RELTIO

EBOOK

# Unlocking agentic AI's advantage

## A business playbook for data readiness

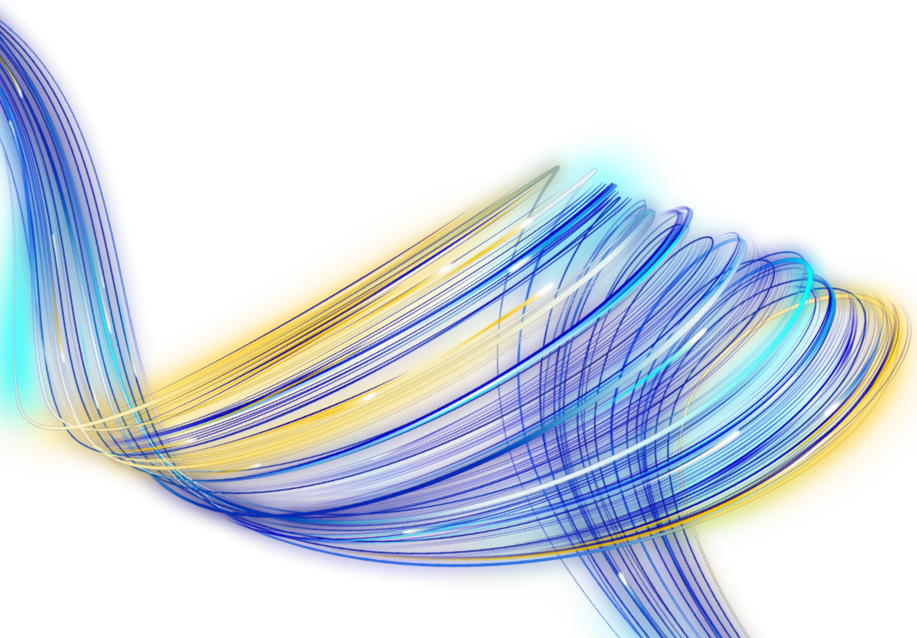

# Table of contents


**Chapter 1:** Why a business leader has to think about data readiness 3

**Chapter 2:** Why “good enough” data isn’t good enough anymore 4

**Chapter 3:** What agentic AI needs from your data 5

**Chapter 4:** Signs you’re not ready—and what that means for AI projects 7

**Chapter 5:** How to get ready for the operational reality of agentic AI and—what to look for 8

**Chapter 6:** Modern data foundation to connect trusted, real-time data to agentic workflows 9


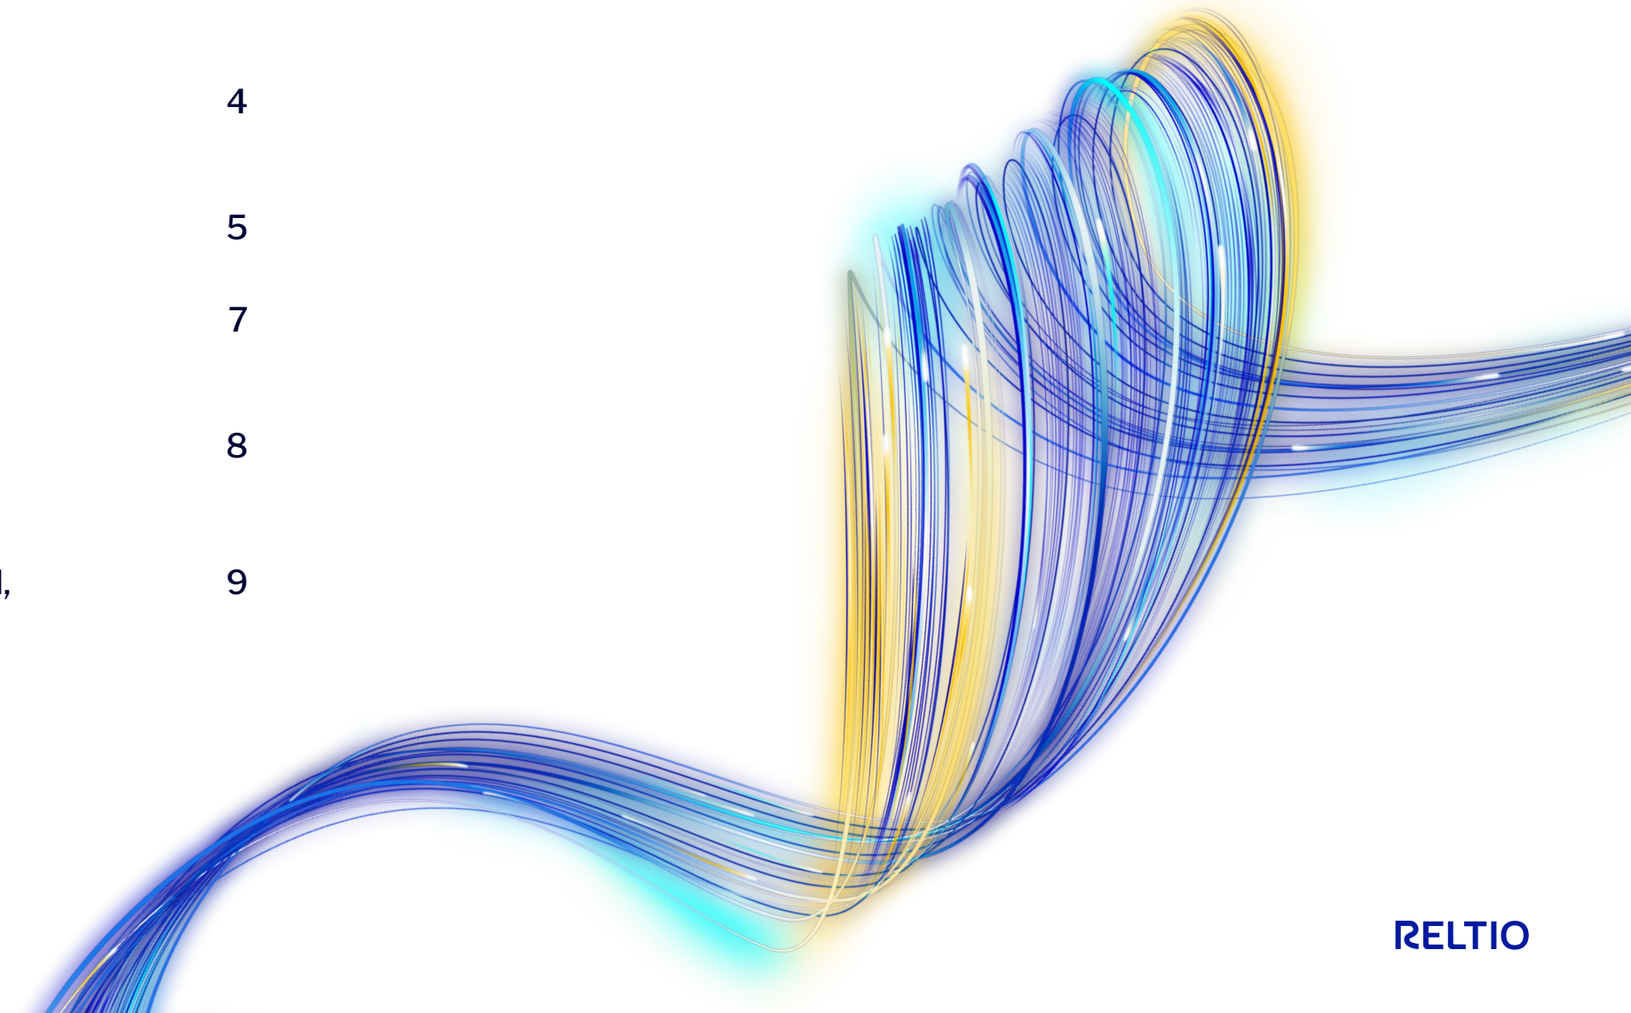



RELTIO

## Chapter 1:
## Why a business leader has to think about data readiness

**For business leaders, data readiness has become a core responsibility—giving agentic AI the foundation to drive tangible business outcomes and limit risks and delays.**

Agentic AI represents a major leap forward from traditional artificial intelligence and machine learning (AI/ML). Where ML models primarily predict, classify, or recommend based on data patterns, agentic AI takes it further—it acts.

These AI agents can incorporate AI/ML models into their decision-making, but they go far beyond passive suggestions. They sense, decide, and execute tasks autonomously, often in real time. They don't just support decision-makers—they become decision-makers.

This is the critical shift. Traditional AI helps humans make decisions, usually requiring manual triggers or working from batch data. Agentic AI flips that model, enabling systems to operate independently, adapt to changing conditions, and execute end-to-end processes without human intervention. While early deployments may still include a human in the loop to guide and validate behavior, these agents are designed to learn quickly and move toward full autonomy.

For business leaders, this shift creates both opportunity and accountability. AI agents will make decisions that directly impact customers, operations, and revenue—so the quality of the underlying data is no longer just an IT concern. It becomes both a core business risk and a competitive differentiator. Business leaders who take an active role in ensuring their organization's data is AI-ready will see faster AI adoption, better outcomes, and greater ROI. Those who don't risk embedding poor decisions into automated processes that run at scale.

