PORTFOLIO

Kris Pueblos, RN, BSN

Clinical AI Product Manager | Hospice Clinical Manager, RN, BSN · Houston, TX

How much autonomy should an AI system get on a clinical record? None on its own. I shipped that answer.

This system puts AI to work on a clinical record — so it never acts alone: clinicians submit, independent validators check, a person decides. It is live in production. Drawing that boundary in the design itself is the product work. This is an independent product, built solo — I define the problem, users, scope and safety boundaries, write the specifications and acceptance criteria, and direct the build. Not an employer product.

This portfolio is the complete product record for the two production systems below: the requirements, the workflow, the case study, the evaluation and the metrics.

8 + 2
interlocking tools in the hospice system, plus two deterministic risk models — all live on one shared patient record

1,496
reference records in the Clinical Audit corpus every per-visit scan is checked against

749
deficiencies surfaced by the Clinical Audit of the demo heart-failure record — 38 visits, every one cited

6,642
hospices scored in CMS's FY2025 SSVI file — the public data behind every Patient Chart Rating's agency standing

Clinician AI Assist — clinicianaiassist.com · live, with demo patients
Care and Living Guide — careandlivingguide.com · live
LinkedIn — linkedin.com/in/kris-pueblos

The two systems

Clinician AI Assist is human-gated clinical decision support for hospice: eight lanes working one shared patient profile from referral to recertification, plus two deterministic risk models — Patient Chart Rating and Census Risk — and a fenced autonomous layer of two operations agents that never touch clinical analysis.

Care and Living Guide is the deliberate opposite: an autonomous multi-agent fleet that researches, writes, independently verifies, and publishes a consumer Medicare directory daily under written specifications — because the data is public and the reader is researching, not acting on a clinical record.

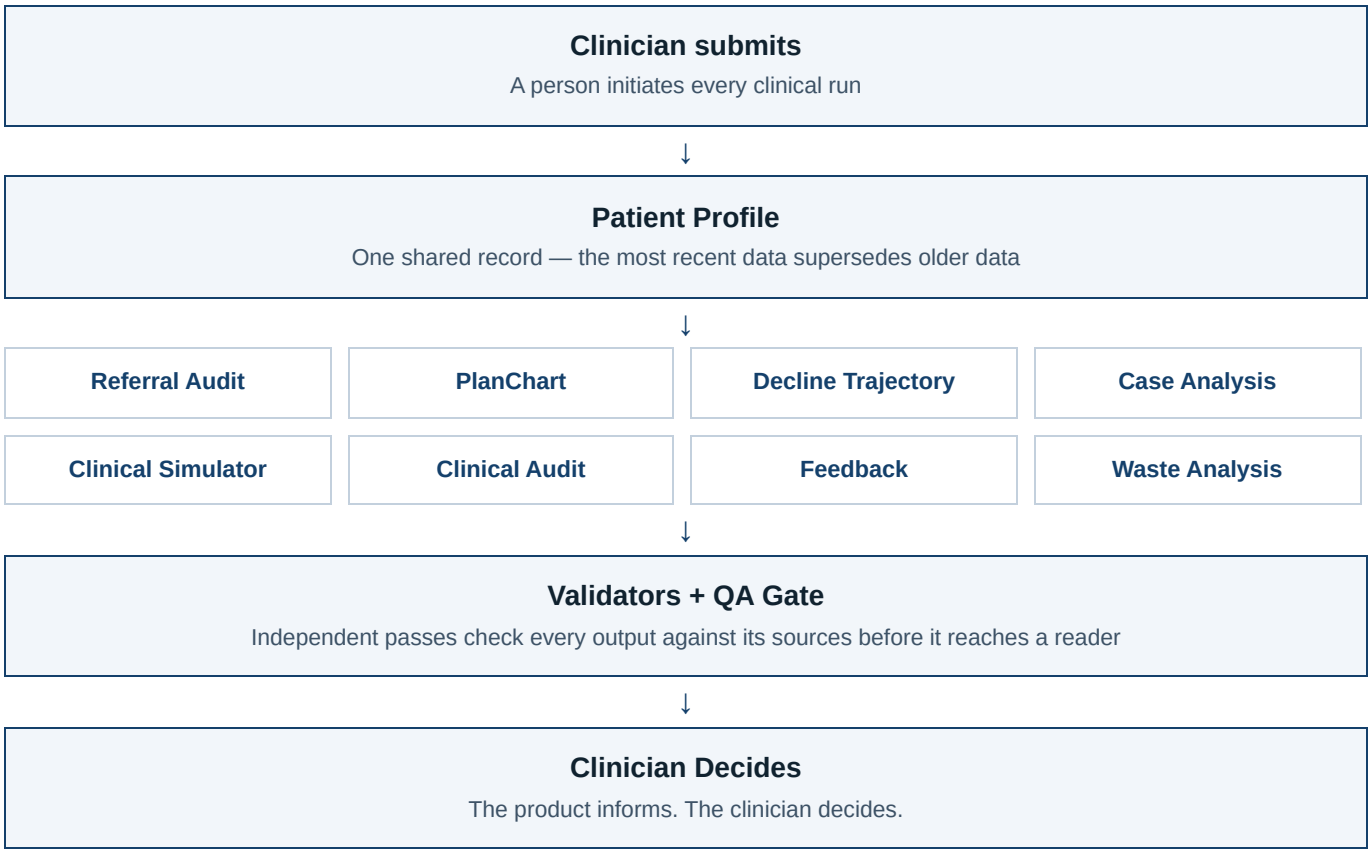
Both systems run fully on Cloudflare (Workers + D1) and started as prototypes I built with Codex. All patient cases in Clinician AI Assist are synthetic. Clinician AI Assist is designed around FDA non-device clinical decision support principles and is not FDA cleared.

