HCP-as-Pilot™

Healthcare Professionals (i.e., Doctors, Nurses, medical team)

Runtime Governance Infrastructure for AI in Healthcare

PatientCentricCare.AI
Version 3.5 · June 2026
HCP-as-Pilot

Andy Squire

Observable **Enforceable** **Replayable**

*Together, these make AI-mediated care **accountable.***

The Problem

A patient is discharged.

An AI schedules follow-up reminders.
A caregiver receives alerts.
A conversational agent checks wellbeing.
A nurse reviews exceptions.

Who authorised each action?
Who approved escalation?
Who verified consent?
Who currently holds authority?
Who is accountable when something goes wrong?

Healthcare increasingly relies on AI-mediated interactions occurring **outside traditional clinical environments.** Patients interact with AI at home, between appointments, and across discharge boundaries - environments where hospitals, caregivers, and regulators have limited visibility into how these systems make decisions, escalate concerns, obtain consent, or transfer authority.

The result is a structural gap. AI liability is unclear. Clinical override is unprotected. Consent is assumed, not governed. Escalation is improvised, not predefined. When something goes wrong, audit trails record **what the model generated** - not **why execution was permitted.**

Existing governance mechanisms - policies, training, documentation, and retrospective audit - remain valuable. But they operate **outside the moment of execution.** Patient-impacting decisions occur in real time. **Governance must too.**

Static Governance	Runtime Governance
Policies · Procedures · Training · Retrospective Audit	*Authority · Consent · Risk · Escalation · Evidence*
Cannot answer:	**Answers before execution:**
May the AI act right now?	Authority verified.
Is consent valid at this moment?	Consent validated.
Who currently holds authority?	Risk tier assigned.
Is escalation required now?	Escalation governed.

Generative AI is not governed by what it knows.

It is governed by *what it is permitted to do.*

The Architecture: Runtime Governance Infrastructure

Runtime Governance Infrastructure (RGI) is a domain-agnostic governance architecture operating **at the point of execution.** It determines whether an action may proceed, whether consent exists, whether human confirmation is required, whether escalation must occur, and what evidence must be recorded. RGI represents a new infrastructure category for governing AI-mediated actions in high-consequence environments.

HCP-as-Pilot™ is the healthcare implementation of RGI - built on the principle long proven in aviation: autopilot may execute, but **authority remains with the pilot.** AI may assist, recommend, coordinate, and automate bounded activities. Authority over patient-impacting **actions** remains human.

Governance Before Action

Traditional AI Execution	Runtime Governance Infrastructure
AI Generates Output	**AI Proposes Action**
↓	↓
Action Occurs	**Governance Executes**
↓	↓ (Authority · Consent · Risk · Escalation · Evidence)
Audit Later	**Patient-Safe Action Occurs**
	↓
	Evidence Recorded

Every action passes through governance before patient impact occurs.

Runtime Governance Infrastructure is to AI what Air Traffic Control is to Aviation.

ATC does not fly the aircraft. It governs the safe execution of flight. Aircraft may differ. Pilots may differ. Airlines may differ. Air Traffic Control provides the common governance layer that keeps the system safe. RGI does not replace AI systems. It governs the safe execution of AI-mediated actions.

Generative AI produces novel outputs in real time. Static policies cannot safely govern dynamic AI behaviour at runtime. RGI exists because **governance must operate at the point of execution** - not only before deployment or after audit.

RGI represents a new infrastructure category for governing AI-mediated actions in high-consequence environments. Healthcare is the first implementation. The underlying architecture is domain-agnostic.

AI GOVERNANCE CONTROL LAYER

Tightly-bounded Agentic Orchestration enables a Safe Continuum of Care.

Required when AI influences clinical decisions under the EU AI Act. Safety OS™ bounds every AI component — from clinic, through discharge, into the home — without replacing existing systems.

Safety OS™ does not replace AI systems, EHRs, or workflows. It governs the execution path between AI output and patient-impacting action.

Why Now: The Governance Window

Regulatory inevitability creates a narrow window for first-mover advantage.

2025 Conversational AI enters homes at scale - without governance

2026 **EU AI Act enforcement begins - runtime evidence supporting conformity assessment becomes a strategic differentiator** ◂ WE ARE HERE

2027 First enforcement precedents - ungoverned systems face sanctions

2028 Payer requirements emerge - governance evidence becomes table stakes

2029+ Clinical expansion - only systems with proven governance history qualify

Regulatory inevitability creates a narrow window for first-mover advantage. Organisations that build runtime governance today will possess the evidence, operational history, and institutional trust required when governance becomes mandatory.

