

The Eight Questions

A one-page brief for the companion scorecard — grade any AI vendor, weighted, including FactoryOS.

FACTORYOS BY COMPTRIO · YOUR DATA STAYS ON YOUR MACHINE. · [COMPTRIO.COM](https://comptrio.com)

One idea ran through the whole series: you don't have to be an AI expert to be a smart buyer of AI — you just need the right questions. These are the eight, turned into a scorecard. Grade your current AI in one column and anything you're evaluating in the other; the sheet weights the answers and shows you the gap. The questions are yours to use on anyone — FactoryOS included.

The criterion: not every question carries equal weight. A vendor can have a beautiful demo and still fail you on the three things a demo can't fake — **where your data lives, where the math runs, and whether compliance is enforced by the architecture**. The scorecard loads the weight there on purpose, so a high score that hides a gap on one of those three can't slip past you.

WHAT THE SCORECARD DOES

For each question, score how well a vendor answers — **Built into the architecture (3), Partial (2), Policy / promise only (1), or No / dodged (0)**. Each score is multiplied by the question's weight; the sheet returns a weighted percentage and a verdict for each vendor, plus the gap between them. It opens pre-filled with an honest read of a typical cloud AI as an example — overwrite it with your own assessment, and re-weight anything to match your priorities.

READING THE VERDICT

- **Strong** (75%+) — holds up where it counts.
- **Mixed** (45–74%) — press hard on the low scores before you commit.
- **Weak** (under 45%) — too much lives in policy and promises.

The three to watch: Q1 (data segregation), Q3 (deterministic math), Q5 (architecture-enforced compliance). A column that's high overall but sits at *Policy-only* or *No* on any of these is the one to chase down first.

USE IT ON EVERYONE

The point of the sequence was never to sell you a vendor — it was to make you the kind of buyer who can grade one. So grade FactoryOS on all eight, the same way you'd grade anyone else. We built it to answer them a particular way; whether it does is yours to check. A buyer who scores every vendor against the same eight questions doesn't get sold to. They decide.

The companion worksheet scores two vendors side by side across all eight questions, weighted, with the verdict and the gap computed live. Nothing leaves your computer — the point of the local, owned AI it's built for.

Source: the eight questions are the recap of Comptrio's Local AI Primer — one criterion per topic, from where AI runs to where its limits are (comptrio.com).

— Shawn Snarski · FactoryOS by Comptrio · your data stays on your machine. · comptrio.com

THE EIGHT QUESTIONS

1. **Where does it run** — can it see only what each seat is allowed, or everything? *Runs on hardware you control; least privilege enforced.*
2. **What does it cost in year three** — renting or owning? *An owned, depreciating machine, not a compounding subscription.*
3. **Where does the math run** — in the model, or an engine that shows its work? *A deterministic engine, with the work and the source.*
4. **Does it read your documents, or map them?** *Builds a knowledge graph and follows cross-references.*
5. **Where does compliance live** — paperwork or architecture? *Enforced by the system, not requested by a policy.*
6. **Where do the gates go** — and who decides? *You set where a human signs off; the dial is yours.*
7. **If you left tomorrow, what would you keep?** *Your data, corrections, workflows, access rules — not a wrapper.*
8. **Where are the limits** — will they hand you the limits sheet? *Honest about what it can't do.*