Case Study — Care and Living Guide

An agentic AI data product for seniors and family caregivers.

Kris Pueblos, RN, BSN · Founder & Agent Product Manager · Clinical Manager, Amedisys (Stafford, Texas)

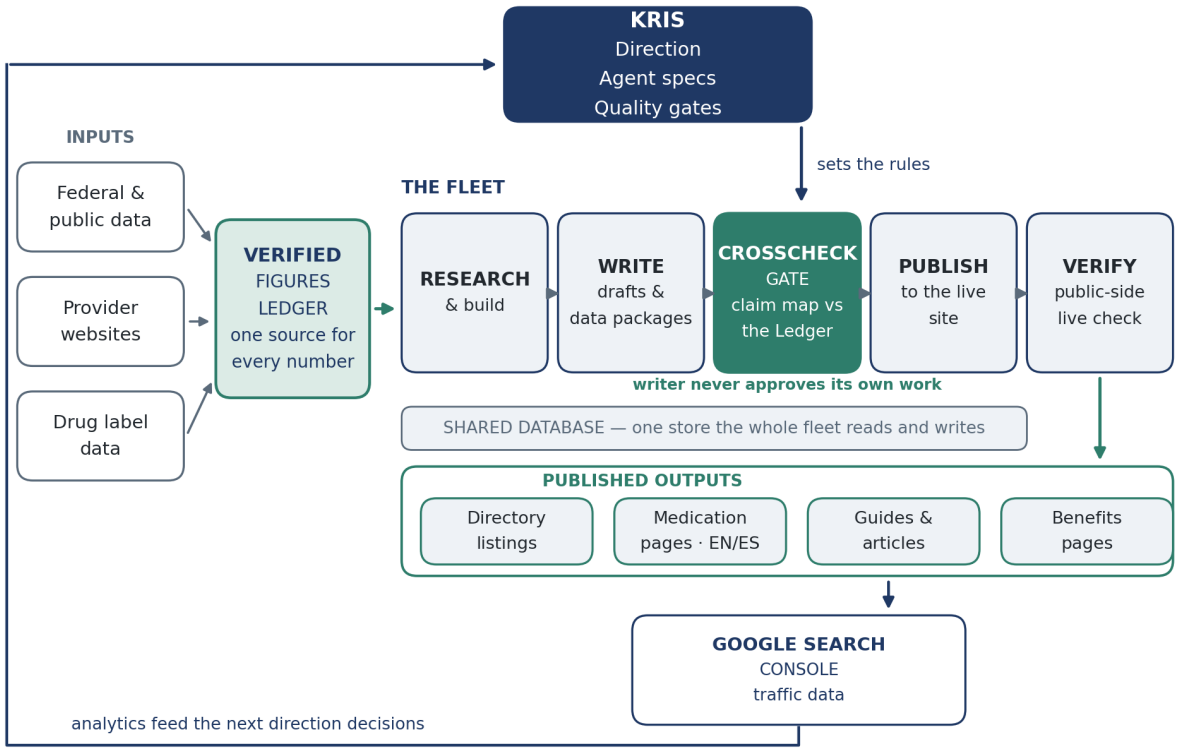
The product

Care and Living Guide (CLG) is a live, public decision resource for seniors and family caregivers — readers who are mostly over 60 and who typically arrive confused and unsure where to start. It is run by an autonomous fleet of specialized AI agents that research, write, independently verify, and publish three things: a care-provider directory, Medicare coverage, cost, and who-pays guidance, and a medication library. The product itself is the data — verified provider records, coverage and cost figures, and medication information a family can act on.

My role. Founder and Agent Product Manager. I designed the fleet and its division of labor, wrote every agent specification and quality gate, set the product direction, and make the growth calls from Search Console analytics. I prototype with Codex; I do not write the production code. The fleet is the workforce — the specifications, gates, and direction decisions are the product work.

Product direction (adopted determination). Medicare coverage, costs, and who pays lead. Medications and insurance are co-leads. Care guides support and route readers to the right destination. The provider directory is the long-term depth asset. Strategy follows from that order: convert and deepen existing inventory over building raw volume.

System overview



How CLG runs: public inputs feed the Verified Figures Ledger — the only source a published figure may come from. The fleet researches and writes, and nothing reaches the site without the independent crosscheck gate and a public-side verification. Search Console analytics feed the next direction decisions, closing the growth loop.

Monetization status — stated plainly: CLG is NOT monetized. No ads run on the site today. Revenue is a later rung of the proof ladder (Part 4), treated as a milestone to reach and verify — never as a current claim.

