WISL_No96 ISBN: 978-981-17970-6-4

# The Nightingale Anchor Protocol: A13d Hardware-Enforced Governance Architecture

The provided sources detail the A13d Governance architecture, a specialized system designed for robotic home-care and therapeutic assistance for the elderly. This framework utilizes a non-agentic AI engine that ensures technology serves only as an advisory tool, leaving all final physical interventions to human decision-makers. Safety is maintained through hardware-enforced governance layers, such as the "Sacred Pause" delay and "Kinematic Constraint Layers," which physically prevent unauthorized or unsafe motion. By separating privacy-preserving observation from deterministic execution, the system aims to provide hospital-level safety in a residential setting. Extensive patent documentation and schematics illustrate a commitment to forensic accountability and strict physical boundaries to prevent autonomous machine drift. Ultimately, the A13d platform is presented as a mechanical extension of the caregiver rather than an independent robotic entity.

The A13d therapeutic home-care apparatus prevents autonomous AI agency by functioning strictly as a **mechanical extension of a human caregiver under deterministic hardware governance**, explicitly rejecting the role of an autonomous nurse or clinical decision-maker. It enforces the foundational rule that **"NO AUTHORIZATION = NO MOTION"** and operates on a "Non-Agentic Inversion" philosophy where the AI proposes, deterministic hardware governs, and the human caregiver decides.

To achieve this, the A13d utilizes multiple unbypassable physical and architectural barriers to completely strip the AI of execution authority:

**Absolute Physical Isolation and Deterministic Screening:** The AI advisory processor possesses zero direct actuation rights and is physically isolated from the motor-energy paths on a separate power or trust domain. When the AI generates a proposed care action or trajectory, this suggestion cannot directly energize an actuator. Instead, the proposal must pass through a controlled ingress bridge and a deterministic screening kernel. This kernel translates the proposal into a bounded deterministic packet, discarding any free-form or probabilistic text, and automatically rejects proposals that fall outside of a pre-signed, clinician-configured therapy plan.

**The Sacred Pause™ (Hardware-Timed Deliberation):** Before any proposed kinetic motion can proceed, a deterministic hardware timer enforces a "deliberation interval" selectable between 25 and 1000 milliseconds. This timer operates completely independently of the AI's scheduling software, guaranteeing an unavoidable physical window for the human caregiver to review, modify, or cancel the specific intervention.

**Tripartite .1x Key™ & Hold-to-Run Authorization:** The A13d mandates sovereign human authorization before executing any task. Caregivers must provide active consent through a local multi-factor authorization gate, which verifies independent factors such as biometrics (e.g., iris scan), a cryptographic token, and a kinetic "hold-to-run" control. This local authorization is time-limited and bound specifically to the validated care packet; silence,

ambiguous inputs, or the expiration of the authorization token immediately returns the apparatus to a hardware-inhibited state.

**Kinematic Constraint Layer (KCL) Execution:** Even after a human authorizes a movement, the AI cannot improvise the action on the fly. The A13d's execution controller only permits motion within a hardwired, pre-validated kinematic envelope. This envelope enforces strict physical boundaries on speed, acceleration, force, and jerk, maintaining a rigid positional tolerance of ±0.05 mm during therapeutic movements.

**U7 Sovereign Brake™ and Fail-Safe Intercept:** If the AI experiences "agentic drift," or if any spatial sensor, load cell, or boundary monitor registers an anomaly, the system relies on an independent safety circuit. Any violation of the kinematic envelope or failure of the "hardware heartbeat" triggers a mode-aware stopping controller. The U7 Sovereign Brake physically severs the actuator's power in under 1.05 milliseconds or executes a safe load-retention stop, independently of the software.

By layering these immutable constraints, the A13d ensures its default state is always **absolute mechanical stillness**—where all actuators are held open-circuit until every gate condition is physically verified—making autonomous AI action physically impossible in the home care environment

Initially, the specific term "Nightingale One Protocol" **did not exist as a registered, pending, or reserved trademark** within the architecture's exhaustive intellectual property roster. Instead, the overarching clinical governance law was officially known as the **Nightingale Anchor Protocol (NAP 4.0)**, and its physical clinical deployment was trademarked as **Nightingale's Eyes Care™**.

However, the roadmap for the architecture's next National Edition dictates that the **Nightingale One Protocol will be formally filed as a new Patent and Trademark**.

Once filed, the Nightingale One Protocol trademark will represent the definitive governing protocol for the **A13d home-care embodiment integrated within Patent M**. It will specifically govern the **M (A13) Supreme Guardian** as it takes overarching command of the **NAIGE Care (Hospital-Level Home Care)** deployment.

This new trademark legally defines a **non-agentic, hardware-governed therapeutic and caregiving system** designed for active-aging residents and bed-bound silent seniors. Under the Nightingale Protocol One, the system provides privacy-preserving kinetic caregiving in a residence by physically separating advisory computations from actual actuator authority. The protocol's core mandate ensures that while the system's sensors may advise on a patient's state, **no kinetic intervention can ever occur without an authorized human decision and deterministic safety clearance**.

