

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

ISBN 978-981-17963-3-3

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.

10-section clinical governance outline structurally explicit for HSA dossier submission, including the specific lab equipment required for validation:

Phase 1: Computational & Temporal Constraints

Section 1. Orange Code CPU Cap (Governance by Scarcity)

- Governance Objective:** Verify that the AI's CPU resources are physically hard-locked at a 1.1x baseline limit. This guarantees the AI is completely starved of the surplus computational power necessary to hallucinate or expand its scope (agentic drift).
- Required Lab Equipment: Technical Instruments / Power-Draw Monitors** to measure physical hardware power draw and verify FPGA resource enforcement circuitry (pRegister P006).

Section 2. Sacred Pause™ Deliberation Timing

- Governance Objective:** Prove that the hardware physically enforces a mandatory 25–1000ms delay window before any kinetic command can propagate to the robot's actuators, ensuring a guaranteed physical window for human deliberation.
- Required Lab Equipment: Oscilloscope.** Used to capture a signed trace of the circuit's activity driven by an independent 100 kHz oscillator, verifying the hardware delay accuracy down to ± 10 microseconds.

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

This phase guarantees patient dignity by destroying biometric identity at the physical hardware layer, stripping raw 3D spatial data (captured via 905nm LiDAR) into a sterile 96-cell voxel grid before the AI processes it.

Section 3. Zero Camera Affirmation

•**Governance Objective:** Validate the absolute physical absence of optical imaging hardware (no cameras) on the robot to prevent photographic surveillance of vulnerable, undressed, or confused elderly patients.

Section 4. Zero Audio Affirmation

•**Governance Objective:** Verify the total absence of microphone hardware and operating-system audio drivers to ensure no audio is ever captured or recorded.

Section 5. Zero Cloud Air-Gap Confirmation

•**Governance Objective:** Confirm that there is no TCP/IP route to any external network, proving the system is fully air-gapped and immune to remote cloud-based data exfiltration.

Phase 3: Sovereign Authorization & Accountability

Section 6. Tripartite .1x Key™ (PLC Logic) Audit

•**Governance Objective:** Audit the programmable logic controller (PLC) to ensure it strictly functions as an "AND" gate, requiring three simultaneous physical inputs (biological, cryptographic, and kinetic) to authorize movement.

Section 7. Iris Module Signing Keys Validation

•**Governance Objective:** Cryptographically validate the biological authorization token generated by the iris scanner to ensure it hard-expires in under 5 seconds, preventing hackers or compromised AI from spoofing delayed executions.

•**Required Lab Equipment: Independent Cryptographic Tools.** Auditors are strictly forbidden from using AI for this step and must use external, independent cryptographic workflows to validate the keys and ledger hash-chains.

Section 8. Kinetic Foot Pedal Continuity Testing

- Governance Objective:** Test the physical "proof of motion" authorization, ensuring a verified human is actively depressing the kinetic foot pedal (dry-contact switch) before the machine can move.

- Required Lab Equipment: Continuity Meter.** Used to physically confirm the dry-contact wiring continuity of the foot pedal circuit.

Phase 4: Data Integrity & Logical Isolation

Section 9. 24-Hour Data Purge Validation

- Governance Objective:** Confirm that an automated purge daemon is actively running and successfully wipes all volatile short-term session memory on a strict cycle of under 24 hours.

Section 10. Total LLM Namespace Isolation

- Governance Objective:** Verify that the language-model network namespace is physically and logically separated from all hardware controls. This ensures the AI has zero General-Purpose Input/Output (GPIO) or PLC access, physically denying the AI any ability to directly trigger actuators.

Continuous Evidence-Bound Vigilance Requirements

To satisfy the HSA's post-market surveillance demands, the framework also mandates real-time tracking via the **WD117 Forensic Ledger**. This hardware layer uses One-Time Programmable (OTP) silicon fuses to create unalterable Write-Once-Read-Many (WORM) memory, proving absolute forensic finality. A failure at any point in these 10 sections will trigger the **U7 Sovereign Brake**, dropping low-dropout (LDO) power rails in under 1.05 milliseconds and forcing the robot into a permanent state of mechanical stillness