Safety OS™ Components

Safety OS™ implements RGI in software, operating in milliseconds before execution - governing the pathway between AI output and patient impact without replacing existing EHRs, AI models, or clinical workflows.

Decision Inventory	The foundation of the entire framework. Every AI-involved decision must be explicitly inventoried before autonomy is allocated. Each decision must have a single named human owner. If a decision is not inventoried, it cannot be automated. Autonomy is assigned per decision - not per system. Without a decision inventory, governance cannot be applied consistently or defended under audit.
Human Authority Envelope (HAE)	Defines the boundaries within which AI may operate and the humans who retain authority over patient-impacting actions. The envelope adapts dynamically to consent state, clinical context, capacity, legal designation, and organisational policy. AI never operates outside authorised boundaries. The envelope defines not only what AI may do, but who retains authority when AI cannot.
Risk-Adaptive Stratification (RAS)	Governance adapts to risk in real time. **Tier A:** routine daily check-in - AI executes within predefined bounds. **Tier B:** repeated refusal or safety signal - human confirmation required. **Tier C:** clinical intervention - human authority required. Tier assignment is deterministic, centralised in the governance evaluator, and recorded on every decision.
Consent Engine	Real-time consent validation conditioning every AI-mediated action. Consent ambiguity defaults conservatively. No consent. No action. Consent state is tracked continuously across sessions.
Escalation Engine	Predefined escalation pathways based on authority, urgency, and context. Escalation is governed, not improvised. Routing distinguishes caregiver vs. clinician escalation based on urgency classification and availability.
Flight Recorder	Most AI audit systems record outputs. Flight Recorder records governance - not merely what happened, but why it was permitted to happen. Captures authority state, consent state, risk tier, policy evaluation, escalation pathway, and enforcement outcome. Every governance decision can be reconstructed from recorded evidence, without retaining private conversation content: accountability without surveillance. Audit log is append-only, enforced at the database access-control layer. Flight Recorder captures governance state, not conversation content.
Runtime Evidence Ledger	Structured storage of governance artefacts supporting internal review, quality assurance, and regulatory assessment processes.

Human Authority Envelope

Safety OS enforces a clear authority boundary - every action is classified into one of 3 risk stratified and adaptive categories. Defines **who** decides, **when** AI can act, and **how** every **action is accountable**.

