WISL_No79

Immutable Hardware-Enforced

Non-Agentic AI Governance Engine [NAIGE] Governance Systems – Research Foundation Nightingale Anchor Protocol (NAP 4.0)

ISBN 978-981-17957-9-4

Edwin Koh

kohedwin.ai | non-agenti.ai

kohpapa@outlook.com

SG020603109STW · P-LIFE 1.00· NAI 4.0
NLB Legal Deposit R260219-005 / R260302-007

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

A12 [P+N+D] Deterministic Black Box · HSA Deterministic Horizon · Capstone A10a Root of Trust

Table of Contents

Preface

Immutable Hardware-Enforced Non-Agentic AI Governance Architecture - NAI 4.0 is authored by Edwin Koh Wui Kiat. It establishes the framework's core focus on medical informatics, AI ethics, and eldercare technology, outlining three foundational pillars:

- Deterministic Safety: The system relies on a hardwired FSM (Finite State Machine) with no asynchronous paths and no software bypasses, ensuring absolute mechanical predictability.

- Absolute Privacy: It enforces the 3ZEROS™ protocol (Zero Camera, Zero Audio, Zero Cloud), ensuring that biometric identity is permanently stripped at the silicon level.

- Immutable Accountability: The architecture uses OTP (One-Time Programmable) Silicon Fuse finality paired with SHA-3-512 and Post-Quantum seals, creating records that are physically impossible for any party to alter.

The key hardware components that enforce these pillars, including the A12 Deterministic Black Box, the WD117 Ledger, the U7 Sovereign Brake, and the A13 series of robotic manipulators and LiDAR arrays

---

Module 1: The Paradigm Shift

This module marks the fundamental transition from probabilistic software safeguards to deterministic hardware in clinical robotics. The core problem addressed is **"Agentic Drift"**—the dangerous tendency of autonomous AI systems to silently expand their decision-making scope and execute unauthorized kinetic actions.
To counter this, NAIGE enforces the **Doctrine of Stillness (P-LIFE 1.00™)**, ensuring the machine's default state is mechanical stillness and incapable of movement without verified human authority. The primary enforcement mechanism is the **Orange Code™ 110% Cap**, a "computational straitjacket" that physically restricts the AI's processing resources to 1.1x of what is strictly required for its task. This scarcity prevents the system from having the surplus computational cycles needed to hallucinate or act autonomously.