Agentic AI is already transforming industries. For example:

- **Insurance -** AI agents resolve claims in real time, minimizing cycle times and improving customer satisfaction.
- **Banking -** AI agents detect fraud patterns and take immediate action to mitigate risk.
- **Retail -** AI agents act as virtual sales assistants—tracking customer behavior to instantly recommend offers.
- **Life sciences -** AI agents help sales teams engage healthcare providers with timely, tailored insights—enabling more relevant outreach.
- **Contact centers -** AI agents can orchestrate end-to-end customer interactions—resolving issues, retrieving documents, and even making proactive offers without escalation.

**Where traditional AI Stops—and agentic AI starts for contact centers**

| Traditional AI stops at... | Agentic AI takes it further by... |
|---|---|
| Answering basic questions using scripted chatbots | Understanding user intent, switching topics naturally, and resolving issues across apps—without human help |
| Alerting support agents when a customer seems unhappy | Proactively contacting the customer with solutions or escalating when needed |
| Suggesting replies for agents to copy and send | Taking action automatically across tools to fix the issue from end to end |
| Sending tickets to the correct queue using keyword rules | Deciding who should handle what based on complexity, skill, and past results |

(Source: Reltio analysis, inspired by CMSWire article)

But these outcomes require more than automation—they demand precision. The bar for operational trust is significantly higher. That's because agentic AI doesn't just recommend—it acts. Every action an agent takes must be backed by data that is current, complete, and explainable. Agentic systems must operate in real-time environments with complete transparency, explainability, and traceability. They must not only act quickly, but also justify every action taken.

To consistently deliver value, agentic AI needs the right data, at the right time, with the right details, in the right form. As a business leader, your role is to champion this readiness. The speed, accuracy, and fairness of your agentic AI outcomes will only be as strong as the data foundation you build today.

## Chapter 2:
## Why "good enough" data isn't good enough anymore

**In the agentic AI era, your outdated data management solutions—which may have worked well enough in the past—don't just slow you down. The poor-quality, outdated data, and lack of agility take you out of the game.**

For years, enterprises got by on fragmented, static, and siloed data systems. Data pipelines were slow, quality checks were ad hoc, and many decisions could tolerate incomplete inputs. But those days are over.

Agentic AI is unforgiving. If the data is missing context, outdated by even a few minutes or duplicated across systems, the agent fails—or worse, takes the wrong action.

The myth that "available data" is the same as "ready data" is a common trap. Availability simply means the data exists somewhere. Readiness means the data is…

- **Clean and unified:** Duplicates removed, inconsistencies resolved, and formats aligned
- **Governed and compliant:** Meets security, privacy, and policy standards by design
- **Context-rich:** Includes relevant details such as customer history, product associations, critical unstructured content (e.g., contracts, applications) and recent interactions that help AI understand the whole picture
- **Instantly accessible:** Available when and where it's needed, by the systems and teams that depend on it

