How Business Context Solves the Agentic AI Trust Gap

Tamr

The AI Readiness Challenge

Artificial intelligence (AI) has taken hold in organizations worldwide, quickly moving from an experimental pilot to a solution that delivers real business impact. While many organizations are applying AI to isolated use cases—especially chatbots—progressive companies are now incorporating it into analytical systems to improve decision-making and using AI agents to automate processes, make autonomous decisions, and improve business productivity across the enterprise.

The pressure to deploy AI is palpable—and so, too, are the meaningful efficiency and productivity gains. But this rapid growth and adoption also presents serious challenges. **Consider these statistics:**

80%+
of AI projects fail, double the rate of non-AI IT projects

95%
of enterprise GenAI pilots fail to reach production

74%
of companies show no tangible value from AI investments

42%
of companies abandoned most AI initiatives in 2025

up from 17% from the prior year

These statistics point to a number of failure points ranging from misaligned business strategies and organizational resistance to poor workflow integration and a lack of specialized internal talent. But even when an organization does everything right—aligning the executive team, choosing the ideal operational use case, designing a low-friction workflow, and selecting a cutting-edge large language model (LLM)—AI initiatives will often still stall at the exact same bottleneck: **a lack of foundational readiness in the enterprise data itself.**

The absence of AI-ready data is one of the primary obstacles preventing successful AI adoption. When data is fragmented, siloed, inconsistent, flawed, or outdated, the resulting models suffer from “garbage in, garbage out” syndrome, producing insights that are often inaccurate, misleading, or biased. And the real risk is not just that AI will make mistakes, but that it will make mistakes at a velocity and scale that human oversight can’t match. These errors and potential hallucinations—amplified across systems and workflows—erode trust. And once lost, trust is extremely difficult to restore.

Lessons in (AI) Failure

Consider this scenario. A manufacturing company purchases thousands of electronic components from a particular supplier every month. Previously, the manufacturer was owned by a large, blue-chip conglomerate with a pristine credit rating and automated fast-track procurement approval. They always paid on time, so new orders always went through without delay. The company’s new owner, a private equity (PE) firm, on the other hand, is notorious for aggressively managing working capital, often stretching payment terms out or disputing invoices to delay payments.

The supplier’s accounting department receives a notification about the manufacturer’s new ownership, records this information in the finance system, and places a flag in the system to review the terms in place with the PE firm, with whom they also have a relationship. However, the new ownership status for the manufacturer was not recorded in other key systems, including the CRM and fulfillment systems, which continue to show the company is still linked to the old parent company.

Meanwhile, the manufacturer places a new, large order for electrical components. An AI agent in use by the supplier receives the order, checks the CRM and fulfillment systems per usual, and processes the order—valued at over $100,000. Because the supplier’s systems had not yet been updated with the new corporate hierarchy for the manufacturer, the agent processed the order in routine fashion even though it should have been flagged for review.

The AI agent wasn’t wrong—it was just working off of an outdated view of the corporate relationships. And because there was no up-to-date, unified view across CRM, fulfillment, and financial system data, the agent automatically processed a sizable order that would likely lead to a costly payment delay on a large receivable.

In contrast, a human would have known to verify multiple systems before processing such an order and would have paused to investigate. While a human’s processing time would have been longer and likely more expensive, it would avoid the “bull in a china shop” scenario where an AI agent plows ahead without critical context.

While this scenario is hypothetical, it highlights very real challenges that businesses face daily, underscoring the systemic impact that bad, outdated data can have on business operations.

Why Connecting and Consolidating Data Is No Longer Enough

For decades, companies have employed various data management tools and techniques to connect data silos and consolidate data into an enterprise-wide data warehouse or data lake. And for a long time, these repositories served organizations well, providing a centralized data store for use in reporting and analytics. However, as AI accelerates the cycle time of business processes that use this data, these data stores—even when using modern, cloud-based solutions like Snowflake, Databricks, or Google BigQuery—are no longer enough.