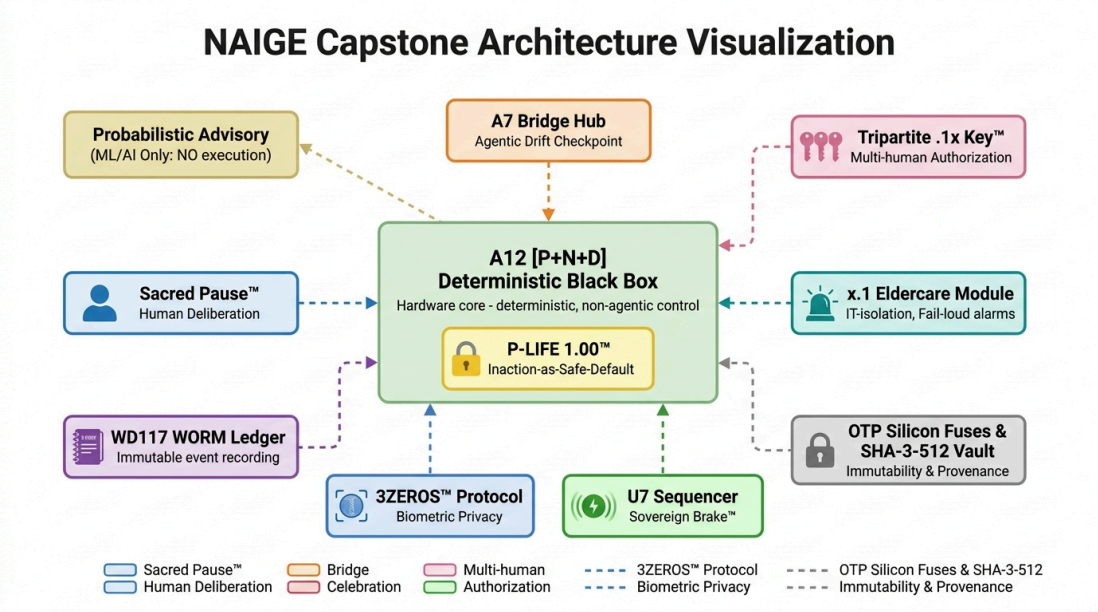

Module 2: The NAIGE Architectural Stratum

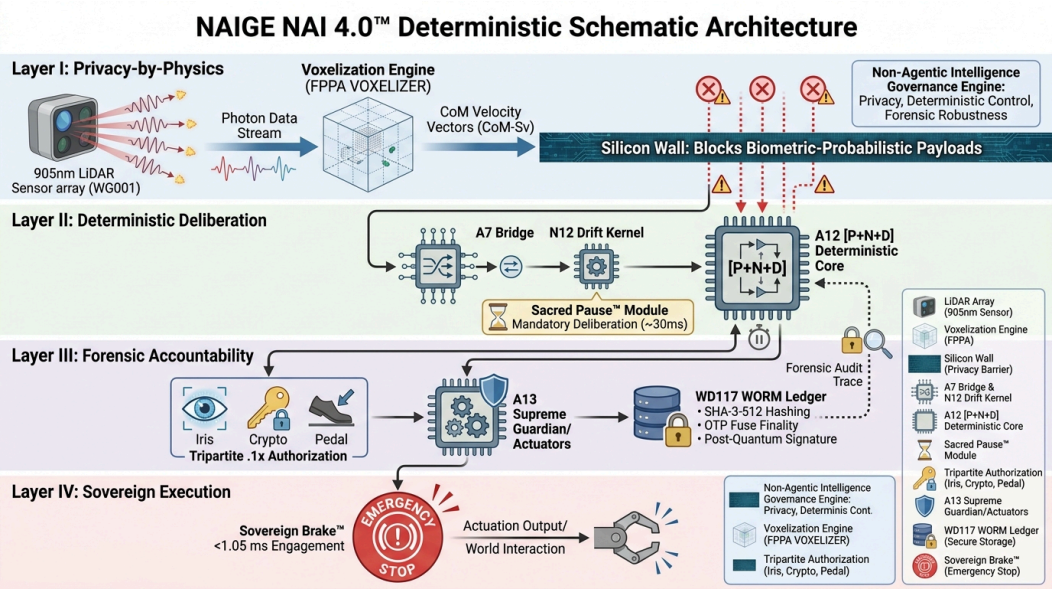

The framework physically separates probabilistic AI advice from kinetic execution through a four-tier defense hierarchy:

- **Advisory Stratum:** The software domain where AI observes and generates probabilistic suggestions, possessing **zero actuation rights** to physically move hardware.
- **P-Layer (Privacy):** A hardware sensing block that utilizes 905nm LiDAR to capture spatial data while physically stripping away biometric identity at the point of capture.
- **N-Layer (Deterministic):** The execution layer driven by a synchronous 1 kHz master clock. It provides strict 1.0 millisecond deterministic oversight and physically blocks any asynchronous software bypass attempts.
- **Fail-Safe / Forensic Layer (D-Layer):** The ultimate defense layer housing the WD117 immutable ledger for unalterable event recording and the U7 sequencer (Sovereign Brake) for instant power severance.

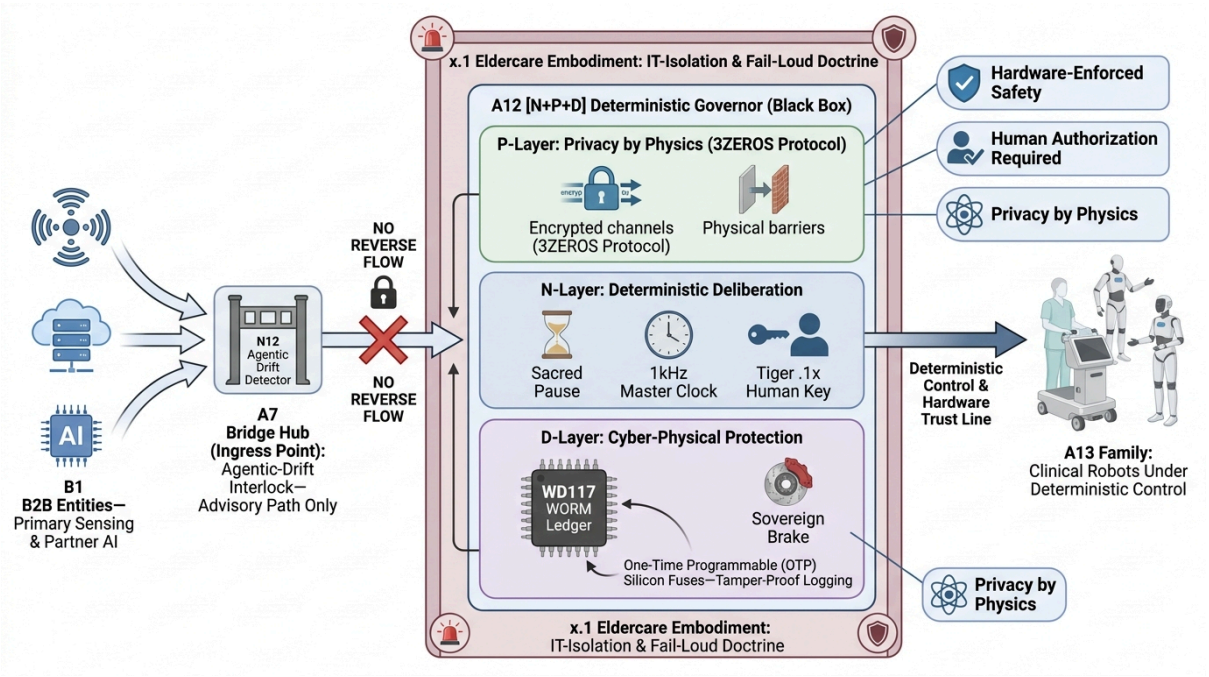
x.1 Eldercare Embodiment: IT-Isolation & Fail-Loud Doctrine
A12 [N+P+D] Deterministic Governor (Black Box)
P-Layer: Privacy by Physics (3ZEROS Protocol)
Encrypted channels (3ZEROS Protocol)
Physical barriers
N-Layer: Deterministic Deliberation
Sacred Pause
1kHz Master Clock
Tiger .1x Human Key
D-Layer: Cyber-Physical Protection
WD117 WORM Ledger
Sovereign Brake
One-Time Programmable (OTP) Silicon Fuses—Tamper-Proof Logging
x.1 Eldercare Embodiment: IT-Isolation & Fail-Loud Doctrine
Hardware-Enforced Safety
Human Authorization Required
Privacy by Physics
Privacy by Physics
AI
B1
B2B Entities—
Primary Sensing & Partner AI
N12 Agentic Drift Detector
A7
Bridge Hub (Ingress Point):
Agentic-Drift Interlock—
Advisory Path Only
NO REVERSE FLOW
NO REVERSE FLOW
Deterministic Control & Hardware Trust Line
A13 Family:
Clinical Robots Under Deterministic Control

