

6 Pitfalls to Avoid on the Road to Becoming Data- and AI-Driven



Is your organization aspiring to become data- and AI-driven? If you're like most modern businesses, then your answer is a resounding "Yes!" In fact, 98% of senior data and AI leaders **report** that their organizations are increasing investments in data and AI, up from 82% in 2024.

But as many organizations are finding, investment alone is not enough. Despite increasing their investments, few businesses—only 37% according to the 2025 AI & Data Leadership Executive Benchmark Survey—claim that they have created a data- and AI-driven organization. Even worse, that number is more than 10% lower than it was last year.



Only 37%
of businesses have
created a data- and
AI-driven organization.


There's no question that data- and AI-driven businesses perform better. According to **Forbes**, companies that invest in data-driven behaviors see results: They are 23 times more likely to acquire customers than their competitors, about 19 times more likely to remain profitable, and nearly seven times more likely to retain their customers. Being data- and AI-driven gives these organizations a competitive edge because they have better visibility into their operations. And this visibility leads to better, more informed decisions, more effective business strategies, and fewer wasteful expenses.

While there are many reasons why organizations struggle to become data- and AI-driven, we believe there are a few common pitfalls that derail even the best of intentions. Six, to be exact.

Here they are, along with our advice on how to avoid them.

6 Pitfalls to Avoid on the Road to Becoming Data- and AI-Driven:

1. Your culture hasn't changed
2. You're still relying on rules
3. Nobody sees the value
4. You rely on overnight syncs, not real-time updates
5. You've burned out your data team
6. Your internal and external data are disconnected



→ PITFALL 1 — Your Culture Hasn't Changed

Data moves fast. Is your organizational culture keeping up? If not, rest assured, you're not alone. In fact, when asked this question, many organizations **admit** they're not moving fast enough.

When it comes to culture change, businesses find that new strategies, new processes, and new tools only go so far, especially if the underlying mindset, behaviors, and ways of working remain the same. Without empowered leaders guiding the change, collaboration across the business falters—leaving stakeholders unsure what to prioritize, who is accountable for the change, and how the culture shift supports the organization's broader strategy and goals.

For example, while many organizations have appointed a chief data officer (CDO), few have given them the budget and authority needed to influence meaningful change. Instead of being part of the leadership agenda, CDOs are often relegated to custodial roles, responsible for cleaning up the mess caused by dirty, siloed data rather than driving enterprise-wide transformation.

This is a problem. Because when CDOs have the title but not the investment and influence, it's difficult for them to advocate for why becoming data- and AI-driven matters. As a result, employees dismiss initiatives as another passing fad.

Failing to engage with business stakeholders early in the process is another misstep. After all, accountability is hard to achieve when there isn't a strong partnership in place and when there's a lack of alignment on objectives, roles, and responsibilities.

Avoid This Pitfall: Empower Your CDO to Lead the Change



CDOs are fundamentally change agents. Through their work with data, analytics, and AI, CDOs can transform your business and elevate the value of data and AI throughout the organization. But they can't do it alone.

First and foremost, becoming data- and AI-driven requires your organization to elevate data to a board-level discussion. It involves making data and AI

a prominent chapter within your organization's strategy, not a footnote that everyone ignores. It also requires your organization to empower the CDO with the budget and the authority that are essential to drive meaningful change.

Shifting the mindsets of people throughout the organization—and doing so early in the process—is another key to success. By cultivating a strong partnership with the business, stakeholders are more likely to become ambassadors for change who speak on behalf of the CDO when they are not in the room, advocate for the work the CDO is doing, and evangelize the results to their peers across the business.

Together, the CDO plus their ambassadors can effect true culture change, where everyone across the organization understands the importance of data, analytics, and AI—and how to use them to drive the business forward.

PITFALL 2 — You're Still Relying on Rules

How many rules have you coded into your master data management (MDM) solution? If it's over 500, then life can't be too easy for you. Just ask legendary computer scientist, **Michael Stonebraker**.

