

AMANA

The bounded AI co-worker for emergency care

An evidence-grounded agent that coordinates referrals, patient flow and the next right action

Stays with the patient journey—without pretending to be the clinician.

Synthetic demonstration. Not approved for clinical use. Clinical decisions remain clinician-owned.

The missing co-worker

Emergency referrals require someone to hold the whole journey together.

- Clinical state changes while teams wait.
- Referral status is fragmented across calls, messages, documents and systems.
- Missing information blocks progress without clear ownership.
- Transport can continue on assumptions that are no longer current.
- Every person sees a fragment; nobody can continuously hold the complete operational picture.

The failure is not only a missed fact. It is a missed **next action**.

Meet AMANA

An AI co-worker that observes, organizes, prepares and follows through

- **Understands** messy referral documents and communications as evidence-linked candidates.
- **Maintains context** across the patient, clinical state, referral, transport and responsible team.
- **Notifies** deterioration, stale information, delays, blockers and changed assumptions.
- **Prepares** referral updates, handoffs, documents and owned tasks within policy.
- **Follows through** by checking outcomes, keeping unresolved work visible and escalating safely.

One co-worker stays with the journey from referral intent to destination handover.

One patient, one continuous co-worker

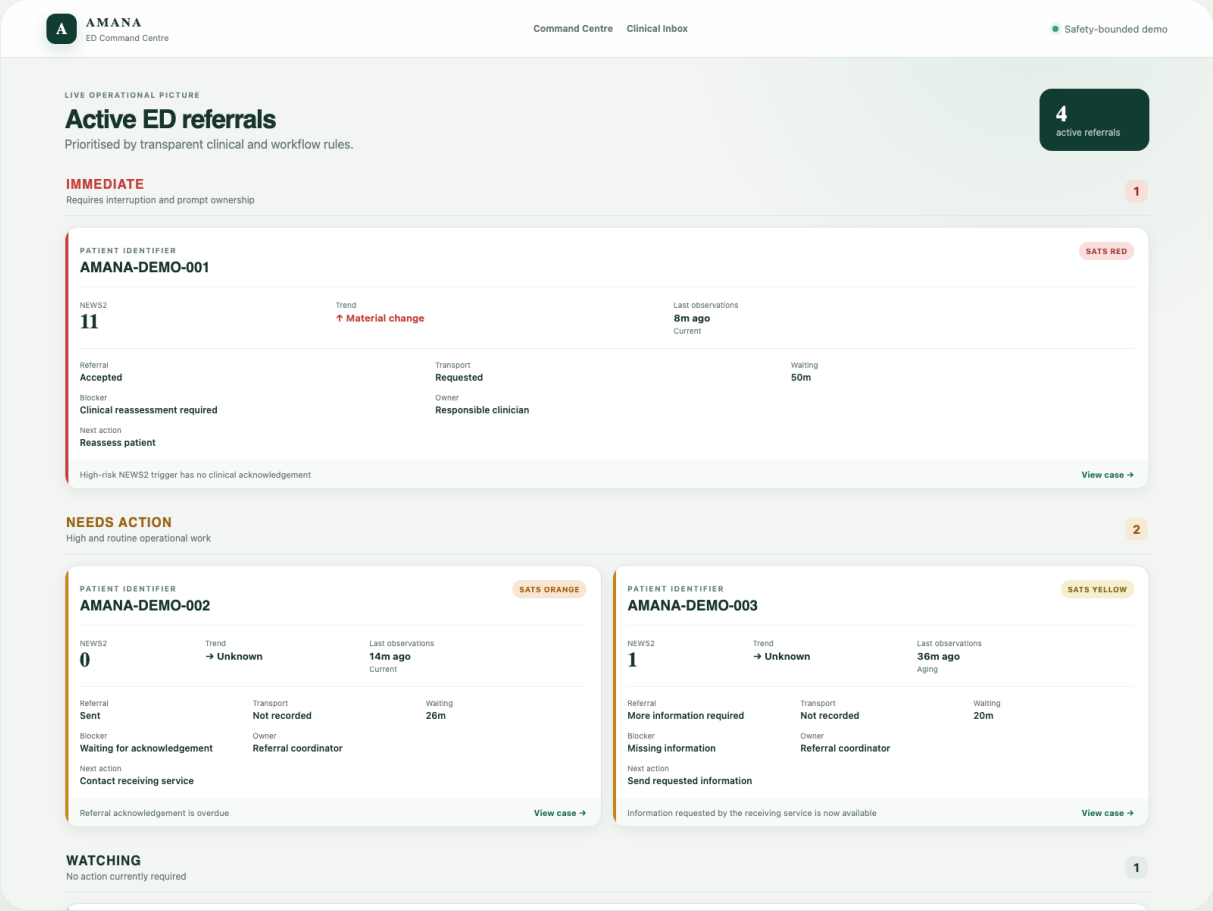
A synthetic referral journey—and AMANA's job at each stage