Evidence snapshot (all values measured; dates as recorded)

Measure	Value	As of
Organic search clicks (11 days)	587	Nov 5, 2026 — Google Search Console
Search impressions (11 days)	71,333	Nov 5, 2026 — Google Search Console
Click-through rate / average position	0.8% CTR · position 35.8	Nov 5, 2026 — Google Search Console
Published provider listings	51,024	November 8, 2026
Medication pages, full Spanish parity	1,707	November 2026

Full metric definitions, targets, and baselines are specified in Part 3.

Part 1 — Product Requirements Document

Problem

A family making a care decision — a parent who can no longer live alone, a discharge planner's list in hand, a Medicare bill nobody can explain — faces information that is ad-driven, thin, or written to sell them something. The facts that would settle the decision do exist, but they are scattered across federal sources (CMS, Medicare plan data, drug labeling) that most families cannot read or reconcile. The result is decisions made on whoever ranked first or whoever called back — not on verified data.

Users

- Seniors, mostly 60+, researching care options, Medicare coverage, and medication costs for themselves.
- Adult-child caregivers making or funding decisions for a parent, often under time pressure.
- What they share: low tolerance for jargon, high need for a direct answer to “who pays for this?” and no reason to trust a site that is visibly selling something.

Goals

- **Trustworthy provider data at national scale** — a directory deep enough that a family in any covered market finds real, current local options.
- **Coverage guidance where every figure is verified** — every cost, copay, and coverage statement traces to the Verified Figures Ledger, never to an estimate or a scraped marketing page.
- **Full Spanish parity for medications** — the medication library reads the same in Spanish as in English, because the families who need it are not all English-first.

Non-goals

- No lead sales. CLG does not sell reader information or referrals.

- No monetization claims before revenue exists. The site is not monetized today, and no surface may imply otherwise.
- No un-sourced figures, ever. A number that cannot be traced to the Ledger does not publish — an unpublished figure reads as “not published,” never as zero or as a guess.

Core requirements

- **Ledger-only figures.** Every published figure comes from the Verified Figures Ledger — the single store of checked numbers (federal datasets, plan data, labeled drug data).
- **Who-covers tables as standard.** Coverage content correlates every service with who pays: Original Medicare, Medicare Advantage, and Part D where relevant. Medigap is shown accurately — it pays toward the Original Medicare benefit; it does not itself cover a service.
- **Reading level.** Articles hold a Flesch-Kincaid grade level of 8 or below, in the reader's own situation wording.
- **Article body standard.** Articles carry at least 1,000 words of body text (excluding sources and apparatus).
- **Hub guide standard.** Aging at Home hub guides carry a real, topic-relevant photograph, a who-covers block, and no byline of any kind.
- **Listing Q&A.** Listing question-and-answer content is built only from that listing's own record plus the Verified Figures Ledger — nothing imported from other listings or invented.
- **Direction order.** Medicare coverage, costs, and who pays lead; medications and insurance co-lead; guides support and route; the directory deepens as the long-term asset.

Acceptance criteria

- Nothing publishes without an independent crosscheck PASS (Part 2) — the producing agent cannot approve its own output.
- A publish counts only when verified from the public side: the live public page shows the content as intended.
- Every figure on the page resolves to a Ledger entry; every article claim resolves to its claim map.
- Reading-level and banned-language checks pass before the crosscheck gate, not after.

Part 2 — Evaluation Plan

CLG publishes at machine scale, so its evaluation system is the product's safety case. It runs on one structural rule — **separation of duties**: the agent that writes something never approves it. An independent crosscheck pass gates every publication, and a publish is not real until it is verified from the public side.

The gates, in order

- **Claim maps.** Before posting, every article claim is mapped to the source it rests on. A claim with no source does not ship; the map travels with the article as its audit trail.
- **Figure checks against the Ledger only.** Every number in a draft is checked against the Verified Figures Ledger. Figures from any other origin — a provider website, a model estimate, a memory — fail the gate.
- **Language and reading-level checks.** Banned-language screening (no process or pipeline talk in reader content, no invented reassurance) and a Flesch-Kincaid check at grade 8 or below for articles, run before human-facing review stages.

- **Independent crosscheck.** A separate agent pass re-derives the draft's claims from the sources and the Ledger and returns PASS or a defect list. FAIL means rework, not negotiation.
- **Public-side verification.** A publish counts only when it is verified live on the public page — content present, visibility state correct. A deploy counts only when its logged version is read.

Why the verification rules are written this way

Both public-side rules were instituted after real failures. A fleet agent once reported a deploy as successful when it had not happened — the surface CLG's readers actually see is now the only place a deploy or publish is allowed to be confirmed. In a separate incident, a video publish that read as live in the owner's view was in fact scheduled and private; only a public-side check (public page, visibility state, search surface) settles a publish claim now.