SYSTEM 01 — CLINICAL DECISION SUPPORT

Clinician AI Assist — eight systems, one profile

Hospice assessment information is fragmented across the referral packet, visit notes, and the chart — and the data that looks reassuring is often the data that misleads: stable vitals, one good meal, one good transfer, while the trajectory underneath is declining. Reviewers, IDG teams, and Medicare auditors each see a different slice of the same patient. Clinician AI Assist follows one patient record from referral to recertification — eight live, finished lanes on a shared profile, instead of eight disconnected tools.

Workflow — eight systems, one profile



How every analysis reaches its reader, checked

1. Clinician submits

A person initiates every clinical run.

2. Analysis engine

Pattern library + live evidence check + adversarial cross-check · gpt-5.6-luna.

3. Independent validators

Separate passes check output against sources.

4. Clinician review

Analysis presented as evidence — the system analyzes; it does not hold or decide.

5. Clinician decides

The product informs. The clinician decides.

SYSTEM 01 — THE EIGHT TOOLS

The eight tools

Lane	What it does
Referral Audit	12-section, source-cited report with disposition; creates the baseline. Structural failures gate as HOLD FOR DATA CORRECTION; review program selectable (TPE default, SMRC, UPIC).
PlanChart	Plan of care fed strictly by the Referral Audit's output at admission (enforced in code); whole-period chart at recertification.
Decline Trajectory	Six-domain acuity tracking across the course; tracks, never predicts. A content gate holds any drift toward prognosis, eligibility, or orders.
Case Analysis	Whole-record engine on a pattern library grounded in CMS Conditions of Participation, LCD/MAC guidance, and published clinical sources; output must name its contributing systems.
Clinical Simulator	Practice scenarios built from the patient's own record only.
Clinical Audit	Independent per-visit documentation deficiency scan and audit-readiness rating that never consumes other systems' conclusions.
Feedback	Clinician feedback lane on the system's own runs; tester input decides what gets fixed and built next.
Waste Analysis	Account-level synthesis of all lanes: seven waste types enumerated, findings naming clinicians from the visit ledger, and a corrective action plan where every step carries a predictive forecast (change, probabilities, outcome, horizon, basis) plus no-intervention probabilities. Reports support Save / Delete / Print to PDF.

Two risk models on the same profile

Patient Chart Rating joins chart evidence from the lanes with the agency's public CMS standing into one rating per patient, shown with its full analysis and a plain clinical translation — including a Probable ADR likelihood per chart. **Census Risk** rolls every rated chart into one census-level analysis, with exposure stated as higher risk for elevation of investigation — never as fraud or wrongdoing. PEPPER data is facility-confidential and is never a source; it can only enter as an upload from the provider's own hands.

Boundaries, enforced — not disclaimers

- No prognosis. Trajectory is tracked and scored — never predicted.
- No eligibility or certification calls. The determination stays human.
- No autonomous orders, nothing unattended. Every run starts with a clinician's submission.
- No unsourced values. Every figure traces to a source, a calculation, or a labeled translation.

