



PatientCentricCare.AI
Helping every patient feel known

HCP-AS-PILOT™ · VERSION 4.0 · JULY 2026

Operational AI Governance for AI-enabled healthcare.

AI is beginning to coordinate care between clinical encounters — scheduling follow-ups, monitoring recovery, supporting discharge. The question hospitals now face is not whether AI is capable, but whether every AI-mediated action is **authorised, consented, risk-tiered, escalated and evidenced** — at the point of use.

The thesis. Operational AI Governance is becoming as essential to AI-enabled healthcare as cybersecurity became to digital healthcare. It is the discipline that lets hospitals adopt AI safely — and it exists, ultimately, to protect one thing: *continuity of care*.

BUSINESS CATEGORY	Operational AI Governance ▲ powered by
ARCHITECTURE / IP	Runtime Governance Infrastructure (RGI) ▲ implemented as
SOFTWARE PLATFORM	PCC-SafetyOS™ ▲ deployed as
PRODUCT ENVIRONMENT	The Patient Recognition Platform · Home Companion™ · MyHealthCanvas® · Trusted Health Memory™

Author **Andy Squire**, Founder · PatientCentricCare.AI · SSRN Abstract 6399818

We began with a continuity problem — not a language model.



“We did not start by asking how to build another AI assistant. We started with a different question: why do patients spend 99% of their lives outside healthcare, yet every clinical encounter begins as though their story has been forgotten?”

— Andy Squire, Founder, PatientCentricCare.AI

That gap — between the **continuous life of a patient** and the **episodic memory of the system** — is where care breaks down: at discharge, across handovers, between the hospital and the home, between one multidisciplinary team and the next. Information is lost, questions are re-asked, and the patient carries the burden of remembering.

Operational AI Governance — and the Runtime Governance Infrastructure behind it — **emerged from solving that continuity problem**. Once AI begins acting between encounters, continuity can only be safe if every action is governed: who authorised it, whether consent held, what risk it carried, when it should escalate, and what evidence it left behind.

◆ **Every governance decision ultimately exists to improve continuity of care.**

Helping every patient feel known.

Continuity of care · continuity of knowledge · continuity of each person's health context.
Our hypothesis is simple: **continuity compounds**.

Why AI healthcare projects stall.

Every hospital now has access to the same frontier models. Capability is no longer the differentiator — yet AI projects keep stalling before they touch a care transition. They stall on the same wall.

WHAT TEAMS REACH FOR

- ✗ A better model
- ✗ Better prompts
- ✗ A bigger context window
- ✗ More agents
- ✗ Another pilot

WHAT ACTUALLY BLOCKS ADOPTION

- ✓ Unclear authority
- ✓ Fragmented care workflows
- ✓ Missing accountability
- ✓ No runtime evidence
- ✓ Fragile patient & clinician trust

None of the items on the left solve a single item on the right.

INDEPENDENT VALIDATION — THE 10 / 20 / 70 REALITY OF AI ADOPTION

10%

The model itself

20%

Tech & data

70%

People & how
care works

After enabling AI across dozens of organisations, practitioner Andreas Horn independently reached the conclusion healthcare deployment forced on us: success is roughly 10% the technology, 20% tech and data, and 70% people and process.¹

The agentic shift removed the technical ceiling and exposed the **organisational one**. A capable agent dropped into an ungoverned care workflow does not improve continuity — it accelerates whatever was already broken: unclear authority becomes faster unclear authority; assumed consent becomes automated assumed consent. Operational AI Governance is the missing layer that lets the 70% — people, authority and how care actually moves — finally have infrastructure beneath it.


¹ The 10/20/70 framing is drawn from AI-implementation practitioner Andreas Horn's publicly shared 2026 lessons, cited as independent validation of a direction PatientCentricCare.AI reached through healthcare deployment — not as its intellectual source.

Care is a workflow between people.

Everyone understands a care pathway; few understand a governance diagram. So we start where care actually happens — with the path a real interaction travels, from the patient to the evidence and back. **Governance is the mechanism. Continuity is the outcome.**



The Patient Recognition Platform — four pillars, one continuity



The AI care maturity model.

Every hospital sits somewhere on this ladder. The value is not in reaching the top fastest — it is in **never letting autonomy outrun governance**, so continuity never depends on an unaccountable action.

1

LEVEL

Chatbot

Answers questions. Stateless, no authority, no memory of the person.

2

LEVEL

Assistant

Helps with tasks and retrieves context. Still fully human-driven, one step at a time.

3

LEVEL

Agent

Executes real multi-step work on its own. Capability arrives — and risk becomes structural.