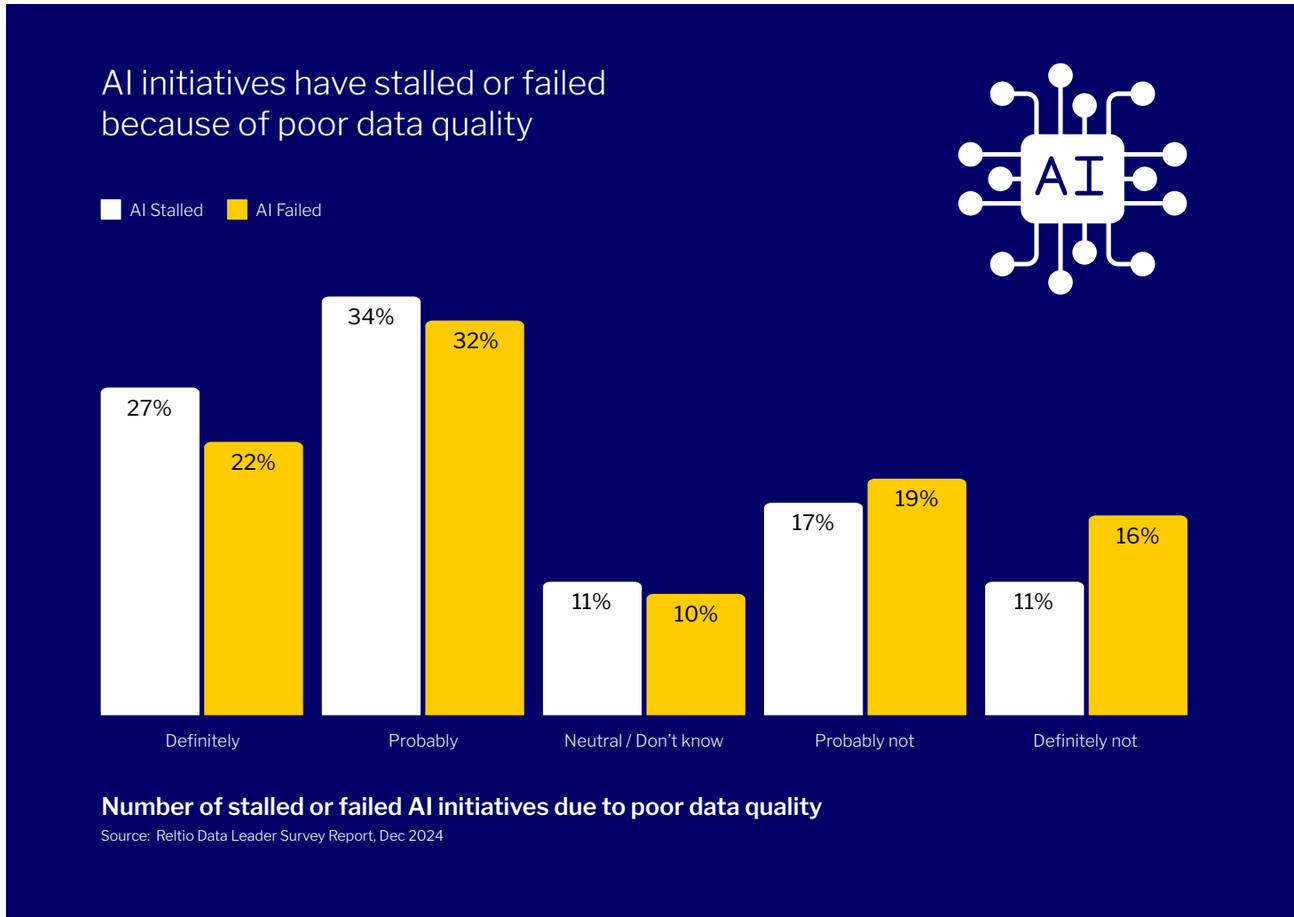


**Number of stalled or failed AI initiatives due to poor data quality**

Source: Reltio Data Leader Survey Report, Dec 2024

The impact of not meeting these standards is real. In agentic environments, low-quality data doesn't just cause technical hiccups—it directly fuels flawed decisions, exposes sensitive processes to compliance violations, and erodes customer trust at scale. Unlike traditional systems, agents act in real time. Therefore, data quality issues can trigger significant operational and reputational consequences.

This isn't just a theoretical risk—it's a widespread reality. According to Forrester, 52% of enterprise IT/technology decision-makers saw data quality and consistency as a significant barrier to adopting AI-driven operations, just behind security concerns. These issues don't just slow progress—they can paralyze transformation altogether.

> **"Lay the groundwork for an agentic AI technology foundation by investing in key enablers—technology infrastructure, data quality, governance frameworks, and workforce readiness."**
>
> **Seizing the Agentic AI Advantage**
> **McKinsey**

## Chapter 3:
## What agentic AI needs from your data

**Smart and trustworthy AI agents need smart and trusted data: clean, current, context-rich, and well-governed.**

So, what does data readiness for agentic AI really look like? It goes well beyond conventional analytics prep. It's about creating a living, breathing, trusted view of the business that agents can interact with safely and intelligently. Specifically, agentic AI requires more....

- **Unified data, not just centralized:** Traditional data warehouses store lots of data but rarely deliver the holistic, real-time views AI agents require. Unified data—free of duplicates and standardized across systems—is essential.
- **Context-rich, connected profiles** including:
  - Full views of entities such as customers, products, and suppliers (See Figure 1)
  - Rich set of attributes for each entity—product name, description, category, and more
  - Relationships between entities, such as vendor A supplies XYZ product
  - Interactions, such as a customer B reached out to the contact center
  - Unstructured content, such as the latest contract on file, call transcripts, documents, and images

Why? This organization of data, which shows its relationships—also known as the semantic layer—gives structure and meaning to your data, defining what things are, how they relate, and how they should be used. **It translates complex data source definitions into business terms, enabling your dashboards, AI agents, and operational apps to read from the same page. Additionally, it allows for dynamically changing multidimensional views based on business context.** It is crucial for AI systems to understand your data, connect the dots, and act intelligently.

