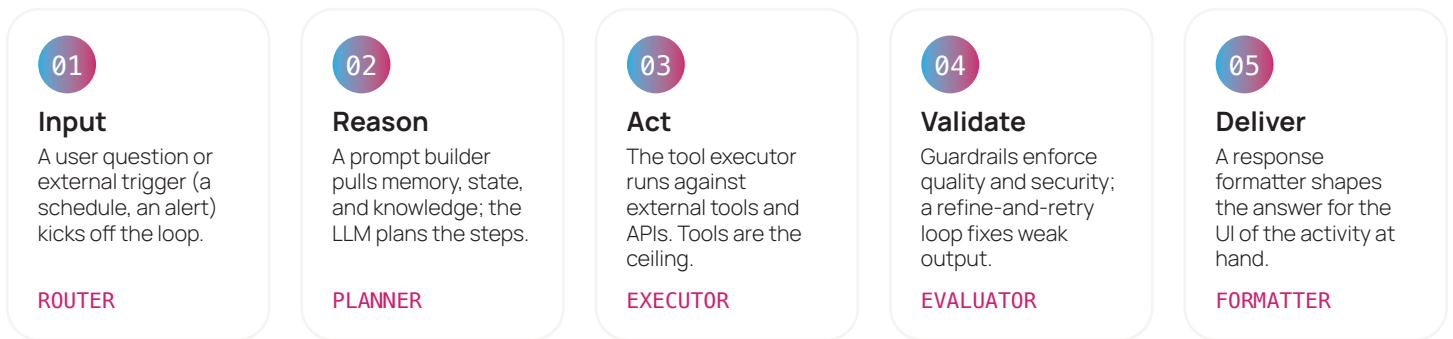


Why AI for BI needs *a great Agent Harness.*

Ask any enterprise BI leader to describe their experience with AI for BI in one word and you'll hear it over and over, **unimpressive!** Business users try these systems, get answers that don't go deep enough, and stop coming back. The problem isn't the intelligence of the models, it's everything around them. The same class of LLM that powers Claude Code and OpenAI Codex is sitting behind most BI agents. **The difference is the harness and overall user experience.**

01 • Anatomy of a great AI harness



02 • What a BI harness must give the LLM

Strip away any of these capabilities, and no amount of model intelligence can close the gap between a chatbot answer and the work of a **skilled human analyst.**

- ✓ **Certified reports across every BI tool** so the agent knows which sources are trusted and can show users the reporting behind every answer.
- ✓ **Semantic models behind those reports** so the agent can query the data through the same definitions the business already trusts.
- ✓ **Cloud database agents** in Snowflake, Databricks, and Fabric to reach certified tables that aren't yet covered by an existing report.
- ✓ **Proactive monitoring of KPIs that matter** to each user, with alerts when something significant changes not just reactive Q&A.

The bar to clear - How a real analyst **actually** answers a hard question.

Why did gross margin drop last quarter?

- GATHER**: **Consult the glossary, then look across every report.**
The analyst checks how the organization defines gross margin and which costs are included, then synthesizes from reports across multiple BI tools and data sources.
- VALIDATE**: **Cross-reference against certified reports the user already trusts.**
Before delivering an answer, the analyst confirms the numbers line up with the dashboards the user sees every day, so the answer feels consistent, not contradictory.
- WATCH**: **Keep monitoring the KPI for days or weeks.**
A good analyst doesn't walk away. They watch what they answered and ping you when the situation normalizes or degrades further. They bring the next insight to you.

03 • Where the obvious approaches fall short

APPROACH A

BI Tool Agents

- ✗ **No cross-report access.**
They only see the semantic model of the dashboard they're attached to, never reports from other BI tools.
- ✗ **No cloud database agents.**
They can't extend analysis beyond the single report they're connected to, so harder questions hit a wall fast.
- ✗ **Stalls on the deeper "why?".**
"What were sales last month?" is the question your dashboard already answers. The next one is where adoption stalls.

APPROACH B

Cloud Database Agents

- ✗ **No access to BI reports.**
They can't see the dashboards users rely on, or the semantic models and calculations embedded in them.
- ✗ **Narrow scope.**
Each agent only sees the tables in its own semantic model. Off-scope questions get no useful answer.
- ✗ **Disconnected from the user's trusted experience.**
Users can't validate answers against a report they trust, so they don't trust the answer.

BOTH APPROACHES SHARE 2 critical gaps

No Business Glossary

Without the data catalog, the agent doesn't know how the business actually defines its metrics. It computes the math, but not the meaning.

No Proactive Monitoring

Both are purely reactive. No business user opens a chatbot and types "did anything important change?"; they expect to be told.

04 • THE SOLUTION FOR ENTERPRISE-READY AI FOR BI

GIVE THE MODEL THE RIGHT HARNESS - METRIC INSIGHTS CONCIERGE

The right context, the right tools, the right framework, pulling from multiple sources, cross-checking against trusted reports, and proactively watching for changes that matter. That's the work an analyst does. That's the bar AI for BI has to clear. That's how Metric Insights Concierge performs.

