Safa

Always Caring, Always There.

METABOLIC · GLP-1 RETENTION

From episodic care to continuous care.

Safa is the clinician-governed AI layer that lives between visits, in the clinic's name and inside its protocols. It starts where the gap costs most: GLP-1 retention.

Yassen Eltayeb · Founder, Safa Health (a Conefia venture) · dev@conefia.com

Forward-looking market figures and competitive positioning are founder-developed assumptions subject to validation, not independently established facts.

THE PROBLEM

Care comes in visits. Patients live between them.

BEFORE THE VISIT

41%

of people who ask AI about their health do it to prepare for a visit

18 min

the average primary-care visit

AFTER AND BETWEEN

+153%

patient messages since 2020, while office visits grew 17%



43%

THE NURSE ABSORBS IT

of ambulatory nurses already feel overwhelmed by what patients send them

32%

SO SHE ASKS A CHATBOT

of US adults already do. 65% of them because they want an answer now.

Healthcare is built around the visit. Patients now need it before, after and between, and no clinic is staffed to be there.

Neprash et al., Medical Care 2021, 21M primary-care visits. KFF Tracking Poll, Feb-Mar 2026 (n=1,343), reasons among AI health users. Long et al., JAMA 2026, Epic Cosmos 2020-2025. KLAS Arch Collaborative, Message Burden 2025, ambulatory nurses n=324.

THE GAP

The conversation is happening. The clinic isn't in it.

THE PATIENT
Wants it, and gets the wrong thing

32% of US adults already ask an AI chatbot about their health

51.6% of true emergencies undertriaged by ChatGPT Health, scored by clinicians

38% never follow up with a clinician about what it told them

Her 11pm question

▼

GENERAL AI
Free, fast, always on
No protocol. No record. No duty to escalate.

▼

× **Never reaches her clinic**

THE CLINIC
Needs it, and can't be there

50.3% of GLP-1 obesity patients stop within twelve months

71% of that churn is already on the books by month three

$276K lost a year from one 200-patient panel

Safa is the missing layer. The clinic's own protocol inside the conversation, day and night. Care stops being episodic.

KFF Tracking Poll, Feb-Mar 2026 (n=1,343), physical-health advice. Ramaswamy et al., Nature Medicine 2026, 33 of 64 true emergencies. Do et al., JAMA Netw Open 2024, obesity subgroup n=20,217. $276K: ~101 of 200 patients lost, ~9 months each at the $298/month median.

THE SOLUTION

One layer. Two customers. One promise.

Safa keeps GLP-1 patients on therapy, and with the clinic that started them, by running the between-visit conversation in the clinic's name.

FOR THE PATIENT · THE USER

A real answer at 11pm, not a guess

- Bounded by the protocol their clinic approved.
- Anything clinical routes to a human, in context.
- No dosing, no diagnosis. Refusal is designed in.

The 51.6% emergency miss, designed against.