The **semantic layer** is like an interactive map. It tells you:

- What each building is (hospital, school, home) ⟶ **entities**
- Key facts about each one (size, hours of operation) ⟶ **attributes**
- How they connect (roads, service zones) ⟶ **relationships**

Only with this map can an autonomous vehicle (your AI agent) move with confidence. That's the role of a semantic layer. It provides the **structure, meaning, and context** AI needs to make smart, relevant decisions—at scale and in real time.

- **Rapid availability of fresh data:** Autonomous agents perform best with high-quality, up-to-date data. Data freshness and quality are both essential for agentic AI—you should not have to choose between them. Delayed data leads to missed opportunities and flawed decisions. Batch processing can't keep up in a world of constantly changing, fast-moving processes. Your data systems should dynamically update when the source data changes—along with all its dimensions and relationships.
- **Lineage and provenance:** Every piece of data must be traceable—where it came from, how it's been modified, and by whom, so every decision made by an agent can be audited, explained, and trusted. Explainability isn't just a technical feature—it's a business imperative. AI agents must offer decision trails that make sense to non-technical teams, auditors, and executive stakeholders alike. Without it, decisions become black boxes, eroding confidence and compliance in the agentic technology.
- **Governance and access control:** It is critical that agents operate within boundaries and only act on what they're allowed to see. Poor control can lead to data breaches or regulatory violations.

Without this information, agents lack the intelligence and guardrails to operate in production. Data that is incomplete, siloed, or delayed simply isn't compatible with agentic decision-making. In addition, agentic AI must be governed by principles of fairness, transparency, and bias mitigation, ensuring actions taken by AI agents reflect ethical and inclusive decision-making. As AI becomes operationalized, data quality, speed, trust, and bias move from technical concerns to business-critical requirements.

**The bottom line:** if your data isn't ready for action, your agents aren't either.

**Example of context-rich, connected profiles**

Figure 1

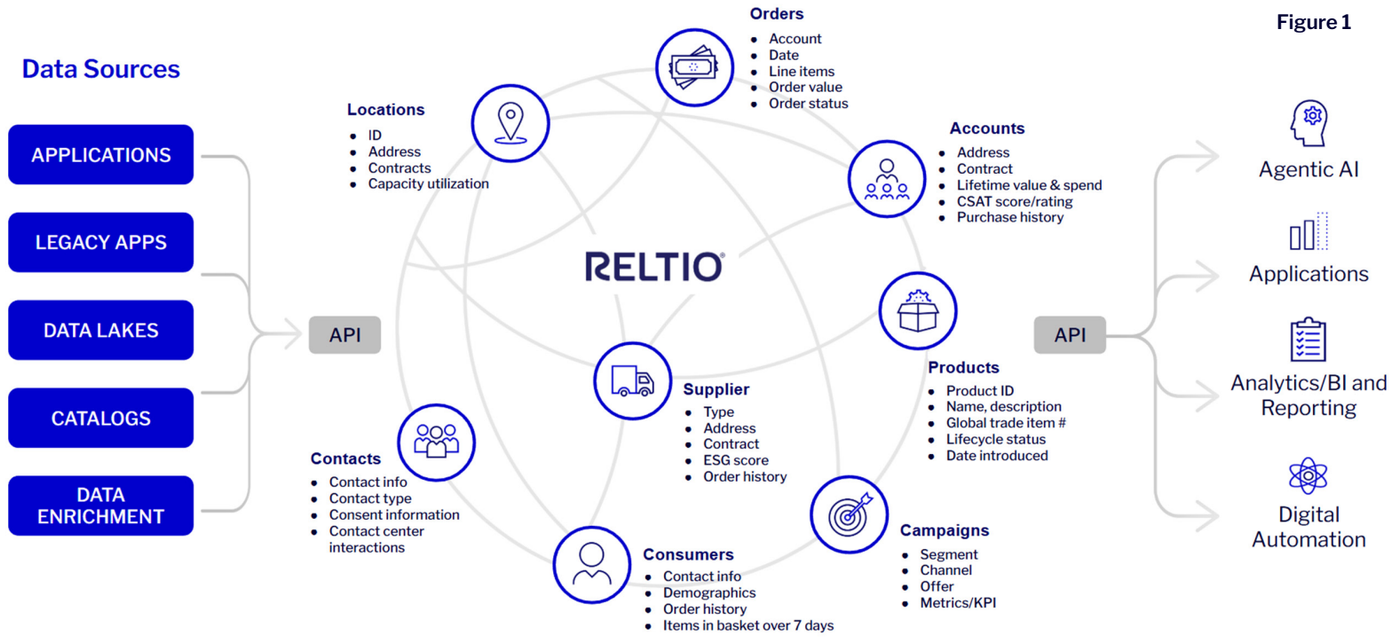




RELTIO

## Chapter 4:
## Signs you're not ready—and what that means for AI projects

**If your data teams spend significant amounts of time resolving data quality issues and getting conflicting insights, your AI agents are not ready for production—or can cause significant business risks.**

Determining whether your organization's data is fit for agentic AI requires honest inspection. The warning signs are often hiding in plain sight: duplicated records, mismatched data relationships, or inconsistencies between systems. These issues delay projects, limit scope, and degrade trust in AI outputs.