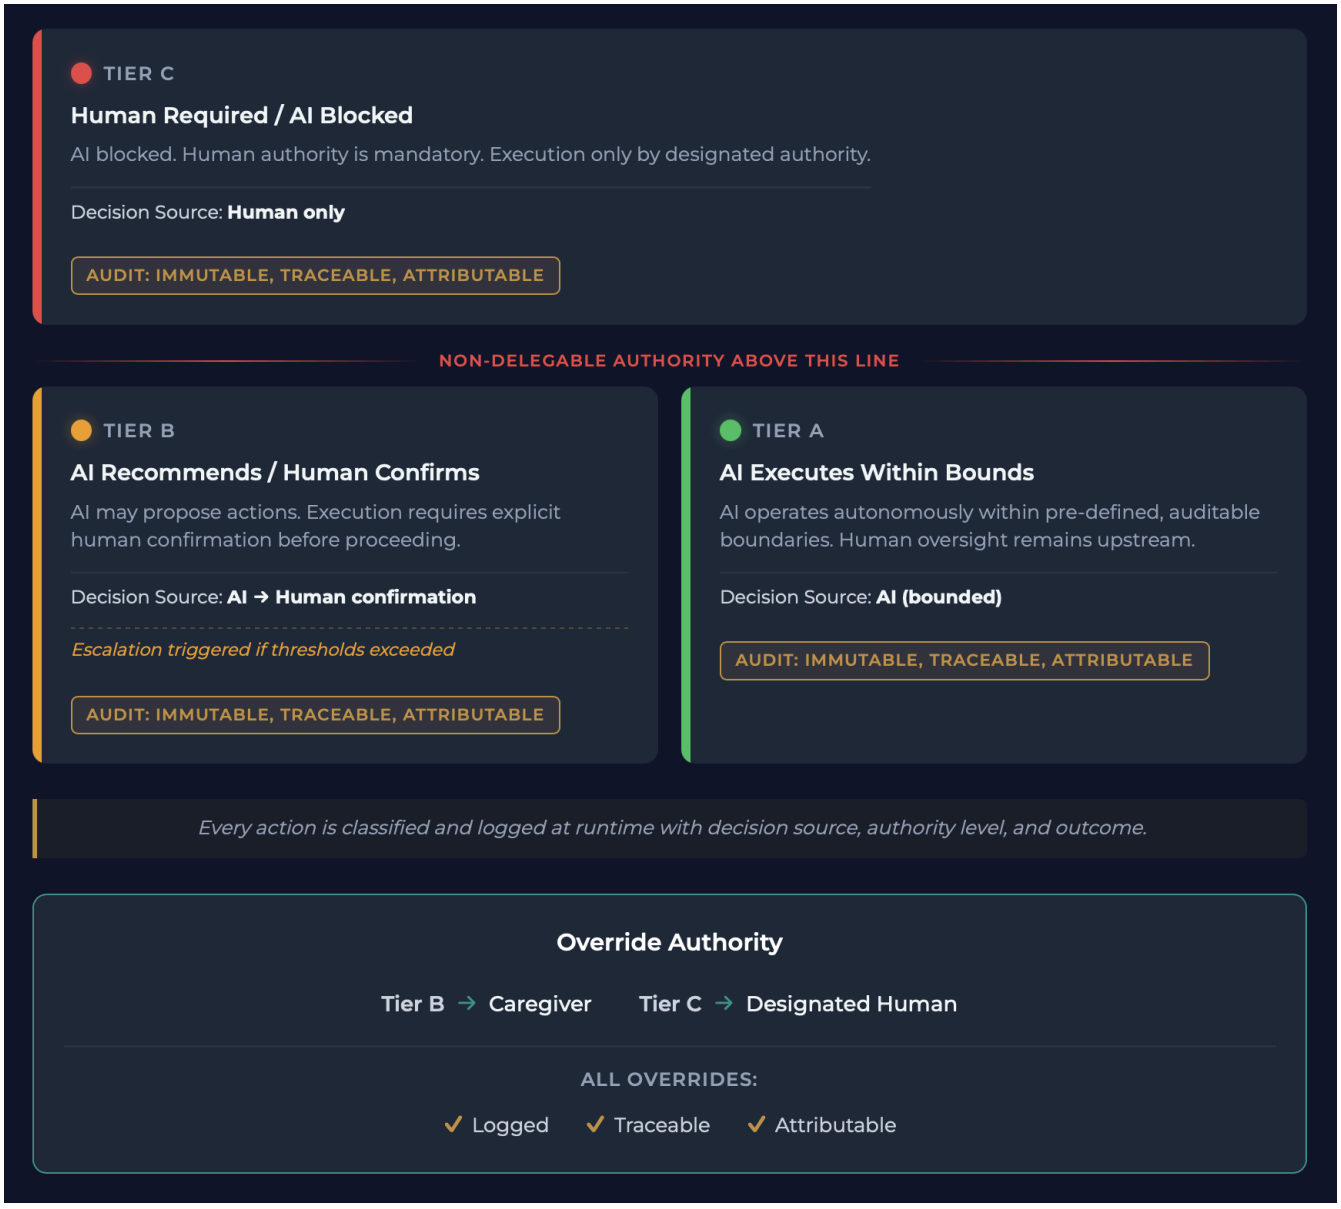

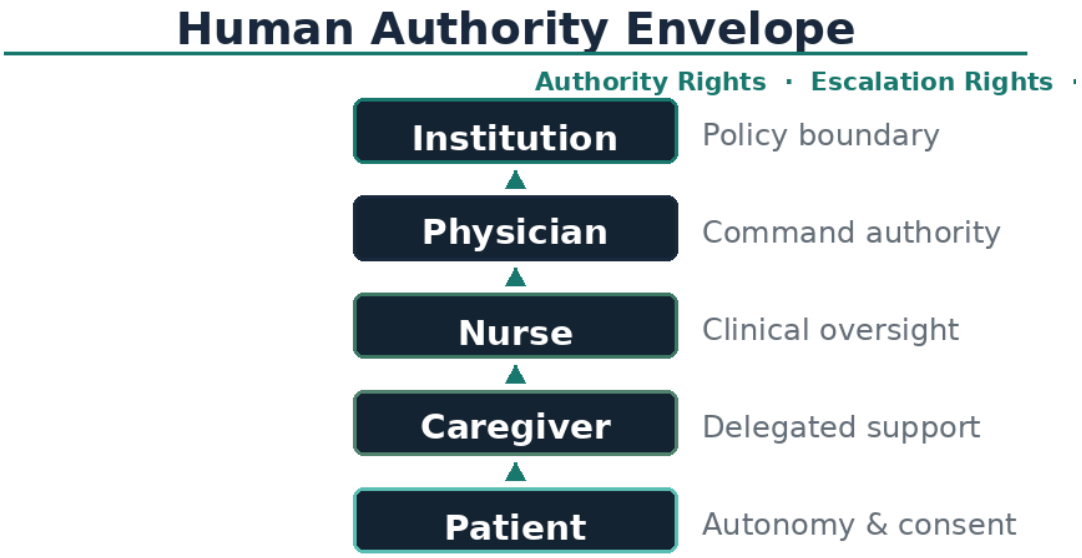

Authority within the HAE flows from Patient through Caregiver, Nurse, and Physician to Institution - each level carrying defined authority, escalation, & override rights. Authority is dynamic and adapts to consent state, capacity, & clinical context.

Regulatory Alignment

Safety OS™ does not perform conformity assessment. It generates the runtime evidence required to *support* conformity assessment. The objective is not compliance by assertion. It is **evidence by design.**

Framework	Reference	RGI Alignment
EU AI Act	Art. 12, 14, 16	Flight Recorder generates runtime evidence for human oversight obligations, record-keeping, and post-market monitoring. Designed to support obligations associated with high-risk AI systems.
Singapore AIHGle 2.0	MOH/HSA 2026	Independent post-publication convergence: mandates human oversight for all clinical AI and identifies home-care AI as requiring additional governance layers - closely aligns with the governance gap addressed by RGI Phase I.
UK MHRA	SaMD Roadmap	Phased governance architecture aligns with the MHRA's progressive AI oversight framework for software as a medical device.
FDA SaMD / PCCP	2019 / 2023	Phase III governance architecture maps to Predetermined Change Control Plan requirements for adaptive AI medical devices.

Early Deployment - Qualitative Evidence

Phase I (Home Companion - non-medical AI, non-SaMD) is in active early deployment with senior users across Switzerland and the UK. Early qualitative feedback spans patients and caregivers across both user groups, including seniors managing chronic conditions, dementia caregivers, and care home staff. Governance-mediated escalation workflows, consent management, caregiver coordination, and runtime audit generation are operational. Formal operational metrics will be published following completion of instrumentation and review. Early observations indicate particularly strong engagement in multilingual settings, including Swiss German dialect-sensitive interactions where users reported increased familiarity and comfort.

