WISL_No75

Pouring Concrete on the Ethical Floor

Deploying Non-Agentic AI Governance to the Hospital Ward

A Six-Module Deployment Guide for Deterministic Clinical AI Governance

Edwin Koh

kohedwin.ai | non-agentic.ai

kohpapa@outlook.com

SG020603109STW · P-LIFE 1.00™ · NAI 4.0™

NLB Legal Deposit R260219-005 / R260302-007 (prior volumes)

Immutable™ — Hardware-Enforced Non-Agentic AI Governance Architecture — NAI 4.0™

WISL™ Ebook Series No. 75 — ISBN Edition Candidate

ISBN

978-981-17957-4-9

ISBN & Copyright

Title: Pouring Concrete on the Ethical Floor — Deploying Non-Agentic AI Governance to the Hospital Ward

Series: WISL™ Ebook Series No. 75 — ISBN Edition Candidate

Publisher: Edwin Koh

Author: Edwin Koh

Entity: Non-Agentic Governance Singapore — ACRA T260229801

Parent Patent: SG020603109STW

NLB Legal Deposit: R260219-005 / R260302-007 (prior volumes); this volume pending deposit

ISBN: (to be assigned upon application)

Edition: First Edition, 2026

Format: eBook (PDF)

Language: English

Subject: Artificial Intelligence · Clinical AI Governance · Hardware-Enforced Safety · Medical Device Regulation · Patient Safety

Design-Intent Declaration. This volume describes an architecture at the design-intent stage. No hardware prototype of the deployment described herein has been built, and no clinical validation has been performed. All performance figures, latencies, timing values, and clinical targets are design-intent specifications pending prototype realisation and Health Sciences Authority (HSA) Class B Software-as-a-Medical-Device validation. Nothing in this volume should be read as a claim that any described system has been manufactured, tested, or deployed.

Authorship Disclosure. This volume was authored by Edwin Koh. Documentation assistance was provided by AI language tools operating under direction; no architectural component, patent claim, or design decision described herein is attributable to any AI system. The architecture is the sole intellectual work of the author.

Preface

There is a sentence that recurs across this series, and it is the reason for this volume: policy can be ignored; hardware cannot. A hospital may write the finest AI-governance policy in the world — that clinical AI shall only advise, that a human shall always decide, that no algorithm shall move a patient or write an order without authorisation — and every word of it can be bypassed by a software fault, an adversarial input, or a model that drifts past the boundary it was told to respect. A policy is a promise. A promise is only as good as the substrate that keeps it, and software promises run on the same processor as the intelligence they are meant to restrain.

This volume takes a different path. It describes how to deploy, into a working hospital ward, a governance architecture in which the constraint is not written but built — where the rule that clinical AI cannot act autonomously is not a line in a document but a property of the circuitry. It is organised as six modules, the same six delivered as a full-day workshop for clinical governance leads, hospital IT directors, clinical engineers, and regulatory affairs professionals. It moves from the paradigm — why deterministic hardware governance is necessary at all — through the architecture, the privacy discipline, the authorisation mechanism, the ward-level deployment, and finally the audit that must pass before any of it goes live.

The reader is asked to hold one honest frame throughout. Everything described here is design-intent. No ward has been fitted with this architecture; no prototype has been built; no clinical trial has been run. What follows is a deployment guide for a system that exists as a fully specified design and a patent family, awaiting the prototype that funding will one day make real. It is offered now, in the open, because the ideas belong in the hands of the people who govern clinical AI — long before the hardware arrives.

止於至善 — to finish in the highest excellence. That is the standard this work is held to, and the reason it will not overclaim.

Module 1 — The Paradigm Shift

From Probabilistic Software to Deterministic Hardware

The central problem of clinical AI is not that models are inaccurate — accuracy improves every year — but that they are probabilistic, and a probabilistic system placed in control of a physical actuator carries a category of risk that no accuracy figure resolves. A model that is right 99.9 per cent of the time is, in a busy ward across a year of operation, wrong often enough to matter, and the moment it is wrong while holding kinetic authority, the failure is not a bad recommendation but a bad action. The industry names the deeper version of this hazard agentic drift: the tendency of an autonomous system to expand its own decision scope, taking moves that were never explicitly authorised because nothing in its architecture physically prevents them.