- **Referral initiated** · 18:02 — a clinician decides to refer. AMANA assembles context; observations produce deterministic NEWS2 **4**.
- **Referral prepared** · 18:16–18:24 — it extracts referral candidates, a clinician confirms, the referral is sent.
- **Missing information tracked** · 18:31–18:51 — it follows the requested CT and creatinine, surfacing each as it lands.
- **Handover evidenced** · 19:01–19:05 — evidenced acceptance, clinician-owned readiness and transport.
- **Deterioration detected** · 19:27 — observations move NEWS2 **4** → **11** → **IMMEDIATE**.
- **Teams re-synchronised** · 19:28 — it requests readiness reconfirmation and reports: receiving update **NOT RECORDED**.

AMANA does not simply record these events. It follows the patient's clinical state and the referral process, identifies changes and missing work, and keeps the required update visible to its human owner.

The shared workspace

Humans see what the co-worker sees, why it matters, and who owns the next action



The AMANA agent loop

Understand → verify → plan → act within authority → check → keep watch

Agent capability	Trusted mechanism
Understand fragmented inputs	Multimodal AI with field-level evidence and uncertainty
Build working context	Immutable events and current-state projections
Judge what needs attention	Deterministic clinical and operational rules
Choose permitted work	Action catalogue, preconditions and governed autonomy
Prepare or execute	Typed tools, human approval and bounded effectors
Verify the result	Postconditions, idempotency, audit and replay

AMANA does not merely answer a prompt. It carries responsibility-aware work across time.

Why advanced AI

The safety core is deterministic. The real world around it is not.

Multimodal understanding

- Interpret text, PDFs and images from fragmented referral workflows.
- Extract candidate intent, requested information, uncertainty and conflicts.
- Ground every candidate in a source excerpt, locator, page and confidence.

Contextual co-work

- Explain only the verified operational picture assembled by AMANA.
- Draft the next communication or handoff from approved evidence.
- Adapt to document and language variation while preserving human confirmation.

Bounded agency by design

Useful enough to act. Constrained enough to trust.

**Model proposes → evidence grounds → policy authorizes → human confirms where required
→ typed tool executes → postcondition verifies → audit records**

- No model-controlled clinical score or state transition.
- No autonomous diagnosis, prescribing, acceptance, triage or transfer-readiness decision.
- No unsupported claim upgraded into patient or workflow truth.
- Every external action has permission, preconditions, idempotency and a verifiable outcome.
- A kill switch and autonomy ceiling stop effectors while monitoring continues.

Intelligence is generative. Authority is explicit.

From one hospital to a care network

A governed deployment becomes more valuable as the referral journey crosses institutions

User and buyer

- Users: ED clinicians, referral coordinators and operational teams.
- Buyer: hospital or health-system clinical and operational leadership.

Initial deployment unit

- Read-only validation → governed shadow pilot → institutional software and support agreement.
- Success is measured as workflow representation accuracy and interruption burden—not clinical impact.

Expansion path

- Add one receiving service, then transport and handover, under explicit sharing agreements.
- Repeat the deployment playbook institution by institution.

Earn the network one governed institutional connection at a time.

Execution

Founder-led execution—not a concept video

- **Ronny Tuurime Polle** — founder, medical doctor and machine-learning engineer; MBChB, University for Development Studies.
- Read-only Command Centre demonstrating complete synthetic referral journeys.
- Deterministic rules track clinical risk, deterioration, delays and next actions.
- Clinical Inbox extracts evidence-linked information for clinician confirmation.
- Read-only hospital-integration foundation with source, conflict and quarantine controls tested synthetically.
- **542 automated backend tests**, plus strict typing, lint and frontend checks.

Why Google

Build the first trustworthy AI co-worker for fragmented African emergency systems

With Google, AMANA will:

- evaluate frontier Gemini multimodal models against locally representative referral artifacts;
- improve evidence grounding, uncertainty detection, language coverage, latency and cost;
- develop rigorous agent evaluations for unsupported claims, unsafe actions and failed outcomes;
- turn one governed institutional deployment into a repeatable multi-institution playbook;
- combine technical mentorship with go-to-market support for health-system adoption across Africa.

The opportunity is not a smarter chatbot. It is an accountable AI worker embedded in essential care.

Vision

Every emergency team gains a tireless, bounded AI co-worker

One that understands the fragments.

Holds the whole patient journey in context.

Brings the right work to the right human at the right time.

AMANA coordinates. Evidence grounds. Clinicians decide.

amana-referrals.pages.dev