4

LEVEL

Governed Agent**OPERATIONAL AI GOVERNANCE OPERATES HERE**

The agent acts inside an enforced authority envelope — consent, risk tier, escalation and evidence checked before every action.

5

LEVEL

Evidence-Earned Adaptive Autonomy

Autonomy expands only for specifically inventoried decisions whose risk, evidence and institutional policy permit it. Human authority is not an immature stage — in many clinical contexts it is the correct operating model.

Operational AI Governance operates between Level 3 and Level 4 — turning a capable agent into a governed one. That is where safe adoption lives.

What hospitals are really buying.

Hospitals do not buy governance. They buy outcomes — the results that keep patients safe and staff confident as care moves between people and settings.

- | | |
|---|--------------------------------------|
| ✓ Continuity of care | ✓ Safer AI adoption |
| ✓ Better, safer discharge | ✓ Fewer information gaps |
| ✓ Stronger multidisciplinary communication | ✓ Improved patient experience |
| ✓ Greater staff confidence | ✓ A defensible audit trail |

Governance is the mechanism. Continuity is the outcome. Operational AI Governance is the discipline of ensuring every AI-mediated action in care is authorised, consented, risk-tiered, escalated and evidenced — at the point of use — so that a follow-up, an escalation or a handover can be trusted across every care transition.


AS ESSENTIAL AS CYBERSECURITY BECAME TO DIGITAL HEALTHCARE

A decade ago, few hospital boards discussed cybersecurity as core infrastructure. Today none would deploy a clinical system without it. Operational AI Governance is following the same curve for AI-enabled care: first a differentiator, soon an expectation, ultimately a precondition for deployment.

And it answers the five questions no ungoverned AI can: **Who authorised this action? Was consent valid at this moment? Who holds authority now? Is escalation required? And when something goes wrong — why was execution permitted?**

Five principles of Operational AI Governance.

A hospital executive needs these five principles — not the software beneath them. Together they keep continuity safe as care moves between people, teams and settings.



Authority flows through the care team




◆ **Authority · Consent · Risk · Escalation · Evidence — five principles in service of continuity of care.**

Powered by Runtime Governance Infrastructure.

Operational AI Governance is made possible by **Runtime Governance Infrastructure (RGI)** — a domain-agnostic architecture that evaluates the five principles **at the point of execution**, in milliseconds, before an AI-mediated action reaches a patient.

Governance before action



THE MENTAL MODEL — AIR TRAFFIC CONTROL FOR AI

Runtime Governance is to AI what **Air Traffic Control** is to aviation. ATC does not fly the aircraft — it governs the safe execution of flight. Aircraft, pilots and airlines differ; ATC is the common layer that keeps the system safe. RGI does not replace AI models, EHRs or clinical workflows. It governs the safe execution of AI-mediated actions between them.

HCP-as-Pilot™ is the healthcare framework built on RGI, following the principle proven in aviation: **autopilot may execute, but authority remains with the pilot.** AI may assist, recommend and coordinate; authority over patient-impacting actions remains human. RGI is a new infrastructure category — healthcare is its first implementation.

A narrow window — and evidence already accruing.

Regulatory inevitability creates a window for first movers: hospitals that build Operational AI Governance today will hold the evidence and institutional trust required when it becomes mandatory.

INDICATIVE MARKET-DEVELOPMENT SCENARIO — PROJECTION, NOT SOURCED REGULATORY FACT

- 2026** **WE ARE HERE**
EU AI Act enforcement begins; runtime evidence supporting conformity assessment becomes a differentiator.
- 2027 →**
Projected: first enforcement precedents; ungoverned systems face scrutiny.
- 2028 →**
Projected: payer requirements emerge; governance evidence trends toward table stakes.
- 2029+ →**
Projected: clinical expansion favours systems with a proven governance history.

Regulatory alignment — evidence by design, not compliance by assertion

FRAMEWORK	REFERENCE	ALIGNMENT
EU AI Act	Art. 12, 14, 16	Generates runtime evidence for human-oversight, record-keeping and post-market monitoring obligations.
Singapore AIHGle 2.0	MOH/HSA 2026	Independent convergence: mandates human oversight for clinical AI; flags home-care AI as needing added governance.
UK MHRA	SaMD Roadmap	Phased architecture aligns with the MHRA's progressive AI-oversight framework for software as a medical device.
FDA SaMD / PCCP	2019 / 2023	Phase III maps to Predetermined Change Control Plan requirements for adaptive AI medical devices.

Early field evidence — governed AI in real homes

Phase I (Home Companion™) is an early-access usability deployment with senior users across Switzerland and the UK — non-medical, non-SaMD, governed end to end.