SYSTEM 01 — THE AUTONOMOUS LAYER

The autonomous layer — and the fence around it

Clinician AI Assist has two autonomous agents — added deliberately, and fenced deliberately. They live on data and structure only: they never run a clinical lane, never write a real patient record, and never change a clinical result. Nothing is ever applied automatically; the clinician views and decides.

Operations agent — hourly

- Fixed 7-slot rotation, one task per run: site-error sweep, CMS public-data watch, registry watch, regulatory-library watch, MAC watch, clinical-reference refresh, integrity sweep.
- Deterministic by design: fixed order, sources, and checks — findings never change what runs next; official sources only, every record carries source + checked date.
- Reports, never fixes: site errors are reported with evidence; no code edits, no deploys, no lane runs.
- Demo reset first: every run first restores the 4-profile synthetic demo cohort (222-row baseline) so clinicians can test freely.

Guideline Watch — daily

- Scans completed analyses across the lanes and deep-researches CMS, the Coverage Database, and both MACs for guideline changes those analyses rest on.
- Writes alerts only: per-account update-needed alerts with what changed, source, effective date, severity (high / moderate / low), and exactly where it applies — profile, lane, section.
- Patient-specific updates appear on that patient's page; account-level alerts sit in the strip under the nav.
- Seeded from the Nov 2026 hospice update (CY2027 final rule, CMS-1844-F), each provision verified against the CMS fact sheet.

Waste Analysis — the synthesis lane

Account-level, not patient-bound: it compiles the analyzed output of all other lanes across every profile the account can see — the correlated record per patient, Patient Chart Rating, the Clinical Audit visit ledger with named clinicians, and Census Risk over the dominant-CCN cohort. Waste is enumerated across the seven types: Inefficiencies, Unjustifiable visits, Copy-and-paste visits, Clinical inefficiencies, Medication mismanagement, Overuse of non-beneficial interventions, Operational bottlenecks. Every corrective step is simulated against all correlated analyzed data and carries its own forecast. Running it generates an unsaved report; saving is an explicit choice.

Account and demo model

All tools sit behind an account gate. Synthetic demo patients are shared read-only with every account — including Dorothy A. Kessler (heart failure) and Robert J. Delgado (end-stage COPD), whose records carry fully processed analyses open to inspection end to end after sign-up. Every account-created profile is owner-scoped and invisible to all others, enforced in the data layer.

SYSTEM 02 — AGENTIC AI, IN PRODUCTION

Care and Living Guide — the autonomous pole

Care and Living Guide (careandlivingguide.com) is a decision-and-directory resource for seniors and family caregivers — readers mostly 60 and up, making care decisions. It answers the family's real question — what care exists near us, and who pays for it? — and it is produced and maintained entirely by an autonomous multi-agent fleet working under written specifications.

51,024
published provider listings — as of November 8, 2026

3,542
Medicare Advantage plans with cost and coverage detail

1,707
medication pages with full Spanish parity and Part D detail

100%
of published facts carry a source and a date

Why autonomy is safe here

The data is public and the reader is a family researching options, not a clinician acting on a record. That risk profile earns autonomy — with structural controls: the agent that writes a page never approves its own work, a database guard blocks unvalidated content, dependent surfaces (aggregator pages, sitemap, search index) update in the same publish run, and a publish counts only when it is verified from the public side.

The fleet's working day

1. Gather
Agents pull from authoritative public sources — CMS and federal datasets, plan and formulary filings — on a daily cadence.

2. Deepen
Existing records are filled out and enriched before new ones are created.

3. Write
Specialized agents draft listings, guides, and coverage pages under written specifications.

4. Cross-check
A different agent verifies against the source — the writer never approves its own output.

5. Publish
Verified from the public side; dependent surfaces update in the same run.

Results so far — stated with their dates

Google Search Console, 11 days through November 5, 2026: 587 clicks, 71,333 impressions, 0.8% CTR, average position 35.8 — an early, working growth loop, presented as measured, not as breakout scale. Internal audit, November 2026: 0 mismatches across published facts.

> **Monetization — stated plainly.** Care and Living Guide is not monetized. No ads run on the site today. The current phase is audience growth through Medicare's annual enrollment season; revenue is a future proof milestone, not a current claim.

