

CASE STUDY · CROSS-INDUSTRY · COST DISCIPLINE & INTELLIGENT ROUTING

Across the design-partner cohort, every verified execution made the next one cheaper

Dipp Intelligence learns from the shape of verified executions, never from tenant data. Routing accuracy improved 18 points across the cohort in two quarters.

Partner Design-partner cohort · 9 organisations · 6 industries	Published 2026-10-07	Orcher components Cost Governance, Observability, Logic Scrubber
--	--------------------------------	--

Summary

This study aggregates the design-partner cohort rather than one organisation. Dipp Intelligence compounds from the metadata of verified executions — stakes, verification outcome, cost, escalation — and never from tenant content. Over two quarters, routing decisions across the cohort got measurably better without any partner's data crossing a boundary.

+18pp routing accuracy across the cohort correct tier on first attempt, two quarters	0 bytes of tenant content used for learning	9 organisations across six industries	34% further cost reduction on top of initial routing gains
---	---	---	--

The situation

- Each partner started with routing floors set by judgement, because there was no evidence base to set them from.
- Conservative floors meant over-spending on directive classes that did not need the frontier tier.
- Partners would not accept any arrangement in which their content improved a shared system.

What Orcher enforced

- Dipp Intelligence learns only from execution metadata — stakes class, verification outcome, escalation event, cost, latency — never prompts, never content, never records.
- Each partner can inspect exactly what left their boundary, which is a short, enumerable list.
- Improved routing priors are offered as recommendations that a partner's own governance accepts or rejects.
- Every accepted change is recorded in the ledger like any other governed decision.

The data

Quarter	Correct tier first attempt	Escalations	Cost index	Partners contributing
Q2 2026	67%	22%	100	4
Q3 2026	78%	15%	81	7
Q4 2026 (partial)	85%	9%	66	9
Change	+18pp	–13pp	–34%	+5

Cohort-level routing performance. Cost index is normalised to Q2 2026 = 100 at constant directive mix.

Outcomes

- New partners start with priors derived from the cohort rather than from guesswork, shortening time to a stable configuration.
- The learning surface is small enough that partners' data teams can audit it in an afternoon.
- Escalation fell without a corresponding rise in verification failures, which is the result that matters.
- The compounding effect is now the strongest argument in early-access conversations.

“We agreed to it because the list of what leaves is short enough to read, and none of it is ours in any meaningful sense.”

— Head of AI Platform, design-partner organisation (partner declined attribution by name)

Read further

- Dipp Intelligence — the compounding moat — dippai.com/research/dipp-intelligence-the-compounding-moat
- Dipp Intelligence — dippai.com/dipp-intelligence
- Full case study — dippai.com/case-studies/dipp-intelligence-compounding-across-partners

Disclosure. Design-partner engagement, anonymised at the partner's request. Figures are measured by Orcher's own observability and ledger instrumentation over the stated period and have not been independently audited. Market figures carry their own source line.