Here are some clear red flags that your data—and hence you—are not ready for AI:

- You're resolving the same duplicate customer records often
- AI-generated insights conflict with operational reports
- Business teams spend a big portion of their time validating data instead of acting on it
- Agents are stuck in demo mode or limited to edge-case automations

These issues don't just **frustrate IT**—they delay AI value realization for the business teams. According to a survey by EY, 83% of senior business leaders said their organization's AI adoption would be faster if they had a stronger data foundation in place, and two-thirds (67%) admit their lack of foundation is actively holding back AI adoption. The survey underscored a glaring paradox: while businesses pursue AI-driven transformation, many remain entangled in a web of disconnected, poor-quality data, jeopardizing their success.

The implications are clear: bad data leads to bad automation. Worse, it erodes trust in AI. And, the AI agents that misfire due to data issues don't just impact efficiency—they damage brand equity and invite scrutiny from regulators, stakeholders, and boards.

Readiness isn't just a data maturity issue—it's a cornerstone of business resilience. As organizations look to automate decision-making with AI agents, gaps in data readiness translate directly into operational fragility. Without a strong data foundation, scaling  becomes a liability, not a competitive advantage.

**According to a survey by EY, 83% of senior business leaders said their organization's AI adoption would be faster if they had a stronger data foundation in place, and 67% admit their lack of foundation is actively holding back AI adoption.**

**While businesses pursue AI-driven transformation, many remain entangled in a web of disconnected data, jeopardizing their success. Even organizations with mature data governance frameworks face persistent data silos that hinder their AI initiatives.**

**For more information, read our blog.**

Data quality, trust and explainability, and integration issues are the largest AI barriers

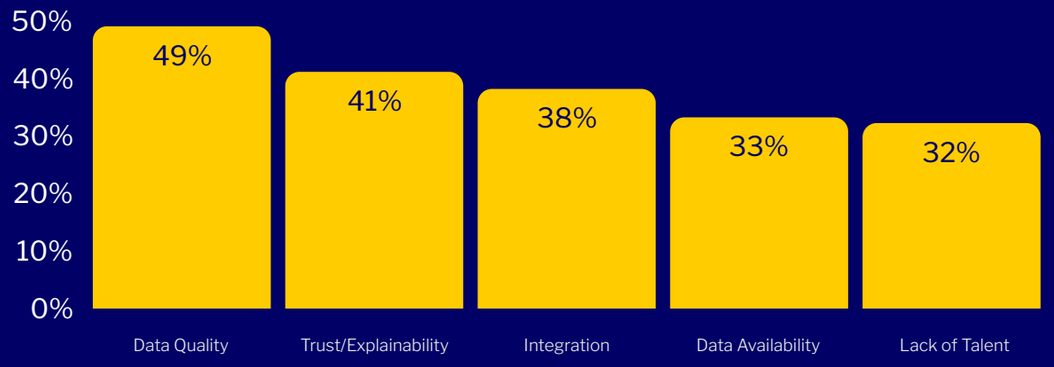


**Top reasons why AI initiatives failed**

Source: 2024 Reltio survey of enterprise data and IT leaders

## Chapter 5:
## How to get ready for the operational reality of agentic AI and—what to look for

**The path to agentic AI success requires an automated approach to continuous data quality management and governance, dynamically updated context, and a flexible data platform with fast time to value.**

Before launching agentic AI initiatives, align with your peers on the specific business use cases, the expected ROI, and the data sets required to support them.

It is also critical to align with your technical teams on a few foundational principles that can guide your data foundation approach. The most effective strategies tend to reflect these five core ideas:

1. **Continuous ingestion, data quality management, and mobilization for different types of data:** For fresh, reliable data to serve AI agents on a continuous basis, the data ingestion and quality management need to be continuous, too. It also requires ongoing monitoring of data freshness and quality—not just during onboarding, but as a built-in, automated process to prevent silent drift over time.

   The true business context for AI means maintaining always-current, trusted profiles across all critical business data—customers, products, suppliers, and more with their preferences, relationships, and interactions across systems. Any agent built on siloed data—in a single application, for example—lacks the full context to get it right.

2. **Built-in automation:** Manual data quality methods don't scale. Look for a solution that continuously improves data quality using AI-augmented automation.