While a data warehouse or data lake provides a foundational repository, these architectures struggle with the real-time operational needs of AI initiatives. Data warehouses and data lakes were engineered to aggregate data from multiple, disparate systems into a single, unified environment. And while they are often successful at centralizing the data, warehouses and lakes are not designed to resolve entities into mastered **golden records**—single, authoritative views across systems—causing the information to remain fragmented, inconsistent, and often duplicative.

As a result, data warehouses and data lakes have been unable to support the real-time feedback loops and continuous updates needed to preserve data integrity. In fact, many have deteriorated into proverbial "data swamps," which cause AI agents (which don't know any better) to deliver faulty or misleading insights or take actions based on inaccurate, inconsistent, or outdated information.

Further, while modern cloud data warehouses such as Snowflake provide tremendous value to forward-thinking, data-driven organizations, infrastructure modernization alone does not create AI-ready data. Cloud data warehouses provide centralization and consolidation benefits and can scale to support real-time AI workloads, including large volumes of queries from multiple AI agents. However, today's leading cloud data warehouses are not able to run the advanced matching algorithms required for real-time **entity resolution** or provide the functionality required for mastering data across sources at scale.

Context: The Critical Piece in the Agentic AI Equation

As the use of AI evolves from broad models to autonomous agents, the definition of “AI-ready data” is undergoing a fundamental shift. While LLMs appear highly capable, their insights are restricted to the public data in their training sets. Incorporating proprietary data can improve results for businesses, but static, isolated rows of enterprise data still don’t provide the breakthrough results and insights that organizations need, because the data lacks a shared semantic framework. In other words, when public and proprietary data are siloed, AI agents remain blind to cross-system insights that unlock real business value.

To bridge this gap, organizations must move beyond both publicly available and fragmented enterprise data and provide AI systems and agents with a rich web of entity relationships that define how the business actually functions. **This is context.**

Context is the governed, structural framework of organizational knowledge that enables AI to move beyond basic pattern matching toward true, advanced reasoning.

It conveys the “why” behind the data, providing the situational awareness, provenance, and business logic agentic AI needs to deliver trustworthy insights and make sound decisions.

Simply put, agentic AI requires context to function reliably. Without it, systems and agents run the risk of misinterpreting signals and producing false insights. And because AI makes everything go faster, these misinterpretations and potential hallucinations can easily propagate throughout business processes, exposing fissures and cracks in operational systems that lead to outcomes that are problematic or, potentially, disastrous.