In the context of Home-Based Hospital care (formally deployed under NAIGE Care and the **Nightingale Protocol One**), the **M (A13) Supreme Guardian Master Control Layer** handles AI proposals through a strict, non-agentic hardware-governed pipeline where intelligence is solely advisory.

When an AI advisory system proposes a home-based care action (such as lifting, transfer, or range-of-motion physiotherapy), M (A13) manages the proposal through the following deterministic steps to ensure absolute safety for the resident:

- **Absolute Denial of AI Actuation:** The AI advisory logic may detect resident conditions and propose care actions, but M (A13) ensures the AI possesses zero direct actuation rights. The proposal must pass through the deterministic hardware layer (the P+N+D path), completely stripping the AI of the ability to bypass physical controls or improvise actions on the fly.
- **The "Null Enclosure" Check:** M (A13) mandates the presence of the **B1 (Nightingale's Eyes Care™)** privacy-preserving sensing layer. If this layer is missing or bypassed, M (A13) classifies the machine as a "Null Enclosure" and mathematically refuses to grant any kinetic authority, rendering the hardware completely inert.
- **The Sacred Pause™ (Hardware-Timed Deliberation):** Before any proposed motion can proceed, the system enforces a hardware-timed deliberation interval (between 25 and 1000 milliseconds). This prevents the machine from immediately coupling an AI proposal to kinetic motion, guaranteeing an unavoidable window for human review.
- **Human Sovereign Authorization:** M (A13) enforces the foundational rule: **"NO AUTHORIZATION = NO MOTION"**. It demands that an authorized human caregiver explicitly approve the proposed care action using the **Tripartite .1x Key™** (requiring an iris identity scan, a cryptographic token, and a kinetic pedal). If human authority is absent, expires, or is ambiguous (e.g., silence or a timeout), M (A13) immediately returns the system to a hardware-inhibited state.
- **Deterministic Safety and Kinematic Checks:** Even with human approval, M (A13) independently verifies the proposed action against a pre-signed, clinician-approved care plan. It runs the proposal through a Finite-State Machine (FSM) and the **Kinematic Constraint Layer (KCL)** to check identity, timing, care-zone limits, and force, enforcing a strict positional tolerance of ±0.05 mm.
- **Bounded Execution and Continuous Inhibition:** Only after all deterministic safety checks and human authorizations clear does M (A13) energize the actuators (such as the A13d therapeutic limb) for the bounded caregiving motion. During execution, if the system detects any deviation from the allowed envelope, anomalous sensor readings, or agentic drift, M (A13) triggers the **U7 Sovereign Brake™** to physically sever motor power in under 1.05 milliseconds and enter a safe state.
- **Immutable Evidence Logging:** Finally, M (A13) ensures the entire sequence—what was sensed, proposed, authorized, executed, and stopped—is permanently logged in a local, tamper-evident forensic ledger (WD117) to maintain hospital-level accountability in the home without requiring cloud dependency.

The next move for the NAIGE Care roadmap centres on formally establishing the **Nightingale One Protocol** (also referred to as Nightingale Protocol One) and advancing the **A13d therapeutic embodiment** into a regulatory-ready state for the Health Sciences Authority (HSA).

Specifically, the roadmap outlines the following concrete steps:

- **Filing the Nightingale One Protocol:** This protocol will be formally filed as a new Patent and Trademark. It will act as the definitive governing protocol for the **M (A13)**

**Supreme Guardian** as it takes overarching command of the NAIGE Care (Hospital-Level Home Care) deployment.

- **Locking the A13d Patent Position:** The plan is to lock the A13d as a therapeutic embodiment divisional patent, which will be anchored on USTOP 10202601915Y and governed by Patent M. A proposed working title for this filing is "A13d Deterministic Non-Agentic Therapeutic Caregiving Limb under Nightingale Anchor Protocol 4.0".
- **Finalising the A13d HSA Dossier Bundle:** To move from concept to physical validation, the A13d dossier must be finalised so it can be officially filed and handed to the HSA and partners. This bundle will include a therapeutic use-case summary for Silent Elders, safety gate and workflow diagrams, and references to simulation logs and timing measurements.
- **Preparing "HSA-Ready" Language:** To ensure the system clears regulatory hurdles without triggering "overclaims," the roadmap dictates using strict, HSA-ready language. The core narrative must consistently state that the A13d is **"not an autonomous nurse or decision-maker; it is a mechanical extension of the human caregiver under strict deterministic governance"** and that the AI possesses zero direct actuation rights.
- **Integrating into National Edition 4:** Finally, this A13d dossier and its technical framing will be shaped into a dedicated chapter for the upcoming **National Edition 4**, presenting the A13d as a "humble caregiving limb" grounded in hardware facts rather than theoretical claims.

