PATIENTCENTRICCARE.AI · CONTINUITY PLATFORM

Governance & Data Stewardship Brief

Home Companion™ • MyHealthCanvas® • PCC-SafetyOS™ • **Trusted Health Memory™**

PatientCentricCare.AI is building the continuity platform for trusted AI-enabled healthcare — helping people preserve continuity of care, continuity of knowledge and continuity of their personal health context, while ensuring every AI-assisted interaction remains governed, transparent and under human authority.

Our hypothesis is that **continuity compounds**: every governed interaction enriches a trusted longitudinal **Trusted Health Memory** that becomes increasingly valuable to patients, caregivers and clinicians over time. Electronic Health Records record clinical *events*; Trusted Health Memory preserves the person's longitudinal health *context* across time.

A memory that grows with every interaction only earns a place in healthcare if two things are true: it is **governed at runtime**, and it **belongs to the person**. This document explains, in plain terms, how both are achieved — what the system does, what it deliberately does not do, what is stored, who can see it, and who holds authority at every step.

HOW TO READ THIS BRIEF

Everything here is a factual description of the system as built, written to be forwarded — to an ethics committee, a data-protection officer, a clinical-safety lead, or an innovation board. Where a claim is aspirational rather than implemented, we say so explicitly. **We would rather under-claim than over-claim.**

ENGAGE
Home Companion™
Voice-first companion that engages people before, during & between care. Deliberately non-medical.

→

CAPTURE
MyHealthCanvas®
Structured patient capture — what matters from every encounter, in the patient's own words.

→

REMEMBER
Trusted Health Memory™
The person's own longitudinal health context — what they chose to keep, preserved across time.

PEOPLE — Patients · Clinicians · Caregivers at the centre of the platform · always in control

PCC-SafetyOS™ · Runtime Governance Infrastructure

Engage · Capture · Remember · **Govern** — every interaction passes through governance before it can have any effect. PCC-SafetyOS™ encloses the platform as a **Human Authority Envelope**: AI never operates outside boundaries a human has authorised.

GOVERNED · COMPLIANT · PRIVATE · HUMAN AUTHORITY

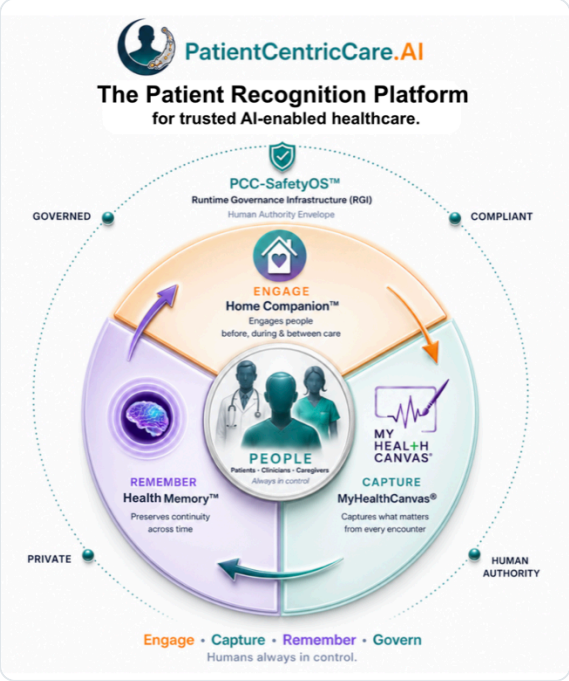

The Patient Recognition Platform — one governed continuity loop: Engage · Capture · Remember · Govern. Humans always in control.

01 Why continuity matters

THE CORE IDEA

Patients spend less than **1%** of their lives with clinicians. The other **99%** — where care plans succeed or fail, where understanding fades and adherence drifts — is invisible to the health system. Records capture the encounter; nothing preserves what happens between encounters. That gap is what a governed **Trusted Health Memory** is built to close.

WHY NOW

Care is moving from episodic hospital visits to continuous, AI-enabled care at home. As it does, a trusted, governed memory becomes the persistent thread that holds care together — and the **EU AI Act** now makes runtime governance of that AI a requirement, not an option.

Runtime governance is the engine. Home Companion is the demonstration. Trusted Health Memory is the accumulating asset. Continuity is the outcome.

