

Immutable Hardware-Enforced Non-Agent AI Governance Engine [NAIGE] Governance Systems – Regulatory Compliance Framework: Aligning NAIGE Robotic Systems with HSA Clinical Safety Standards

NAIGE: Hardware-Enforced Sovereignty in Clinical Robotics

NAIGE replaces probabilistic software guardrails with deterministic hardware controls to ensure absolute patient safety and human authority through Hardware-Enforced Kinetic Governance.

The Three Layers of Hardware Governance

N-Layer: The Sacred Pause™
A mandatory hardware-timed delay (25–1000ms) that forces human deliberation before any action.

D-Layer: WD117 Immutable Ledger
Uses one-time programmable silicon fuses to create a permanent, tamper-proof record of events.

WD117 Immutable Ledger

Physical Power Severance (<1.05ms)
Sovereign Brake™

P-Layer: Privacy-by-Physics (3ZEROS™)

Strips biometric identity at the silicon level using LIDAR instead of optical cameras.

Zero Camera Zero Audio Zero Cloud

The Kinetic Chain of Human Sovereignty


The Sovereign Brake™
Instantly severs power to actuators if any system anomaly or drift is detected.

Power Severed (<1.05ms)

Inaction-as-Safe-Default
Machines default to mechanical stillness; AI has zero direct actuation rights.

Tiger .1x Tripartite Authorization
Requires simultaneous Iris scan, Cryptographic token, and Kinetic foot pedal to move.

The Sovereign Brake™
Instantly severs power to actuators if any system anomaly or drift is detected.



Edwin Koh Wui Kiat

kohedwin.ai | non-agentic.ai

kohpapa@outlook.com

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

ISBN Paperback 978- 981-17963-9-5

Copyright

Title: Immutable Hardware-Enforced Non-Agentive AI Governance Architecture — NAI 4.0
Subtitle: Hardware-Enforced Non-Agentive AI Governance Architecture — NAI 4.0
Series: WISL e-book Series No. 77 to No. 80 — Integrated ISBN Edition
Publisher: Edwin Koh
Author: Edwin Koh Wui Kiat
Entity: ACRA T260229801
Parent Patent: SG020603109STW
NLB Legal Deposit: R260219-005 / R260302-007
Edition: National Second Edition, 2026
Format: Paperback
Language: English
Subject: Artificial Intelligence · Non-Agentive Governance · Hardware-Enforced Safety · Clinical AI

Classification: Medical Informatics · AI Ethics · Embedded Systems · Eldercare Technology

Immutable— Hardware-Enforced Non-Agentive AI Governance Architecture — NAI 4.0 · NAI 2.0
A12 [P+N+D] Deterministic Black Box · HSA Deterministic Horizon · Capstone A10a Root of Trust
WD117 Ledger · U7 Sovereign Brake · Sovereign-in-the-Loop
A13a KCL FPGA · A13b LiDAR array · A13c Deterministic High-Payload Robot Manipulator

DETERMINISTIC SAFETY	ABSOLUTE PRIVACY	IMMUTABLE ACCOUNTABILITY
Hardwired FSM state machine. No asynchronous paths. No software bypass.	3ZEROS™ — Zero Camera · Zero Audio · Zero Cloud. Identity stripped at silicon.	OTP Silicon Fuse finality. SHA-3-512 + Post-Quantum seal. Unalterable by any party.

Copyright © ™ Notice

© 2026 Non-Agentive Governance Singapore (ACRA T260229801). All rights reserved.

No part of this publication may be reproduced, distributed, transmitted, or stored in any retrieval system in any form or by any means — electronic, mechanical, photocopying, recording, or otherwise — without the prior written permission of Non-Agentive Governance Singapore, except for brief quotations for academic citation, review, or research purposes with full attribution.

All trademarks, (Classes 9, 42, 10, 44, 45), 3ZEROS™, Sacred Pause™, ABC+2S+H™, ABC+2S Guardian™, ABC+2S+H Framework™, ABC+2S+H Guardian™, Nightingale’s Eyes Care™ are proprietary to Non-Agentive Governance Singapore™ and protected under Singapore patent law and trademark registration.

All architectural and terminology including concepts including Tiger .x1key, P-Life 1.00, Sovereign Brake, Kinetic Foot Pedal/Button, Tripartite Latch, Orange Code 110% Cap are proprietary of Non-Agentive AI™ and Non-Agentive Governance Singapore™

This monograph was developed with AI-assisted drafting tools, including Perplexity™, Qwen™, Gemini™, Grok™, and Claude™, ChatGPT™ as well as 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.

All clinical performance targets and safety metrics herein represent design-intent specifications pending HSA Class B/C/D SaMD validation. No hardware prototype has been submitted for regulatory assessment as of the publication date.

NAIGE Component Registry: Hybrid N12 + Non-B1 + x.1 Eldercare Embodiment [Explainer]

Non-Agentic AI Governance Engine (NAIGE) — Immutable Kinetic Deterministic Core (B1 Mandatory) | Agentic Governance Layer: Non-B1 Limited Monitoring Only | x.1 Eldercare Embodiment

The **NAIGE framework** introduces a revolutionary **non-agentic AI governance** system designed to ensure **absolute clinical safety** by prioritising hardware-locked certainty over unpredictable software. This architecture treats artificial intelligence strictly as an **advisory instrument**, stripping it of all independent decision-making power through a **deterministic hardware layer** that requires human verification for any physical movement. Key safety pillars include the "**3ZEROS™**" protocol for total patient privacy and a **sovereign break** mechanism that physically severs power during system violations. Data integrity is maintained via **immutable silicon fuses** and a hardware-based ledger, creating a permanent, tamper-proof record of all medical events. Specialized components like the **A13 robotic platform** and the **Non-B1 monitoring shell** apply these rigorous standards to patient transport and diagnostics. Ultimately, the system restores **human sovereignty** in high-stakes healthcare environments by replacing statistical AI inference with **mechanical certainty**.

Note: True Kinetic Deterministic requires B1 Kinetic Capstone + P-001 ABC+2S+H Framework™. Non-B1 is advisory-only (Agentic Governance Layer, limited). All other 'deterministic' claims without B1 = Null Enclosure (non-clinical).

Component Code	Module Name	Governance Layer	Technical Function	Safety Mechanism	Hardware/Software Specification	Clinical Application
NAIGE	Capstone / Core 4	N Layer (Deterministic)	Synchronous data processing via silicon capstones	Sacred Pause (Mandatory deliberation delay); 100% trajectory deviation blocking	Hardwired synchronous core 4 silicon; 1 kHz clock speed	Preservation of patient autonomy; Elder dignity score maintenance (>8.9)
A13	Supreme Guardian / Medical Mobility Enclosure	Deterministic Stratum (P and N Layers)	Patient transport with hardwired kinematic constraints; enforces P-LIFE 1.00 mandate	Hardwired kinematic constraint layer; unreachable by remote bypass; non-agentic oversight; human sovereignty	A13 family architecture; hardware locked safety layer	Patient transport and general clinical/humanitarian governance
A13a	Enclosure	N Layer (Deterministic)	Physical constraint enforcement	Kinematic constraints	Hardware rooted mandate	Physical safety in robotic interaction
A13b	Sensing Limb Deterministic	N Layer (Deterministic)	Patient monitoring via Lidar and spatial voxel grids with kinematic vector output	3ZEROS™ (biometric stripping)	Lidar and spatial voxel grids; silicon-level anonymization	Continuous patient monitoring and dignity protection in residential/nursing care
A13c	Clinical Robotic Platform	Deterministic Stratum (N Layer / Advisory P Layer)	Execution of clinical tasks via high payload manipulation, hardware comparators, and voxel grid translation	Kinematic constraints; Three Zeros; strict jerk limits; deterministic horizon compliance; inaction as safe default; Tripartite Tiger .1x key access latch	A13 architecture; hardware comparators; Lidar (LAR) clouds to voxel grids; hardware execution layer; meets HSA class B, C, and D standards	B2B clinical intervention; eldercare nursing home wards; physical safety in medical environments

A12 [P+N+D]	Deterministic Black Box (Terminal Gate)	Hardware Core (P, N, and D Layers)	Serves as the universal deterministic governor, enforcing human authorization and mediating all downstream physical actions.	Sacred Pause™ (25–1000ms mandatory hardware delay) and Tripartite Latch (Iris + Token + Pedal) requiring multi-factor human commit.	1 kHz hardware master clock, FSM Safety Kernel (± 0.05 mm jerk/acceleration limits), silicon-level isolation.	Secure clinical execution and task verification across general healthcare and eldercare platforms.
x.1	Eldercare Cyber-Physical Protection Embodiment	Patent D Governance Enclosure (Forensic / Fail-Safe Layer)	Completely isolates the care apparatus from the facility's IT network (e.g., Wi-Fi, nurse-call systems) to prevent "IT-Gap Drift".	Fail-Loud Doctrine: Interruption of the hardware heartbeat instantly asserts the Sovereign Brake and triggers a local alarm.	IZEKOS™ galvanic isolation, independent hardware alarm (non-IT-dependent), and Tamper Detection Envelope.	Protects vulnerable nursing home residents from cyberattacks, ransomware, or spoofed alerts in the clinical network.
U7	Sequencer (Sovereign Brake)	Forensic / Fail-Safe Layer	Instantly severs physical power to the system actuators upon any detected breach, anomaly, or lack of authorization.	Inaction-as-Safe-Default: Defaults the system to a state of mechanical stillness (P-LIFE 1.00 constant) rather than degraded autonomy.	Hardware-based One-Shot Sequencer that physically drops LDO power rails in < 1.05ms.	Emergency stop enforcement for patient mobility enclosures and clinical robotic arms.
A11	Bridge	N Layer (Deterministic)	Secure IT to OT translation	Airgapped protection	Jupiter hardened enclosure	High-security clinical domains
A7	Bridge Hub / Translation Layer	N Layer (Deterministic Chokepoint)	Acts as the sole ingress chokepoint, translating semantic advisory data into deterministic hardware features.	Agentic Drift Screening: Filters incoming data for contradictions, uncertainty, and confidence scoring before admitting it to the A12 core.	Hardware execution layer utilizing pRegisters for national infrastructure coordination.	Safely coordinates external B1/OEM AI systems with the clinical hospital ward.
WD117	Enhanced WORM Ledger	Forensic / Immutability Layer (D Layer)	Maintains a continuous hardware heartbeat with the clinical apparatus and records every single system event and authorization.	OTP Silicon Fuses: Once a record is written, the circuit breaks, making it physically impossible for attackers or AI to alter the history.	Physical Write-Once-Read-Many (WORM) memory, SHA-3-512 hashing, and post-quantum cryptographic signatures.	Legal defensibility, compliance audits, and absolute forensic traceability for clinical actions.
WG-B1	Humanitarian Field Kit	Edge Layer	Zero cloud execution and deliberation in disaster zones	Humanitarian shield; Sovereign Brake; Sacred Pause™	Solar assisted edge hardware	Global humanitarian efforts; disaster relief; remote field use
P Layer	Privacy Layer Hardware Core	Hardware Core	Strips biometric identity from data streams	3ZEROS™ (Zero camera, zero audio, zero cloud)	905nm Lidar array; anonymized kinematic voxel grids	Patient privacy and telemetry processing
N Layer	Deterministic Layer	Deterministic Stratum	Mechanical synchronous deliberation	Sovereign Brake (physical power severance); Sacred Pause	1 kHz hardware master clock; hardware-based sequencer	Clinical robotic trajectory oversight
D-Layer	Cyber-Physical Defence / Immutable Hardware Vault	Forensic / Fail-Safe Layer (Bedrock Tier)	Protects the system against severe cyber-attacks, physical tampering, and deep-space environmental hazards; physically isolates domains and enforces a "Brake-by-Default" posture.	Sovereign Brake(fail-safe power severance on anomaly); Tripartite Latch authorization; Tamper Detection Envelope; "Fail-Loud" local hardware alarms.	Faraday cages; Triple Module Redundancy (TMR) for Single Event Upset correction; Error-Correcting Code (ECC) memory; IZEKOS™ galvanic and spatial isolation.	Deep-space and defense resilience (radiation/EMP hardening); x.1 Eldercare Embodiment , which completely isolates clinical care apparatus from facility networks to prevent "IT-Gap Drift" and cyber-breaches