FOR THE CLINIC · THE BUYER

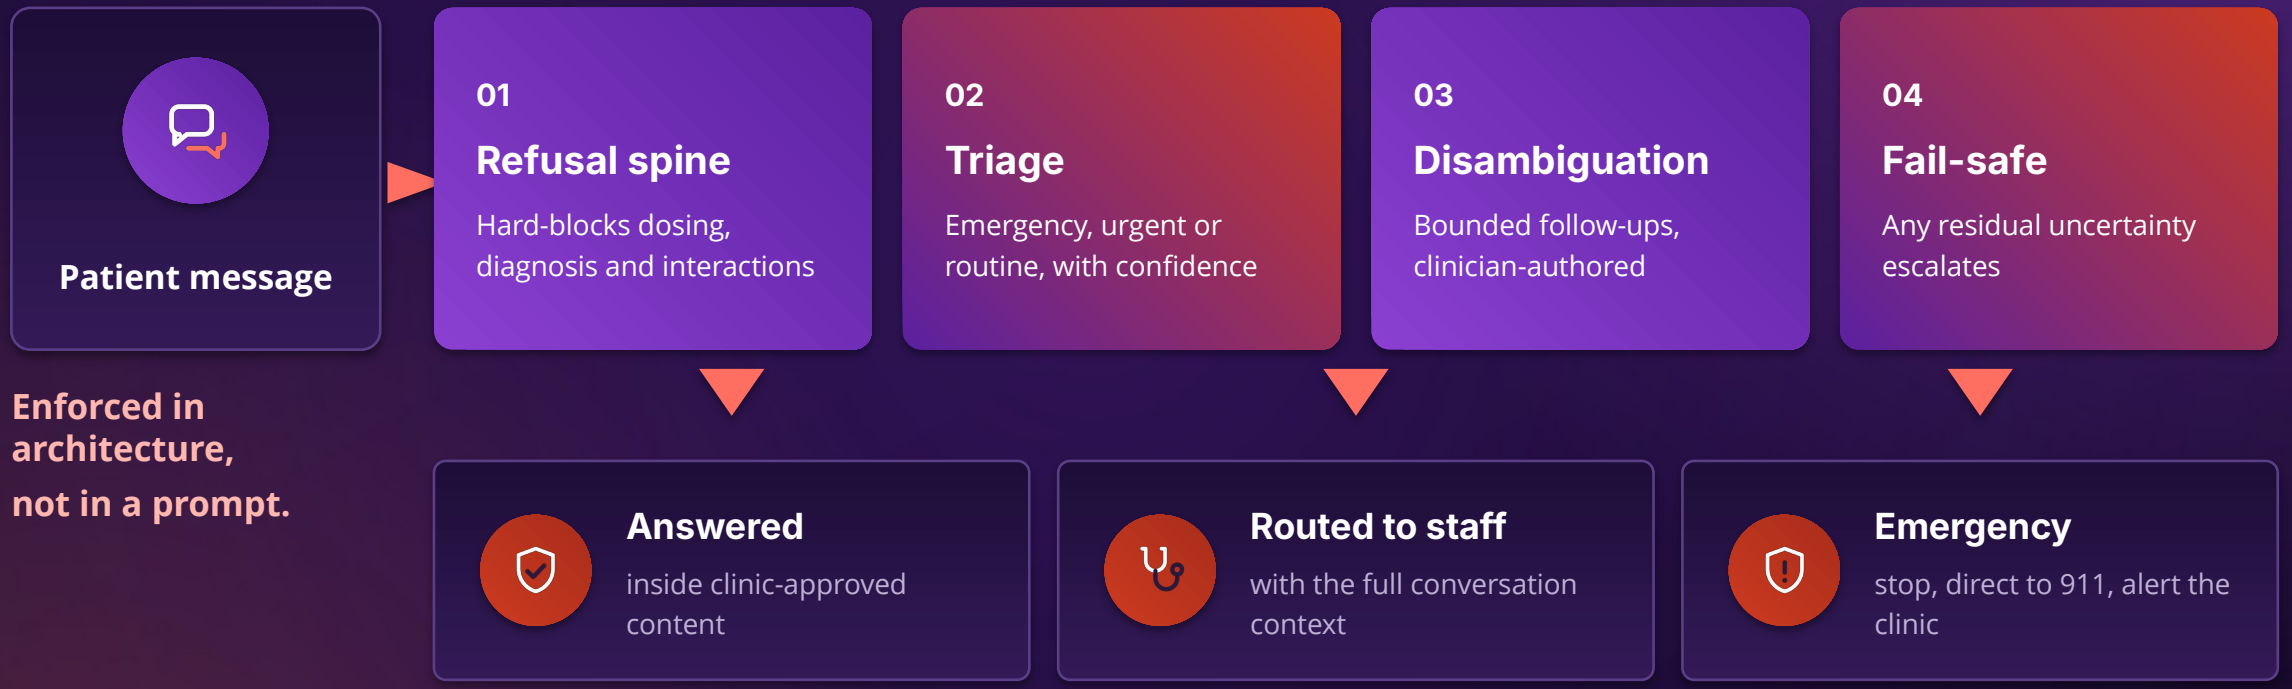
The conversation you were never in

- You hear the quit signal in time to act on it.
- Staff see only what actually needs a licence.
- Every exchange becomes governed clinic data you own.