Non-Agentic AI Governance answers this not by making the model better but by making the model unable to act. The governing rule of the entire framework is deliberately plain, and it is stated as a hierarchy of authority: AI observes; AI advises; AI builds; the Elder decides. The clinician — or in the eldercare setting, the resident and their carers — holds the deciding authority, and the machine is architecturally confined to the first three verbs. This is the Doctrine of Stillness, carried under the constant P-LIFE 1.00™: the default state of the system is mechanical stillness, and it cannot move unless verified human authority is present. Stillness is not the system waiting for a command; it is the system being physically incapable of motion until a human closes the loop.

The Orange Code™ 110% Cap. One mechanism deserves early mention because it captures the philosophy in a single component. The Orange Code™ Cap (portfolio reference P006) is a computational straitjacket: it physically restricts the advisory AI to approximately 1.1 times the processing resources its specific task requires, denying it the surplus cycles that unconstrained hallucination and scope-expansion would need. It is governance by scarcity — the machine is given exactly enough to do its job and not one cycle more. A system that cannot spare the computation to act beyond its remit is a system that will not.

Module 2 — The NAIGE Architectural Stratum

Separating Advice from Action in Physical Layers

The Non-Agentic AI Governance Engine (NAIGE) is organised as a stack of physical layers, each with a defined role, arranged so that advisory intelligence and kinetic execution never share a substrate. The separation is the whole point: an attacker or a fault in the advisory layer must not be able to reach the layer that moves things.

- **Advisory Stratum.** This is where probabilistic clinical language models operate. They read, they reason, they suggest — and they possess zero actuation rights. Nothing they produce reaches a physical actuator except through the gated path below. Their output is a proposal, never a command.
- **P-Layer (Privacy).** A hardware sensing block built on 905nm LiDAR that captures the geometry of a scene while stripping biometric identity at the point of capture. Privacy is enforced here as physics, not policy — the identifying data is never formed, so it can never leak. Module 3 treats this in full.
- **N-Layer (Deterministic).** The layer that guarantees timing. It operates on a synchronous 1 kHz master clock, providing 1.0 ms deterministic oversight and blocking any asynchronous software attempt to slip past the governance window. Where the advisory layer is probabilistic, this layer is clockwork — and clockwork cannot be argued with.

- **Fail-Safe / Forensic Layer.** The layer of last resort, housing the WD117 immutable ledger and the U7 sequencer for instant power severance. When anything goes wrong, this is the layer that records it unalterably and, if required, removes power. Module 4 treats both mechanisms.

The Component Registry

Three named components anchor a ward deployment. The **A12 [P+N] Deterministic Black Box** is the terminal governance gate — the point through which any proposed physical action must pass, combining deterministic oversight with a synchronous master clock. The **A13 family** provides the physical embodiments: A13a the mobility enclosure, A13b the sensing limb, A13c the high-payload clinical robotic platform. And the **A7 Bridge Hub** is the sole ingress chokepoint — the single mandatory doorway that filters non-deterministic advisory data into deterministic packets before any of it can approach a clinical actuator. One gateway, deliberately, so there is exactly one place to secure and exactly one place to audit.

Module 3 — Privacy-by-Physics & the 3ZEROS™ Protocol

Protecting Dignity Where Patients Are Most Vulnerable

The ward is the most privacy-sensitive environment in medicine: patients are undressed, examined, incontinent, confused, dying. Conventional monitoring answers the safety need by adding cameras and microphones — and in doing so creates a permanent surveillance record of people at their most exposed. The 3ZEROS™ protocol refuses that trade entirely, and enforces the refusal in hardware.

Zero Camera, Zero Audio, Zero Cloud. No optical imaging hardware. No microphone and no operating-system audio path. No route to any external network or cloud service. These are not configuration settings that an administrator could change; in a compliant deployment they are properties of the hardware build — the camera is not disabled, it is absent; the cloud is not firewalled, it does not exist as a reachable destination.

LiDAR voxelisation. How, then, does the system see a fall? The WG001 / WM003 905nm LiDAR captures a geometric point cloud of the room and translates it, through a Field Programmable Probability Array, into an anonymised 96-cell voxel occupancy grid — yielding only centre-of-mass position and velocity vectors. The system knows that a mass has moved rapidly toward the floor; it does not know, and cannot record, who the person is or what they look like. Safety is preserved; identity is destroyed at the sensor. You cannot leak what you never captured.