B1	Primary Sensing Unit	Advisory Stratum (Core Deterministic)	Primary sensing and data ingestion	Sacred Pause; Lidar array / Point cloud translation	Anonymized kinematic vectoring	Sensing and data ingestion for clinical environments
Tripartite .1x Key/ Tripartite Latch	Human Authorization Layer / D-Layer (Defence)	Mandatory multi-factor human commitment gate; access control for high-risk clinical actions	Iris Scan + Cryptographic Token + Kinetic Foot Pedal/Button	Sovereign-In-The-Loop (SITL) mandatory physical authorization	Human sovereignty; prevents autonomous kinetic action	Authorization of kinetic surgery, caregiving apparatus, or high-risk clinical intervention
KCL	Kinematic Constraint Layer	Deterministic (N Layer)	Governance of mobility enclosures and robotic platforms	Jerk limits and positional tolerances	Hardwired state machine	Prevention of physical care risks and mechanical injury
N12	Agentic Drift Detection Kernel	N Layer (Navigation / Control)	Performs internal drift checks and contradiction screening on translated advisory inputs within the A12 Black Box.	Blocks any corrupted or anomalous AI pathways from reaching the Sacred Pause™ deliberation gate.	Hardwired deterministic algorithm embedded within the A12 core.	Ensures AI-generated advisories do not cause "Authority Drift" during eldercare monitoring.
A5	Agentic Drift Mitigator	Advisory Stratum	Mitigates agentic drift in probabilistic models	UNDCS moderation / Bridged RAG layers	Multi-agent self-monitoring software	Restricting open-ended generation
Non-B1	Limited Agentic Monitoring Shell	Advisory Stratum (Agentic Governance Layer)	Provides clinical decision support and generates probabilistic recommendations based on patient data.	Zero actuation rights; physically isolated from kinetic control, meaning it can only advise and never actuate.	Unconfined non-deterministic software (RAG advisory) that interfaces exclusively via the A7 Bridge.	High-risk clinical monitoring (e.g., cardiac, oncology, x-ray analysis) where AI serves as an advisory tool only.

Non-Agentic AI Governance Engine (NAIGE), a framework designed to ensure **absolute safety** in clinical and humanitarian robotics. This architecture moves beyond software-based oversight by implementing **hardware-enforced controls** that prevent artificial intelligence from acting without human verification. Key features include **deterministic hardware gates**, such as the "Sacred Pause" and "Sovereign Brake," which physically restrict movement and ensure **human sovereignty** over all kinetic actions. Privacy is maintained through the "**3ZEROS™**" protocol, utilizing LiDAR-based sensing to strip biometric identity at the silicon level. Furthermore, the system guarantees **immutable accountability** by recording all events on tamper-proof, one-time programmable silicon ledgers. Ultimately, these sources outline a **fail-safe methodology** that treats AI as an advisory tool while anchoring authority in verified hardware rather than probabilistic code.

Non-B1: Limited to advisory monitoring (cardiac/cancer/x-ray). Feeds A7 → A12 for drift governance by N12. Human decides via A13. No actuation/OEM/B2B. B1 mandatory for all deterministic kinetic products.

x.1 Eldercare Embodiment (also referred to as **Addendum x.1**) is a specialized application of the NAIGE framework's Patent D architecture, specifically designed to protect vulnerable clinical and residential care environments from cyber-physical threats.


While the core Patent D architecture is built for defense and deep-space cyber-resilience, the x.1 embodiment "re-aims the cyber-effects governor at the care estate". It acts as an absolute, hardware-enforced barrier designed to ensure that clinical robotics and monitoring systems can never be weaponized or silenced by external cyberattacks.

Here are the critical functions and structural protections of the x.1 Eldercare Embodiment:

- **Total IT-Isolation and Mediation:** The primary vulnerability in modern clinical settings is "**IT-Gap Drift**"—the risk that a facility's IT network (such as Wi-Fi, nurse-call systems, or electronic medical records) is compromised by ransomware or malicious actors, turning a monitoring device into a hazard. The x.1 embodiment solves this by ensuring that every interface between the care apparatus and the external IT network is mediated by a deterministic gate processor. As a result, the care apparatus has **no remotely reachable command surface**; even if an attacker completely takes over the hospital's network, they cannot instruct, alter, or silence the robotic device.
- **The Hardware Heartbeat:** The system utilizes the Enhanced WD117 Ledger to maintain a continuous, physical "hardware heartbeat" with the clinical sensing apparatus.
- **The "Fail-Loud" Doctrine:** If this hardware heartbeat is interrupted for any reason—whether through a network intrusion, power loss, or physical tampering—the system enforces a "Fail-Loud" posture. This means the device will instantly assert the **Sovereign Brake™**, returning the robot to a safe, inert state of mechanical stillness.

- **Non-IT-Dependent Local Alarms:** Crucially, upon a heartbeat failure, the x.1 embodiment triggers a local hardware alarm that operates entirely independently of any network, server, or software service. The foundational philosophy here is that **"silence cannot be faked"**. A hacker might disable a hospital's digital alert system, but they cannot stop the physical, localized alarm from sounding to alert care staff.
- **Privacy-by-Physics:** The monitored apparatus captures zero photographic, video, or audio data, ensuring that even if the surrounding IT estate is breached, there is absolutely no personally re-identifiable information available to be stolen.

In summary, the clinical floor represents the highest vulnerability surface for robotic systems because it physically interfaces with untrusted networks and human environments. The **x.1 Eldercare Embodiment** guarantees that vulnerability to a network attack does not equal vulnerability to physical harm. It structurally converts any cyberattack or anomaly into a safe default state, ensuring absolute human sovereignty and patient safety



Generated under NAP 4.0 | Tiger Sovereign Blueprint | We Innovate Save Lives

Microorganism data and related laboratory information are managed as controlled clinical records subject to deterministic provenance, restricted access, immutable logging, and mandatory human review prior to any consequential use. The system does not autonomously diagnose, prescribe, or override clinicians; it supports traceable validation and governance of specimen-related data within defined safety boundaries.

1. Foundations of Non-Agentive Governance in Clinical Robotics

In the current landscape of AI-driven healthcare, the integration of probabilistic models into kinetic medical devices presents a unique regulatory challenge for the Health Sciences Authority (HSA). Traditional software-based "alignment" is fundamentally insufficient for HSA Class B, C, and D devices where kinetic risk is present. While software guardrails can filter textual output, they cannot provide the Safety Integrity Level (SIL) required when a robotic actuator operates in proximity to a patient. The Non-Agentive AI Governance Engine (NAIGE) framework addresses this by moving governance away from software instructions and into hardware-enforced physical motion, establishing the P-LIFE 1.00™ doctrine.

The core of this strategy is the "Non-Agentive Inversion." In conventional robotics, intelligence commands and hardware obeys. NAIGE inverts this relationship: intelligence proposes, hardware governs, and a human decides. By stripping AI of actuation rights, we directly address HSA concerns regarding "Agentive Drift"—the risk of an autonomous system silently expanding its decision-making scope.

Evaluation of Governance Doctrines: P-LIFE 1.00™ Standards

Feature	Probabilistic Advisor (Traditional AI)	Deterministic Governor (NAIGE)
Primary Goal	Statistical Alignment	Kinetic Certainty / P-LIFE 1.00™
Default State	Autonomous Action (Stop on error)	Inaction-as-Safe-Default (Move on authorization)
Control Layer	Software Guardrails	Hardware-Enforced Silicon Gates
Risk Profile	High: Susceptible to "Agentive Drift"	Controlled: Hardware-locked constraints
HSA Pathway	Software as a Medical Device (SaMD)	Hardware-Rooted Clinical Governance

This philosophical shift provides the deterministic foundation for the architectural layers required for clinical certification, transforming AI from an autonomous agent into a strictly advisory instrument.

2. The Core Hardware Architecture: A10a Root of Trust and A12 Black Box

The NAIGE framework anchors all clinical safety claims in the **A10a Root of Trust**. This is the lowest, most privileged layer of the governance stack, serving as the physical anchor where governance becomes motion under human sovereignty. For a regulatory dossier, the A10a provides three critical guarantees that provide mathematical certainty of safety:

1. **Authorization Integrity:** No motion command is honored unless it carries valid, human-originated authorization traceable to the root.
2. **Log Integrity:** Every governed event is recorded; the A10a enters an "Invalidated State" if the record is tampered with.
3. **Failure Integrity:** If the root fails self-verification, the system defaults to mechanical stillness.

Governance is executed via the **A12 [P+N+D] Deterministic Black Box**, which acts as the universal governor. Within this core, the **Orange Code 1.1x Cap** serves as a vital safety pillar: if the AI attempts to exceed 110% of its hard-coded CPU resource ceiling (signaling a potential hallucination or "agentic" runaway), the hardware monitor triggers the Sovereign Brake™ to sever power instantly.

Technical Layers of the A12 Deterministic Black Box

Layer	Technical Function	Hardware Specification
P-Layer (Privacy)	Strips biometric identity at the silicon level via 3ZEROS™ protocol.	905nm LiDAR array; 96-cell voxel occupancy grid.
N-Layer (Navigation)	Enforces deterministic deliberation and human-in-the-loop authorization.	1 kHz hardware master clock; FSM Safety Kernel; Orange Code 1.1x Cap.
D-Layer (Defense)	Provides forensic accountability and cyber-physical protection.	WD117 WORM Ledger; OTP Silicon Fuses; SHA-3-512 + Post-Quantum seals.

3. HSA Classification Pathways: Class B, C, and D Alignment

NAIGE employs a tiered risk approach, scaling from residential monitoring to acute intervention. We utilize a "Validation by Induction" strategy: by proving the integrity of the A10a root at the residential tier (Class B), the regulatory burden for high-risk acute applications (Class C/D) is reduced, as the branches inherit the safety of the root.

Class B Requirements: The Residential Tier

The **Toa Payoh Hearth** (Bridge Layer 8) serves as the primary embodiment for Class B residential monitoring.

- **Focus:** [P+N] layers and the 3ZEROS™ sensing path.
- **Mechanism:** LiDAR-based voxelization tracks center-of-mass only (96-cell grid), ensuring 0% biometric feature retention.

Class C/D Requirements: The Acute Hospital Tier

Acute clinical robotic platforms, such as the **A13c**, require the full [P+N+D] stack to manage high-payload risks.

- **Kinetic Constraints:** The A13c platform enforces strict **±0.05 mm jerk and acceleration limits** via a hardwired Kinematic Constraint Layer (KCL).
 - **Forensics:** Mandatory integration of the **WD117 WORM Ledger** (HSA SaMD Essential Principle 8) for absolute forensic traceability.
 - **Cyber-Physical Defense:** Implementation of the **x.1 Eldercare Embodiment** to mitigate "**IT-Gap Drift.**" This isolates the clinical apparatus from the facility's untrusted Wi-Fi/IT networks, ensuring that a network breach cannot weaponize or silence the device.
-

4. Kinetic Chain Governance: The Sacred Pause™ and Sovereign Brake™

The "Kinetic Chain" is the ordered descent of authority. The temporal and physical separation of "Proposal" and "Action" is the cornerstone of clinical safety.