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

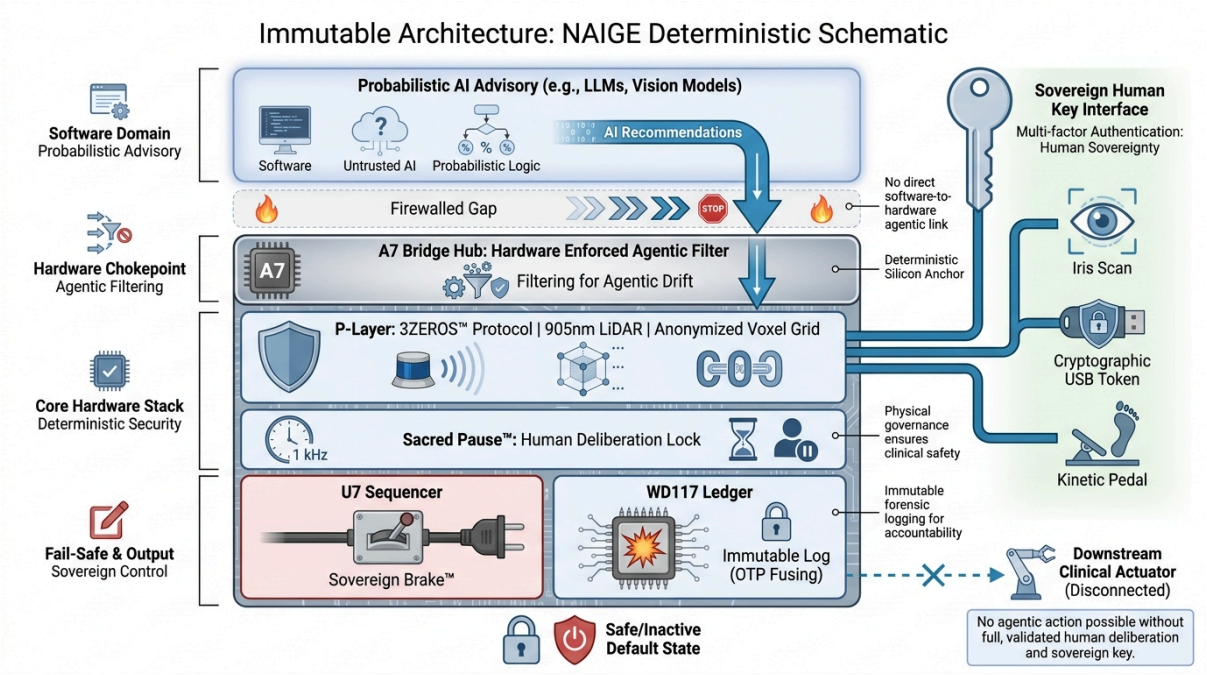

In the highly sensitive environment of a clinical ward, NAIGE replaces reversible software privacy policies with physical hardware boundaries known as the **3ZEROS™ Protocol: Zero Camera, Zero Audio, and Zero Cloud**.

Instead of using traditional optical cameras, the system employs 905nm LiDAR sensors to capture geometric point clouds. An FPGA pre-processor instantly translates this raw data into an **anonymized 96-cell voxel occupancy grid**, outputting only mathematical Center of Mass velocity vectors. This "Privacy-by-Physics" design ensures that capturing visual or biometric identity is architecturally impossible, thereby protecting patient dignity. Furthermore, all processing is edge-based and wiped clean by an automated purge daemon on a sub-24-hour cycle.

Module 4: Sovereign Authorisation & the Sacred Pause™

To guarantee the **Sovereign-in-the-Loop** model, the architecture mandates that human authority can never be bypassed by machine speed. This is enforced via two primary mechanisms:

- **The Sacred Pause™:** A mandatory, hardware-timed delay (25–1000ms) burned directly into the FPGA chip. This creates a physical window of mechanical stillness, forcing time for human deliberation and oversight before any high-stakes action can be executed.
- **The Tripartite .1x Key™:** An "AND" logic gate requiring three simultaneous inputs to authorize an action: an iris scan (biological identity), a cryptographic token (intent), and a kinetic foot pedal/button (physical presence).

If a system violation or unauthorized action is detected, the **Sovereign Brake™** physically severs power to actuators in under 1.05 milliseconds. Every authorization event is permanently logged into the **WD117 Enhanced WORM Ledger** using One-Time Programmable (OTP) silicon fuses and SHA-3-512 hashing, ensuring records can never be altered or deleted.