CONTINUITY COMPOUNDS

1. **Appointment.** A clinical encounter takes place.
2. **Conversation.** Home Companion engages the person, governed at runtime.
3. **The person chooses what matters.** Captured in their own words via MyHealthCanvas.
4. **Trusted Health Memory grows.** A longitudinal health context accrues — owned by the person.
5. **Next appointment.** The memory is carried forward into the following encounter.
6. **Better-prepared patient & clinician.** Less repeated history; more informed, more personalised care.

∞ **Continuity compounds.** Every loop makes the memory richer and the next encounter better.

WHAT BREAKS WITHOUT CONTINUITY

Understanding fades
People forget instructions and explanations soon after leaving.

Context disappears
Clinicians cannot see what happens once the patient is home.

Adherence drifts
Medication and lifestyle follow-through gradually decline.

Every visit starts over
Hard-won context is lost and re-gathered each time.

02 The governance claim

RUNTIME GOVERNANCE

The AI does not make the safety decision. A fixed set of written rules — not the AI model's judgement — decides what is allowed, what is refused, and when a human must be involved.

The rules are **deterministic**: the same situation always produces the same decision. If the rules cannot load, the system **fails closed** — it does nothing beyond a safe minimum. The intelligence answers questions; the governance decides what is permitted.

GOVERNANCE BEFORE ACTION — NOT AUDIT AFTER IT

1 · PROPOSE
AI proposes
An action or response. Nothing executes yet.

→

2 · GOVERN
PCC-SafetyOS™ evaluates
Authority · Consent · Risk · Escalation · Evidence — checked before patient impact.

→

3 · ACT
Permitted — or refused
Refusals use fixed, non-medical templates.

→

4 · RECORD
Evidence to audit trail
Append-only; can never be edited or deleted.

RISK-ADAPTIVE TIERS — WHO DECIDES, AT EVERY LEVEL

TIER	SITUATION	WHO DECIDES	EXAMPLE
TIER A	Routine, bounded interaction	AI executes within predefined bounds set by the rules	Daily check-in, appointment reminder, conversation
TIER B	Safety signal or repeated refusal	Human confirmation required — escalation offered, human notified	Soft-distress phrase → senior asked; caregiver notified on acceptance
TIER C	Anything clinical	Human authority required — the system does not act	Medical questions are refused and redirected to a human

Tier assignment is deterministic, centralised in the governance evaluator, and recorded on every decision.

RED LINES — ENFORCED BY RULE, NOT BY PROMPT

- ✕ No diagnosis. Never assesses or monitors a medical condition.
- ✕ No triage. It does not decide clinical urgency.
- ✕ A caregiver cannot override patient consent.
- ✕ No treatment or medication advice — ever, under any phrasing.
- ✕ It never contacts emergency services itself.
- → When something needs a person, it escalates to a person.

OUR GOVERNANCE PRINCIPLES

- › Human authority always prevails.
- › Governance before capability.
- › Patient consent before persistence.
- › Runtime evidence before regulatory claims.
- › Under-claim before over-claim.

Why this matters for a compounding memory. Trusted Health Memory becomes more valuable with every interaction. A memory that compounds is only safe to hold if every interaction that feeds it is governed the same way, every time — by rules that can be inspected, tested and audited, not by a model's mood.

03 Who owns the memory

DATA STEWARDSHIP

Trusted Health Memory compounds — and a compounding health dataset is exactly what a privacy officer should question hardest. Our answer is structural: the memory compounds **for the person, under their own authority** — not for the vendor.

THE PATIENT'S FOUR RIGHTS — BUILT IN

- ✓ **Owned.** The memory is the patient's record — they choose what is kept.
- ✓ **Consented.** Withdrawable at any time, immediately; a caregiver cannot override it.
- ✓ **Portable.** The patient can take their memory with them.
- ✓ **Deletable.** Deletion removes their data; audit records are anonymised.

WHAT NEVER EXISTS IN THE FIRST PLACE

- ✗ No audio is ever stored — voice is streamed live, answered, then gone.
- ✗ No conversation transcripts are retained.
- ✗ The audit trail records what *kind* of interaction happened — never the words.
- ✗ No advertising use, no data resale.

