

An amplifier, not a brain

An AI agent doesn't reason its way past your data — it **amplifies** it. Point it at clean, governed, consistent data and it looks brilliant. Point it at fragmented, duplicated, stale data and it scales the confusion — **fluently**. It doesn't fix what's underneath; it operationalizes whatever's there, at machine speed.



THE 3 CHECKS BEFORE YOU DEPLOY AN AGENT — identity · completeness · source of truth

- I

Identity
“Can you reliably tell when two records are the same entity?”
→ If not, every ranking double-counts, fragments the history, and distorts priority.
- 2

Completeness
“Are the specific fields the agent reads actually populated?”
→ If a field is blank the agent doesn't stop — it reasons from partial visibility, confidently.
- 3

Source of truth
“When ERP, CRM and support disagree, is there one answer?”
→ If not, the agent inherits the disagreement and picks one — and the moment users catch it, trust collapses.

WHAT THE AGENT ACTUALLY SEES — one account · three records

“Vihren Labs” ①
linked to the invoices

“Vihren Labs” ②
holds the support history

“Vihren Labs” ③
has the executive contact

Same company, three records — plus industry blank on 40% of accounts and stale activity dates. The agent can't tell they're one entity, so it **ranks confidently on a fiction** and writes a fluent paragraph explaining why.

Why it fails **quietly**: bad data used to *look* broken — the empty dashboard, the failed join. An agent gives no such warning; it writes a coherent, persuasive, wrong answer. And you can't lean on humans quietly patching it — the reps know “the real Vihren Labs record,” Finance keeps a reconciliation sheet. **An agent takes the data literally.**

Monday morning

YOUR FIRST MOVE

Before you shortlist a single agent, pick the **one entity** any AI decision would touch first — customer, supplier, product, or contract — and **pull the duplicates**. Count how many times your highest-value entity exists more than once, and check whether the fields an agent would read are populated.

That one count tells you more about your AI readiness than any vendor demo. Self-check: would you trust your own data enough to let a machine decide from it — continuously, in real time?

OPERATOR INSIGHT · LAYER 2

AI doesn't create truth; it amplifies the truth your systems already contain.

“Your AI maturity will never exceed your operational maturity.”

Free

Score your own data — the AI-Ready Operations Assessment runs the identity, completeness & source-of-truth checks on your systems. vihrenlabs.gumroad.com/l/rmkje