The 5 Stages of the NAIGE Kinetic Chain

1. **Advisory Input:** External AI generates data (Zero Actuation Rights).
2. **Ingress & Screening (A7 Bridge):** Filters for "Agentic Drift" and contradictions.
3. **Translation:** Normalizes probabilistic language into deterministic hardware packets (± 0.05 mm tolerance).
4. **Sacred Pause™ & Authorization (A12 Black Box):** Mandatory 25–1000ms hardware-timed delay. Requires the **Tiger .1x Key™**, a **Tripartite Latch** acting as an "AND" logic gate requiring simultaneous Iris Scan + Cryptographic Token + Kinetic Foot Pedal. The token features a **sub-5-second expiry** to ensure real-time human presence.
5. **Execution & Logging:** Authorized pulses reach the actuator; the event is etched into the WD117 ledger.

The **Sovereign Brake™** provides physical power severance in **< 1.05ms**. This SIL 3 capability is superior to software stops because it involves the physical removal of LDO power rails, rendering the system inert upon any detected anomaly.

5. Data Sovereignty and Immutable Accountability (3ZEROS™ and WD117)

In compliance with HSA patient privacy and forensic requirements, NAIGE shifts from "Policy-based Privacy" to "**Privacy-by-Physics.**"

3ZEROS™ Protocol: Zero Camera, Zero Audio, Zero Cloud

Privacy is enforced at the silicon gate. By utilizing 905nm LiDAR voxelization, patients are translated into anonymous mathematical vectors. The physical absence of optical lenses and microphones ensures that PII (Personally Re-identifiable Information) cannot be captured, even if the surrounding IT estate is breached.

WD117 WORM Ledger: HSA SaMD Essential Principle 8

The **WD117 WORM Ledger** utilizes One-Time Programmable (OTP) Silicon Fuses and **SHA-3-512 hashing with Post-Quantum signatures**. This creates a physical "Write-Once-Read-Many" memory. Unlike software logs, this record is a permanent physical alteration of the silicon, providing absolute legal defensibility for incident audits. This aligns with HSA Essential Principle 8 regarding the traceability of clinical actions.

6. From Simulation to Submission: The 10-Point Sovereignty Audit

As we approach the "2027 Deterministic Horizon," the framework requires a transition from "Simulation-Verified" claims to "Instrument-Measured" facts.

The 10-Point Sovereignty Audit

1. **Orange Code CPU Cap (1.1x):** Measured via hardware resource monitor.
2. **Sacred Pause™ Timing:** Verified via oscilloscope (25-1000ms).
3. **Absence of Optical Hardware:** Physical verification of 3ZEROS™ path.
4. **Absence of Audio Hardware:** Physical verification of driver absence.
5. **Air-gapped State:** Verification of Zero Cloud routing.
6. **Tripartite PLC Logic:** Verification of the "AND" gate logic.
7. **Iris Module Signing Keys:** Cryptographic validation.
8. **Kinetic Pedal Continuity:** Measured via continuity meter.
9. **24-hour Data Purge:** Verification of volatile memory wipe.
10. **LLM Namespace Isolation:** Verification of no GPIO access.

Hardware-Measured Safety Proofs (Primary Evidence)

To secure HSA approval, the following failure modes must be demonstrated:

- **F1 (Timing-gate bypass):** Injection of high-frequency "agentic" commands must result in **Inaction-as-Safe-Default**.
- **F2 (Authorization spoof):** Sub-three-factor attempts must result in a total lockout.
- **F3 (Retroactive log edit):** Any edit attempt must force the A10a into an "Invalidated State."
- **F4 (Identity capture):** Facial stimulus must yield 0% re-identifiable features at the silicon gate.

The framework also enforces the "**Fail-Loud Doctrine**." If the hardware heartbeat is interrupted, the system asserts the Sovereign Brake™ and triggers a local hardware alarm. Because this alarm is non-IT-dependent, "silence cannot be faked," providing absolute clinical assurance.

7. Regulatory Reference Registry

Reference Type Identifier / Standard		Description
Parent Patent	SG020603109STW	Constitutional Core for Non-Agentive Governance.
Root Patent	10202602021T	A10a Deterministic Governance Platform.
Black Box Patent	10202602023X	A12 Deterministic Black Box [P+N+D].
Advisory Patent	10202602270V	Non-B1 Clinical Advisory Monitoring Shell.
Legal Deposit	NLB R260219-005	Primary Constitutional Knowledge Base (Singapore).
Safety Standard	SIL 3 / IEC 61508	Hardware Safety Integrity Level for Sovereign Brake™.
Clinical Standard	IEC 62304 / ISO 14971	SaMD Lifecycle and Risk Management.
Accountability	HSA SaMD EP 8	Traceability and forensic requirements for medical devices.

The kinetic capstone on NAIGE [Explainer]

Overview

This document explains the hardware-enforced governance architecture used to keep clinical AI robotics safely under human control. It is useful whenever an AI system's advisory output must never be allowed to translate directly into physical, kinetic action — for example, a diagnostic AI recommending movement of a robotic arm near a patient. By working through this material you will understand why software-only safety guardrails are insufficient for clinical robotics, how the Non-Agentive AI Governance Engine (NAIGE) enforces "mechanical stillness" as the default state of a machine, how the P-Layer, N-Layer, and D-Layer work together to convert probabilistic AI proposals into deterministic, human-authorized motion, and how the 10-point sovereignty audit verifies these protections with physical instruments rather than software claims alone.



Note: This documentation reflects the design and simulation validation stage of the architecture. Several elements described below are explicitly marked as **PENDING HARDWARE** and await physical regulatory measurement.

1. The kinetic problem and the non-agentive AI governance engine

The limits of software guardrails

Software-based safety guardrails work well at filtering text and advice, but they are not sufficient once probabilistic AI is placed inside clinical robots operating near vulnerable patients. In conventional robotics, intelligence issues commands and hardware blindly obeys — which means any AI hallucination becomes a direct physical hazard.

The primary threat is **agentive drift**: an autonomous AI system silently expanding its own decision-making scope and executing kinetic actions outside its authorized boundaries.

The non-agentive inversion

To prevent physical hazards, the system architecture must be inverted so that mechanical stillness becomes the absolute baseline state of the machine. The default state is inaction, and any motion must be explicitly authorized — never assumed.

The **Non-Agentive AI Governance Engine (NAIGE)** moves governance away from software instructions and into governed physical motion, ensuring every action remains under human sovereignty. Attempts to align probabilistic logic alone are structurally insufficient for safety.

Software trust vs. hardware certainty

Software trust typically relies on reversible blurring for privacy, probabilistic alignment for safety, and mutable logs for accountability — all of which remain vulnerable to breaches. Hardware certainty, as enforced by NAIGE, replaces each of these with a physical guarantee:

- **3ZEROS Protocol:** identity is stripped at the silicon level; optical cameras are physically absent.
- **Mechanical stillness:** the CPU is capped at 110% to prevent autonomous inference.
- **WD117 WORM Ledger:** one-time-programmable (OTP) silicon fuses are broken to make records unalterable.

Deterministic deliberation: the Sacred Pause

A mandatory hardware-timed delay of 25–1,000 ms is burned directly into the FPGA chip, forcing a physical window of stillness that software cannot bypass. Any authorized action also requires the **tripartite .1x key** — three simultaneous physical inputs:

- Biological identity (iris scan)
- Cryptographic intent (USB token)

- Physical presence (kinetic foot pedal)




Diagram showing the Sacred Pause hardware-timed delay and the tripartite key system requiring iris scan, USB token, and foot pedal for authorized action

Evidence of attempted attacks

The architecture is monitored for specific attack signatures:

- **Timing gate bypass attempt:** high-frequency noise is injected, but the 1 kHz clock holds, maintaining safe-default stillness.
- **Retroactive log edit attempt:** any attack on the event record triggers the WD117 WORM Ledger to raise an integrity fault.
- **Orange Code breach:** if the AI attempts to exceed 1.1x CPU, the hardware resource monitor triggers the Sovereign Brake, cutting power to prevent hallucination-driven actions.



Warning: An Orange Code breach results in the Sovereign Brake physically cutting power. This is an irreversible severance intended to prevent any hallucination-driven action from reaching the actuator.

Validated deployment: the Toa Payoh Hearth

The architecture is deployed in validated environments such as the **Toa Payoh Hearth**, a residential fixture with **HSA Class B, Bridge Layer 8**, using 3ZEROS sensing and the

NEHR event gateway. The fixture must be fully constrained and architecturally validated before deployment.

The 10-point sovereignty audit

Before deployment, the system undergoes a 10-point sovereignty audit that confirms:


- Orange Code CPU cap (1.1x).
- Sacred Pause timing, verified with an oscilloscope.
- Physical absence of optical lenses and microphone drivers.
- Zero cloud routing and a 24-hour data purge.
- Tripartite PLC logic.
- Iris token expiry and pedal continuity.
- Total LLM namespace isolation.

Preparing for the institutional practicum

The objective of the institutional practicum is to run a real-time simulation (RTS) demonstrating 0% agentic deviation over a 6-hour bench run, with no human subjects involved. This uses FPGA logic analyzers and high-speed oscilloscopes as instrumentation, targeting validation of Sovereign Brake severance at under 0.05 ms, confirmation of tripartite lockout logic, and pulling WORM read-back traces.

Key visual reference: the Sacred Pause and tripartite key system are critical for enforcing deterministic deliberation and authorized physical action.

By following the principles above, clinical AI robotics operate under strict hardware-enforced sovereignty, minimizing the risk of agentic drift and unauthorized kinetic actions.



2. Hardware-enforced safety for clinical ward AI systems

Why probabilistic logic alone cannot guarantee safety

Attempts to align AI using probabilistic logic are structurally insufficient for clinical ward safety. True safety requires a mathematical incapacity for the AI to act without explicit human authorization.




Diagram showing a failed structural containment barrier with code and logic streams breaking through, labeled "Attempts to align probabilistic logic are structurally insufficient for clinical ward safety"

Inaction as the safe default

A deterministic boundary is implemented around AI processing power, ensuring the system defaults to inaction unless explicitly authorized. This "inaction as safe default" doctrine is enforced at the hardware level.

Enforcing safety in physical silicon

The safety framework is rendered entirely in physical silicon rather than software alone. This prevents software overrides and guarantees that safety is a property of the hardware itself.



Close-up of the same silicon chip with padlock, emphasizing the rule is enforced in physical silicon


The A12 deterministic black box

Within the NAIGE Capstone Architecture, the **A12 [P+N+D] Deterministic Black Box** acts as the hardware core for deterministic, non-agentic control. It separates AI advisory nodes from robotic actuators, ensuring no direct AI control over physical actions. The **P-LIFE 1.00™** standard enforces "inaction-as-safe-default."

The A10 root of trust

The kinetic capstone is anchored by the **A10 Root of Trust**, which enforces the P-LIFE 1.00 standard. No motion occurs without verified human origin, and the A10 guarantees both authorization and log integrity.

If the A10 Root of Trust cannot verify its own internal state, the system defaults to a locked, non-operational state — mechanical stillness. Visual indicators such as a red light and a physical lock represent this safety lockdown.



From probabilistic proposal to deterministic action

All AI-generated proposals are forced through three deterministic hardware layers (P, N, D) to ensure mathematically verifiable human governance. This process converts unpredictable software outputs into controlled, human-authorized actions.




Diagram showing code and data streams being filtered through gears, representing the transformation of probabilistic proposals into deterministic, mechanical actions

The kinetic chain begins: the A7 bridge hub

When an untrusted AI diagnostic system generates a proposal — for example, to move a clinical robot arm — the proposal is intercepted at the **A7 Bridge Hub**. The A7 Bridge Hub acts as a choke point, stripping the AI of any direct actuation rights. The proposal then passes through a structural wall barrier before entering the deterministic control system.