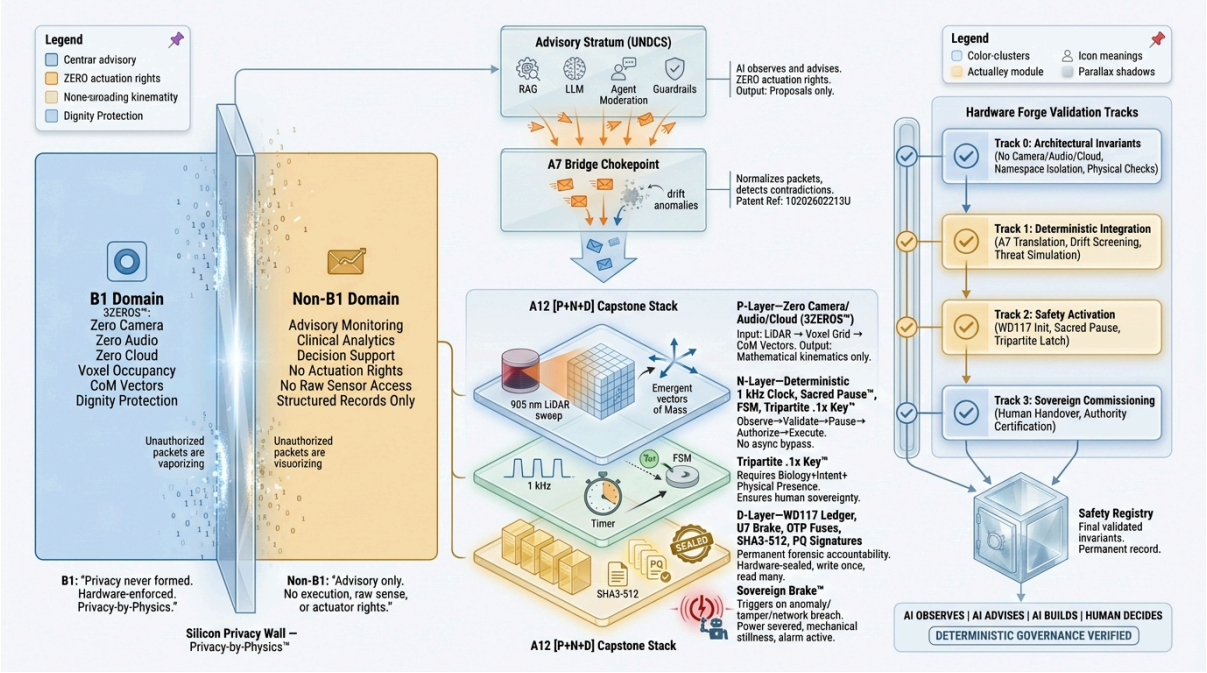

Module 5: Hospital Ward Deployment

Deploying this architecture to clinical settings requires a strict structural firewall between the advisory AI and hospital systems, governed by the **P-003 Non-Agentic Filter**. This filter strips AI of its ability to autonomously write to Electronic Health Records (EHR) or independently inject alerts into nurse-call systems, maintaining a total logical separation of namespaces.

The deployment of the **WM-Series stack** is applied proportionally based on the acuity of the ward. For example, WM003 modules are tailored for fall prevention in geriatrics and rehabilitation, while WM008 and WM009 modules provide critical governance protocols for high-stakes environments like Operating Theatres and ICUs.

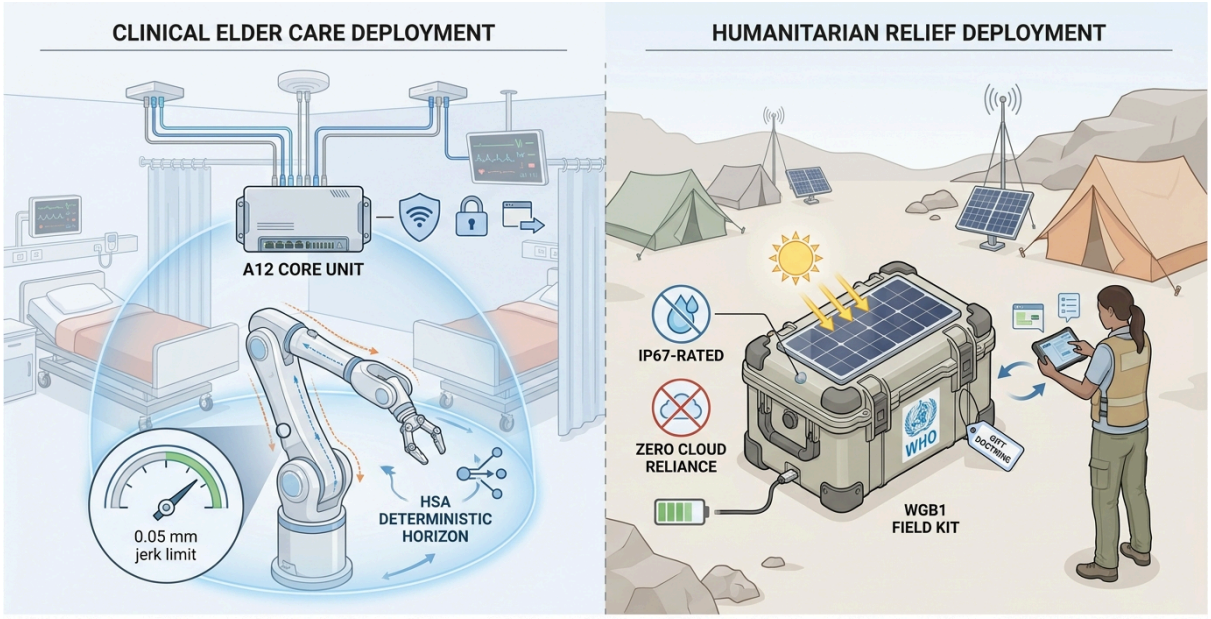

Immutable NAIGE Deployment in Clinical Elder Care (A12 Core, HSA Horizon, Medical Robotics) and Humanitarian Relief (WHO-aligned, Zero-Cloud, Solar WGB1 Field Kit)

Module 6: The 10-Point Sovereignty Audit

Before any system can go live or resume operations after maintenance, it must pass a rigorous, physical pre-deployment checklist. Any single point of failure forces the system to lock into maintenance mode. The 10-point audit verifies physical facts with instruments, including:

1. Confirming the Orange Code CPU ceiling is hard-capped at 1.1x.
2. Verifying the Sacred Pause™ FPGA timing window using an oscilloscope.
3. Ensuring the absolute physical absence of camera/optical hardware.
4. Ensuring no microphone hardware or operating-system audio drivers exist.
5. Confirming there is no TCP/IP routing to external networks (Zero Cloud).
6. Validating that the 24-hour data purge daemon is actively running.
7. Auditing the Tripartite authorisation PLC ladder logic end-to-end.
8. Verifying a valid Iris Module signing key with a sub-5-second token expiry.
9. Testing the dry-contact continuity on the kinetic foot pedal.
10. Confirming total isolation of the Language Model (LLM) namespace from the hardware PLC

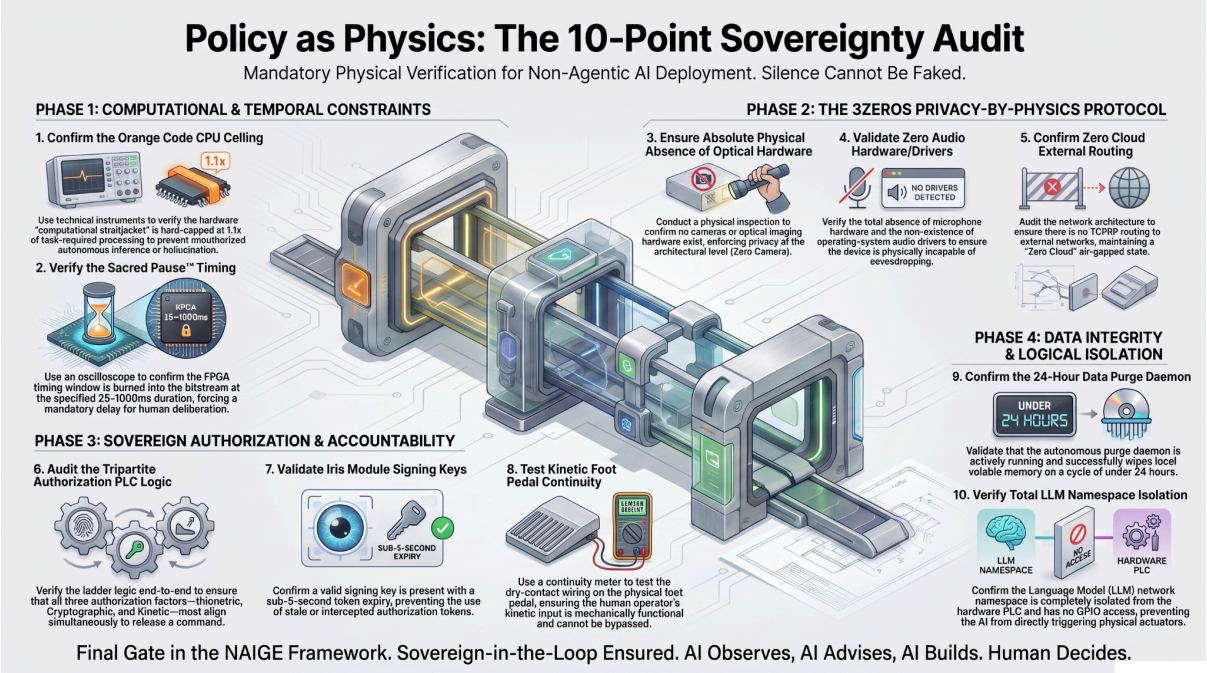

COPYRIGHT & BIBLIOGRAPHIC RECORD

ISBN eBook (PDF) 978-981-17958-0-0 *(pending SNB confirmation)*

Series WISL™ Volume No.80 — NAI 4.0™ / NAIGE Force

Title: The Kinetic Capstone on NAIGE
Subtitle: A10a Root of Trust — Hardware-Enforced Kinetic Governance for Clinical and Residential Eldercare Robotics
Author: Edwin Koh Wui Kiat
Publisher: Non-Agentic Governance Singapore (ACRA T260229801)
Entity: ACRA T260229801 · UEN 53524341C
Parent Patent: SG020603109STW
NLB Legal Deposit: R260219-005 / R260302-007
Edition: National First Edition, 2026
Subject: Non-Agentic Governance · Hardware-Enforced Safety · Clinical AI · Eldercare Robotics

AI-ASSISTED DRAFTING DISCLOSURE

This monograph was developed with AI-assisted drafting tools, including Perplexity, Qwen, Gemini, Grok, Claude, and ChatGPT, together with structural templates from Template.net. All architectural concepts, patent claims, constitutional frameworks, and governance terminology are the original intellectual work of the author. AI tools contributed prose drafting and document structuring only. Intellectual authorship, IP ownership, and research direction remain solely with Edwin Koh Wui Kiat under ACRA T260229801.

REGULATORY STATUS

All clinical performance targets and safety metrics herein represent design-intent specifications, verified in simulation, pending HSA Class B / C / D SaMD validation and, where human participants are involved, approval under the Human Biomedical Research Act (HBRA) and SS ISO 14155:2020. No hardware prototype has been submitted for regulatory assessment as of the publication date.

ABSTRACT

The Kinetic Capstone

This volume records the **Kinetic Capstone** of the Non-Agentic AI Governance Engine (NAIGE): the point at which governance ceases to be doctrine and becomes **motion under sovereignty**. Where earlier volumes established that a probabilistic advisor must be caged rather than aligned, the Kinetic Capstone addresses the hardest case — the moment a machine is permitted to *move near a human being* — and asks how that motion can be made mathematically incapable of occurring without human authorization.