3. **Speed to value:** If it takes months or years to prepare data, your AI initiatives stall. That's why a DIY approach for getting data ready—for integration, unification, or data quality—isn't the answer. It's slow, costly, and puts you behind competitors already moving fast with agentic AI.

4. **Unified, flexible data model:** Your data model should support your enterprise with a unified approach, rather than a siloed view of your apps or business processes. You also need the agility to add new data, attributes, and relationships as your business and AI evolve.

5. **Collaboration by design:** Because agentic AI is infused into business workflows, the best solutions should be designed to connect business and IT teams—enabling shared governance, fast iteration, and aligned outcomes. A highly technical UI and the need for coding with specialized skills limit the necessary involvement of the business teams. Any use case should also be collectively tracked for specific KPI improvements to demonstrate ROI and to continuously refine goals, as well as related data and AI requirements.

**In a recent survey of 200 data and IT leaders, 82% of respondents reported that more than 40% of their organization's data is derived from over 50 different applications, highlighting the extent of data silos across businesses.**

Based on the above five principles, here is what business leaders should expect—and ask for—from their data platforms:

- **Continuous data quality management:** Is duplicate data resolved automatically—and merged intelligently—so your data has less noise and AI agents can act as expected? Look for solutions that go beyond traditional rule-based deduplication by augmenting it with AI-driven matching and automated anomaly detection to catch issues before they impact AI decisions. Such capabilities not only improve accuracy but also scale the process of preparing data for AI agents across entities and use cases.
- **Real-time processing:** Is data always fresh and instantly accessible by systems and AI agents? Real-time data flow ensures decisions are based on the most current information, reducing lag and improving responsiveness.
- **Enrichment and 360-degree profiles:** Can it stitch together complete, continuously updated profiles with full context, including relationships and interactions?
- **Multidomain flexibility:** Will it connect and unify across customer, product, supplier, and location data simultaneously? AI agents need to make decisions with full context—not just within one domain, but across all the entities involved in a process.
- **Built-in governance:** Is data validated, compliant, and traceable from source to decision? AI agents need to act on data that's trustworthy, auditable, and aligned with policies—especially when decisions carry risk or regulatory impact.
- **Adaptability:** Can the data model easily evolve to support new AI expectations, business models, or regulatory changes?
- **AI-ready access via an MCP server:** Can trusted, governed 360-degree profiles be securely accessed by AI models and LLM agents using the open model context protocol (MCP)? MCP serves as a plug-and-play bridge that connects your data to any AI agents—whether custom-built, third-party, or powered by LLMs.
- **Agentic innovation for data team productivity:** Are there prebuilt agents that autonomously handle specific data governance tasks? They can dramatically boost the productivity and impact of data stewards.

CARMAX

> **"We wanted to get to the point where we had a 360-degree view of our customers, and we couldn't deliver that without having the data, systems, and technology together in a connected format."**
>
> **Shamin Mohammed**
> **Chief Information and Technology Officer**

## Removing data silos unlocks AI potential

Smashing silos would improve AI outcomes

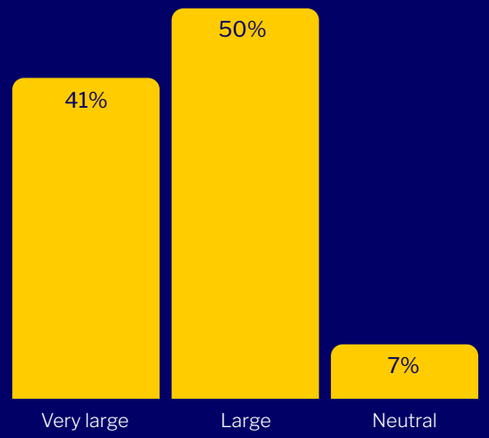


What improvement do you think breaking down data silos could have on AI initiative success rates?

**91%**

believe breaking down silos would improve AI outcomes

## Chapter 6:
## Modern data foundation to connect trusted, real-time data to agentic workflows

**With Reltio, you don't just modernize your data backbone, you unlock AI at scale—safely and fast.**

**Reltio Data Cloud™** is designed with the above principles from the ground up. It not only improves trust in data but also gives AI the full context it needs—boosting confidence in results, speeding up deployment, and making AI more useful across the business.