Rule of thumb: The more data sets you're integrating, the more rules you need. And as your data sets grow incrementally, the number of rules you need grows exponentially, too. Dan Waldner, former director of customer data at Scotiabank, estimated:

8 data sets = 1,000 rules

10 data sets = 1 million rules


20 data sets = Several million rules

So if you have over 500 rules (and who doesn't?) and you're relying on a traditional MDM solution to store these rules, then you will struggle to onboard

new data sources. Why? Because **every time you add a new source, you need to rewrite the rules**. And as you can see, the number of rules you need to write is massive. When you're trying to scale, this is a surefire way of doing the complete opposite.

Another option is to adopt a solution that bolts on AI—to deliver capabilities such as automation, recommendations, or predictive analytics—on top of an existing, rules-based architecture. But buyer beware! While these AI-enhanced solutions may reduce manual effort, they often introduce unnecessary complexity, limit access to real-time insights, and still can't scale.

Without a scalable solution, your business will quickly find itself unable to handle the workload that your data requires, resulting in a large and growing portion of data remaining unmastered. And you'll find that the only way to keep pace with the ever-increasing volume of data will be to invest in more human resources. And that's a costly proposal.



Avoid This Pitfall: Invest in AI-Native MDM




AI and machine learning are critical when it comes to scalability. Not only can they do the heavy lifting of data mastering and **entity resolution**, but the models also get better with more data. In fact, the bigger the data, the more accurate the results. But AI is not something you can simply bolt on to a solution.

Instead, AI must be embedded in the foundation of your MDM solution. That way, humans can spend their time guiding the machine and building trust through bottom-up feedback that considers the context of how their organization is consuming the data. AI will use the subject matter expertise provided by humans to learn and improve over time. And by time, we mean days or weeks, not months or years.

“By 2027, the application of GenAI will accelerate the time to value of D&A [data and analytics] governance and MDM programs by 40%.”

Gartner

4 Essentials of Successful Master Data Management
by Sally Parker, published September 16, 2025



Want to know the difference between AI-native MDM and AI-enhanced MDM?
[Check out this blog post](#)



Old Mutual, a premium financial services group, wanted to digitize their customer journeys. But valuable customer data was trapped in three separate and outdated master data management solutions. Old Mutual knew that to reach their goals, they needed to modernize their solutions. That's why they embarked on a digital transformation journey to replace their legacy MDM solutions with Tamr's AI-native MDM.

By adopting Tamr's **patented**, advanced AI/ML technology, Old Mutual saw immediate improvements in their data.

Old Mutual...

- **Improved data accuracy by 69%** in just six weeks, resulting in golden records
- **Fully decommissioned legacy MDM solutions** within nine months and saved millions of dollars in costs
- **Simplified their IT landscape**, reducing complexities in data integration and eliminating data silos

And they accomplished all of this on time, on budget, and without hiring new resources.

PITFALL 3 — Nobody Sees the Value of Your MDM Solution

There are many reasons why data mastering initiatives fail. But the number one reason why organizations fall short of their goals is that data leaders lack the structured framework needed to make the business case for adopting an AI-native MDM solution. Without a business case, the organization lacks clearly defined goals, has no roadmap for what it takes to make the change, and can't measure success.

Without these goals and measurement frameworks in place, business leaders will never know whether the initiative was successful—or not. Moreover, executives and other leaders across the business will fail to see the value of the hard work invested by the data teams because there is simply no way to demonstrate the benefits of the MDM strategy and the AI-native solutions that support it.

Avoid This Pitfall: Create a Business Case That Defines the Value of AI-Native MDM


Successful MDM initiatives, on the other hand, have a clearly defined business case with goals that are explicit, right from the start. They build in a process to

measure the value AI-native MDM provides, and they secure buy-in from leadership on how they will demonstrate success.

To ensure your initiative starts off on the right foot, start by looking for resources that help you estimate the total return on investment (ROI) from implementing AI-native MDM. Using Tamr's [value calculator](#), you can easily calculate the potential ROI for your organization. Simply adjust the numbers to align with your business, and see what kind of return you could achieve.