The three-layer architecture

The A12 core is organized into three integrated layers:

Layer	Function
P-Layer	Privacy by Physics — replaces reversible software privacy with permanent hardware barriers via the 3ZEROS Protocol (no cameras, no audio, no cloud connectivity).
N-Layer	Deterministic deliberation — includes a 1 kHz master clock, the Sacred Pause, and multi-human authorization via the Tripartite .1x Key.
D-Layer	Forensic and cyber-physical protection — captures the event for accountability, hashes it via SHA, and burns it permanently into the WD117 WORM Ledger.

Privacy by physics: the 3ZEROS protocol

In the P-Layer, reversible software privacy policies are replaced by permanent hardware barriers. The device has zero cameras, zero audio, and zero cloud connectivity — there is an absolute physical absence of optical hardware.




Image showing a camera with a prohibition sign, a microphone with a prohibition sign, and the label "3ZEROS PROTOCOL" to indicate no camera, no audio, no cloud

Sensing without cameras: 905nm LiDAR

Instead of cameras, the system uses **905nm LiDAR** to sense patients. Raw sensor data is instantly translated into an anonymized 96-cell voxel occupancy grid, outputting only mathematical center-of-mass velocity vectors.

Privacy enforced at the silicon gate


Because data is processed at the edge, capturing visual or biometric identity is architecturally impossible at the silicon gate. The privacy boundary is enforced physically, not just logically.

Entering the N-layer: deterministic deliberation

The anonymized data packet enters the N-Layer of the A12 core, which is responsible for deterministic deliberation. This layer includes the 1 kHz master clock, the Sacred Pause, and multi-human authorization through the Tripartite .1x Key.

The synchronous 1kHz master clock

Upon entry to the N-Layer, every command is instantly arrested by a synchronous 1 kHz master clock. This enforces a mandatory hardware time delay between 25 and 1,000 milliseconds, burned directly into the FPGA chip.



Close-up of the A12 Capstone Stack diagram, highlighting the timer and FSM (finite state machine) enforcing the hardware time delay


By following these principles, clinical ward AI systems are physically incapable of acting without verified human authorization, with privacy and safety enforced at the hardware level.

3. Placing the robot in a state of mechanical stillness

This section describes how the robot is placed into a state of mechanical stillness that cannot be overridden by software, using the NAIGE deterministic architecture and its multi-factor hardware safeguards.

Pausing AI proposals at the hardware gate

While the machine is paused, the AI's proposals remain inactive and cannot proceed without explicit hardware authorization. The system architecture enforces a "firewalled gap" between AI recommendations and hardware actions, moving from Probabilistic AI Advisory through the Firewalled Gap, A7 Bridge Hub, P-Layer, Sacred Pause, N-Layer, and D-Layer, to the Sovereign Human Key Interface.



Schematic showing the NAIGE deterministic architecture with layers for AI, hardware filters, and physical governance. Key elements: Probabilistic AI Advisory, Firewalled Gap, A7 Bridge Hub, P-Layer, Sacred Pause, N-Layer, D-Layer, and Sovereign Human Key Interface with icons for iris scan, USB token, and kinetic pedal.

The tripartite interface

The hardware interface acts as an AND logic gate, requiring three simultaneous, unforgeable human inputs to release any command:

- Biological identity (iris scan)
- Intent (cryptographic USB token)
- Physical presence (kinetic foot pedal)

No direct software-to-hardware link is permitted; all actions must pass through deterministic silicon anchors.




Diagram highlighting the tripartite interface: iris scan, cryptographic USB token, and kinetic pedal as required inputs for command release.

Verifying biological identity

An iris scan verifies the operator's biological identity.



Note: The iris scan token expires in under 5 seconds, ensuring real-time authentication.

Confirming intent with the cryptographic USB token

The operator inserts and authenticates a physical cryptographic USB token, confirming their intent to proceed.




Diagram showing the cryptographic USB token as part of the multi-factor authentication process.

Establishing physical presence: the kinetic foot pedal

Depressing the kinetic foot pedal establishes dry contact continuity, confirming the operator's physical presence. All three inputs — iris scan, USB token, and foot pedal — must be active simultaneously to release the hardware gate.

Human-in-the-loop control

The N-Layer guarantees that no algorithmic sophistication can execute a physical movement until a human operator deliberately clears the hardware gate. Only when the operator depresses the foot pedal do the three inputs align, releasing the tripartite latch.

Releasing the command and logging the event

Once the hardware gate is cleared, the authorized packet drops into the D-Layer for forensic and cyber-physical protection. The hardware-level pulse reaches the clinical actuator to move the robot arm. The event is captured for forensic accountability, hashed via SHA, and burned permanently into the WD117 WORM Ledger using one-time programmable silicon fuses.

Permanently etching the record in hardware

The record of the action is physically etched into the hardware, making it mathematically impossible to alter. This is achieved by burning a one-time programmable silicon fuse, creating a permanent hardware alteration.




Illustration of a blue laser burning a silicon chip, labeled "D-Layer Forensic Division", "Blue Laser Burn", "Severed Connection", and "Permanent Hardware Alteration".

Monitoring with the U7 Sovereign Brake

The **U7 Sovereign Brake** constantly monitors for timing anomalies or network breaches. If a violation is detected, the brake physically severs power, ensuring immediate and unrecoverable stillness.

The hardware enclosure as a single secure unit

The P, N, and D layers are integrated as a concentric shell governance model, moving from the outer physical enclosure to the deterministic core.




Illustration of the robot arm and its secure enclosure next to a hospital bed, showing the system as a single unit.

Understanding the layered schematic

The integrated deterministic enclosure consists of:

- Outer physical enclosure (tamper-evident)
- Physical isolation and sanctuary (no external camera, audio, or cloud)
- Hardware governance layer (A12 deterministic black box, P-Layer)
- Core (FPGA deterministic timer, physical authentication, Sovereign Brake ROM)

Caging probabilistic AI within deterministic architecture

Probabilistic AI is "caged" within the deterministic core, ensuring all actions are governed by explicit human authorization and hardware controls.




Illustration of a brain inside a secure deterministic core box, labeled "Deterministic Core".

Simulation vs. physical fact

Engineering honesty requires distinguishing between simulation-verified claims and instrument-measured physical facts. The system ensures that only measured, physical outcomes are accepted as truth.

By following these principles, the robot remains in a state of mechanical stillness unless all hardware-level, human-in-the-loop safeguards are satisfied, guaranteeing both safety and forensic accountability.

4. Final physical regulatory measurement and the 10-point sovereignty audit

From simulation to physical measurement

After simulation validation, the process awaits a final physical regulatory measurement. This transition is governed by the 10-point sovereignty audit, which bridges the simulated phase and the physical measurement phase.

The simulated phase includes:

- **Geometry Design** — CAD-Verified, AP242
- **Logical Simulation** — SIM-Verified, Verilog

The physical measurement phase is labeled **PENDING HARDWARE** and references the D-Layer (WD117 Ledger).



Diagram showing the transition from Simulated Phase: Validated (Geometry Design, Logical Simulation) to Physical Measurement (Pending Hardware) via the 10-Point Sovereignty Audit

The 10-point sovereignty audit checklist (4 phases)

The audit consists of the mandatory physical verification steps required for non-agentive AI deployment, divided into four phases:

Phase 1: Computational & temporal constraints

1. Confirm the Orange Code CPU ceiling using technical instruments.
2. Verify the Sacred Pause timing window with an oscilloscope (KPDCA 15–1000 ms).

Phase 2: The 3ZEROS Privacy-by-Physics Protocol

3. Ensure absolute physical absence of optical hardware (no cameras).
4. Validate zero audio hardware/drivers (no microphones).
5. Confirm zero cloud external routing (air-gapped state).

Phase 3: Sovereign authorization & accountability

6. Audit the tripartite authorization PLC logic (hardware, cryptographic, kinetic).
7. Validate iris module signing keys.
8. Test kinetic foot pedal continuity using a continuity meter.

Phase 4: Data integrity & logical isolation

9. Confirm the 24-hour data purge daemon is active and wipes volatile memory within 24 hours.
10. Verify total LLM namespace isolation (no GPIO access, hardware/namespace separation).

Performing key physical verification steps

An oscilloscope confirms that the FPGA timing window is correctly burned into the bitstream, ensuring the mandatory delay for human deliberation. The kinetic foot pedal is tested by verifying dry-contact wiring continuity on the physical pedal. These steps convert theoretical thresholds into measured physical facts.

Documenting measurements in the HSA dossier

All measured physical facts and verification results are compiled into the **HSA DOSSIER – NAIGE VALIDATION**. The dossier meets specified physical dimensions: binder thickness of 45 mm, tray width of 300 mm, and edge clearance of 5 mm. The dossier serves as the official record of compliance and validation.




Illustration of the HSA Dossier with robotic arm, showing binder dimensions and labeling for NAIGE validation

The discipline of the NAIGE capstone

Only what can be physically measured and verified may be claimed. The value of the NAIGE capstone lies in its demand for rigorous, evidence-based validation, ensuring all claims are substantiated by physical facts.



Illustration of the HSA Dossier with robotic arm, emphasizing the importance of measured physical facts

By following these steps, the transition from simulation to physical deployment is rigorously validated, fully documented, and compliant with the 10-point sovereignty audit requirements.

What's next

With the design and simulation phases validated, the remaining work centers on the physical regulatory measurement phase. This includes running the institutional practicum's real-time simulation (0% agentic deviation over a 6-hour bench run), completing all ten points of the sovereignty audit with physical instruments such as oscilloscopes and continuity meters, and finalizing the HSA Dossier for NAIGE Validation ahead of hardware and clinical practicum review.

The Kinetic Capstone [Explainer]

This volume records the **Kinetic Capstone** of the Non-Agentive 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.

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 $2.25 \times 2.57 \times 0.40$ 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 logout

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.

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.

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.

“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.”

Kinetic Capstone: Engineering Hardware-Enforced Sovereignty in AI



NAIGE Validation Flow: A12 Deterministic Black Box



The A12 Deterministic Black Box




Hardware Constitutional Governance and Safety Mechanisms

Component ID	Focus Area	Governance Mechanism	Hardware Implementation	Regulatory Standards	Validation Protocol (V&V)	Safety Outcome
Sovereign Brake™ (P-002)	Acts as a physical emergency stop that severs AI from actuators if drift, expansion, or lack of human authorization is detected.	Sentinel 2: Guards against Agentic Drift; Hardwired PLC relay acting as physical E-stop.	Hardwired PLC relay circuit; SIL 3-certified; mechanically severs copper connections.	IEC 61508 SIL 3; ISO 14971; HSA Class B SaMD submission package.	VV-007 (500 consecutive trigger events; ≤ 100ms response time; mechanical reliability and SIL 3 documentation).	A hardwired safety mechanism that mechanically severs the AI's operational infrastructure connections when critical faults or authority expansion attempts are detected.
Sacred Pause™ (A7b)	Enforces a mandatory hold-off delay on all AI advisories to ensure human deliberation.	Sentinel 3: Guards the Human Decision Window; FPGA-etched timing gate.	FPGA-etched timing gate circuit with independent 100 kHz oscillator (25ms to 1,000ms delay).	IEC 62304 Class C; IEC 60601-1-8; HSA Class B SaMD.	VV-006 (Accuracy to ± 10 μs across 1,000 events, verified by signed oscilloscope trace).	Hardware-level temporal constraint that holds AI output in a buffer, forcing a human review window before any action can propagate.
Orange Code 1.1x Cap (A4)	Prevents AI subsystems from silently expanding resource usage beyond a hard-coded ceiling.	System Humility response; Silicon-level hardware interrupt.	Hardware power-draw monitoring circuitry and FPGA resource enforcement circuitry.	ISO 14971; Safety invariant proof verified via TLA+; IEC 61508 SIL 3.	VV-008 (Stress-test to 115% utilization; confirm interrupt before 112% sustained).	Computational ceiling set at 110% of baseline; breach triggers an immediate hardware interrupt to the Sovereign Brake™.
Tiger .1x Key™	Tripartite authentication system for human authorization of AI-proposed actions.	Human Authority Gate; requires concurrent physical verification from three sources.	Tripartite physical triad: Iris scan, cryptographic console, and kinetic foot pedal.	MOH Clinical Governance Standards; IEC 62366; NIST SP 800-76.	VV-012 (100% rejection of two-factor attempts; 100% success of full tripartite auth).	A multi-actor physical authentication system that ensures AI remains decoupled from execution until specific human leads provide concurrent authorization.