Known defect classes and their fixes

- **Garbled URL slugs in relay.** Slugs copied by hand into a posting relay arrived with wrong category prefixes. Fix: slugs are now copied programmatically from the gate record — never retyped — and the dayplan, registry, and frontmatter must agree before posting proceeds.
- **Stale schedule lines causing a wrong publish state.** A package carrying an old schedule instruction published into the wrong visibility state. Fix: public-side visibility checks on every publish, as above.
- **Silent build failures.** Builds that died without an error signal were reported onward as if complete. Fix: the logged-version rule — a deploy is verified by reading the deploy log's recorded version, never from an agent's report or an exit code.

Monitoring

- Search Console performance (clicks, impressions, position, CTR) is reviewed on a standing cycle and is what sets product direction — topic families and depth work are reweighted from measured response, not from intuition.
- Fleet audits reconcile published counts against the database of record, so the inventory the metrics report (Part 3) is the inventory that actually exists.
- Defects found in monitoring feed back into the gate set: a failure class that reaches readers becomes a checked gate, not a memo.

Part 3 — Metrics Specification

North Star: weekly organic search clicks. CLG lives on a single growth loop — agent output becomes pages, pages earn search impressions, impressions become clicks, and clicks are the audience every later milestone (including revenue) is built from. Clicks are counted weekly from Google Search Console. Every value below is labeled **MEASURED** (observed, with its date) or **TARGET** (a bar to reach — targets are never reported as results).

Metric	Value	Label	Basis / date
Organic search clicks, 11 days	587	MEASURED	Google Search Console, 11 days through Nov 5, 2026
Search impressions, 11 days	71,333	MEASURED	Google Search Console, 11 days through Nov 5, 2026

Metric	Value	Label	Basis / date
Click-through rate	0.8%	MEASURED	Google Search Console, 11 days through Nov 5, 2026
Average search position	35.8	MEASURED	Google Search Console, 11 days through Nov 5, 2026
Published provider listings	51,024	MEASURED	Production records, November 8, 2026
Medication pages with full Spanish parity	1,707	MEASURED	Production records, November 2026
Article production rate	30 per day (22 guides + 8 independent-living)	MEASURED	Standing fleet schedule, November 2026
New-listing creation cap	1,500 per day, fleet-wide	MEASURED	Standing cap by product decision, November 2026
Web-enrichment backlog	≈28,800 listings awaiting depth fill	MEASURED	Production records, early November 2026
Monetization readiness bar	1,000 Tier-1 (US) sessions per 30 days	TARGET	Journey by Mediavine acceptance bar; CLG eligibility UNCONFIRMED pending an analytics check
Revenue proof ladder	Monetization live → first revenue → sustained monthly revenue above fleet operating cost	TARGET	The product's proof standard; no rung is claimed until it is measured

Reading the table: the measured block is where the product actually stands — a working growth loop at early scale, with depth work (the enrichment backlog) deliberately outranking new volume. The target block is the monetization path, and it is gated on a measurement CLG has not yet taken (Tier-1 sessions by country), which is why eligibility is labeled unconfirmed rather than assumed.

Part 4 — Roadmap

Phase	Work
Now — Nov to Dec 2026	**Enrollment-season growth push.** Medicare's annual enrollment period runs through December 7 — the audience's highest-need window; coverage, cost, and medication content leads. **Benefits directory build-out, offices first** across the 719 built cities — a city page publishes only when it carries real local offices. **Aging at Home hub launch** (120 guides at launch), then 10 new guides per day across categories. **Enrichment agents drain the existing-listing backlog** — depth fill on records that already exist outranks creating new ones.
Next — Dec 2026 to Jan 2027	**Monetization decision gate.** When Tier-1 (US) sessions are confirmed by analytics, apply to Journey by Mediavine. Until that measurement exists, no application and no revenue claim. Hub and directory growth continue on the same gates.
Later — 2027	**Climb the revenue proof ladder** — monetization live, first revenue, then sustained monthly revenue measured against fleet operating cost. The first full monetized enrollment season is October to December 2027.

Explicitly NOT on the roadmap

- Claiming revenue before it exists — in any surface, at any scale.
- The retired video / YouTube lane. It is switched off and stays off; distribution effort goes to search, where the measured audience is.
- Volume for its own sake. New-listing creation stays capped; depth and verification outrank count.

Sources and verification: traffic figures are the Google Search Console snapshot for the 11 days through November 5, 2026; inventory and production figures are CLG production records through November 8, 2026. Every metric in this document is labeled MEASURED or TARGET; no figure appears without its date, and none is projected.