

Article Master Data in Mid-Sized Manufacturing: The Quiet Problem That Now Has a Deadline

July 2026

Most mid-sized manufacturers run on article master data that has degraded silently over decades of growth. The cost never appears as a line item – it hides in inventory, purchasing, reporting, and every manual workaround in between. Three regulatory and technology deadlines, and the arrival of AI on both sides of the equation, are now turning a chronic annoyance into an acute management issue.

Executive Summary

In many family-run, engineering-driven manufacturers, article master data – the records describing products, materials, customers, and suppliers – has never had an owner. The companies grew, organically and through acquisitions, faster than their data discipline. Each new site, system, and product line added records; nobody was responsible for consistency. Because the resulting cost is distributed across departments and buried in daily routines, it was never prioritized.

Key issues

- Inconsistent and incomplete records accumulate for years – duplicates, spelling variants, empty attributes – with no validation at entry and no owner
- Users build manual “bridges” (Excel mappings, re-keying, local tools) so that system A can talk to system B – each bridge is a permanent efficiency tax and a key-person risk
- After international growth, there is no group-wide article master: the same part carries different numbers in every plant
- Duplicate parts inflate inventory and fragment purchasing volumes
- Management reporting requires a monthly reconciliation marathon because every source uses different dimensions
- Enterprise MDM suites – the textbook answer – are too heavy, too long, and too expensive for mid-sized structures, so nothing happens and the gap compounds

What has changed

AI has altered the economics of the fix. Large language models can clean, deduplicate, and enrich hundreds of thousands of records at a fraction of the historical manual effort, and AI-assisted software development delivers a tailored central data repository in weeks rather than quarters. The outcome that once required an enterprise program is now within reach of a mid-sized IT organization – provided the organizational side (ownership, processes, adoption) is built alongside the technology.

A Problem Nobody Owns

Ask a mid-sized manufacturer who owns product costing, and you get a name. Ask who owns the article master, and you usually get a pause. Master data sits between IT, production, purchasing, sales, and finance – each department maintains *its* fields, in *its* systems, for *its* process step. Everyone depends on it; nobody is accountable for it as a whole.

This is not negligence. It is the natural result of how successful mid-sized companies grow. A niche product becomes a world-market position. A second plant opens abroad. An acquisition brings its own ERP. Sales adds a CRM, engineering a PLM, the shop floor an MES, e-commerce a webshop. Every system holds its own copy of the article master, slightly different from the others. Twenty years later the company runs excellently – and on data that no single person can fully explain.

How the Damage Shows Up

The pattern is remarkably consistent across mid-sized manufacturers. Six symptoms appear again and again:

1. Inconsistent, incomplete records

Decades of unvalidated free-text entry produce spelling variants, duplicates, and gaps. We have seen a company whose systems contained roughly sixty different spellings of the city of Brussels. Important attributes – material class, customs code, weight, minimum order quantity – remain empty for years because filling them was never anyone's job. Manual cleanup at this scale is hopeless, which is precisely why it never happened.

2. Manual bridges between systems

When systems disagree, pragmatic employees help themselves: an Excel file that maps article numbers between ERP and webshop, a re-keying routine between CRM and ERP, an Access database that translates plant A codes into plant B codes. Each bridge solves a local problem and creates a permanent one: someone must maintain it, forever. In mature landscapes we find dozens of such bridges – an invisible, distributed headcount tax, typically known in full by exactly one long-tenured employee.

3. No group-wide article master

The same physical part exists as three different article numbers in three plants. Intercompany orders are re-typed by hand. Group purchasing cannot see total spend per supplier or material, and therefore cannot bundle volumes. Group reporting maps everything manually – every month. For companies that became international niche leaders through acquisitions, this is the default state, not the exception.

4. Duplicates that cost real money

When search fails, engineers create the part again rather than find it. Industry analyses attribute 5–7% of MRO spend to duplicate purchasing driven by poor material data,¹ and

¹Industry analyses of material master and MRO data (e.g., CODA Technology Solutions, Sparetech) attribute 5–7% of MRO spend to duplicate purchasing driven by poor data.

duplicate records inflate inventory directly: the same item is stocked twice under two numbers while purchasing negotiates two smaller volumes at worse prices.

5. Reporting nobody quite trusts

When every source system carries its own dimensions and hierarchies, business intelligence becomes reconciliation. Finance spends the first week of each month explaining differences rather than analyzing results, and management meetings open with the question of whose number is right. New BI tools do not fix this – they display the inconsistency faster.