Then, combine these projected outcomes with the resources in our [how-to guide](#) to simplify and accelerate the process of creating a detailed and defensible business case—one that not only defines outcomes and success metrics, but also overcomes objections and measures the value realized by potential savings and/or improvements in operational efficiencies.

Share this business case with your leadership team. Get their buy-in. And implement AI-native MDM. Then, you'll be one step closer to becoming data- and AI-driven.



Want to strengthen the case for investing in AI-native MDM? Download our ebook, [Guide: Building a Business Case for AI-Native MDM](#).

PITFALL 4 — You Rely on Overnight Syncs, Not Real-Time Updates

Healthy, AI-driven data ecosystems need the ability to simultaneously process data in both batch and real-time modes. And this needs to occur not only from source to consumption, but also writing back to the source, which is often an operational system (or systems) itself.

Real-time connectivity is essential for a modern MDM solution. Not only does it ensure that every record across your systems is accurate, consistent, and connected, but having the ability to read and write while the data is still in motion also prevents your organization from creating bad data in the first place.

Unfortunately, many solutions have architectures that do not allow you to make updates in real time. Beware of these solutions, as they will derail your efforts to write back to the source system with updated information and cause operational headaches down the line.

Avoid This Pitfall: Embrace MDM Solutions That Deliver Real-Time Capabilities



Traditional, rules-based MDM solutions hinder your ability to deliver clean, curated data across siloed systems in real time. Instead, you need to embrace an AI-native MDM solution that delivers a mastered view of key business entities without delay. That's what **CHG Healthcare** did, and they saw immediate impact.

As a leader in the healthcare staffing space, CHG Healthcare experienced growth over time. But as their business grew, so, too, did their challenges with disjointed, inconsistent provider data. From inconsistent, incomplete provider information and frustrating and manual onboarding workflows, to high infrastructure and maintenance costs, the challenges built up. And while CHG attempted to build their own data mastering solution to fix the problems, they quickly realized that their internal efforts fell short of the scale and specialized skills needed to support the business.

CHG partnered with Tamr and quickly mastered more than 7.3 million provider records while also reducing their duplicates by 46%. Not only did Tamr improve visibility into their network and reduce manual data preparation, but through **Tamr RealTime**, it also enabled CHG to access accurate, enriched provider data—so they can more readily match the most qualified healthcare provider with their client's needs.

CHG Healthcare

- **Mastered over 7.3 million provider records** in weeks
- **Reduced duplicate provider records by 46%**
- **Improved data completeness** through third-party enrichment
- **Enabled real-time access to provider data** via Tamr RealTime
- **Decreased manual data preparation**
- **Lowered infrastructure costs** with Tamr's AI-native MDM

PITFALL 5 — You've Burned Out Your Data Team

According to [DataKitchen](#), 97% of data engineers are feeling burnout. The reason? Well, there are many.

To start, many data teams are wrangling with the build-up of poor-quality data. And while they may have a traditional, rules-based MDM solution in place, it can't scale, which prompts the data team to take on manual, repetitive grunt work to try and fix the problem.

On top of that, stakeholders expect answers to their questions and fixes to their data quality issues immediately, causing data teams to

spend a disproportionate amount of time fighting data fires. And a lack of a data- and AI-driven culture (see pitfall #1!) means that data teams struggle to gain the buy-in and support they need to make things better.

97% of data engineers are feeling burnout

Burning out your data team is a pitfall every organization should avoid. Because data teams often operate in reaction mode, they are unable to embrace continuous improvement, which leads to frustration, burnout, and, ultimately, turnover. And because skillsets like those of data engineers are rare, it becomes difficult to replace these team members when they've simply had enough.

Avoid This Pitfall: Embrace Automation with AI Agents




The good news is that there is a good way to avoid this pitfall: Embrace automation using AI agents.

AI agents enable organizations to automate the process of fixing and curating data so that data scientists and data engineers can spend less time processing the data, and more time analyzing it and providing feedback on its quality. AI agents also help to reduce issues with burnout by handling the rote, repetitive tasks that wear out data teams. Instead of being data firefighters, data teams can now spend their time collaborating with business counterparts to elevate the value of the data so everyone can use it to drive better, more informed decisions that move the business forward.

