

The Data Quality Playbook: Agentic AI with Review Gates

Why AI projects die on dirty data, how agentic cleansing works without losing control — and the metrics that prove it.

The one-paragraph version: Gartner predicts that through 2026, organizations will abandon **60% of AI projects** that are not supported by AI-ready data. Duplicates, format drift, and missing identifiers are not an IT annoyance — they are the reason automation initiatives stall. Agentic AI can clean master data semantically (recognizing that "Müller AG", "Mueller AG" and "Müller AG, Grenchen" are one customer), but only a human review gate makes the result trustworthy: **AI proposes, humans approve.**

The four failure patterns of dirty master data

- **Duplicates no rule engine catches** — name variants, umlauts, abbreviations, relocated addresses
- **Format drift** — phone numbers, dates, and IDs written differently after every system change
- **Missing identifiers** — records that cannot be joined across ERP, CRM, and spreadsheets
- **Disagreeing systems** — ERP, CRM, and Excel each "know" a different address or price

Why agents — and why context windows matter

Traditional data-quality tools execute hand-written rules; agents read data the way a careful employee would. Current Claude models hold 200K–1M tokens of context (roughly 300–2,500 pages, or tens of thousands of records), so one agent can consider your schema, business rules, and large samples simultaneously. But published research (NVIDIA RULER, Chroma's 2025 context-rot study) shows effective context is smaller than advertised — so a serious setup combines long context for reasoning with direct MCP database queries and chunked batch runs. Never paste a database into a prompt.

The four-step operating model

1. **Audit:** agents profile ERP, CRM, and files; scored report of duplicates, gaps, inconsistencies — each with business impact.
2. **Cleansing with review gates:** corrections and merges proposed in batches; a human approves; nothing is written back without sign-off.
3. **Validation at the source:** new records checked on entry — plausibility, completeness, cross-system consistency.
4. **Monitoring:** a data-health dashboard and anomaly alerts, so quality stops degrading silently.

The metrics that make it verifiable

- Duplicate rate** — measured before the first cleansing batch, tracked after each one
- Error rate on critical fields** (identifiers, addresses, payment terms)
- Hours of manual cleanup per month** — the number your team feels
- Records failing validation at entry** — the leading indicator of future mess

Define the baseline in the audit; watch the numbers move after every batch. If they don't move, you see that too — that is the point.

The market context (so you can compare)

Gartner renamed the category "Augmented Data Quality Solutions" in 2024; the 2025 Magic Quadrant names Ataccama, Informatica, and Qlik as Leaders. These are excellent enterprise platforms with enterprise timelines and budgets. The SME-scale alternative: Claude + MCP against the systems you already run (bexio, ABACUS, Microsoft 365) — weeks instead of quarters, fixed prices, human-in-the-loop by design.

This guide is the condensed edition of our published in-depth article — sources and detail: eflury.com/en/services/data-quality/. As of July 2026. Not legal advice.

Questions about applying this in your business?

Free 30-minute call: eflury.com/en/contact/ · me@eflury.com · +41 79 910 77 87