

This plan is deployed, not proposed. Every cost figure below is printed from the live engine's run ledger — measured, not modelled.

1 · LEAD MAGNET LADDER each rung earns its place with a mechanism

- **UNAWARE — UK Agency Client-Churn Index (quarterly)**
Aggregated from the engine's exhaust: CS-leadership moves and named account losses across trade press. Owned data no one else compiles; every agency MD reads it asking "am I in it?". Marginal cost ~ £0.
- **PROBLEM-AWARE — CARE Client-Health Self-Audit**
Score your top five accounts by hand on C-A-R-E. The buyer discovers mid-worksheet they can't answer without data — the gap is the demo. Deliberately manual, so it funnels into the product, never replaces it.
- **SOLUTION-AWARE — Signal Teardown → Paid Pilot**
"What the market can see about your agency": a personalised scan of their public footprint, bridging straight into Kaizan's existing paid-pilot motion (£2,995 / 90 days) — the ask sales already makes.

2 · CHANNEL MIX one hypothesis each, tested not assumed

- **Signal-led outbound (email + LinkedIn DM)**
Timing beats persona: a first touch triggered by a verified event converts ≥3× a cold list. The engine's core motion.
- **Rep content engine — assigned by market signal**
The radar doesn't just find prospects; it assigns each rep non-overlapping weekly angles from live market events. Reps read as industry insiders changing the space, not account reps.
- **Churn Index as PR / inbound**
A quarterly data release is inherently newsworthy to the same trade press the engine reads. The magnet and the radar are one machine.
- **Paid — deliberately not yet**
Paid amplifies a proven message. Enters only after an organic message clears its MVT. Seed money never amplifies a guess.

3 · AI ARCHITECTURE — KAIZAN'S PRODUCT, POINTED OUTWARD (the part most people skip)

Kaizan watches client relationships for risk and white space (CARE). This engine runs the same discipline on the market: public signals in, human-approved revenue motion out.



Two surfaces, one board. The dashboard is for review and approval; MCP is for depth — each rep's Claude carries an installable skill (official skills released by admin; reps build and share their own). A rep asks "should we go after Boots?"; their Claude researches freely, proposes better adjacencies ("No7 is the sharper wedge"), and writes the sourced conclusion onto the board — where the same suppression check, scoring and human gate apply. Claude proposes; the board disposes. Sending stays human.

Automated: hunting, citation-verified structuring, enrichment, scoring, drafting, CRM hygiene, the cost ledger. Human: approval, sending, the relationship.

Deferred, with triggers: HeyReach sequencing (after reply-rate proof) · Clay (at team scale, as the collaboration surface) · paid (after message proof) · mobile approve-queue (when reps exist).

4 · MINIMAL VIABLE TESTS — HYPOTHESES WITH KILL CRITERIA, NOT CAMPAIGNS

	HYPOTHESIS	MECHANIC	SUCCESS METRIC	KILL CRITERIA	COST
MVT-1 Outbound	Signal-triggered first touch converts ≥3× a cold list — timing beats persona	Radar sweep → score → human gate → 60–90 word first touch	≥8% reply in 4 weeks; ≥2 demos per 50 sends	<3% reply after 100 sends across two signal types → re-cut the ICP. The engine survives; the hypothesis dies.	~£0.15 / lead + review minutes
MVT-2 Content	Market-signal-assigned posts make reps read as insiders, not account reps	Engine assigns 2 angles / rep / week, non-repeating across the team	≥5 ICP-title conversations per month from content	<2 after 6 weeks → change format, keep cadence	~£0 (engine exhaust)
MVT-3 Churn Index	Owned churn data pulls the exact buyer inbound	Quarterly release, gated download, PR to the same trade press the engine reads	Download → demo ≥3%	One cycle without a sourced demo → demote to blog content	~1 day / quarter

5 · SEED-STAGE ECONOMICS — MEASURED ON THE LIVE ENGINE, 3–4 JULY 2026

£0.008 score a lead (ICP + signals)	£0.005 draft email + LinkedIn note	~£0.10 full radar sweep (two search angles)	<£0.15 per scored, drafted, human-reviewed lead	£59k–£108k Kaizan ACV (Team → Growth tier)	days CAC payback, even at 0.5% lead→customer
---	--	---	--	--	--

6 · IF HIRED — THE FIRST FOURTEEN DAYS

DAYS 1–3 Audit what exists: HubSpot state, Apollo plan, domain health. Repoint the engine at production; warm a sending domain.	WEEK 1 First 50 signal-led sends through the human gate (MVT-1 live). Baseline: reply rate by signal type, cost per lead.	WEEK 2 MVT-2 live: content cycle from market signals, with Leah on format. Churn-Index scaffold begun. Kill/scale call on MVT-1.
---	---	--

HELD IN RESERVE (ROADMAP)

Skill registry (official + rep-built) · churn-index automation · HeyReach sequencing · Clay handoff at team scale · PLG signals when self-serve lands · mobile approve-queue