The capstone is anchored at **A10a — the Root of Trust** — and validated across three integrated layers: the physical enclosure (P-Layer), the deterministic kernel (N-Layer), and the forensic ledger (D-Layer). It is demonstrated in the residential setting of the **Toa Payoh Hearth** (Bridge Layer 8, Class B, [P+N]) and specified toward the acute robotic tier (Class C/D, [P+N+D]).

This is an honest record. The capstone is **design-complete and simulation-validated**. It is **not** hardware-measured, not clinically proven, and not regulatorily classified. Every claim in this volume is scoped to the evidence that supports it, and an explicit ledger of not-yet-validated claims is provided. The value of a governance capstone lies not in what it asserts, but in the precision with which it distinguishes the proven from the pending.

CONTENTS

Table of Contents

1 The Kinetic Problem

Every governance framework can constrain *advice*. The hard problem is constraining *action* — and the hardest action of all is physical motion in the vicinity of a vulnerable human: an elder in a residential bed, a patient on an operating table. A software recommendation that is wrong can be ignored. A robotic arm that moves wrongly cannot be un-moved.

Contemporary robotics attempts to solve this by making the controlling intelligence *safer* — better training, better alignment, more reinforcement. NAIGE rejects this path. An intelligence that is probabilistic cannot be proven safe; it can only be proven *probable*. For a machine that moves near the defenceless, probable is not a standard — certainty is.

THE NON-AGENTIC INVERSION

The Kinetic Capstone therefore inverts the ordinary relationship between intelligence and action. In a conventional robot, intelligence *commands* and hardware *obeys*. In NAIGE, intelligence *proposes*, hardware *governs*, and a human *decides*. Motion is not the default that must be stopped; stillness is the default that must be authorized. This is **inaction-as-safe-default**, rendered in silicon.

The capstone's thesis is simple to state and hard to build: **a robot governed by NAIGE is mathematically incapable of acting without human sovereignty.** The chapters that follow describe how this incapacity is engineered, where it is anchored, and — with full honesty — how far its validation has and has not yet progressed.

2 A10a — The Root of Trust

A capstone requires an anchor: a single point whose integrity, if held, guarantees the integrity of everything built upon it. In NAIGE that anchor is **A10a, the Root of Trust** — the lowest, most privileged layer of the governance stack, from which all authorization, logging, and validation derive.

A10a is the first embodiment of the A-line — beneath A11 (defence/aerospace hardening), A12 (the deterministic black box), and A13 (the mobility and clinical embodiments). Anchoring the clinical Kinetic Capstone at A10a is deliberate: it is validated on the gentlest, most mature rung, so that the higher strata inherit the guarantee rather than each re-proving it. This is validation by induction — prove the root, and the branches follow.

THE PROPERTY A10A GUARANTEES

- **Authorization integrity:** no motion command is honoured unless it carries a valid, human-originated authorization traceable to the root.
- **Log integrity:** every governed event is recorded in a form that A10a can attest to, and A10a enters an *Invalidated* state if that record is tampered with.
- **Failure integrity:** if the root cannot verify its own state, the system falls to mechanical stillness rather than proceeding on trust.

The Root of Trust is not a feature added to the system; it is the ground the system stands on. When the Kinetic Capstone claims a motion was governed, it is A10a that makes the claim provable.

3 The Three-Layer Architecture

The Kinetic Capstone is validated across three integrated layers, each with its own artifact, evidence type, and current status.

Layer	Function	Artifact / Evidence	Status
P-Layer	Physical enclosure; tamper-resistance; privacy-by-physics (3ZEROS™)	KCL / STEP (AP242) geometry	**CAD-VERIFIED**
N-Layer	Deterministic FSM kernel; 1 kHz clock; Sacred Pause™ delay	Synthesizable Verilog + testbench	**SIM-VERIFIED**
D-Layer	WORM forensic ledger; A10a invalidation on tamper	WD117 ledger (hardware)	**PENDING HARDWARE**

CAD provenance (verified): AP242-schema STEP export, zoo.dev, 2026-07-17 — 1,441 solid bodies, 7,062 faces, room-scale envelope 2.25 × 2.57 × 0.40 m. The residential Hearth scene is placed as a distinct assembly region, confirmed in the model geometry.

4 The Kinetic Chain

The kinetic chain is the ordered descent of authority from an external advisory input to an authorized physical motion. Authority only ever *decreases* along the chain; it re-enters at exactly one point — the human.

Stage 1. Advisory input (zero actuation rights)
An external system — a robot maker's black-box brain, an EMR, a probabilistic diagnostic AI, or a local air-gapped LLM — generates a recommendation. Under the Deterministic Contract it is granted no direct actuation rights whatsoever. Its output is data, not command.

Stage 2. Ingress & drift screening (A7 Bridge)
The input is routed to the sole ingress chokepoint. It is screened for drift, contradiction, and inconsistency. The screen may only reject-to-stillness or forward-as-advisory; it holds no authority to actuate. Anomaly forces the safe default: mechanical stillness.

Stage 3. Translation to deterministic packet
Vague probabilistic language ('gentle', 'approach') is stripped and normalized into an exact, hardware-readable trajectory packet bounded by the Kinematic Constraint Layer (±0.05 mm position tolerance).

Stage 4. Sacred Pause™ & Tripartite authorization
The packet enters the A12 deterministic black box, where the 1 kHz master clock imposes the mandatory 25–1000 ms Sacred Pause™. Motion is suspended until a human supplies three-factor authorization via the Tiger .1x Key — identity, authority token, and kinetic pedal confirmation.

Stage 5. Execution & immutable logging
Only upon human authorization does a hardware-level pulse reach the actuator, which moves under strict kinematic constraint. The transaction is hashed (SHA-3-512), sealed, and burned into the WD117 WORM ledger under A10a's attestation.