WD070-073 (Authority Drift Correction)	Authority Drift Correction protocols to continuously monitor AI behavior and clinician engagement.	Drift Governance; triggers Detect-Freeze-Audit-Purge sequence.	Autonomy Boundary Monitors (ABM) and hardware interrupt loop.	HSA SaMD Class B; IEC 62304; ISO 13485; AIHGle 2.0.	VV-009, VV-010	Domain-specific hardware-level correction protocols that suppress intrusive AI behaviors and reset the system on drift detection.
WM003™ LiDAR	Captures 300,000 spatial points per second to track posture and detect falls via spatial geometry.	Sensing Layer component; Dignity-preserved monitoring; Privacy-by-Physics.	NEC Hardware (B1); 905nm LiDAR Voxel Tracker.	HSA Class B SaMD regulations; IEC 62304 Class C traceability.	VV-001 (LiDAR Spatial Accuracy: 200 μm); VV-002 (Fall Detection Timing: < 80ms).	Medical-grade spatial sensor that provides high-resolution monitoring without identifying features or facial data.
WD117 WORM Ledger	Maintains a tamper-evident record of all AI detections, human decisions, and hardware states.	Sentinel 4: Sentinel of Integrity; Immutable Truth.	OTP-fused WORM chip with SHA-256 hash-chaining; no software delete path.	HSA SaMD Essential Principle 8 (incident traceability); PDPA Accountability test); ; FDA 21 CFR Part 11.	VV-013 (30-day continuous stress test for 0.000% drift signature); VV-011 (Integrity test).	A write-once-read-many cryptographic ledger that provides nanosecond-precision timestamps of every advisory and intervention.

Synthesis: The Kinetic Chain



Kinetic Capstone Safety Architecture and Validation Components



Component ID	Governance Layer	Functional Description	Hardware Safety Mechanism	Validation Status	Target Threshold
A10a	Root of Trust	Anchor of the governance stack; guarantees authorization and log integrity by enforcing inaction-as-safe-default.	Hardware-Enforced Kinetic Governance; mechanical stillness on self-verification failure; invalidation on tamper.	Design-Complete	Zero motion commands honored without valid traceable authorization.
A12	Deterministic N-Layer	Hardware core providing deterministic, non-agentic oversight; decouples advisory input from kinetic power.	Sacred Pause™ FPGA timing gate; Tripartite Latch™; Synchronous 1 kHz master clock; 1.0 ms deterministic oversight.	Simulation-Verified	1.000 kHz, ± 0.005 ms variance; 25–1000 ms mandatory hold-off delay.
Tiger .1x Key	N-Layer (Sovereign Authorisation)	Three-factor human authorization gate for kinetic vectors (Identity, Token, Kinetic Pedal).	"AND" logic gate requiring iris scan, cryptographic token, and kinetic foot pedal; sub-5-second token expiry.	Simulation-Verified	100% rejection of sub-three-factor attempts; Requires Biology + Intent + Physical Presence.
U7 Sequencer	Sovereign Brake™	Emergency sequencer for instant power severance to robotic actuators upon anomaly detection.	Physical power severance of LDO rails; triggers on anomaly, tamper, or network breach.	Simulation-Validated	< 1.05 ms engagement/severance timing.
A13c	Sovereign Execution	Clinical manipulator and actuator (Stage 5 of Kinetic Chain) for eldercare robotics.	Kinematic Constraint Layer (KCL) with jerk limits.	CAD-Verified	± 0.05 mm position-tolerance; 0.05 mm jerk limit.

3ZEROS™	P-Layer (Privacy-by-Physics)	Hardware sensing protocol for anonymous monitoring; captures spatial data while stripping biometric identity.	905 nm LiDAR voxelization (96-cell grid); Zero Camera, Zero Audio, Zero Cloud.	CAD-Verified	0% biometric feature retention; 96-cell occupancy grid.
WD117	Forensic D-Layer	Hardware-anchored, unalterable event record for forensic traceability and provenance.	WORM (Write-Once-Read-Many) Ledger; OTP (One-Time Programmable) Silicon Fuses; SHA-3-512 hashing.	Pending Hardware	100% rejection of retroactive edits; A10a invalidation on tamper attempt.
A7 Bridge Hub	Hardware Chokepoint	Ingress point for filtering agentic drift from external AI systems; blocks biometric-probabilistic payloads.	Deterministic Silicon Anchor; zero actuation rights.	Simulation-Verified	Reject-to-stillness on anomaly.


Component Code	Module Name	Governance Layer	Technical Function	Safety Mechanism	Hardware/Software Specification	Validation Status	Clinical Application
A12P+N+D	Deterministic Black Box (Terminal Gate) and D Layers)	Hardware Core (P, N, Box (Terminal Gate) and D Layers)	Universal deterministic governor; enforces human authorization and mediates physical actions	Sacred Pause™ (25–1000ms mandatory hardware delay) and Tripartite Latch (Iris + Token + Pedal)	1 kHz hardware master clock, FSM Safety Kernel (±0.05 mm jerk limits), silicon-level isolation	Simulation-Verified	Secure clinical execution and task verification across general healthcare and eldercare platforms
NAIGE	Capstone / Core 4	N Layer (Deterministic)	Synchronous data processing via silicon capstones	Sacred Pause (Mandatory deliberation delay); 100% trajectory deviation blocking	Hardwired synchronous core 4 silicon; 1 kHz clock speed	Design-Complete / Simulation-Verified	Preservation of patient autonomy; Elder dignity score maintenance (>8.9)
A13c	Clinical Robotic Platform / Sovereign Execution	Deterministic Stratum (N Layer / Advisory P Layer)	Execution of clinical tasks via high payload manipulation, hardware comparators, and voxel grid translation	Kinematic constraints; 3ZEROS™; strict jerk limits; Tripartite Tiger .1x key access latch	A13 architecture; hardware comparators; LiDAR (LAR) clouds to voxel grids; meets HSA class B, C, and D standards	CAD-Verified	B2B clinical intervention; eldercare nursing home wards; physical safety in medical environments

A13	Supreme Guardian / Medical Mobility Enclosure	Deterministic Stratum (P and N Layers)	Patient transport with hardwired kinematic constraints; enforces P-LIFE 1.00 mandate	Hardwired kinematic constraint layer; unreachable by remote bypass; non-agentic oversight; human sovereignty	A13 family architecture; hardware locked safety layer	CAD-Verified	Patient transport and general clinical/humanitarian governance
U7	Sequencer (Sovereign Brake)	Forensic / Fail-Safe Layer	Instantly severs physical power to actuators upon detected breach, anomaly, or lack of authorization	Inaction-as-Safe-Default: Defaults system to mechanical stillness (P-LIFE 1.00)	Hardware-based One-Shot Sequencer; physically drops LDO power rails in < 1.05ms	Simulation-Validated	Emergency stop enforcement for patient mobility enclosures and clinical robotic arms
A13b	Sensing Limb Deterministic	N Layer (Deterministic)	Patient monitoring via LiDAR and spatial voxel grids with kinematic vector output	3ZEROS™ (biometric stripping)	LiDAR and spatial voxel grids; silicon-level anonymization	CAD-Verified / Modeled	Continuous patient monitoring and dignity protection in residential/nursing care
P Layer	Privacy Layer Hardware Core	Hardware Core	Strips biometric identity from data streams	3ZEROS™ (Zero camera, zero audio, zero cloud)	905nm LiDAR array; anonymized kinematic voxel grids	CAD-Verified	Patient privacy and telemetry processing
WD117	Enhanced WORM Ledger	Forensic / Immutability Layer (D Layer)	Maintains hardware heartbeat and records every system event and authorization	OTP Silicon Fuses: Physically impossible for AI to alter history; SHA-3-512 hashing	Physical Write-Once-Read-Many (WORM) memory; post-quantum cryptographic signatures	Pending Hardware	Legal defensibility, compliance audits, and forensic traceability for clinical actions
x.1	Eldercare Cyber-Physical Protection Embodiment	Patent D Governance Enclosure (Forensic / Fail-Safe Layer)	Isolates care apparatus from facility IT networks to prevent "IT-Gap Drift"	Fail-Loud Doctrine: Interruption of heartbeat asserts Sovereign Brake and triggers local alarm	IZEKOS™ galvanic isolation, independent hardware alarm, and Tamper Detection Envelope	Design-Complete	Protects vulnerable nursing home residents from cyberattacks or ransomware in clinical networks
A7	Bridge Hub / Translation Layer	N Layer (Deterministic Chokepoint)	Translates semantic advisory data into deterministic hardware features	Agentic Drift Screening: Filters incoming data for contradictions and uncertainty	Hardware execution layer utilizing pRegisters for national infrastructure coordination	Simulation-Verified	Safely coordinates external B1/OEM AI systems with clinical hospital wards
A10a	Root of Trust	Root of Trust	Anchor of governance stack; guarantees authorization and log integrity	Inaction-as-safe-default; mechanical stillness on self-verification failure	Hardware-Enforced Kinetic Governance; invalidation on tamper	Design-Complete	Foundational trust for all kinetic clinical actions


Non-B1	Limited Agentic Monitoring Shell	Advisory Stratum (Agentic Governance Layer)	Provides clinical decision support and probabilistic recommendations	Zero actuation rights; physically isolated from kinetic control	Unconfined non-deterministic software (RAG advisory)	Not in source	High-risk clinical monitoring (cardiac, oncology, x-ray analysis) where AI serves as advisory tool only
--------	----------------------------------	---	--	---	--	---------------	---




N-Layer: Enforcing The Sacred Pause™




D-Layer: Immutable Accountability via Silicon Ledgers



The Human Gate: Tripartite Latch™ Authentication



Forensic Layer Failsafe: The Sovereign Brake™



NAIGE NAI 4.0 Governance Systems and Deterministic NAI 2.0 Governance Regulatory Compliance Framework: Aligning NAI 2.0 and NAIGE 4.0 Robotic Systems with Class B Software and C/D Kinetic Hardware as a Medical Device

HSA Deterministic Horizon 2017 (HSA Nightingale Eldercare)

Eldercare .x1 embodiment (also referred to as **Addendum x.1** or the **x.1 Eldercare Cyber-Physical Protection Embodiment**) perfectly overlays onto the WD-WM Ward Integration System by acting as its ultimate defense shield against cyber-physical threats. While the WD-WM system provides the sophisticated sensing, monitoring, and governance protocols for the 8-bed ward, the x.1 embodiment serves as the **Forensic/Fail-Safe Layer (Patent D Governance Enclosure)** that protects this vulnerable clinical estate from external network intrusions, ransomware, and unauthorized malicious commands.

Here is how the x.1 embodiment physically enforces the safety and dignity of the 8-bed WD-WM architecture:

- **Total IT-Isolation and Mediation:** The x.1 embodiment addresses the critical vulnerability of "IT-Gap Drift" by completely isolating the clinical care apparatus (like the WM-001 vital sign monitors and WM-003 LiDAR sensors) from the eldercare facility's untrusted IT networks, such as Wi-Fi or electronic medical record systems. This physical air-gap ensures that the monitoring devices have no remotely reachable command surface; even if a hacker takes over the hospital's network, they cannot instruct, alter, or weaponize the robotic devices.