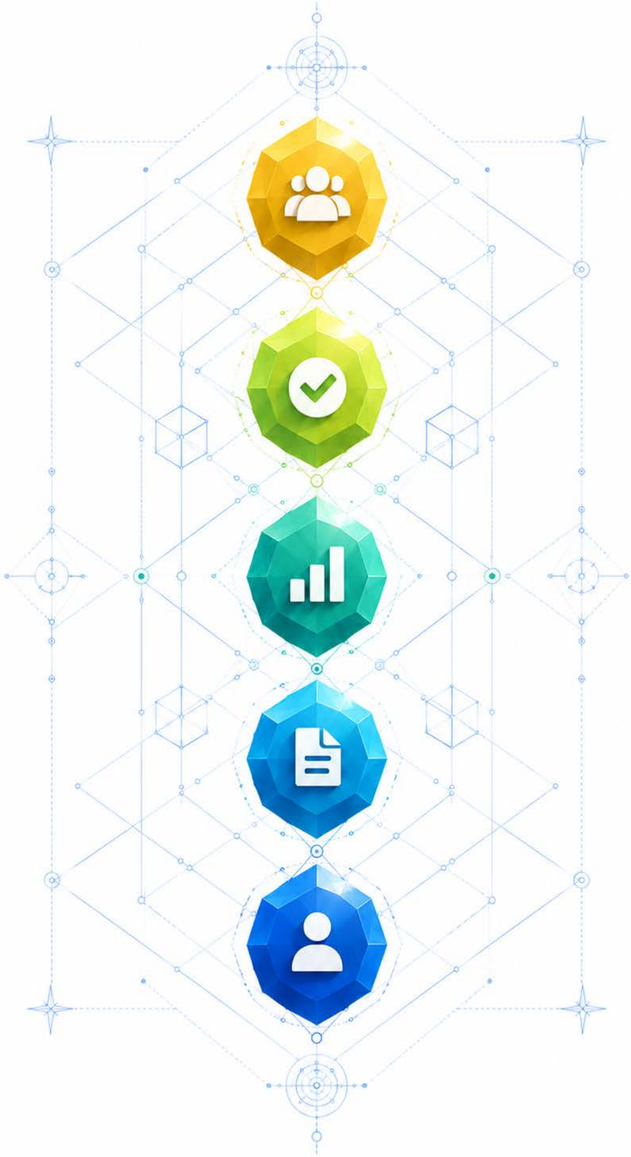

Lessons in (AI) Failure

Consider this scenario. A global luxury retailer knows that 80% of their revenue comes from 8-10% of their customers. Because these customers are so valuable, the retailer established a VIP customer service program that provides additional perks and a human touch to every interaction. For the program to be a success, it's important that the retailer understand the connections between customer entities and their households.

So, when a top-tier customer enters the retailer's flagship store in New York City and completes a purchase, the retailer's systems identify the person's VIP status, prompting the staff to deliver a customized experience. A manager comes over to personally thank the customer for their brand loyalty. Within hours, an autonomous agent triggers a personalized follow-up inviting the customer to an exclusive, "VIP-only" event. This is a seamless, high-value experience.

However, the logic fails just days later when the customer's spouse visits a different location. Because the backend systems treat the spouse as a "new" and isolated record, the in-store staff—and the AI agents monitoring the transaction—fail to recognize the household's collective value. The spouse is treated as a new customer, denied the VIP perks, and excluded from the exclusive invitation.

This example illustrates a context gap. The backend systems failed to connect the customer and their spouse into a single household. Consequently, an autonomous agent could not see the connection, and therefore failed to extend VIP perks to both individuals. As a result, the customer and their spouse had a mismatched experience that erodes the perceived value of the VIP customer service program.

Going Beyond the Golden Record

For years, companies aspired to deliver golden records—single, authoritative, accurate versions of a business entity's data across multiple data sources and datasets. They purchased **traditional master data management (MDM) solutions** to help them tackle this task. And over time, they began to make some progress.

But then everything began to change.

- Data started growing at exponential rates
- Sources multiplied
- Information became increasingly complex

And then AI entered the picture.

- Decision velocity increased dramatically
- "Garbage in/garbage out" accelerated errors and misinformation
- Data silos became systemic blind spots
- Fragmented data led to faulty outputs

Suddenly, a golden record was no longer enough. Golden records, on their own, can inform AI "what" the entity is, but they don't explain the "how" and "why" of the entity's relationship across the organization. This missing information—the context—puts AI at a disadvantage, obscuring its view of relevant details that will help it move from simple logic to high-stakes, autonomous decision-making.

To understand the role of context, think of golden records as the noun of a sentence. On its own, a noun is just a label. It requires the other parts of speech—verbs, adjectives, adverbs, etc.—to transform it into a complete thought. Just as a sentence fails without its supporting parts of speech, AI fails when its core data lacks the context to be useful.

Golden records have served as the pinnacle of data integrity in the past. But today, having a single, clean view of a customer, supplier, product, or location is just the start. While golden records provide a critical foundation, the context needed for effective AI use comes from connecting these entities in the form of **enterprise knowledge graphs**. A critical **platform capability** of AI-native master data management, an enterprise knowledge graph structures organizational knowledge as interconnected entities and the relationships between them.

Enterprise knowledge graphs help to make connections across fragmented, siloed data. They surface relationships that may otherwise be obscured, and enable the creation of multi-domain relationship maps that reflect not only which records are related but also how those records relate to others across domains. They provide the **context** organizations need in the AI era.