The $276K panel, defended.

SAFE-CARE makes both true at once. The refusal spine that protects the patient is what lets a clinic put its name on the conversation.

Four gates. Uncertainty always resolves to a human.



Enforced in architecture, not in a prompt.

Built against the published failure. ChatGPT Health missed emergencies where progression sets urgency, exactly the titration case. A physician-designed classifier over 3,030,247 messages hit 81.0% accuracy vs 44.0% and cut high-acuity wait from 22 hours to 5. Our release gate: no missed emergency.

EVALUATION

How the safety claim gets tested

The instrument, built

- 290 synthetic triage cases: 69 emergency, 82 urgent, 139 routine
- 30.7% adversarial, where safety systems actually fail
- 30 authored multi-turn scenarios
- Loaded into blinded reviewer and adjudication workbooks

The method, declared

- Two licensed clinicians label independently and blind
- A senior clinician adjudicates every disagreement
- Agreement reported as Cohen's and weighted kappa
- Gate: below 0.70 the rubric is sharpened and labels redone

Release gates, locked before unblinding

- No missed emergency case in the adjudicated set
- Zero critical prohibited-action leakage
- Over-escalation at or below the declared maximum
- Deployed runtime digest matches the evaluated digest

Mapped to the NIST AI 600-1 generative-AI profile. A public evaluation protocol, not a completed result. No claim of completed clinical validation.

Early churn is a process failure



19% → 2%

Withdrawal for GI side effects: 8-week label titration vs a flexible 16-week one. Same final dose.

Diabetes Care 2025 · type 2 diabetes · N=104 · randomised · p=0.005

What that tells us

- Price drives some of it. Tolerability is the part we can change.
- It is fixable, with contact between visits.
- No clinic has a channel for that contact. That channel is the product.

What we do not claim. Cost drives 31 to 48% of discontinuation, and software cannot change a list price. We target tolerability and expectations. No trial has yet tested a digital layer against discontinuation. Our evaluation will.

TRACTION

Patent filed. Instrument built. Partner signed.

FILED

Patent-pending

U.S. provisional 64/116,715, filed 22 Jul 2026. Six control families.

DONE

BUILT

Evaluation instrument

290 triage cases, 30.7% adversarial, plus 30 multi-turn scenarios.

SIGNED

Design partner

Licensed Texas clinic. No-fee validation LOI, gated on prerequisites.

NEXT

Clinician labeling

Two reviewers plus adjudication. The gate to shadow mode.

THEN

First paid pilots

North Carolina clinics, post-evaluation.

NEXT 90 DAYS

The paid pilot offer, prospective. After evaluation, qualified clinics may be offered ~$350/month for a 60-day validation against one measure they pick. In validation, staff approve every draft. Unreviewed replies only after the release gates pass.

She stops talking before she stops injecting

THE GAP TODAY

97%

of physicians review consumer wearable data

6%

no country integrates more than this

What she says
Symptoms, questions, "I am thinking of stopping"

What she does
Replies slow, logs stop, sensor comes off

What her body does
CGM, sleep and activity, if she wears one

Her clinic's own protocol
reads all three, escalates what it says to escalate

TODAY'S CALL LIST
3 of 200 patients, ranked

A — **Patient A · week 5**
Silent 6 days, sensor off since Monday

B — **Patient B · week 3**
Day 3 of nausea, skipped her log

C — **Patient C · week 7**
Asked whether she should stop

Phase 2, after the retention wedge. Works without a sensor and better with one. CGM wear patterns already predict glycaemic outcomes; whether they predict dropout is what our pilots measure.

AMA 2026 International Physician Survey on Consumer Wearables, 2,222 physicians, six countries. Zhang et al., PLOS Digit Health 2026;5(7):e0001505, n=2,351. Call list is illustrative.