6. Knowledge in heads, not systems

The mappings, exceptions, and correction rules that hold all of this together live in personal files and personal memory. With the demographic wave now rolling through the German Mittelstand, this knowledge is measurably walking out the door. The successor inherits the bridges without the manual.

What It Actually Costs

Because the cost is distributed, it is systematically underestimated. Gartner puts the average cost of poor data quality at USD 12.9 million per organization per year;² research published via MIT Sloan Management Review estimates that bad data costs most companies 15–25% of revenue.³ Even if the true number for a given mid-sized manufacturer is a fraction of these figures, it is a seven-digit annual amount hiding in inventory carrying costs, duplicate purchasing, manual reconciliation hours, delayed quotations, and failed automation initiatives.

Why This Now Has a Deadline

For twenty years, poor master data was a chronic condition a healthy company could live with. Between 2026 and 2028, four developments turn it acute:

- **ERP deadlines.** Mainstream maintenance for SAP ECC ends on 31 December 2027, with extended maintenance until 2030 at a premium.⁴ Non-SAP legacy systems face similar generational upgrades. Migrating dirty data multiplies project cost and risk – every migration consultant's first finding is the state of the master data. Cleaning *before* migrating is dramatically cheaper than cleaning during or after.
- **Regulatory data demands.** Germany's B2B e-invoicing mandate phases in from 2027; the EU carbon border mechanism (CBAM) entered its definitive phase in January 2026; Digital Product Passports arrive from 2027 onward.⁵ All of them demand structured, complete, consistent article and supplier attributes that in most mid-sized companies currently exist nowhere.

²Gartner, Data Quality research: poor data quality costs organizations an average of USD 12.9 million per year. [gartner.com/en/data-analytics/topics/data-quality](https://www.gartner.com/en/data-analytics/topics/data-quality)

³Thomas C. Redman, research published via MIT Sloan Management Review / Harvard Business Review: poor data quality costs most companies an estimated 15–25% of revenue.

⁴SAP: mainstream maintenance for SAP ERP (ECC / Business Suite 7) ends 31 December 2027; optional extended maintenance until end of 2030 at a premium.

⁵Germany (Growth Opportunities Act): mandatory B2B e-invoice issuance from 2027 for companies above EUR 800k turnover, from 2028 for all. EU: CBAM definitive phase since January 2026; Digital Product Passports phased in from 2027 under ESPR.

- **Customer pressure.** OEMs and large customers increasingly require EDI connectivity, electronic catalogs, and product carbon footprint data – requests that land directly on the article master.
- **AI ambition.** Boards want AI in quoting, forecasting, and service. In KPMG's manufacturing survey, 56% of manufacturers name data challenges as the primary barrier to AI adoption.⁶ Every AI use case stands on master data the organization must be able to trust. AI starts with the article master.

Why the Textbook Answer Never Fit the Mittelstand

The classic response to this problem is an enterprise master data management suite. For a corporation with a data governance department, that can work. For a mid-sized manufacturer it means a 12–24 month program, six-to-seven-figure licensing, dedicated stewardship roles, and consultants who assume an organization that does not exist at this size. The predictable outcome: the business case never closes, the initiative never starts, and the problem compounds for another five years.

The requirement is different: the *outcome* of MDM – one trusted, central article master feeding every system and report – without the enterprise apparatus.

What Changed: AI Rewrote the Economics

Two things that were prohibitively expensive three years ago are now fast and affordable.

AI-assisted data cleansing. Large language models are exceptionally good at exactly the work that made master data projects fail: recognizing that sixty spellings are one city, that two differently described parts are the same item, that a missing material class can be inferred from the description. What used to be months of manual review by teams becomes a machine-assisted process with human approval of edge cases – runnable once as a cleanup and then continuously as a quality guardrail on new entries.

AI-assisted application development. A central master data repository – tailored tables, validation rules, role-based editing, import from Excel, APIs for the systems around it – no longer requires a development program. With AI-assisted development, a working, secure web application on a standard cloud stack is a matter of weeks, tailored to the company's entities instead of forcing the company into a product's data model.

In a recent engagement, we consolidated an international automotive supplier's fragmented Excel-based mappings and reporting dimensions into exactly such a central repository – a lightweight web application on a cloud SQL backend, with structured table definitions and direct copy-paste import from Excel. Internal tools and Power BI reports were remapped to the repository as the single source. The core system was built in days, not months; the Excel silos disappeared.