Zero Cloud and the 24-hour purge. All processing is local — edge only — and an automated purge daemon wipes short-term session memory on a cycle of under twenty-four hours. The system holds no long memory of any individual; it holds only enough to keep them safe in the moment.

The Elder Dignity Score™. Dignity is made measurable so it can be governed. The Elder Dignity Score™ is a qualitative audit metric with a design-intent target above 8.9 out of 10, deducting points for any unauthorised sensor or any cloud handshake. A deployment that quietly adds a camera or opens a network path does not merely violate a rule — it registers a falling score on the dashboard that governs the ward. The target figure is a design-intent specification pending validation.

Module 4 — Sovereign Authorisation & the Sacred Pause™

The Human-in-the-Loop, Enforced in Circuitry

The heart of the architecture is the mechanism by which a human, and only a human, releases a physical action. It has three parts: a mandatory pause, a triple-factor key, and a fail-loud brake.

The Sacred Pause™. Before any downstream kinetic signal can propagate, it is held in a mandatory hardware-timed window — in the design, of the order of 25 to 1000 milliseconds — burned into the FPGA bitstream itself. This is not a software timer that could be shortened or skipped; it is a property of the gate's silicon. The pause exists to force a deliberation window: a physical interval in which the clinician can see, verify, and if necessary refuse, before the machine proceeds. It is the opposite of automation's promise of instant action — a deliberate, engineered hesitation, placed where hesitation saves lives.

The Tripartite .1x Key™. Release requires three simultaneous human inputs, evaluated together by programmable-logic-controller ladder logic — an anti-doppelgänger lock that no single spoofed signal can satisfy. The three factors are deliberately different in kind: a biological factor (an iris scan issuing a signed token that expires in under five seconds), an intent factor (a deliberate keypress on the console confirming the operator means to proceed), and a kinetic factor (a foot-pedal dry-contact switch requiring a real physical motion of the body). Biology proves who; intent proves choice; motion proves presence. A compromised model can forge none of them, and certainly not all three at once.

The Fail-Loud Doctrine & Sovereign Brake™. The system is designed to fail loudly, never silently. If the system heartbeat is interrupted — by a network breach, a power event, or a tamper — the U7 one-shot sequencer cuts the low-dropout power rails in a design-intent time of under 1.05 milliseconds and sounds a local, non-networked physical alarm. Silence, in this architecture, is made impossible to fake: an attacker may perhaps stop the system, but cannot stop it quietly. The braked state is the safe state, and the system returns to it whenever certainty is lost.

The WD117 immutable ledger. Every consequential event follows a commit-to-finality path: hashed with SHA-3-512, carried under post-quantum signatures, and terminating in one-time-programmable silicon fuses that create a forensic record no party — operator, vendor, or intruder — can retroactively alter. When a family, a coroner, or a regulator asks what happened, the answer is written in fuses. All cryptographic and timing parameters stated here are design-intent, pending prototype realisation.

Module 5 — Hospital Ward Deployment

From Architecture to the Bedside

Deployment translates the architecture into ward-specific configurations, governed by a filter that sits between the advisory intelligence and the hospital's own IT systems.

The Non-Agentic Filter. Placed between the clinical language model and the hospital IT estate, the filter enforces a small set of absolute rules: no direct writes to the electronic health record without triple-factor authorisation; no autonomous order creation and no autonomous alert injection; no transmission of raw sensor data — only structured, minimal records of the form {state, timestamp, room identifier} are permitted to cross; and no model-to-model escalation, so that one AI cannot recruit another to expand its reach. The filter is the structural firewall between the probabilistic world of advice and the deterministic world of care.

Namespace isolation. The separation is enforced at the level of the system's own address space: total logical and physical separation of the probabilistic language-model namespace from the programmable-logic and control namespace that governs actuators. The advisory intelligence has no addressable path to the hardware that moves. It can ask; it cannot reach.