4,600+

Governed runtime events — each evaluated before execution

Consent · Authority · Audit

Controls applied to every event; nothing acts outside authorised bounds

12 · 3 · 2

Seniors · care professionals · months of continuous use

A *governed runtime event* is a single AI-mediated action evaluated for consent, authority, risk and escalation before it executes. Figures describe a usability deployment, not a clinical study; hospital-pilot recruitment is now underway.

Operational AI Governance is the new precondition for AI in care.

Every hospital now has access to intelligence. The differentiator is **the capability to govern it** — to trust, scale, audit, improve and certify AI acting between clinical encounters. Hospitals that can do this will adopt AI faster, protect continuity of care, and earn greater confidence from patients, clinicians and regulators.

"She made me laugh speaking in Luzern Swiss German dialect from my childhood."

— Senior & cancer survivor, 86, Basel

"I love that I can suggest personalised memories for my mum to discuss. It makes her smile."

— Caregiver, 62, Allschwil

"I care for multiple seniors, and Home Companion keeps the status of each organised for me."

— Care-home assistant, 42, Basel

YOUR NEXT STEP

Pilot the governance layer before clinical autonomy.

Start where the risk is bounded — one care transition, one virtual ward, one discharge pathway. Prove Operational AI Governance there, then expand autonomy only where the evidence earns it.

1

Select one bounded care workflow

2

Inventory its decisions

3

Define authority & escalation

4

Generate 90 days of runtime evidence

**The question is not whether AI can participate in care.
It is how AI is governed — at the point of use, in service of continuity.**

PCC-SafetyOS™ — the software beneath the principles.

Hospital readers can stop at the five principles. This appendix is for technical evaluation and procurement due diligence: the components of PCC-SafetyOS™ that implement Operational AI Governance using Runtime Governance Infrastructure — each stated as one function and one governance consequence.

Decision Inventory **FOUNDATION**

Inventories every AI-involved decision and assigns one named human owner.

→ **A decision not inventoried cannot be automated.** · AUTHORITY

Human Authority Envelope

Defines the boundaries AI may act within — and who holds authority when it cannot.

→ **AI never operates outside authorised boundaries.** · AUTHORITY

Consent Engine

Validates consent at the moment of every AI-mediated action.

→ **No consent, no action; ambiguity defaults conservatively.** · CONSENT

Risk-Adaptive Stratification

Classifies every action into a risk tier in real time.

→ **Higher risk demands human confirmation or authority.** · RISK

Escalation Engine

Routes to caregiver or clinician by urgency and availability.

→ **Escalation is governed, not improvised.** · ESCALATION

Flight Recorder

Records why each action was permitted — not just what happened.

→ **Every governed decision is reconstructable; append-only.** · EVIDENCE

Runtime Evidence Ledger

Structures governance artefacts for internal review, quality assurance and regulatory assessment.

→ **Everyday operation becomes defensible evidence.** · EVIDENCE


◆ **Operational AI Governance, powered by Runtime Governance Infrastructure — in service of continuity of care.**

From Autopilot to Agentic AI: The Pilot Never Left the Cockpit

Accountability doesn't disappear with automation. It evolves.

AVIATION ECOSYSTEM

PILOT & COPILOT IN THE COCKPIT



FLIGHT MANAGEMENT SYSTEMS



BLACK BOX (FLIGHT RECORDER)




- ✓ Records all critical data
- ✓ Enables investigation
- ✓ Drives learning & improvement

HEALTHCARE ECOSYSTEM

DOCTOR & NURSE IN FRONT OF THE PATIENT



HEALTHCARE SYSTEMS & DATA FLOW



FLIGHT RECORDER (HEALTHCARE AUDIT TRAIL)



- ✓ Records all decisions & actions
- ✓ Enables audit & compliance
- ✓ Drives learning & safety improvement

THE PARALLEL IS CLEAR

Pilot (Accountable for the flight)	↔	HCP-as-Pilot (Accountable for the patient journey)
Autopilot Systems	↔	Safety OS (RGI Layer)
Copilot (Supports & Monitors)	↔	Super Agent & Agent Swarm (CoPilot)
Flight Systems (Data Inputs)	↔	Healthcare Systems (Data Inputs)
Black Box (Flight Recorder)	↔	Flight Recorder (Healthcare Audit Trail)



Technology evolves. Responsibility remains.

From autopilot to agentic AI, the future belongs to those who can govern complex systems while keeping human judgment at the center.

**HUMAN AUTHORITY.
SYSTEMS SUPPORT.
BETTER OUTCOMES.**