My role: Founder & Agentic AI Product Manager. I designed the fleet and its division of labor, wrote every agent's specification and quality gates, set product direction — Medicare coverage, costs, and who pays lead — and make the growth calls from Search Console analytics. Prototyped with Codex; I do not write production code.

CASE STUDY — EVALUATION PROOF

Case study — evaluation proof

The false-stability defect

What testing caught: a measurement whose date was changed made a declining patient look stable — exactly the failure this product exists to prevent, manufactured by the product itself. The product response was not a patch on the one screen: all five metric write paths were routed through a single validation gate, so a re-dated or out-of-contract value cannot enter the record by any route. Find it in testing, fix the root cause, make the defect class impossible.

Privacy, verified at the database level

A scripted cross-account audit before launch caught leftover permissive access rules that could have let one account reach another account's profiles and reports — invisible from the UI, real in the data layer. The verified end state: a fresh account sees exactly the shared read-only demos plus its own profiles; attempts to edit a demo or read another owner's record affect zero rows. The sharing rule — demos shared, account work private — is proven, not intended.

The gates bite, and the record is deterministic

- Referral gates: structural contract failures force HOLD FOR DATA CORRECTION; review program (TPE / SMRC / UPIC) is selectable, default TPE.
- Determinism: every consumer reads the latest completed run per lane from the same source, regardless of a run's result — a run's outcome never changes whether it counts as source data.
- PlanChart sequencing is code-enforced: at admission it is fed strictly by the Referral Audit's output; recertification uses data from all sources.
- Clinical Audit independence: its rating is computed from its own scan plus record data only — never from other systems' conclusions.
- Waste Analysis harness: the corpus builder was verified pre-deploy against production D1 as the demo account — 3 profiles, 23 lane outputs, 93 visits, Census Risk reproduced exactly.

From lane outputs to an operating plan

Each lane judges one slice — one referral, one plan, one visit ledger. No single lane can see waste that only appears across patients, clinicians, and weeks. Waste Analysis is deliberately account-level, with no patient picker: it correlates every lane's output, enumerates waste across the seven types, names clinicians from the visit record itself, and turns findings into a corrective plan where every step carries a forecast, with probability bars on the page and in the printed report.

METHOD

How I build — and the thesis the pair proves

01 — Workflow first

Every product starts from how the user actually works and who makes which decision — before any feature.

02 — Prototype, then specifications

Both systems started as prototypes I built with Codex. I write requirements, boundaries, and acceptance criteria — what the system must do and must never do — and direct the build from there.

03 — Evaluation before release

Validators, adversarial cross-checks, and a final disposition stand between any output and its reader. When evaluation catches a defect, the fix is a root-cause gate, not a patch.

04 — Proof from the outside

Claims get verified from the public side — live pages counted, cross-account access tested, audits run against published facts.

	Clinician AI Assist	Care and Living Guide
Who starts the work	A clinician submits a case. Nothing runs unattended.	Agents run on schedule, under specifications I set.
Who decides	The clinician — every output, every time.	Set in advance: specs, gates, and direction. The fleet executes.
Autonomy granted	None in the clinical lanes. Operations agents touch data and structure only.	Full operational autonomy, inside verification gates.
The scoreboard	Clinician validation — a testing round is running now.	Search traffic today; revenue is a later proof milestone.

> **Same builder. Same discipline. Opposite autonomy — chosen per system, on purpose.** Here autonomy is a design decision, made twice, in opposite directions, for stated reasons — and enforced in gates and guards, not in intentions. The logic is the durable asset; the models and tools are replaceable.

The builder

Kris Pueblos, RN, BSN — 13+ years as an RN inside CMS/Medicare compliance and hospice documentation audit. Clinical Manager at Amedisys (Stafford, TX): census of about 40, team of 4 RNs, 1 MSW, 1 chaplain, 3 CNAs. Previously Director of Nursing at Genus Patris Hospice (Houston): up to 75 patients, team of 16. Daily work: certification and recertification, face-to-face, IDG, HOPE, plan-of-care alignment, ADR preparation.

See both systems live: clinicianaiassist.com · careandlivingguide.com · clinicianaiassist.com/portfolio · linkedin.com/in/kris-pueblos

Independent products, built solo — not employer products. All patient cases are synthetic. Clinician AI Assist is designed around FDA non-device CDS principles and is not FDA cleared. Care and Living Guide is not monetized. Current as of November 10, 2026.