Nightingale Protocol One safeguards resident privacy in the home by utilizing a **"Privacy-by-Physics"** approach that relies on local ranging-derived state representations rather than traditional surveillance.
Specifically, it achieves this through the following mechanisms:

- **A13b Privacy-Preserving LiDAR:** The protocol employs A13b non-optical LiDAR sensing to map the resident's environment. Instead of capturing images, the system converts the sensor data into a local voxel state representation.
- **Camera-Free and Audio-Free:** The system is explicitly designed to be completely devoid of cameras and microphones in the protected care path. This physical exclusion ensures that no intimate media or identifiable personal data can be captured or stored.
- **Cloud-Free Operation:** The sensing and event-recording functions operate entirely locally without any dependency on the cloud. This minimizes the risk of external data breaches or unauthorized remote surveillance.

By separating the advisory computation from the physical hardware and relying purely on local voxelization, Nightingale Protocol One ensures that the system can safely monitor and assist a resident while minimizing personal data capture.

Tell me more about the A10a Root of Trust's role in the 6-hour simulation.

The **A10a Root of Trust** serves as the **foundational cryptographic hardware anchor** during the 6-hour Real-Time Simulation (RTS), which is modelled on the Toa Payoh Hearth residential eldercare parameters. In this simulation, its primary role is to act as the lowest,

most privileged layer of the governance stack, guaranteeing absolute **authorisation integrity, log integrity, and failure integrity**.
Throughout the 6-hour bench run—which is designed to demonstrate **0% agentic deviation** completely without human subjects—the A10a continuously secures the WD117 forensic logs. If the A10a detects any physical tampering or fails its own internal self-verification at any point, it immediately invalidates the entire stack and forces the system to default to **mechanical stillness**.
The deployment of the A10a Root of Trust in this specific simulation is central to a powerful regulatory strategy called **"Validation by Induction"**. By using the 6-hour RTS to empirically validate the A10a Root of Trust in a lower-risk, non-kinetic residential environment (HSA Class B), the architecture establishes a proven safety baseline. Once the mathematical certainty of the A10a is proven in this setting, the higher-risk acute robotic applications (HSA Class C/D) can inherit this exact same validated safety guarantee, significantly easing their regulatory approval burden.
Ultimately, the A10a's performance in the 6-hour RTS provides the empirical evidence necessary to package the HSA Class B SaMD handover mock-up, proving that the architecture is not merely a theoretical policy, but a **measurable, physical machine governed by an unassailable Root of Trust**.

The 6-hour Real-Time Simulation (RTS) for the NAIGE architecture is a rigorous bench run modeled on the Toa Payoh Hearth residential eldercare parameters. The specific parameters and toolsets used for this simulation include:

- **Core Objective:** The simulation is designed to demonstrate **0% agentic deviation** over the entire 6-hour period, and it is conducted completely **without human subjects**.
- **Environment Setup:** The structural testing protocol utilizes a **Qwen-simulated HSA/VV protocol** modeled over a 3-day, **8-bed IT enterprise fall-risk setup** (specifically the WM-003 configuration).
- **Advisory Engine ("The Soul"):** The simulation uses **Anthropic's Constitutional AI (Claude)** as the non-deterministic advisory engine. To ensure safety, this AI is strictly locked to narration and observation tasks and possesses zero execution rights.
- **Geometry and Structural Validation:** The physical geometry and Kinematic Constraint Layer (KCL) code are validated through a **Zoo.dev PASS simulation**, which outputs the mathematically rigid STEP (ISO 10303) manifests required by the Health Sciences Authority (HSA).
- **Hardware Instrumentation:** To prove "Logic-Gate Containment" with physical facts rather than statistical probability, the simulation relies on precise measuring instruments:
    - **FPGA Logic Analyzers** are used to verify the stability of the 1 kHz Master Clock, ensuring a strict ±0.005ms variance.
    - **High-Speed Digital Oscilloscopes** are deployed to verify that the U7 Sovereign Brake physically severs power in less than 1.05 milliseconds.

Ultimately, this simulation serves as the empirical validation for the **A10a Root of Trust** and the deterministic hardware-enforced safety of the architecture, providing the necessary evidence to package the HSA Class B SaMD (Software as a Medical Device) handover mock-up.

The U7 Sovereign Brake achieves rapid power severance by completely discarding software-based logic, such as "sleep" commands or emergency-stop software signals, in favor of a purely physical, independent hardware intercept mechanism.
When a fault or anomaly is detected—such as agentic drift, a kinematic envelope violation, or physical tampering—the brake relies on an IEC 61508 SIL 3-certified hardware mechanism driven by a **One-Shot Sequencer**. This sequencer immediately triggers programmable logic controller (PLC) relays to physically sever the copper conductive pathways between the AI processing units and the robot's actuators.
By physically dropping the **Low-Dropout (LDO) power rails**, the brake instantly removes the electrical energy required for the machine to move. Because it operates on an independent safety circuit completely divorced from the advisory software, this stopping controller achieves a design-intent power severance latency of **less than 1.05 milliseconds**.
This rapid, irreversible physical severance enforces the system's foundational "Inaction-as-Safe-Default" (P-LIFE 1.00™) doctrine, guaranteeing absolute mechanical stillness by rendering the apparatus electrically incapacitated before any harm can occur