

ASSESSMENT — FINDINGS & RECOMMENDATIONS

AI Value & Accountability Assessment

Prepared for Dana Whitfield, CIO — Northwind Mutual Health (*illustrative*)

ILLUSTRATIVE SAMPLE — methodology demonstration only. Fictional client and data; contains no real or confidential information from any prior employer. Shown to demonstrate deliverable quality and approach.

Executive Summary

Over 18 months, Northwind invested ~\$3.0M in AI across claims, member service, and underwriting. The investment was real; the accountability around it was not. We found spend disconnected from value, 14 unsanctioned AI agents with no owners, a vendor renewal carrying multi-year lock-in, and 'governance' that existed on paper but not in enforcement.

~\$1.1M of the \$3.0M was producing measurable value; ~\$1.9M was unmeasured or redundant. The changes below let Northwind defend the program to the board, reclaim ~\$700K/year, and avoid a seven-figure vendor lock-in.

Where the AI budget actually went




Figure 1 — Of \$3.0M in AI spend, only ~\$1.1M was tied to measurable value.

What Was Wrong

Spend with no line to value (High). AI budget approved per-initiative with no shared definition of 'return.' A \$3.0M program and no number the CFO would accept.

The accountability gap is real (Critical). 14 AI agents/tools in production that IT never approved — built on corporate and personal accounts, no owners, several touching member PHI.

Vendor lock-in hiding in the renewal (High). The pending renewal assigned derived-data rights to the vendor and carried a 36-month exit penalty — unflagged by procurement or engineering.

Governance theater (Medium-High). An 'AI Usage Policy' existed with no enforcement — no approval gate, no monitoring, no consequence.

Cost trajectory unmanaged (Medium). Compute/token consumption grew ~22%/quarter with no budget owner — a curve that doubles the bill in under a year.




Figure 2 — Unmanaged AI consumption was on track to double the bill within a year.

Why It Happened

Not incompetence — velocity outran controls. Adoption was incentivized; no one owned the *system* (intake, ownership, value definition, enforcement). A classic process gap: the org optimized for adoption and never built the operating layer underneath.

How We Fixed It

ROI framework + board narrative. Mapped each initiative to tangible and intangible returns. Result: a defensible number and a one-page board story leadership can stand behind.

Agent inventory + ownership model. Catalogued all 14 agents; assigned owners or decommission dates. 6 retired, 8 brought under governance.

Governance with enforcement. A lightweight operating layer: intake/approval gate, monitoring, ownership tied to budget, steering-committee cadence — control that scales without killing experimentation.

Vendor diligence → leverage. Flagged the data-ownership and exit clauses; delivered a renegotiate recommendation with three specific leverage points and a switch-cost model.

12-month roadmap, value-prioritized. Double down on the ~\$1.1M that works, reclaim the ~\$700K that doesn't, gate the rest; cost controls cap the curve.




Figure 3 — Every one of the 14 shadow agents was retired or brought under governance.

Illustrative Impact

- **Board:** program defended and re-funded with a credible ROI position.
- **Financial:** ~\$700K/yr reclaimed; seven-figure lock-in avoided; spend curve capped.
- **Risk:** PHI-touching shadow agents brought under control before audit exposure.
- **CIO:** becomes the leader who got AI under control.

Methodology: operator-led discovery → value mapping (monetized) → agent inventory → root-cause analysis → governance design → vendor diligence (lock-in, data terms, evaluation method, exit) → board-ready synthesis. Lean Six Sigma rigor applied to the operating layer.