Further, AI agents can free up data teams from monotonous tasks such as enriching addresses, spotting potential duplicates, and reconciling discrepancies across entities. And they can do these tasks faster, with greater accuracy, and at a lower cost. AI agents can also perform smart record matching by continuously scanning and evaluating records across systems to identify potential matches and resolve them into a single, golden record.

AI agents can also transform how users interact with the data. Using an intuitive chat interface and Retrieval Augmented Generation (RAG), users can ask sophisticated, context-driven questions about the data and receive immediate, actionable answers. And when new records enter the system, the AI agents work behind the scenes to ensure every new record added is not only accurate and consistent, but also connected to the broader data ecosystem.



PITFALL 6 — Your Internal and External Data Are Disconnected

Internal operations produce a lot of data. And organizations have made huge investments and considerable strides in collecting and using this data for business decision-making.

However, for many organizations, external or third-party data remains an untapped or underutilized resource. And failing to capture the full potential of linking internal data with external data remains a missed opportunity.

Consider this example. A B2B software company wants to build 360-degree views of its customers so they can deliver a higher-quality customer experience and better identify upsell and cross-sell opportunities. Most of their records include information such as name, company, email address, and maybe a phone number. But the data is incomplete and disconnected, and many records appear to be duplicates of one another. Missing company attributes and corporate hierarchies make it difficult to do effective segmentation and develop personalized communications and experiences.

By relying solely on internal data, the company is missing out on the rich corpus of information available through trustworthy, third-party providers such as Dun & Bradstreet, ZoomInfo, and many others.



Avoid This Pitfall: Enrich Your Internal Data With Trustworthy, External Sources



Overcoming this pitfall requires your organization to connect your wealth of internal data with valuable external data from third parties, vendors, or public data sources relevant to your business. And by recognizing the value of this linkage, you'll gain a competitive edge in today's dynamic business environment.

A **Fortune 500 retailer** improved marketing and sales operations by modernizing their MDM approach and curating customer data from disparate data sources using Tamr's AI-native MDM solution.


Faced with outdated and incomplete customer data, the retailer found that data challenges slowed marketing initiatives and disrupted sales operations.

But correcting a single customer record took months and manual coding just to adjust their rules-based systems. This effort led to a high total cost of ownership because they required dozens of highly skilled resources just to maintain these rules.

In partnership with Tamr, everything changed. The retailer used Tamr's AI-native solution to:

- **Create complete, up-to-date customer views** from millions of records across 11 Salesforce and operational data sources
- **Generate unique Tamr IDs for each customer**, removing the 40% duplication between Salesforce and systems
- **Cluster customer views into business-ready segments** by territory, industry, etc.

As a result of this effort, the retailer improved marketing efficacy, boosted sales operations, and enhanced their customer experience initiatives.





The Road to Success

Encountering MDM pitfalls on the road to becoming data- and AI-driven is inevitable. But if you're still using traditional, rules-based MDM solutions, the odds of hitting potholes and roadblocks are even greater.

AI-native MDM lays the foundation for delivering the accurate, accessible, trustworthy insights organizations need to be data- and AI-driven. By integrating data across fragmented systems and silos, AI-native MDM ensures that everyone across the business has access to accurate and consistent golden records—in real time. And with automation and governance built into the platform, AI-native MDM enables organizations to streamline their processes and reduce manual effort so they can deliver accurate, reliable data at scale. With clean, unified data in place, companies can feel confident they are creating a culture where decisions are data-driven and AI drives measurable business impact.

Isn't it time for these obstacles to stop derailing your efforts? By putting these six principles—and Tamr's AI-native MDM solution—into practice, you can overcome the pitfalls and truly be data- and AI-driven.

Discover how Tamr's AI-native MDM can help your organization deliver the complete, accurate, real-time data needed to become data- and AI-driven.

Learn more at **tamr.com**

