer service queries will be resolved by AI agents in less than five years.

But here's the catch: agents are only as good as the data they run on.

Why Data Still Trips Up AI

No matter the cutting-edge nature of the AI tool or its sky-high promises, one constant remains when it comes to the data they're operating on: garbage in, garbage out.

Companies racing against competitors to deploy AI agents without taking a step back to evaluate the sources they're operating on face a major risk—if those agents rely on fragmented or inaccurate data, they won't perform as expected. Even the most capable AI systems can't deliver results if they're built on bad information.

According to MIT Technology Review Insights, 78% of global companies are not ready to deploy AI agents and [LLMs](#). What's stopping them? Their data is not

prepared to support AI. At the core of AI's success is unified, accurate and real-time customer data.

When AI agents are powered by bad, disjointed data, the consequences can be costly. Last year, Air Canada was forced to reimburse a customer when its [chatbot](#) promised a discount that didn't exist. And, in April, a tech company suffered fallout after a customer service agent's mistake resulted in a wave of canceled subscriptions.

These types of mishaps can threaten customer loyalty and result in churn. AI agents are only as smart and useful as the [data](#) on which they're built. In order to trust your AI agent, you have to trust your data foundation.

Identity Resolution, Reimagined for Agents

The most essential—and most overlooked—piece of making agentic AI work is identity resolution. Without a clear, accurate view of who the customer is across historically disconnected and fragmented systems, agents are flying blind.

That's changing. AI agents can now take on identity resolution as part of their function, matching records in real time, continuously refining connections and operating without brittle rule-based systems. Rather than depending on static, one-size-fits-all profiles, agentic identity resolution builds a living picture of the customer, improving with each interaction and fostering enhanced [productivity](#) and accuracy.

This means fewer errors, less time-consuming manual data prep and faster time-to-insight for every downstream system.

Getting the Data Foundation Right

Before AI agents can operate effectively, the underlying data must be:

Unified: Data from every touchpoint, ranging from eCommerce and [CRM](#) to customer support, should be stitched together into a single, accessible layer that's usable for marketing and engineering teams alike.

Accurate: Identity resolution must reconcile inconsistencies or duplicates across multiple channels and touchpoints to build a reliable profile.

Contextual: Different use cases need different views. Marketing might need probabilistic profiles for broad targeting, while support needs deterministic, single-session accuracy.

Governed: Access controls, human oversight, feedback loops and consent tracking are table stakes for compliant and trustworthy AI – especially in the wake of evolving privacy regulations.

A modern lake house architecture, paired with AI-native tools for identity resolution and customer profile building, can drastically reduce the manual effort required and make real-time, AI-powered decisions viable.

Data as Competitive Differentiator

Often, data quality is treated like plumbing, which is necessary but invisible. But in the age of AI agents, it becomes a competitive asset.

High-quality, agent-ready data enables better personalization, faster experimentation, and safer automation. It allows AI to act with confidence, knowing who it's interacting with, what they want, and how to best respond efficiently and effectively.

When done right, data doesn't just support AI - it elevates it.

What's Next

Agent-based AI is already reshaping expectations for responsiveness, personalization, and automation. But the true breakthrough isn't in the models, it's in the data. Companies that invest in a high-quality data foundation will now be the ones that make AI useful, reliable, and transformative not only for their operations but also for the end customer experience. That's the difference between a flashy interface and a top-notch algorithm and an impactful, scalable solution. Before you build your next agent, build the data foundation it needs.