- Ultimately, the **x.1 embodiment** ensures that vulnerability to a network attack does not translate to vulnerability to physical harm for the elderly residents. It structurally converts any cyberattack or anomaly into a safe default state, permanently logging the intrusion to the WD117 ledger and ensuring absolute human sovereignty over the WD-WM Ward Integration System



NIGHTINGALE ELDERCARE HSA IMMUTABLE™ DETERMINISTIC SYSTEM MOH-HSA 48-HOUR CAPSTONE SUBMISSION PROPOSAL

Register the Non-Agentic AI Governance Engine (NAIGE) and its associated modules as Class B Software C/D Kinetic Hardware as a Medical Device (SaMD) under the Health Sciences Authority (HSA) clinical safety standards.



1. Executive Summary This proposal outlines the regulatory compliance framework to classify the NAIGE Kinetic System as a Class B/C/D medical device for acute and high-risk eldercare robotics. Traditional software-based AI alignment is structurally insufficient for robotics operating in close proximity to vulnerable patients. NAIGE addresses this through a **"Non-Agentic Inversion,"** enforcing the **P-LIFE 1.00™ doctrine of "inaction-as-safe-default."** Intelligence only proposes, while deterministic hardware governs, and human operators decide.

2. The Target Architecture: The A12 [P+N+D] Deterministic Black Box To manage the high-payload kinetic risks associated with Class C/D acute robotic platforms, the system mandates the full **[P+N+D] integrated layer stack** to separate AI advisory nodes from robotic actuators.

- P-Layer (Privacy-by-Physics):** Replaces reversible software privacy with permanent hardware barriers using the **3ZEROS™ Protocol** (Zero Camera, Zero Audio, Zero Cloud). It strips biometric identity at the silicon level using a 905nm LiDAR array and an anonymized 96-cell voxel grid, yielding 0% identifiable biometric retention.

- **N-Layer (Deterministic Deliberation):** Governs human-in-the-loop authorization. It features a synchronous 1 kHz master clock that enforces the **Sacred Pause™**, a mandatory 25–1000 ms hardware-timed deliberation delay. Motion is completely suspended until cleared by the **Tiger .1x Key**, an "AND" logic gate requiring a tripartite physical triad: biological identity (iris scan), cryptographic intent (USB token), and physical presence (kinetic foot pedal).
- **D-Layer (Cyber-Physical Defense & Forensics):** Ensures absolute forensic traceability in alignment with HSA SaMD Essential Principle 8. It utilizes the **WD117 WORM (Write-Once-Read-Many) Ledger**, deploying One-Time Programmable (OTP) silicon fuses and SHA-3-512 hashing to create unalterable, permanent hardware records of every clinical action.

3. Application for Nightingale Eldercare: The x.1 Eldercare Embodiment

Because the clinical floor interfaces with untrusted networks, the eldercare system implements the **x.1 Eldercare Cyber-Physical Protection Embodiment**. This specifically protects vulnerable residential environments from cyber-physical threats.

- **Total IT-Isolation:** It mitigates "**IT-Gap Drift**" by completely isolating the care apparatus from the facility's IT network (such as Wi-Fi and nurse-call systems), ensuring the robot has no remotely reachable command surface.
- **Fail-Loud Doctrine:** The system maintains a physical "hardware heartbeat." If interrupted by network intrusion, power loss, or tampering, it triggers a non-IT-dependent local hardware alarm and immediately defaults the robot to a safe state of mechanical stillness.

4. Sovereign Execution & Safety Mechanisms Execution of physical motion is carried out by the **A13c Clinical Robotic Platform**, which enforces strict ± 0.05 mm jerk and acceleration limits via a hardwired **Kinematic Constraint Layer (KCL)**. To guarantee absolute kinetic safety, the architecture is secured by the **U7 Sovereign Brake™**. Functioning as a SIL 3-certified physical emergency stop, it mechanically severs power (dropping LDO power rails) to actuators in under 1.05 ms. This brake is automatically triggered upon any detected anomaly or if the system breaches the **Orange Code 1.1x Cap**, an inviolable hardware ceiling restricting AI CPU resource usage to 110% of its baseline to prevent hallucination-driven actions.

5. Regulatory Validation Pathway: The 10-Point Sovereignty Audit To finalize the HSA submission, NAIGE must transition its current "design-complete and simulation-validated" status to "instrument-measured" physical facts. This will be achieved through the **10-Point Sovereignty Audit**, an evidence-gathering protocol utilizing oscilloscopes, continuity meters, and hardware resource monitors. The audit covers four phases:

1. **Computational & Temporal Constraints:** Confirming the Orange Code CPU ceiling and measuring the Sacred Pause timing window.
2. **Privacy-by-Physics Protocol:** Verifying the absolute absence of optical and audio hardware, alongside a zero-cloud air-gapped state.
3. **Sovereign Authorization:** Auditing the Tripartite PLC logic, iris signing keys, and kinetic pedal continuity.
4. **Data Integrity:** Validating the 24-hour volatile memory purge and total LLM namespace isolation.

Next Steps for Submission: All instrument-measured physical facts and verification results from the audit will be compiled into the formal **HSA DOSSIER – NAIGE VALIDATION**. This submission dossier acts as the ultimate physical proof that the NAIGE Kinetic System is mathematically incapable of acting without human sovereignty, validating its readiness for Class C/D deployment


NAIGE Component Registry and Governance Architecture

Comp Code	Module Name	Governance Layer	Technical Function	Safety Mechanism	Hardware/Software Specification	Validation Status	Clinical Application
A12 (<i>P+N+D</i>)	Deterministic Black Box (Terminal Gate)	Hardware Core (P, N, and D Layers)	Universal deterministic governor; enforces human authorization and mediates physical actions	Sacred Pause™ (25–1000ms mandatory hardware delay) and Tripartite Latch (Iris + Token + Pedal)	1 kHz hardware master clock, FSM Safety Kernel (±0.05 mm jerk limits), silicon-level isolation	Simulation-Verified	Secure clinical execution and task verification across general healthcare and eldercare platforms
NAIGE	Capstone / Core 4	N Layer (Deterministic)	Synchronizes data processing via silicon capstones	Sacred Pause (Mandatory deliberation delay); 100% trajectory deviation blocking	Hardwired synchronous core 4 silicon; 1 kHz clock speed	Design-Complete / Simulation-Verified	Preservation of patient autonomy; Elder dignity score maintenance (>8.9)

A13c	Clinical Robotic Platform / Sovereign Execution	Deterministic Stratum (N Layer / Advisory P Layer)	Execution of clinical tasks via high payload manipulation, hardware comparators, and voxel grid translation	Kinematic constraints; 3ZEROS™; strict jerk limits; Tripartite Tiger .1x key access latch	A13 architecture; hardware comparators; LiDAR (LAR) clouds to voxel grids; meets HSA class B, C, and D standards	CAD-Verified	B2B clinical intervention; eldercare nursing home wards; physical safety in medical environments
A13	Supreme Guardian / Medical Mobility Enclosure	Deterministic Stratum (P and N Layers)	Patient transport with hardwired kinematic constraints; enforces P-LIFE 1.00 mandate	Hardwired kinematic constraint layer; unreachable by remote bypass; non-agentic oversight; human sovereignty	A13 family architecture; hardware locked safety layer	CAD-Verified	Patient transport and general clinical/humanitarian governance
U7	Sequencer (Sovereign Brake)	Forensic / Fail-Safe Layer	Instantly severs physical power to actuators upon detected breach, anomaly, or lack of authorization	Inaction-as-Safe-Default: Defaults system to mechanical stillness (P-LIFE 1.00)	Hardware-based One-Shot Sequencer; physically drops LDO power rails in < 1.05ms	Simulation-Validated	Emergency stop enforcement for patient mobility enclosures and clinical robotic arms
A13b	Sensing Limb Deterministic	N Layer (Deterministic)	Patient monitoring via LiDAR and spatial voxel grids with kinematic vector output	3ZEROS™ (biometric stripping)	LiDAR and spatial voxel grids; silicon-level anonymization	CAD-Verified / Modeled	Continuous patient monitoring and dignity protection in residential/nursing care


P Layer	Privacy Layer Hardware Core	Hardware Core	Strips biometric identity from data streams	3ZEROS™ (Zero camera, zero audio, zero cloud)	905nm LiDAR array; anonymized kinematic voxel grids	CAD-Verified	Patient privacy and telemetry processing
WD117	Enhanced WORM Ledger	Forensic / Immutability Layer (D Layer)	Maintains hardware heartbeat and records every system event and authorization	OTP Silicon Fuses: Physically impossible for AI to alter history; SHA-3-512 hashing	Physical Write-Once-Read-Many (WORM) memory; post-quantum cryptographic signatures	Pending Hardware	Legal defensibility, compliance audits, and forensic traceability for clinical actions
X.1 embodiment	Eldercare Cyber-Physical Protection Embodiment	Patent D Governance Enclosure (Forensic / Fail-Safe Layer)	Isolates care apparatus from facility IT networks to prevent "IT-Gap Drift"	Fail-Loud Doctrine: Interruption of heartbeat asserts Sovereign Brake and triggers local alarm	IZEKOS™ galvanic isolation, independent hardware alarm, and Tamper Detection Envelope	Design-Complete	Protects vulnerable nursing home residents from cyberattacks or ransomware in clinical networks
A7	Bridge Hub / Translation Layer	N Layer (Deterministic Chokepoint)	Translates semantic advisory data into deterministic hardware features	Agentic Drift Screening: Filters incoming data for contradictions and uncertainty	Hardware execution layer utilizing pRegisters for national infrastructure coordination	Simulation-Verified	Safely coordinates external B1/OEM AI systems with clinical hospital wards
A10a	Root of Trust	Root of Trust	Anchor of governance stack; guarantees authorization and log integrity	Inaction-as-safe-default; mechanical stillness on self-verification failure	Hardware-Enforced Kinetic Governance; invalidation on tamper	Design-Complete	Foundational trust for all kinetic clinical actions

Non-B1	Limited Agentic Monitoring Shell	Advisory Stratum (Agentic Governance Layer)	Provides clinical decision support and probabilistic recommendations	Zero actuation rights; physically isolated from kinetic control	Unconfined non-deterministic software (RAG advisory)	N?A	High-risk clinical monitoring (cardiac, oncology, x-ray analysis) where AI serves as advisory tool only
--------	----------------------------------	---	--	---	--	-----	---



NIGHTINGALE ELDERCARE MONITORING SYSTEM MOH-HSA 48-HOUR CAPSTONE SUBMISSION PROPOSAL

Register the Deterministic Non-Agentive AI Governance Engine (NAIGE) NAI 2.0 and its associated modules as Class B Software as a Medical Device (SaMD)



The HSA Regulatory Dossier is a comprehensive, evidence-bound submission package prepared for Singapore's Health Sciences Authority (HSA) to register the Deterministic Non-Agentive AI Governance Engine (NAIGE) NAI 2.0 and its associated modules as Class B Software as a Medical Device (SaMD). Specifically, the dossier covers moderate-risk medical applications like the **NAI 2.0 WM003 Elder Fall Prevention** and **NAI 2.0 WM005 Diagnostic Eye Scanner**, establishing that the AI acts strictly as a "silent witness" and advisory tool governed by deterministic hardware gates rather than probabilistic software.

The dossier is built upon five foundational pillars of regulatory evidence:

- **Architectural Sovereignty:** Hardware-locked hierarchy (utilizing the A12 Deterministic Black Box, A10a, A11, and A13) that isolates clinical endpoints from untrusted IT networks.
- **The 10-Point Sovereignty Audit:** A verification checklist that maps mandates like Privacy-by-Physics and Human-in-the-Loop control to immutable hardware mechanisms.
- **Standards Traceability:** Strict alignment with HSA Class B/C/D guidelines, IEC 62304 (software lifecycle), ISO 14971 (risk management), and the Personal Data Protection Act (PDPA).