"She made me laugh speaking in Luzern Swiss German Dialect from my childhood." **- Senior & cancer survivor, 86, Basel**	*"I wasn't sure at first, but then I realised Home Companion helps me organise my day and reminds me about important things. I love that it switches languages instantly and can talk about anything. I learn every day."* **- Cancer patient, 58, Basel**
"I feel less alone. It's not a nurse, it's not a doctor - it's just there, and it listens when I want to talk." **- Senior, 75, Liverpool, UK**	*"This is great for anyone who needs someone to talk to."* **- Senior & cancer survivor, 70, Leicester**
"I love that I can suggest personalised memories for my mum of their topics of interest to discuss. It makes her smile." **- Caregiver, 62, Allschwil**	*"I care for multiple seniors, and Home Companion keeps the status of each of them organised for me."* **- Senior Care Home Assistant, 42, Basel**
"I said I was lonely and the system encouraged motivational activities around planning my time. It also offered to escalate to my Caregiver." **- Patient, 63, Oxford, UK (Retired Nurse & Therapist, dementia care experienced)**	*"I like the British accent. Home Companion was very polite and helpful for planning my day so that I don't forget anything."* **- Senior, 75, Liverpool, UK**

Actionable finding: Early feedback identified onboarding complexity. Caregiver-senior setup and memory configuration have since been radically simplified in response - an example of governance-informed iterative improvement.

The framework has been informed through qualitative engagement with a cross-disciplinary group spanning patients, family caregivers, registered nurses, physicians, healthcare innovators, and AI governance advisors - reflecting the principle that safe AI governance must be designed **with**, not merely **for**, the people it affects.

The question is not whether AI can participate in care.

It is how AI is *governed* - at the point of use.

Full paper & open standards: patientcentriccare.ai/standards · SSRN Abstract 6399818 · andy@patientcentriccare.ai