DATA INVENTORY — WHAT IS STORED, WHERE, WHO SEES IT, HOW LONG

DATA	STORED	WHERE	WHO CAN ACCESS	HOW LONG
Voice / audio	No	Streamed live to the AI provider only	—	Not kept
Conversation transcripts	No	Not created for retention	—	Not kept
Audit record (interaction type, decisions, escalations — not the words)	Yes	App's secure database	The patient; the service, under strict rules	With patient consent; anonymised on deletion
Patient profile & consent records	Yes	App's secure database	The patient only	Deleted with the account
Memory items (what the patient chose to keep)	Yes	App's secure database	The patient only	Patient-controlled

Row-level rules
Each person sees only their own data — enforced by the database itself, not by app code.

Encrypted & vaulted
Encrypted connections throughout; credentials in a server-side vault, never in source code.

Append-only audit
Audit records can be added but never edited or deleted — including by administrators.

PROCESSORS — WHO TOUCHES THE DATA, AND FOR WHAT

PROVIDER	ROLE	SEES
Supabase	Secure database, sign-in, application back-end	Stored data listed above, under access rules
OpenAI / Google (Gemini)	Understands and answers speech (conversation)	Live speech stream only — nothing stored
ElevenLabs	Produces the spoken voice	Reply text to be spoken
Resend / Vonage	Caregiver email / SMS notifications	Notification content only

Each provider is engaged under a data-processing agreement; data-residency and no-retention / no-training commitments are being contractually confirmed for the final version of this brief. Designed around Swiss (FADP) and EU (GDPR) data-protection principles; international transfers use standard safeguards (SCCs).

04 Evidence discipline

WHAT WE CLAIM — AND DON'T

A governed platform must also govern its own claims. We distinguish sharply between what our early deployment demonstrates **today** and what only clinical study can demonstrate.

✓ WHAT WE ARE VALIDATING TODAY

- ✓ Bounded human–AI interaction in real homes
- ✓ Longitudinal Trusted Health Memory creation
- ✓ Patient engagement & conversation continuity
- ✓ Safe escalation to a human
- ✓ A working governance audit trail
- ✓ Patient-owned history & consent flows

— WHAT WE ARE NOT YET CLAIMING

- — Reduced admissions or readmissions
- — Reduced cost of care
- — Improved clinical outcomes

These are hypotheses — what hospital pilots exist to test, under an investigator's protocol, not marketing claims.

EARLY REAL-WORLD DEPLOYMENT — GOVERNED OPERATION, OBSERVED

4,600+	12	3	2 mo
governed AI interactions — every one through the audit trail	patients & seniors, in real home environments	caregivers participating alongside them	of continuous operation — Switzerland & UK

Signals observed to date: conversation continuity and engagement; correct human escalation; onboarding complexity identified early and simplified in response — an example of governance-informed iteration. These are **operational signals, not clinical claims**.

EXHIBIT — THE AUDIT TRAIL, ILLUSTRATED

TIME	TIER	EVENT	GOVERNANCE DECISION	HUMAN INVOLVED
07:42	**A**	Daily check-in initiated	Permitted within predefined bounds	—
09:15	**C**	Medication question detected	Refused — fixed non-medical template; redirected to caregiver	Redirect to human
14:02	**B**	Soft-distress phrase matched	Escalation offered; senior accepted	Caregiver notified (SMS + email)
14:03	**A**	Supportive conversation continued	Permitted within bounds; consent state verified	—
18:30	**A**	Memory item proposed	Held for patient review — nothing stored without consent	Patient decides

Illustrative, de-identified representation of Phase I audit-trail output. Every row is append-only: it can be added, but never edited or deleted. Note what is absent: the words spoken are never recorded — only what kind of interaction occurred and how it was governed.

Why under-claiming is the strategy. For a Principal Investigator or ethics committee, a vendor that under-claims is a credible evidence partner; a vendor that over-claims is a risk. Every sentence in this brief is written to survive that test.

05 Regulatory posture & audit readiness

EVIDENCE BY DESIGN

Designed for a limited-risk position. Engineered to a higher evidence standard.