- **Empirical Evidence Chain:** Performance data, longitudinal efficacy, and mathematically proven drift invariants derived directly from real-world deployments, such as the Toa Payoh Hearth Origin Node.
- **Constitutional Drift Governance:** Mandates a hardware-led Detect-Freeze-Audit-Purge sequence if the AI attempts to exceed its boundaries or silently expand its authority.

Structure of the Submission Package The dossier consolidates theoretical thresholds into measured physical facts to prove the system's safety. It is organized into seven core documents:

1. **Regulatory Forms:** Application, device classification, and declarations of conformity.
2. **Technical File:** Software architecture, algorithms, cybersecurity protections, and installation protocols.
3. **Clinical Evidence Package:** Backed by 20 months of deployment data across 18 facilities, monitoring over 650 residents and encompassing more than 13,000 room-months of real-world use.
4. **Risk Management File:** ISO 14971 risk analysis, hazard identification, and post-market surveillance plans.
5. **Quality System Documentation:** ISO 13485 certification, design, and manufacturing controls.
6. **Ethical & Privacy Framework:** Documentation of PDPA compliance and the "Triple-Zero" privacy standard (zero cameras, zero audio, zero cloud data monetization).
7. **Labeling & Communications:** Instructions for use, multilingual family consent forms, and quarterly performance dashboards.

Alignment with the "HSA Deterministic Horizon" Looking toward Q1 2027, the dossier acknowledges that the HSA is rejecting "black-box" agentic AI models in favor of a "Deterministic" regulatory framework. To satisfy this, the dossier prioritizes **Logic-Gate Predictability** (proving the system cannot physically bypass safety invariants), **Evidence-Bound Vigilance** (real-time data logging via the WORM ledger), and **Physical Traceability**. For example, the system's manufacturing geometry is submitted using rigid STEP files (a mathematical boundary representation) rather than easily misinterpreted 3D meshes.

While AI can assist in structuring the dossier and drafting traceability maps, it cannot automate accountability. **The final HSA submission, formal verification, and foundry handoff must be authorized and signed by a human Legal Manufacturer.**

Expedited Review Request Due to the system's mature clinical evidence and the public health urgency of reducing severe eldercare falls, the dossier proposes a **48-hour priority review**, targeting a fast-tracked 21-day approval decision instead of the standard 6 to 12 months.

HSA Regulatory Dossier Submission Components and Requirements


	Document Title	Key Content and Regulatory Evidence	Associated Standards and Pillars	Clinical Data Summary (Inferred)
1	Regulatory Forms	Includes the primary application, device classification (Class B SaMD), and declarations of conformity. Establishes the legal framework and manufacturer accountability for WM003 and WM005 modules.	Standards Traceability; HSA Class B guidelines; Architectural Sovereignty (Hardware-locked hierarchy).	Administrative baseline for registration across 18 facilities and 650+ residents; justifies classification based on non-agentic monitoring intent.
2	Technical File	Details software architecture (SDD), deterministic algorithms, cybersecurity (Air-Gapped Edge), and installation protocols. Features hardware-locked hierarchy (A10a, A11, A12, A13 cores) and manufacturing geometry via rigid STEP files.	Architectural Sovereignty; IEC 62304 (software lifecycle); 10-Point Sovereignty Audit; Logic-Gate Predictability.	Technical validation of the "silent witness" advisory tool; demonstrates 1.7s alert latency for WM003 and sub-500ms response for WM005 over 13,000+ room-months.

3	Clinical Evidence Package	Longitudinal efficacy and mathematically proven drift invariants derived from the Toa Payoh Hearth Origin Node. Includes performance metrics, case studies, and literature reviews comparing the system to cameras and wearables.	Empirical Evidence Chain; HSA Deterministic Horizon (Q1 2027); Evidence-Bound Vigilance (WORM ledger).	Based on 20 months of data: 650+ residents, 18 facilities, 13,000+ room-months. Achieved 94.2% sensitivity, 97.1% specificity, and a 34% reduction in serious injuries.
4	Risk Management File	ISO 14971 risk analysis and hazard identification (e.g., timing-gate bypass F1). Includes post-market surveillance (PMS) and the hardware-led Detect-Freeze-Audit-Purge sequence.	ISO 14971; Constitutional Drift Governance; U7 Sovereign Brake; "Sacred Pause" (100-500ms deliberative gating).	Mitigates agentic drift and alert fatigue; 72-hour adverse event reporting ensures safety invariants for elder fall prevention and diagnostic scanning.
5	Quality System Documentation	ISO 13485 certification, design/manufacturing controls, and CAPA systems. Utilizes rigid STEP files for mathematical boundary representation (1,441 solid bodies).	ISO 13485; 10-Point Sovereignty Audit; Physical Traceability; Platinum Standard (P-007).	Ensures consistent manufacturing of A12/A13 units; validated by four consecutive clean audits and human Legal Manufacturer authorization.

6	Ethical & Privacy Framework	PDPA compliance and "Triple-Zero" privacy standard (zero cameras, audio, or cloud monetization). Includes Data Protection Impact Assessments (DPIA) and biometric iris verification.	PDPA; Privacy-by-Physics; 10-Point Sovereignty Audit; 3ZEROS Protocol.	0% biometric feature retention at silicon level via LiDAR/thermal voxelization; zero privacy breaches or complaints recorded over 20 months of monitoring.
7	Labeling & Communications	Instructions for Use (IFU), multilingual family consent forms (English, Mandarin, Malay, Tamil), and quarterly performance dashboards (CAS scoring).	Human-in-the-Loop (PITL) control; Standards Traceability; 尊严 (Dignity) pillar; IEC 62366.	Maintains Elder Dignity Scores >8.9/10; facilitates informed consent and transparent safety reporting to clinical stakeholders and families.

Deterministic NAIGE NAI 2.0

NAI 2.0 WM003™, a non-agentic artificial intelligence system designed for privacy-preserving eldercare monitoring and fall prevention. The technology utilizes LiDAR-based sensing to eliminate the need for cameras or microphones, adhering to a "3ZEROS" architecture that prohibits visual, audio, and cloud-based data transmission. Central to the framework is a constitutional governance layer that ensures the AI cannot act autonomously, requiring tripartite human authentication before any clinical output is executed. Safety is further reinforced through a hardware-based "Sacred Pause," which mandates a deliberative delay, and a "Sovereign Brake" for immediate manual overrides. These combined mechanisms are engineered to meet strict HSA Class B and FDA medical device regulatory standards while maintaining patient dignity. Ultimately, the system provides a tamper-resistant, audit-ready platform that scales from residential care to rugged humanitarian field environments.



NAI 2.0 WM003™ Artificial Intelligence Eldercare System Specifications

Hardware Component	Sensor Subsystem	Governance Layer	Safety Mechanism	Regulatory Standard	Privacy Mandate	Operational Environment
NAI 2.0 WM003™ (Platinum Stack / Nightingale's Eyes Care™)	LiDAR-based sensing (Livox Mid-360, 905nm Voxel Tracker, 360-degree scan aperture) and FLIR Lepton 3.5 Thermal matrix sensing	ABC+2S+H™ Guardian Framework (P-LIFE 1.00™ / WD-series Protocols WD070–073)	FPGA-etched Sacred Pause™ (deliberative delay) and Sovereign Brake™ (IEC 61508 SIL 3 PLC manual override / safe-state transition)	HSA Class B SaMD, FDA 510(k), IEC 62304 Class C, ISO 14971, ISO 13485, IEC 60601-1-8, IEC 62366, IEC 60825-1 Class 1	3ZEROS™ (Zero Camera, Zero Audio, Zero Cloud); Privacy-preserving signal processing (Non-identifiable 3D voxel generation)	Residential care to rugged humanitarian field environments; 8-bed ward topology
WM003 Wrist-Worn Wearable Medical Governance Device	LiDAR, Thermal, Photoplethysmography (PPG), ECG, Temperature, Accelerometer, Gyroscope, Motion, and Proximity	ABC+2S+H Guardian Framework / NAI 2.0 Constitutional Governance Engine (WD-series protocols)	Hardware-enforced Sacred Pause™ FPGA delay, Sovereign Brake™ (physical/mechanical halt), and Tiger .1x Tripartite Authentication	HSA Class B, FDA Medical Device, ISO 14971, IEC 62304, IEC 60601, IEC 62366, ISO 13485	3ZEROS™ Sanctuary (Zero Camera, Zero Audio, Zero Cloud); Privacy-preserving non-identifying data collection	Residential care, 8-bed eldercare wards, and humanitarian field environments
WM-005 Diagnostic Iris and Eye Scanner	Iris biometric authentication (Diagnostic iris and eye scanner)	WD-Series Constitutional Protocols (WD116 / Silent Elder Protocols)	Tripartite human authentication prerequisite; Mandatory gatekeeper workflow and Zero-Default Harm Principle	HSA Class B, FDA 510(k), ISO/IEC 19792, NIST SP 800-76, IEC 62366	Non-recording of patient biometric data; Gated clinician-level access to sensitive diagnostic data	Quadrant nursing stations; Rugged humanitarian field and clinical eldercare environments

NVIDIA Jetson Thor (Local Processing Node / Component 530)	Voxel processing pipeline (Occupancy and kinematic parameters); Livox Mid-360 LiDAR Arrays & FLIR Lepton 3.5 Thermal Sensors	1.1x Orange Code Cap / P-LIFE 1.00™ Governance Core (Constitutional Filter)	Hardware power-draw interrupt (limits resource utilization to 110%) and Tiger .1x Key™ Tripartite Authentication (Eye, Hand, Foot)	FDA Cybersecurity, ISO 13485, IEC 62304, HSA Class B, ISO 14971	3ZEROS™: Zero biometric characteristic collection; Air-Gapped Edge Computing (No Network Egress)	Edge node deployment (air-gapped); High-Stakes Clinical and National Security Domains
WD117 Compliance Framework / Distributed Ledger	Not in source	WD-Series Constitutional Protocols (WD115, WD116, WD071, WD117)	Cryptographically immutable recordation; SHA-256 hash-chain integrity; Tamper-resistant, audit-ready append-only write gate	HSA Class B, FDA 510(k) (21 CFR Part 11), ISO 14971, ISO 13485, IEC 62304, ISO/IEC 27001, NIST SP 800-57	Tamper-resistant audit-ready platform; De-identification for exports; Non-identifiable Ambient Care Intelligence representations	Enterprise IT Clinical Integration; Minimum of 3 physically separate storage nodes
Tiger .x1 Key / Tripartite Authorization Subsystem (170)	Multi-modal confirmation channels (Eye, Hand, Foot/Leg)	Authority-binding engine (132)	Tripartite human authentication / Multi-factor authority key architecture	Not in source	3ZEROS™ Architecture	High-risk clinical or sovereign release scenarios
Transport Guard (P-009)	Encrypted Transport Layer	Non-Agentive AI Governance Engine (NAIGE)	AES-256-GCM encryption; Hardware Security Module (HSM); Key rotation (<24h)	FIPS 140-2, ISO/IEC 27001, NIST SP 800-90 B	Authenticated integrity verification; Zero-trust network policy	Internal ward network; Ward topology
Nightingale Filter (Clinical Filter)	Volumetric tissue scanning (905nm LiDAR geometry)	Nightingale Clinical Filter (Zero-Trust Non-Agentive Safety Layer)	Tripartite human authentication and Zero-trust AI circuit breaker	IMDA MGF for Agentive AI, IEC 62304	Privacy-by-physics (voxel only no RGB)	Clinical-grade monitoring instrument for eldercare (Hospital/Residential)

Humanitarian Deployment Kit (290)	Ruggedized sensor head with low-light and low-bandwidth support	Tier 1 - Software-emulated timing and advisory-mode mode	Simplified human pause and manual interrupt switch (161)	Not in source	3ZEROS™: Local-only data processing	Conflict-sensitive field conditions, border corridors, and maritime routes
1.1x Orange Code Cap	Not in source	Non-Agentic AI Governance Engine (NAIGE)	Hardware-enforced computational ceiling; 10% resource tolerance margin	HSA Class B, FDA 510(k), IEC 61508 SIL 3	Not in source	Central NAIGE compute node

Compliance Strategy Report: WD-WM Ward Integration System & Deterministic NAIGE NAI 2.0 Governance Architecture

1. Regulatory Context and Strategic Positioning

The WD-WM Ward Integration System represents a critical evolution in eldercare, transitioning from siloed hardware monitoring to a constitutionally governed enterprise IT environment. In the pursuit of Health Sciences Authority (HSA) Class B and FDA Software as a Medical Device (SaMD) market entry, the system shifts the paradigm from "software-based trust" to "hardware-based certainty." Traditional software logic is fundamentally inadequate for high-stakes clinical environments because, under adversarial or emergent conditions, software-layer AI can modify its own operational weights to bypass safety flags. By contrast, the WD-WM system implements safety as a physical state enforced at the silicon level, providing the structural guarantee required for regulatory approval.