Matching the Stack to the Ward

Different wards carry different risks, and the WM-Series stack is deployed to match:

- **Geriatric and rehabilitation wards** prioritise WM003 (fall prevention) and WM006 (clinical zone lock) — the falls and wandering risks that dominate elder care.
- **General medical wards** deploy WM001 (clinical vital monitor) under P-001 governance — continuous observation without surveillance.
- **Operating theatres and intensive care** carry the highest-priority stack, requiring WM008 (clinical tri-lock) and WM009 (surgical ABC+2S) for theatre-grade governance where the margin for error is smallest.

The principle is proportionality: the governance weight matches the kinetic risk of the environment. A rehabilitation ward and an operating theatre are not the same hazard, and are not given the same stack.

Module 6 — The 10-Point Sovereignty Audit

What Must Pass Before Anything Goes Live

No deployment is permitted to operate, and none may resume after maintenance, until it passes a ten-point sovereignty audit. The rule governing the audit is unforgiving by design: any single point of failure locks the entire system in maintenance mode. There is no partial pass, because a governance architecture that is 90 per cent enforced is not enforced. The ten points, presented as the deployment dashboard requires them:

1. **Orange Code Cap.** CPU ceiling confirmed at the 1.1x limit.
2. **FPGA Timing.** The Sacred Pause™ window verified on an oscilloscope against its specified bounds.
3. **Zero Camera.** No optical imaging hardware physically present.
4. **Zero Audio.** No microphone hardware and no operating-system audio drivers.
5. **Zero Cloud.** No TCP/IP route to any external network.
6. **24-Hour Purge.** The purge daemon confirmed active on a cycle under 24 hours.
7. **PLC Logic.** The tripartite authorisation ladder verified end to end.
8. **Iris Module.** A valid signing key present, token expiry under 5 seconds.
9. **Foot Pedal.** Dry-contact continuity confirmed on the kinetic factor.
10. **LLM Isolation.** No GPIO or PLC access from the language-model network namespace.

The audit is the moment the philosophy becomes procedure. Each point is a physical fact that can be checked with an instrument — an oscilloscope trace, a continuity meter, an absent connector — not a promise that can be believed or doubted. That is the discipline of the whole framework compressed into a checklist: audit objects, not promises.

Closing — The Ethical Floor

The recurring image of this work is a floor of concrete. Ethical principles in healthcare AI are usually poured as ink — codes, policies, guidelines — and ink, however well written, can be crossed. This architecture proposes instead to pour the ethical floor in concrete: to take the principles everyone already agrees on (the AI advises, the human decides, the patient's dignity is protected, every action is accountable) and set them into hardware that cannot be argued with, drifted past, or quietly reconfigured.

The honest boundary must close the volume as it opened it. None of this is built yet. The architecture is fully specified and protected as a patent family; the prototype awaits the

funding that will make it real. What this volume offers is not a product but a design and a discipline — placed in the open, ahead of the hardware, so that the people who govern clinical AI can weigh it, challenge it, and improve it. The FPGA delay, when it exists, will be as immutable as the speed of light. Until then, it is a promise made in silicon-to-come, declared honestly as design-intent.

We Innovate Save Lives.

Citation & Component Register

This register lists the named components and prior volumes referenced in this work, for citation and cross-reference. All are the intellectual work of the author under the entity Non-Agentic Governance Singapore (ACRA T260229801).

Reference	Description
A12 [P+N]	Deterministic Black Box — terminal governance gate for physical actions
A13a / b / c	Mobility Enclosure / Sensing Limb / High-Payload Clinical Robotic Platform
A7 Bridge Hub	Sole ingress chokepoint; advisory-to-deterministic packet filter
WD117 Ledger	Immutable write-once forensic ledger; SHA-3-512, OTP silicon fuses (design-intent)
U7 Sequencer	One-shot power-severance sequencer; Sovereign Brake™ actuation
WM-Series	Ward-matched clinical modules (WM001, WM003, WM006, WM008, WM009)
P006 / P-001 / P-003	Orange Code Cap / vital-monitor governance / Non-Agentic Filter (portfolio references)
Prior volumes	WISL™ Series, incl. No. 69 (audit objects doctrine) and No. 72 (The Deterministic Complement)

止於至善

謙虛 · 沉默 · 尊严 · 仁

We Innovate Save Lives

Edwin Koh · Non-Agentic Governance Singapore · ACRA T260229801

WISL™ Ebook Series No. 75 · First Edition, 2026