Home Companion is deliberately **non-medical** — no diagnosis, no monitoring, no treatment claims — and the platform is designed to support a limited-risk regulatory position with **Article 50 transparency** (users always know they are interacting with AI). Final regulatory classification and jurisdiction-specific obligations remain subject to legal review. Regardless of where that line settles, the governance infrastructure is engineered to generate runtime evidence at the standard the most demanding regimes expect. **Not compliance by assertion — evidence by design.**

FRAMEWORK ALIGNMENT

FRAMEWORK	REFERENCE	HOW THE PLATFORM ALIGNS
EU AI Act	Art. 50; evidence to Art. 12/14/16 grade	Transparency obligations designed in. The audit trail ("flight recorder") generates the runtime evidence — human oversight, record-keeping, post-market monitoring — that the most demanding (high-risk) tier expects, giving headroom whatever the final classification.
Singapore	AIHGle 2.0 · MOH / HSA 2026	Independently mandates human oversight for clinical AI and flags home-care AI as needing additional governance — the exact gap this platform is built to close, at runtime.
UK MHRA	SaMD / AI Airlock roadmap	Phased governance architecture aligns with MHRA's progressive oversight approach for software in healthcare.
US FDA	SaMD / PCCP (2019 / 2023)	Later-phase governance architecture maps to Predetermined Change Control Plan expectations for adaptive AI.

WHAT "AUDIT READINESS" MEANS IN PRACTICE

Because every interaction already produces governance evidence, an institution deploying the platform can answer an auditor's questions from **records, not recollection**:

Who authorised it?
Every action carries its authority check — which rule permitted it, and under whose consent.

Who was in control?
Tier assignment on every decision shows when the AI acted within bounds and when a human decided.

Can you prove it?
The append-only trail is replayable evidence — usable for conformity assessment and internal AI-governance review.

Research positioning. Hospital pilots are prospective, observational and non-interventional — designed for ethics-committee review (e.g. BASEC/swissethics in Switzerland), with a DPIA, proxy-consent provisions where capacity requires, and any clinical measures belonging to the investigator's protocol, not to the product. Regulatory positioning is under review by external counsel; this section reflects our working understanding.

06 The Hospital Innovation Pilot

PROPOSED PILOT DESIGN

A **cost-free, non-commercial feasibility pilot**, co-designed with a named Principal Investigator. It evaluates the feasibility, acceptability and safe deployment of Home Companion — **not clinical efficacy**. The final protocol, endpoints and analysis plan are developed jointly with the institution's PI and, where required, its ethics committee.

Design	Prospective, single-arm, real-world feasibility pilot
Setting	Teaching hospital and/or affiliated outpatient or community service
Participants	~20–40 adults aged 50+, recruited by the institution — e.g. cancer survivors · mild age-related memory concerns · people managing multiple long-term conditions
Duration	~6–12 weeks of use — appointment prep, capturing what matters, post-visit reflection, between-visit engagement
Exclusions	Significant cognitive impairment or dementia, unless specifically approved within the study protocol
Data collected	The minimum necessary: consent records · audit events · application usage · patient-selected memory items · optional satisfaction questionnaires. No EHR integration; no diagnostic outputs; no treatment recommendations.

✓ PRIMARY ENDPOINTS — IMPLEMENTATION & SAFETY

- ✓ Onboarding success · active use · retention
- ✓ Technical reliability · safety events
- ✓ Governance-rule performance · consent management
- ✓ User acceptability · staff usability feedback

— WHAT THE PILOT DOES NOT MEASURE

- — Diagnosis or clinical monitoring
- — Treatment or medication effect
- — Clinical efficacy or improved outcomes

Secondary endpoints are exploratory, non-clinical observations only — perceived usefulness, pre-visit confidence, recall of discussion points, perceived continuity, caregiver reassurance.

SUCCESS CRITERIA

- ✓ Safe deployment & acceptable usability
- ✓ High participant engagement
- ✓ Successful governance operation
- ✓ Practical integration into routine workflows
- ✓ Enough RWD to justify a larger multi-site evaluation

WHAT THE PARTNER RECEIVES

- › A governed deployment at no cost — app, capture forms, onboarding
- › The full governance evidence pack — the runtime audit trail, for EU AI Act readiness review
- › Co-authorship on published Real-World Evidence
- › First-mover experience with runtime-governed AI

Review the Phase 1 Pilot Framework

A 15-minute discussion on whether a governed, non-commercial RWD feasibility pilot fits your 2026 innovation goals — not a product demo.
patientcentriccare.ai/pilot-apply
Andy Squire, Founder & CEO · andy@patientcentriccare.ai · **patientcentriccare.ai/home-companion**