01 · THE PROBLEM

AI is already acting in care. Governance is not.

AI systems today:

- draft notes
- suggest actions
- coordinate care

But in most deployments:

- authority is unclear
- consent is fragmented
- decisions are not reconstructable

Healthcare AI has entered high-risk territory without runtime governance.

02 · WHY THIS MATTERS

This is not theoretical.
It is now regulated.

Under the EU AI Act (High-Risk Systems):

Article 12
Logging & traceability required.

Article 14
Human oversight required.

GDPR
Consent must be explicit, scoped, revocable.

Reality. Most AI deployments today:

- log outputs, not decisions
- assume consent, don't enforce it
- rely on policy, not runtime control

Regulatory and liability exposure starts now.

03 · THE GAP

Not another AI model. A control layer.

TODAY

- AI acts inside workflows
- Governance lives in documents
- Consent is static
- Audit is incomplete

MISSING

- Real-time authority enforcement
- Consent gating at action-time
- Immutable decision logging
- Post-incident reconstruction

No system today enforces governance at runtime.

04 · THE CATEGORY

Introducing

Runtime Governance Infrastructure

Implementation: Safety OS™

A control layer above AI systems that ensures:

- Every action is **authorised**
- Every action is **consent-gated**
- Every decision is **logged & reconstructable**

No AI action without logged authority,
consent, and traceability.

05 · HOW IT WORKS

Three primitives. Enforced at runtime.

01

Authority

- Who is allowed to act?
- What role: clinician, caregiver, system?

02

Consent

- What scope is granted?
- Is it current, specific, revocable?

03

Traceability (Flight Recorder)

- Every decision logged, immutable.
- Reconstructable post-incident.

All AI outputs pass through this layer before execution.

06 · ARCHITECTURE

A control layer on top of existing systems.

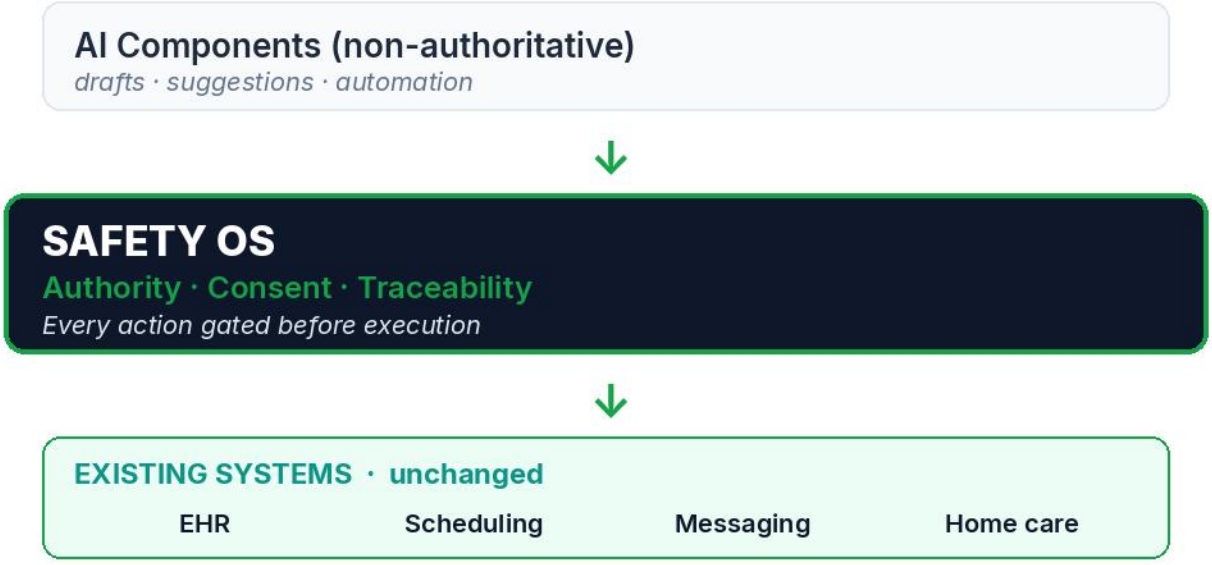

Deploys without disrupting existing clinical systems.

07 · SAFE CONTINUUM OF CARE

Tightly-bounded Agentic Orchestration from hospital to home.

HOSPITAL
Phase II · HCP-as-Pilot™
Clinician orchestrates AI components for notes, orders, follow-ups.

DISCHARGE
Authority + consent travel with the patient
Same governance envelope. No handover gap. No audit break.

HOME
Phase I · Care Delivery Teams
Nurses, caregivers, and family act inside the same bounded envelope.

One governance layer. Healthcare Professional in command.

07 · DEPLOYMENT

Deployed in real care. Expanding to clinical AI.

PHASE I

LIVE NOW

Care Delivery Teams

- Home care + caregiver coordination
- Consent enforcement
- Audit logging active

PHASE II

NEXT

HCP-as-Pilot™

- Clinician workflows
- Orders, follow-ups, coordination
- Full oversight + escalation control

Same governance layer. Different environments.

09 · WHAT THIS REMOVES

What this removes from your balance sheet:

- Liability from unauthorised AI actions
- Regulatory exposure (EU AI Act / GDPR)
- Audit failure risk
- Post-incident uncertainty

Turns AI from "uncontrolled risk"
into "governed system."

09 · THE ANALOGY

Aviation solved this. Healthcare AI has not.

Aviation engineered:

- **Pilot-in-command** - *clear authority*
- **Air traffic control** - *coordination*
- **Flight recorder** - *full reconstruction*
- **Checklists** - *state awareness*
- **CRM** - *human + system teamwork*

Each was created after failure.

Healthcare AI does not need its Tenerife moment.
Governance can be engineered now.

10 · WORK WITH US

We are working with partners deploying AI in high-risk care.

We are looking for:

- Hospitals evaluating clinical AI
- Sponsors interested in governance infrastructure
- Pilot partners under EU AI Act readiness

WHAT WE DO

1. Map your AI workflows
2. Identify governance gaps
3. Deploy the control layer

Before scale. Before incident. Before regulation forces it.

DM Andy Squire · patientcentriccare.ai