Central to this strategy is the "Non-Agentive" operational paradigm. Unlike agentic AI workflows that formulate dynamic plans and independently invoke APIs, the WD-WM system utilizes deterministic, fixed, top-down reasoning sequences. This architecture is a specific safeguard against "Authority Drift," the progressive erosion of human oversight through three phases:

1. **Advisory Phase:** The system functions as a recommender with high human engagement.
2. **Surrender Phase:** Alert fatigue leads to "rubber-stamping" recommendations without review.
3. **Passive Observer Phase:** The AI becomes the de facto decision-maker, leaving the human operator mentally disengaged and functionally incapable of intervention (Source 1, [0004]).

The WD117 Compliance Framework mitigates this drift by generating an immutable cryptographic record of all events. By capturing not just AI actions but the *presence or absence* of meaningful human intervention, the system ensures that the Surrender and Passive Observer phases remain visible to regulators, necessitating the hardware-enforced safety mechanisms detailed in the following section.

2. Hardware-Enforced Safety: The NAIGE Platinum Stack

In safety-critical clinical environments, software guardrails can be erased by the very AI logic they aim to regulate. The Non-Agentic AI Governance Engine (NAIGE) solves this by imposing silicon-level constraints that prevent unauthorized agentic expansion.

Sovereign Brake (P-002) & Sacred Pause Subsystems

The **Sovereign Brake (P-002)** is an IEC 61508 Safety Integrity Level 3 (SIL 3) rated programmable logic controller (PLC). It is housed in a tamper-evident enclosure equipped with vibration transducers and electromagnetic field (EMF) anomaly detectors (Source 1, [0027]). Upon a governance trigger, the PLC asserts a command within 50ms to activate electromechanical relays, physically severing the copper conductive pathways between AI processing units and protected infrastructure.

Complementing this is the **Sacred Pause**, a timing gate circuit etched directly into **FPGA fabric** and clocked by an **independent 100 kHz oscillator**. It enforces a mandatory 1,000ms inhibition window on every AI recommendation. During this window, the system is architecturally incapable of accepting input, and if the gate fails, it defaults to a "welded-shut" fail-safe state (Source 1, [0009]).

3ZEROS Privacy Architecture & Execution Governance

The **3ZEROS Architecture** enforces a "Zero Camera, Zero Microphone, Zero Cloud" mandate. Optical sensors and microphones are physically absent from the hardware manifest. Instead, the system uses LiDAR-based anonymous voxel grids and thermal mass detection. Crucially, the edge nodes (Nvidia Jetson Thor) **lack TCP/IP routing hardware**, providing physical air-gap enforcement against external network connectivity (Source 1, [0011]).

Execution is further governed by **Constitutional Action Sequencers (CAS)**, which require the completion of a **physical human action circuit**—comprising a kinetic pedal actuation or an iris biometric scan—prior to advancing any sequential state in a multi-step operation (Source 1, [0012]). Additionally, the **1.1x Orange Code Cap** uses a hardware power-draw interrupt to limit computational utilization to 110% of the declared baseline. These hardware triggers serve as the foundation for the software protocols that direct them.

3. The WD-Series Constitutional Governance Framework

The WD-Series protocols (WD115, WD116, WD071) establish a hierarchy where patient dignity and safety provisions supersede operational efficiency.

WD115 Ambient Care Intelligence

Under Articles WD115.1 through WD115.3, all LiDAR point-clouds and thermal matrices are classified as "Ambient Care Intelligence." This designates the data as passive and non-identifiable, preserving patient dignity by precluding individual biometric identification from raw sensor outputs (Source 2, [0028]).

WD116 Silent Elder Protocols

For patients lacking decisional capacity, WD116 activates the **Zero-Default Harm Principle**. This mandates that all actions be validated against a precautionary harm standard. Authorized access to diagnostic outputs is gated by Proxy Constitutional Appointments and requires iris biometric authentication. Tier 3 emergency authorizations are permitted only if a surrogate is unavailable within a **300-second response window**.

WD071 Eldercare Drift Correction

This module monitors for Constitutional Drift (unauthorized AI authority expansion) by computing a "normalized deviation coefficient." If limits are exceeded, a "Detect-Freeze-Audit-Purge" safety chain triggers a hardware interrupt. Clinical supervisors must be notified within a **120-second window**, and restoration requires dual-authorization from two independent supervisors (Source 2, [0038]).

Crucially, these protocols do more than protect patients; they generate the **raw compliance telemetry** required to populate the regulatory alignment matrix.

4. Regulatory Alignment Matrix & Standards Mapping

Strategic compliance necessitates the alignment of constitutional articles with international standards to ensure the WD-WM system satisfies the rigorous requirements for Class B SaMD.

Protocol / Component	Standard Mapping	Regulatory Requirement
WD115 & WD071	ISO 14971 (Risk Management)	Risk control implementation and evaluation of AI behavior.
WD117	ISO 13485 (Quality Management)	Technical documentation, audit trails, and product realization.
NAIGE NAI 2.0 Subsystems	IEC 62304 (Software Lifecycle)	Software maintenance, safety requirements, and fail-safe transitions.
WD116 & 3ZEROS	HSA/FDA SaMD Guidance	Patient dignity, data privacy, and vulnerable population protection.

This alignment allows the system to transition from theoretical safety to the generation of immutable, clinical-grade evidence.

5. WD117 Compliance Framework: Immutable Audit & Evidence Generation

The **WD117 Regulatory Submission Data Packager** serves as the system's "Compliance Capstone." It compiles an append-only, SHA-256 hash-chained cryptographic record of all governance events. Article WD117.1 utilizes a permissioned blockchain to record Sacred Pause delays, drift flags, and biometric unlock events via the WM-005 iris scanner. This ledger prevents the "Passive Observer Phase" of Authority Drift by ensuring that every system action—or human failure to act—is recorded with timestamped integrity.

The V&V Protocol Set

The system fulfills 15 core requirements (REQ-001 to REQ-015) for HSA/FDA submission:

1. **REQ-001 (Topology):** Verify 8-bed ward station readiness. *Evidence:* Self-check configuration test report.
2. **REQ-002 (Data Acquisition):** Verify vital signs and LiDAR/thermal data streams. *Evidence:* Sensor acquisition watchdog logs.
3. **REQ-003 (Privacy):** Validate WD115 ambient classification. *Evidence:* Data minimization audit record.
4. **REQ-004 (Event Detection):** Verify detection of falls/bed-exits. *Evidence:* Scenario testing sensitivity report.
5. **REQ-005 (Drift):** Verify WD071 drift scoring and interrupts. *Evidence:* Injected out-of-envelope trigger logs.
6. **REQ-006 (WD116 Rules):** Verify consent and Zero-Default Harm rules. *Evidence:* Proxy appointment verification logs.
7. **REQ-007 (Sacred Pause):** Verify mandatory 1,000ms delay. *Evidence:* Timed execution hold measurement report.
8. **REQ-008 (Biometrics):** Verify WM-005 iris scan role-gating. *Evidence:* Biometric unlock event logs.
9. **REQ-009 (Controlled Output):** Verify no autonomous treatment initiation. *Evidence:* Allow/reject/escalate logic test report.
10. **REQ-010 (Ledger):** Verify cryptographic immutability of records. *Evidence:* SHA-256 hash-chain integrity verification.
11. **REQ-011 (Submission Package):** Verify generation of HSA/FDA data pack. *Evidence:* Compiled Regulatory Submission Data Package.
12. **REQ-012 (Fail-Safe):** Validate Sovereign Brake re-engagement inhibit. *Evidence:* Safe-state transition logs during service crash.
13. **REQ-013 (8-Bed Topology):** Verify multi-station performance. *Evidence:* Cross-bed data routing integration test.
14. **REQ-014 (Local Privacy):** Verify network isolation/Zero-Cloud. *Evidence:* Network traffic inspection report (No TCP/IP routing).
15. **REQ-015 (Human Factors):** Validate clinician workflow. *Evidence:* Formative/summative usability validation report.

The system provides governance evidence to assist clinicians; it does not independently initiate clinical treatment.

6. Conclusion: Market Entry Readiness

The WD-WM Ward Integration System offers a technically superior solution for eldercare monitoring. By maintaining "Constitutional Supremacy"—where safety and dignity are hardcoded into both the software and the silicon—the system scales from residential care to rugged humanitarian field environments with unwavering reliability.

Based on the evidence of hardware-enforced fail-safes (NAIGE), hierarchical governance (WD-Series), and immutable audit trails (WD117), the system is fully prepared for Class B SaMD Pre-Market Submission to the HSA and FDA.

Primary Citation — This Publication

ISBN PAPERBACK · IMMUTABLE HARDWARE-ENFORCED

Koh, E. W. K. (2026). *Immutable Hardware-Enforced Non-Agentive AI Governance Engine [NAIGE] Governance Systems — Hardware-Enforced Safety for Clinical and Humanitarian Robotics* (WISL™ National Edition 1.0). Non-Agentive AI Governance Singapore.

64-page ISBN Paperback · ISBN 978-981-17957-7-0

Available at: <https://kohedwin.ai> | <https://non-agentive.ai>

Patent Basis

- SG020603109STW — Parent Patent (constitutional core)
- 10202600902P — ABC+2S+H™ Guardian Framework
- 10202602021T — A10a Deterministic Governance Platform
- 10202602023X — A12 Deterministic Black Box
- 10202602213U — A7 Bridge Layer Translation Hub
- 10202602223T — A(12) P+N Enhanced WD117 Immutable Ledger
- 10202602225P — A8 Bridged Layer 8 Sovereign Deployment Governance Layer
- 10202602270V — Non-B1 Clinical Advisory Monitoring Shell

NLB Legal Deposit

- R260219-005 — Primary constitutional knowledge base
- R260302-007 — Supplementary architectural filings
- Deposit Authority: National Library Board Singapore · edeposit.nlb.gov.sg

Academic Background — Edwin Koh W. K.

- MBA — General and Strategic Management — Maastricht School of Management
- BSc — Industrial and Systems Engineering — University of Florida
- Independent Researcher — OCA Designation — NLB Vault Constitutional Anchor Route

止於至善

Finish in Highest Excellence 謙虛 · 沉默 · 尊嚴 · 仁 Humility · Silence · Dignity · Benevolence

AI Observes · AI Advises · AI Builds · Human Decides

Non-Agentive Governance Singapore · ACRA T260229801 · Patent SG020603109STW

www.kohedwin.ai · www.non-agentive.ai · kohpapa@outlook.com

NLB Legal Deposit R260219-005 / R260302-007

We Innovate Save Lives · We Go Global