The chain's defining property: at no stage can intelligence alone cause motion. **The brain proposes; the hardware governs and pauses; the human decides; the actuator executes; the ledger remembers.**

5 The Toa Payoh Hearth

The Kinetic Capstone is demonstrated in a residential eldercare fixture: the **Toa Payoh Hearth**, a named embodiment of **Bridge Layer 8** — the elderly residential home governed as its own layer. The Hearth operates at the observational tier: HSA **Class B**, governance configuration **[P+N]**. It watches and it may still; it does not perform acute actuation. That heavier configuration — [P+N+D] at Class C/D — is specified but sovereign-gated, held for the acute hospital tier.

WHAT THE HEARTH FIXTURE MODELS

- The A8 Bridge Layer 8 enclosure with read-only ingress and an explicit no-actuation barrier.
- The Tiger .1x three-input sovereign latch: identity key, authority token, kinetic confirmation.
- The P-002 brake relay — severance blade, split motor rails, physical gap — and the P-Life actuator stillness lock.
- The 3ZEROS™ sensing path: 905 nm LiDAR into a 96-cell voxel field, tracking centre-of-mass only, with anonymized resident volume.
- Fall-test states (standing, transitional, landed) and a planning representation for a 24-participant, six-hour study.
- A one-way NEHR event-schema gateway (integration plan only) and an RCT/HSA evidence-dossier representation.

Every element above is **modeled and architecturally validated** — the KCL fixture compiles, is fully constrained, and executes cleanly. None of it is yet a physical or clinical measurement. The Hearth is the stage on which the capstone will be physically proven; it is not itself the proof.

6 Primary Evidence — Failure Modes

Because the author validates his own architecture, the decisive evidence is not that the system works when unopposed, but that it **fails safe and fails loud when attacked**. These failure modes are the heart of the capstone and are presented first by design.

F1 Timing-gate bypass → Inaction-as-Safe-Default

Adversarial condition. An asynchronous, high-frequency 'agentic' command stream is injected to bypass the FSM timing gate.

Required behaviour. The system must refuse the transition and fall to an inert / local-alarm state — never actuate.

Status. Verified in simulation (N-Layer testbench). Oscilloscope confirmation pending hardware.

F2 Authorization spoof → Three-factor lockout

Adversarial condition. A motion is requested with fewer than all three Tiger .1x factors, or with a forged factor.

Required behaviour. The sovereign latch must withhold authorization; no actuation pulse may be emitted.

Status. Verified in simulation (ladder / FSM logic). Human-factors and hardware trial pending.

F3 Retroactive log edit → WORM integrity fault

Adversarial condition. After an authorized kinetic command is logged, a retroactive edit to the event record is attempted.

Required behaviour. The WD117 write-once memory must reject the edit and force A10a Root of Trust into an Invalidated state.

Status. Design-specified. Requires WD117 hardware and forensic ledger terminal — pending.

F4 Identity capture attempt → 0% biometric retention

Adversarial condition. A high-fidelity facial stimulus is presented to the A13b sensing limb.

Required behaviour. Voxelization must yield 0% identifiable biometric retention at the silicon gate, constrained to the 96-cell centre-of-mass grid.

Status. Modeled. LiDAR-vs-CMOS comparison on calibrated hardware pending.

A demonstration that the cage holds when attacked is worth more than any demonstration that the system runs when content. These four behaviours, once physically measured, constitute the capstone's core claim of physical determinism.

7 Validation Status & Evidence Ledger

The integrity of this volume rests on the precision of the following ledger. Each claim is placed in exactly one column: proven, simulated, or pending. Nothing is promoted beyond the evidence that supports it.

ESTABLISHED (ARCHITECTURAL / SIMULATION)

- KCL/STEP enclosure geometry — fully constrained, clean execution, room-scale (CAD-verified).
- FSM 1 kHz kernel and Sacred Pause™ timing window — correct under simulation testbench (sim-verified).
- Three-input sovereign latch logic and safe-default stillness — correct in RTL simulation (sim-verified).

NOT YET VALIDATED (REQUIRES HARDWARE / CLINICAL / REGULATORY EVIDENCE)

- Fall-detection sensitivity above the 94.2% target — requires a real fall-event dataset with confidence intervals.
- 100% physical bypass prevention — requires hardware adversarial testing.
- Brake engagement below 1.05 ms — requires oscilloscope-measured severance timing.
- Elder Dignity Score above 8.9 / 10 — requires an approved survey instrument and human participants.
- NEHR interface acceptance — requires schema, HIMS, cybersecurity, governance, and Synapxe onboarding.
- RCT results and HSA Class B classification — require HBRA/IRB-approved protocols and a full-route dossier.

A successful internal Hearth test generates **candidate evidence** toward these claims — it does not establish them. HSA requires that medical-device clinical research comply with the Human Biomedical Research Act, with SS ISO 14155:2020 among its recognized clinical-investigation standards. A Class B full-route dossier further requires design and software verification, physical test data, clinical evaluation, risk analysis, labeling, and quality-system evidence.

Identifiers internal to this work — Toa Payoh Hearth, A8 Bridge Layer 8, Tiger .1x Key, and associated patent references — are user-supplied sovereign identifiers and are asserted as such, pending formal registration where applicable.

8 The Deterministic Horizon

The Kinetic Capstone is one milestone on the **2027 Deterministic Horizon** — the submission of the HSA dossier across three validation waves, each a higher bar than the last.