⁶KPMG, Intelligent Manufacturing (2025): 56% of manufacturers identify data challenges as their primary barrier to AI adoption; 52% report insufficient integration between systems.

A Pragmatic Approach

Technology alone does not keep data clean – organizations do. A workable approach for mid-sized structures combines four elements, delivered as one engagement:

- **See the landscape first.** A short assessment maps systems, data flows, and manual bridges along one or two end-to-end processes, quantifies the pain in euros and hours, and defines the target picture. No tooling decisions before this exists.
- **Build the lightweight hub.** A central repository for the article master (and adjacent domains as needed), built AI-assisted on a standard cloud stack, connected to the surrounding systems – replacing the bridge landscape instead of adding another tool to it.
- **Clean with AI, approve with humans.** LLM-based deduplication, normalization, and enrichment of the existing records – first as a one-off deep clean, then as a recurring quality process on new entries.
- **Anchor ownership.** Data ownership roles, entry standards, and simple governance routines – sized for a mid-sized company, not a corporation – plus the change management that makes the organization actually use the new single source. This is where such projects are won or lost.

The sequence matters: assess, build, clean, anchor. Companies that start with a tool purchase reliably end up automating their existing confusion.

Practical First Steps

A large program is not required to begin. Three questions locate the pain precisely:

1. **Where do employees re-key or map data between systems by hand?** List the bridges. Each one is a cost and a risk with a name attached.
2. **Which article attributes does the next deadline require – and do they exist?** Check e-invoicing, CBAM, or the upcoming ERP migration against today's data.
3. **Whose number is right?** Take one management report and trace why the figures differ per source. The path usually leads straight to the master data.

From there, a two-to-three-week assessment produces a quantified picture and a concrete roadmap – typically with a first visible result (one cleaned domain, one retired bridge landscape) within a quarter.

FAQ: Master Data in Mid-Sized Manufacturing

Why is article master data such a common problem in mid-sized manufacturers?

Because these companies grew faster than their data discipline, and master data never had an owner. Each department maintains its own fields in its own systems; consistency across the whole was nobody's job. The cost is real but distributed, so it never made it onto a management agenda.

Is this an IT problem?

Only partly. IT can build the repository and run the cleansing, but the root cause is organizational: missing ownership, missing entry standards, and processes that tolerate

local workarounds. That is why a sustainable fix combines technology with process and change work.

Can't we just clean the data manually?

At the accumulated scale – tens or hundreds of thousands of records, decades of variants – manual cleanup is economically hopeless, which is why it has not happened. LLM-based cleansing with human approval of edge cases changes the economics fundamentally.

Do we need an enterprise MDM suite?

For most mid-sized manufacturers, no. Enterprise suites assume governance structures and budgets that do not exist at this size. A lightweight central repository, built AI-assisted and tailored to your entities, delivers the outcome – one trusted source – at a fraction of the cost and time.

How long does an initiative like this take?

A meaningful first stage – assessment, central repository, AI-assisted cleanup of a priority domain, remapped reporting – typically fits into three to four months. Full multi-site rollouts are staged from there, guided by where the pain is largest.

What is the connection to our ERP migration?

The strongest one there is. Data migration is the most expensive and risky part of any ERP program, and its cost scales with the dirtiness of the data. Cleaning and centralizing the master data before the migration reduces project cost, shortens the timeline, and prevents importing twenty years of inconsistency into a new system.

What does AI have to do with master data?

Two things. AI is the reason the fix is now affordable – LLMs handle deduplication, normalization, and enrichment at scale. And AI is the reason the fix is now urgent – every AI use case the business wants stands on master data it must be able to trust.

Editorial notes for web publication (remove before publishing)

- Suggested meta description: Article master data quietly erodes margins in mid-sized manufacturing – duplicates, manual bridges, unreliable reporting. Why 2026–28 deadlines make it urgent and how AI changes the fix.
- Suggested category: Steering Foundation (alt.: Articles)
- Suggested tags: mid-sized companies, SME, data engineering, ERP, AI, Business Intelligence, finance transformation
- Internal cross-links: ERP Transformation and Planning: Why Forecasting Breaks Down; Why Most AI in Procurement Fails; Planning in Mid-Sized Companies: Why It Fails Under Pressure
- Reference case is anonymized (“international automotive supplier”) per decision 2026-07-03.