

MANUFACTURING · DATA PLATFORM

One governed context layer for every agent on a manufacturer's data platform

Procurement, finance and support agents each carried their own copy of supplier contracts, terms and rules, so a change to payment terms reached them at different times. Wexa served all of them from one governed context layer through a single MCP tool.

0 vs 5

manual edits to propagate a contract-term change

5.6 s vs 87 s

to propagate it to every agent

-61%

tokens per answer (107,698 → 41,560)

20 / 20

cross-agent answers consistent before and after the change

CLIENT

Manufacturer running AI agents on a leading cloud data platform (with the platform vendor)

REGION

US

WHEN

2026

ENGAGEMENT

Partner build

WEXA COMPONENTS

Wexa · data-platform connector · context graph · MCP get_context · per-agent policies

The challenge

- Context defined three times by three teams
- Rules such as "support never sees unit prices" enforced per agent
- No single trace of what context an agent used for an answer

What Wexa built

- ✓ Platform tables and contract PDFs linked in the context graph through the data-platform connector
- ✓ One MCP tool, `get_context`, registered by every agent; each agent has its own key and clearances
- ✓ Field masking and object hiding enforced on every read; only the active contract version served
- ✓ A trace per answer: which row and clause, which version, which policy allowed or refused each field
- ✓ A fourth (quality) agent added live in under **3 seconds**

How it runs on Wexa




FIGURE How this use case runs on Wexa: client systems connected, ontology discovered, context graph in the Wexa memory layer, agents and process flows, and governance. Highlighted components are central to this use case.

Product in use




FIGURE Live answers from two agents with the trace drawer showing evidence rows, contract clauses and the policy applied (product screen).

Results

MEASURE	BEFORE / WITHOUT WEXA	WITH WEXA
Engineering effort to define and keep each agent team's context in sync (3 teams)	3 FTE	0.5 FTE
Annual cost of the process	\$450K	\$75K
Time to implement the AI solution	6 months	2 weeks

MEASURE	BEFORE / WITHOUT WEXA	WITH WEXA
AI build cost (one-time)	\$480K	\$15K
AI operating cost per year (models, infrastructure, upkeep, platform)	\$312K	\$171K
Total AI cost in year 1	\$792K	\$186K

WITH WEXA

\$204K

net annual value to the client after Wexa's
own operating cost

13x

faster to implement

45%

lower AI operating cost

\$606K

less AI spend in year 1

- ✓ Head-to-head against the platform's native setup on the same agents and questions: median answer 11.4 s vs 15.5 s
- ✓ Governance acceptance tests 72 / 72; a specific trace found among 873 in 2.2 s
- ✓ Context call latency p50 0.71 s, p95 1.52 s