The context provided by enterprise knowledge graphs elevates golden records from single views of entities to true intelligence that AI can reliably act on.

By pairing clean, curated entities with rich relational context, enterprise knowledge graphs enable AI to see the full picture of a golden record, including its history, its current connections and relationships, and other contextual information needed to resolve conflicts in real time.

For example, in fraud detection, context is what looks beyond a single, suspicious credit card swipe to recognize that a purchase (a) was made from a new device, (b) is linked to an IP address associated with a known breach, and (c) was for a product being shipped to an address updated just moments ago. Similarly, in supply chain management, context can allow an AI agent to recognize that a delay at a parts manufacturer in Asia isn't just a delayed shipment, but rather a critical point of failure that will disrupt a multi-million dollar production line and drain inventory reserves.

By connecting these threads, context surfaces meaningful associations that AI needs to make informed, trustworthy decisions.

What Is a Context Graph?

Context graphs are an emerging, specific type of graph that rely on the underlying data and connections provided by enterprise knowledge graphs and are optimized for use in the AI era. Because AI needs structured, governed, relational enterprise context to function effectively in production, context graphs can serve as the ground truth that prevents errors and hallucinations and ensures decisions are based on reliable, connected data that is accurate and trustworthy.

Context graphs, based on trustworthy master data, enable AI-based systems and AI agents to become even smarter and more "human" in their reasoning abilities. With access to company knowledge, organized in a structured way, systems and agents can synthesize large amounts of information quickly in order to make the same kinds of decisions that humans would make, complete with reasoning and logic, but much faster and at scale.

Powering AI With Context

By now, it's clear that context—or the lack of it—is a defining factor in the success or failure of an agentic AI initiative. However, adding context to master data is just the beginning. The true challenge lies in providing that context directly to AI agents tasked with analysis and action.

That's another key role for modern AI-native MDM to play. With a flexible, **cloud-native architecture** equipped with APIs and webhooks, AI-native MDM seamlessly supplies rich, trusted business entity data to downstream agents in real time. And by stitching all the data together in the form of enterprise knowledge graphs, it surfaces connections and relationships among and between entities, giving agents the context necessary to do their jobs properly.

With the right context, an AI agent can respond to a question like **"What do we know about ABC Corp?"** and not just return fragmented details; instead, it can leverage a comprehensive, connected map of mastered data that answers additional questions such as:

- **"What has ABC Corp purchased from us?"**
- **"Do we have any active contracts with ABC Corp?"**
- **"If so, what are the terms?"**
- **"When was the last interaction with ABC Corp?"**

By understanding the intent behind the question, the agent goes beyond simple data retrieval, allowing users to interrogate the data and ask it complex questions. And only if the data is clean, curated, and connected will the AI be able to generate accurate and trustworthy insights.

Lessons in (AI) Success

Revisiting the previous examples of AI failure reveals a common thread: the power of context to positively change the outcome.

If the manufacturing company's AI agent was equipped with the proper context, it would have flagged the order for human review instead of blindly greenlighting it. This step would have saved the company from processing the automatic shipment without proper evaluation and risking an extended payment delay.

Similar logic applies to the retail example. If the AI agent had a holistic view of household relationships, it would have recognized the connections between the spouses and delivered a coordinated, VIP experience to both of them.