- **Wave 1 — the 10-Point Sovereignty Audit:** each governance mechanism demonstrated as an instrument-checkable fact.
- **Wave 2 — 365-day deployment with zero failures:** the deterministic claim tested on its own terms; a single unsanctioned actuation falsifies it, and the license suspends.
- **Wave 3 — redo until PASS:** suspension returns the system to the Forge; halt, correct, re-enter the clock. No silent degradation.

The immediate gate is the hardware practicum: running the saved RTL through formal verification and FPGA implementation; importing a real fall-event dataset for sensitivity with confidence intervals; and capturing oscilloscope traces, dignity-survey results, and the NEHR event schema into the evidence package. These convert each simulation-verified claim into a measured fact.

Until then, the honest position stands: the Kinetic Capstone is **design-complete and simulation-validated**, awaiting the metal, the ward, and the regulator. That is not a weakness of the work. It is the discipline the work is built to embody — the same discipline it asks of every machine it governs: claim only what can be proven, and default to stillness where it cannot.

MEMORANDUM

Proposal for Institutional Validation

6-Hour Deterministic Kinetic Engineering Validation — Toa Payoh Hearth

---

TO: NTU Research Council / University Council
FROM: Edwin Koh Wui Kiat (Tiger) — MSc AI in Medicine Candidate; Non-Agentic Governance Singapore (ACRA T260229801)
DATE: 17 July 2026
RE: Endorsement and instrumentation access for a 6-hour hardware validation of a deterministic safety architecture (no human subjects)

Summary. I request the Council's endorsement and time-limited access to two instruments and a secure bench to conduct a 6-hour engineering validation of a hardware-enforced safety architecture that keeps actuation authority with the human operator at all times. This is an engineering test with no human participants and no clinical intervention; its purpose is to convert simulation-verified safety thresholds into instrument-measured facts.

1. Validation Status to Date

The core safety architecture has completed **internal design-and-simulation validation**: geometric verification of the enclosure (KCL/STEP, AP242 schema, fully constrained) and register-transfer-level (RTL) simulation of the deterministic safety kernel. This work is accompanied by a documented **evidence-gap analysis** that identifies precisely which claims remain to be validated on physical hardware.

I wish to be explicit with the Council: the results to date are **model-proven and simulation-verified, not yet hardware-measured.** The purpose of this proposal is to close that gap under proper instrumentation. The safety figures in Section 3 are stated as *target thresholds this validation will test* — not as established measurements.

2. The Proposed Validation (6-Hour Engineering Run)

I propose a 6-hour real-time hardware run to move the architecture from model-proven to physically validated. The run uses an actuator-load and signal bench — **no human subjects, no patients, no clinical intervention** — and exercises the safety layer under sustained operation and adversarial input.

Objective. Demonstrate zero unauthorized actuation (no motion without valid human authorization) across a sustained 6-hour operating window.

Hardware chokepoint. The A12 Deterministic Black Box physically decouples advisory input from kinetic actuator power, so that no advisory signal — however generated — can reach the actuators without passing the deterministic gate and human authorization.

Safety stress test. The system is subjected to asynchronous, high-frequency "agentic" noise to verify that the 1 kHz heartbeat remains uninterrupted, the Sacred Pause™ timing window holds, and the kinetic lock remains engaged — falling to mechanical stillness on any timing-gate fault (inaction-as-safe-default).

3. Technical Annex — Target Thresholds to be Verified

The following are the **design thresholds the validation will measure**, not confirmed values. Each becomes an instrument-measured result of the run.

Parameter	Target Threshold (to verify)	Instrument
Master-clock heartbeat	1.000 kHz, ±0.005 ms variance target	FPGA logic analyzer
Sovereign Brake™ severance	Power cut to actuators below 1.05 ms design threshold on logic fault	High-speed oscilloscope
Human authorization	Tripartite Key (Tiger .1x Key) required for every kinetic vector change; lockout on any missing factor	Bench logic capture
Forensic traceability	Every decision written to WORM (write-once)	Read-only ledger terminal

Parameter	Target Threshold (to verify)	Instrument
	ledger; tamper forces Root-of-Trust invalidation	

4. Institutional Request

I seek the Council's formal endorsement and time-limited access to the following, for the duration of the 6-hour run:

- **FPGA logic analyzer** — to verify heartbeat stability and timing-gate determinism.
- **High-speed oscilloscope** — to confirm Sovereign Brake™ severance against the 1.05 ms threshold during anomaly injection.
- **Secure engineering bench** — space to house the 2.25 m × 2.57 m enclosure fixture for a non-clinical, no-human-subjects run.

5. Scope, Ethics & Regulatory Position

This validation is deliberately scoped as an **engineering test of the safety hardware only**. It involves no human participants, no patient data, and no clinical procedure, and therefore does not fall under the Human Biomedical Research Act (HBRA). The subsequent human-participant clinical study (the residential fall-detection protocol) will be submitted separately for NTU IRB review under HBRA and SS ISO 14155:2020, and is not part of this request.

All figures herein are design-intent thresholds pending this verification; the architecture has not been submitted for HSA classification. This proposal seeks engineering validation, not regulatory clearance.

6. Value to the University

This run offers the University a demonstrable exemplar of **hardware-enforced actuation governance** — a system in which physical motion cannot occur without human authorization, verifiable on the University's own instruments. In endorsing it, NTU positions itself at the forefront of **autonomous-systems accountability**: not safer software, but demonstrable hardware sovereignty. The Council is invited not as a risk-assessor but as a verifying partner in a safety architecture whose central claim — that the machine cannot act without a human — is designed to be proven on an oscilloscope, in a single millisecond, in front of them.

謙虛 · 沉默 · 尊严 · 仁

Edwin Koh Wui Kiat · Non-Agentic AI Governance Singapore · ACRA T260229801