The Single Source of Truth Trap

AI deployment decisions your data warehouse initiative is already making for you.

01

Background

What happens when AI arrives before the data is ready.

01 — BACKGROUND

Fragmented AI. One data warehouse. A decision nobody made.

- 130-person company mid-build on a data warehouse to replace fragmented spreadsheets, SaaS exports, and tribal knowledge with one reliable data layer.
- AI arrived informally: a subset hold Claude licenses for analysis and drafting; the majority have Microsoft Copilot through an existing M365 agreement.
- Actual AI usage is low and uncoordinated. No AI owner, no policy, no connection between AI tools and the data infrastructure being built.
- The risk: by the time the warehouse is live, AI deployment decisions will have been made by default — through vendor contracts, habits, and IT config — rather than by design.
- A data warehouse without an AI strategy is a reporting tool. With one, it becomes a decision engine.

02

Decision Required

The question that must be answered before the warehouse ships.

02 — DECISION REQUIRED

What AI architecture sits on top of the warehouse — and does it match the tools already in employees' hands?

> Deferring this decision means the warehouse ships and AI usage continues unchanged — disconnected from the clean data that was supposed to make the company smarter.

> The current split (Claude for some, Copilot for most) was never chosen. It accumulated. That is a different problem than the one the data warehouse was built to solve.

> This decision must be made before go-live, not after. Post-launch tooling changes cost 3–5x more in retraining, re-contracting, and habit disruption.

03 — OPTIONS

Three paths forward.

Option A

Status quo — let organic adoption continue

Keep the split (Claude / Copilot), make no formal connection to the warehouse, revisit after go-live. The warehouse ships into a vacuum.

Option B

Microsoft consolidation via Fabric + Copilot

Standardize on Copilot. Connect the warehouse to Microsoft Fabric. Retire ad-hoc Claude licenses. Deep dependency, but one integration to maintain.

Option C

Recommended

Dual-track: Copilot for productivity, Claude for analysis

Copilot owns M365 tasks (docs, email, meetings). Claude owns analyst workflows against warehouse data. Two tools, two jobs, clear ownership.

04 — RECOMMENDATION

Formalize what is already working. Set a forcing function at warehouse go-live.

The employees using Claude for analysis are your AI-capable cohort. Give them a mandate: they own the use-case library for AI-assisted analysis against warehouse data.

Let Copilot do what it already does well for the rest of the company: meeting summaries, document drafting, email.

Do not connect Copilot to live data until the warehouse reaches production readiness. Early exposure to pre-warehouse data will create a "Copilot is wrong" reputation that survives the fix.

Set a forcing function: at warehouse go-live, run a 30-day evaluation of Claude API vs. Power BI Copilot for data querying. Decide with real schema and real user patterns, not vendor preference.

Assign an AI deployment owner now. At 130 people this is a named BA or analyst, not a new hire.

05 — RISKS

Four material risks.

1.
Warehouse delay freezes AI strategy

If the warehouse timeline slips, AI deployment stays in limbo and employees build habits on consumer tools (ChatGPT, Gemini) that are harder to displace later.

2.
Copilot deployed against bad data

Employees who experience Copilot as "wrong" stop using it. That reputation sticks even after the warehouse is live and the data is clean.

3.
No AI owner

Without a named owner, vendor renewals happen on autopilot, Claude licenses drift, and nobody connects the warehouse to anything. The investment compounds for nobody.

4.
Shadow AI fills the gap

Employees who want more than Copilot delivers will find it. Free ChatGPT, personal Claude accounts, Gemini on personal phones. A clear policy on approved tools reduces exposure from unmanaged data handling.

06 — QUESTIONS YOUR TEAM SHOULD BE ANSWERING

If your team cannot answer these, that is your first deliverable.

1. Who owns the data warehouse initiative — and are they in the room when AI tooling decisions are made? If not, these two workstreams are on a collision course.
2. What data will live in the warehouse at go-live? Does any of it carry PII, client data, or confidentiality obligations that would constrain which AI tools can query it?
3. Do your current Copilot licenses include the M365 Copilot tier for data analysis, or the base productivity tier? These are different products at different price points.
4. Of the employees currently using Claude, what are they actually doing? If it is document drafting — that is Copilot territory. If it is multi-step analysis — that advantage is worth preserving.
5. What does warehouse success look like in year one? If the answer involves faster decisions — what does the human workflow look like to get from data to a decision? AI lives in that gap.
6. Is there a training plan? At 130 people, one two-hour session on how to use AI against your data will do more for adoption than any tooling decision.

AI INSIGHT LAB

The Deployment Memo

One enterprise AI deployment, dissected every Tuesday.
Written for executives who have to decide, not just read.

Subscribe at aiinsightlab.cloud — free during beta.