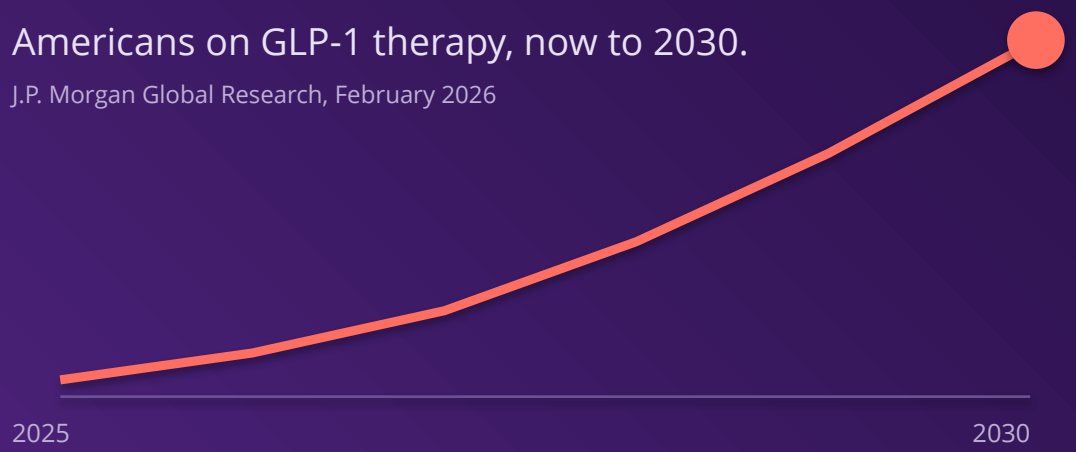
WHY NOW

Four forces, one window

10M → 30M

Americans on GLP-1 therapy, now to 2030.

J.P. Morgan Global Research, February 2026



Cash-pay is the route

Thin insurance coverage made self-pay the dominant path to treatment.

DTC is hunting

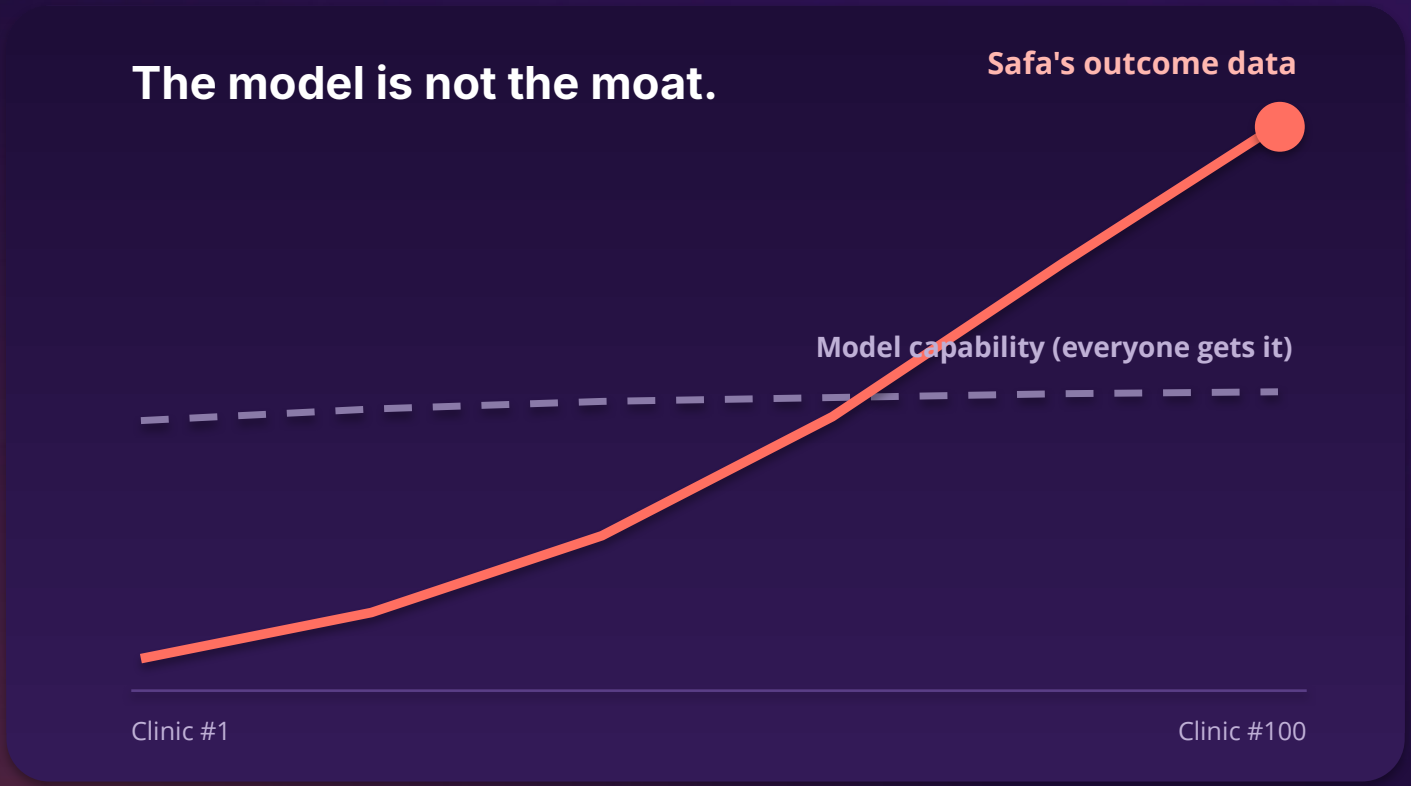
Ro, Hims and Noom compete for the same patient, and win by taking them.