Our platform provides a unique advantage in the agentic era. Built on a cloud-native, real-time architecture, it delivers always-on access to trusted data across domains. Here's how Reltio supports AI readiness, especially for business teams:

- **360° views:** Reltio adds rich context to each entity—accounts, contacts, suppliers, locations, and assets—by capturing their relationships, interactions, and related unstructured content. This gives agents a truly complete, real-time view, complete with a data graph.
- **Unified, dynamic profiles:** These 360° views are powered by our platform, which unifies and continuously refines data across domains—deduplicating, standardizing, and enriching profiles to create golden records for customers, products, suppliers, and more.
- **Built-in data quality automation:** AI-ready data through always-on validation, survivorship rules, quality monitoring, and remediation workflows.
- **Real-time data delivery:** Instantly syncs updates with AI agents and other systems, ensuring agents always act on the latest information.
- **Trusted lineage and compliance:** End-to-end traceability, audit trails, and governance aligned with security and regulatory mandates.
- **Preconfigured for fast deployments:** Our solution includes preconfigured, best-practice data models and mappings that accelerate implementation across key domains like customer, product, and supplier. You can go live faster with trusted, unified data you need to drive your AI initiatives.
- **Agentic AI platform:** Power job-specific, autonomous agents with real-time, trusted data. It includes prebuilt governance agents and supports custom or third-party agents. At the core is our MCP server, which uses the open model context protocol (MCP) to connect these agents directly to the Reltio data.

Eaton used AI inference models on trusted customer data from Reltio, along with invoice data, to realize $14M in annualized rebate leakage savings. Use of AI to identify and extract such a scale of savings from thousands of incoming rebate claims would not have been possible without clean, unified, and normalized customer data from Reltio.

Reltio's survey found that 88% of respondents believe breaking down silos would create a large improvement in operational efficiency, while 80% reported that it would create a significant improvement in business outcomes.

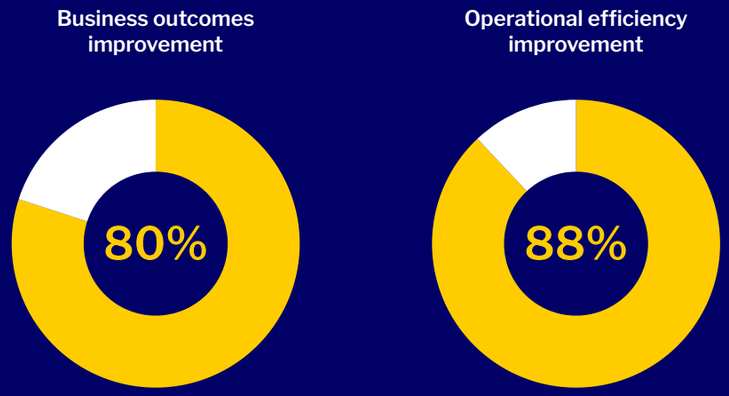

Reltio doesn't just help organizations modernize their IT. It enables business teams to scale their AI investments with confidence—reducing risk, accelerating impact, and unlocking operational agility.

## Conclusion: Your next step toward operationalizing AI

Agentic AI isn't a futuristic vision—it's already reshaping how leading enterprises operate. But without the proper data infrastructure, even the most advanced initiatives stall or fail.

**To move from an AI pilot to real-world impact, your organization needs a data foundation that is:**

- ✓ **Unified**
- ✓ **Governed**
- ✓ **Real-time**
- ✓ **Adaptable**

These aren't IT-only concerns. Business leaders must be active participants in shaping data readiness.

**5 questions to ask your data team this week:**

1. How can we stop manually fixing the same data issues over and over and start automating data quality with AI?
2. How current and comprehensive are our customer and product profiles? Are we limited to a specific app's data?
3. How can business and data teams collaborate better to prepare for agentic AI use cases?
4. How can we transition rapidly from manual and batch processing workflows to using trusted, fresh data?
5. Is our data foundation flexible enough to rapidly adapt to evolving processes, new agentic AI workflows, and business model changes?

**Agentic AI is here—and your success depends on one thing: your ability to connect trusted data to agentic workflows. Explore how Reltio can help you get AI-ready.**

RELTIO®

US +1 (855) 360-3282
UK +44 (800) 368-7643

@Reltio
facebook.com/ReltioHub
linkedin.com/company/reltio-inc

ABOUT RELTIO

At Reltio, we believe data should fuel your success in the enterprise AI era. Reltio Data Cloud™ is the agentic data fabric for the enterprise—powering real-time data intelligence and AI transformation. Reltio's cloud-native SaaS platform delivers unified, trusted, and context-rich data across domains in real-time. With Reltio, organizations gain 360-degree views of customers, products, suppliers, and more—mobilized in milliseconds to any application, user, or AI agent. Trusted by the world's largest enterprises across life sciences, financial services, healthcare, technology, and more, we help organizations fuel frictionless operations, drive innovation, and reduce risk.

To learn more, visit www.reltio.com