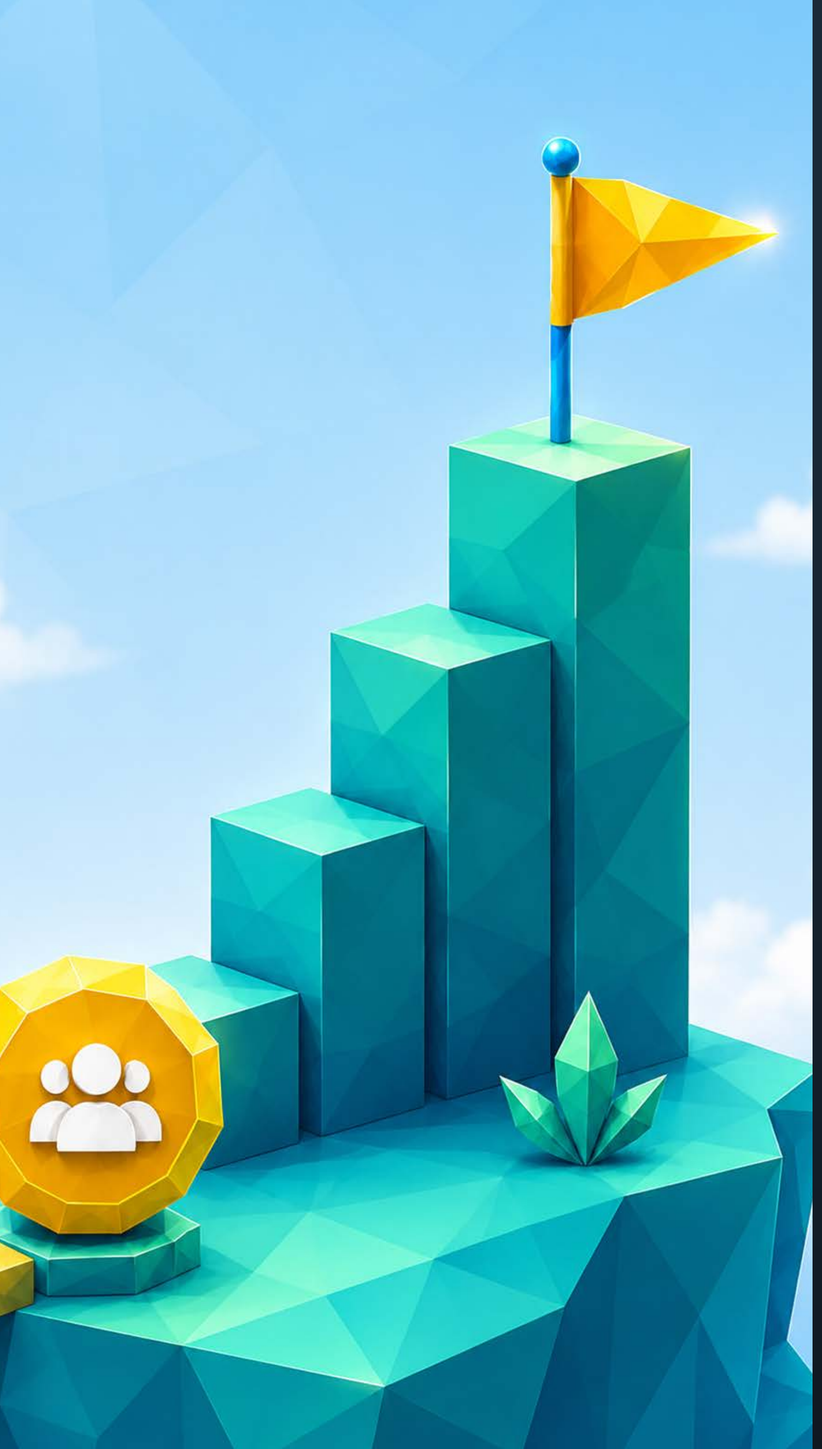

Activating the Next Generation of AI With AI-Native MDM

AI is fundamentally reshaping how organizations interact with and consume enterprise data. To turn this shift into competitive advantage—while also protecting the business from operational and reputational risk—organizations must anchor their AI strategy in high-quality data.

Tamr has spent more than a decade focused on building and using AI and machine learning to tackle the hard problem of delivering clean, contextual, trustworthy data cost-effectively and at scale. Everything we provide—from architecture and workflows to user interfaces and platform capabilities—is built around AI.

As the AI-native MDM standard, Tamr transforms fragmented data into golden records, connects data across domains, and produces enterprise knowledge graphs—critical ingredients for the valuable context AI systems and agents need to deliver trustworthy results.

Learn more at **tamr.com**

Tamr