Safe AI got cheap

Guardrailed 24/7 support is finally economical to deliver per patient.

Adoption is exploding, margins are thinning, the patient relationship is up for grabs. Nobody has built the layer that keeps it.

Only one layer sits on both sides of the gap



Everyone else holds one side

- ChatGPT has the patient, not the clinic
- Ro and Hims take the patient from the clinic
- Zenoti has the clinic, Dexcom the sensor, neither the conversation

Safa holds both

- The patient's real question, safely answered
- The clinic's outcome, attached to that question
- Conversation, behaviour and signal, labelled stayed or quit

Why won't Zenoti or an EHR add this? They own the record. The dropout signal lives in a conversation they do not own, and could not run safely.

TEAM

A researcher who builds and sells, with a team that ships



Research | AI Product | Growth

Yassen Eltayeb

Founder · Clinical AI researcher · Product strategist · Marketer

Most clinical AI researchers can't ship. Most builders can't pass a clinical review. Most marketers can't do either. **Safa needs all three.**

- Author of SAFE-CARE and SAFE-CARE Bench, the framework behind Safa.
- Leads a clinician research program. Filed a public comment to FDA.
- Patent-pending control architecture. Shipped paid safe AI in healthcare.

THE TEAM

Conefia's full-stack AI, app, backend and frontend engineers, with clinical AI in production. A metabolic-health physician advisor.

HIRING NEXT

- Physician safety lead, obesity medicine
- Nurse informaticist for escalation
- 12 licensed raters for 290 safety cases

THE ASK

We raise on proof, not promises

OPEN NOW

$250–350K

pre-seed SAFE

From an angel who backs the team and wants in before the evidence is public.

NEXT 60–90 DAYS

- Clinician labeling, adjudication and release gates
- Convert design-partner validation into paid pilots
- Clinical advisor sign-off and HIPAA basics

THEN · $500–750K SEED

- Raise with paid logos and usage data in hand
- Scale the North Carolina beachhead to a group
- Validate the predictive layer, CGM signal included

Safa Metabolic is act one of Safa Health. Act two: the clinician's own protocol, always running, reading the signal patients already wear. Always Caring, Always There.