SOVEREIGN STEWARD™

NAI 2.0™ Clinical Governance Monograph

No. 6

How to Conduct Studies in NAI 2.0 Clinical Governance

PITL Research · Formal Methods · Multi-Sensor Fusion · Silent Elder

Edwin Koh Wui Kiat

Founding Father · Non-Agentic AI 2.0™

ACRA T26022980I · SG020603109STW · 2026

□□ · □□ · □□ · □ · □□□□

Table of Contents

Abstract
Chapter 1: Research Framework — Two Directions, One Constitution
Chapter 2: Direction 1 — Hardware-Enforced PITL Studies (LKCMedicine / NHG)
Chapter 3: The Three Measurable Outcomes Across All Study Designs
Chapter 4: Direction 2 — Constitutional Drift and Formal Methods
Chapter 5: The 10-Point Sovereignty Audit as a Formal Research Framework
Chapter 6: Multi-Sensor Data Fusion — Monitoring Heart Function and Physical Distress
Chapter 7: Giving Voice to the Silent Elder — Multi-Sensor Fusion in Practice
Chapter 8: Research Governance — Ethics, Data, and Constitutional Constraints
Appendix A: Study Design Summary Table
Appendix B: SVA Formal Properties — 10-Point Sovereignty Audit
References

Abstract

This monograph — NAI 2.0™ Clinical Governance No. 6 — is the research methodology companion to the CG Series. Where CG No. 1 established the engineering foundations of the Sacred Pause™, CG No. 2 mapped global regulatory compliance, CG No. 3 documented the patent architecture, CG No. 4 established the academic framework, and CG No. 5 provided the Q&A examination canon, CG No. 6 descends to the level of research practice: how, precisely, do you conduct studies on a constitutionally constrained clinical AI system?

The monograph addresses two research directions confirmed by Low Sani Elias (NTU LKCMedicine, 2 April 2026): Direction 1 — observational and simulation studies on how hardware-enforced PITL (3ZEROS™, Sacred Pause™, tripartite authentication) affects workflow, alert fatigue, and clinical safety; and Direction 2 — modelling and testing of constitutional drift and the 10-Point Sovereignty Audit using AI and formal methods groups.

A third domain is also covered: the clinical monitoring of cardiac and physiological signals through multi-sensor data fusion — specifically how the NAI 2.0™ framework traces heart function and physical distress through HRV wearable biometrics and WM003™ LiDAR geometry, giving a constitutional voice to the silent elder.

> "These studies do not merely evaluate a clinical AI. They evaluate the constitutional boundary between the AI's observation and the human's decision — the most important boundary in medicine."

Chapter 1: Research Framework — Two Directions, One Constitution

1.1 The NTU Collaboration Context

On 2 April 2026, Low Sani Elias (NTU LKCMedicine) confirmed institutional interest in two research directions for academic-non-commercial collaboration with Non-Agentic AI Governance Singapore (ACRA T260229801). Both directions treat the NAI 2.0™ constitutional architecture not merely as a clinical tool — but as a subject of rigorous, formal scientific inquiry.

Direction	Collaborating Groups	Research Domain	Primary Question
Direction 1	LKCMedicine / NHG Clinical Research	Hardware-Enforced PITL Studies	How do 3ZEROS™, Sacred Pause™, and tripartite authentication affect workflow, alert fatigue, and clinical safety?
Direction 2	AI / Formal Methods Groups	Constitutional Drift Modelling	Can constitutional drift and the 10-Point Sovereignty Audit be formally verified as a framework for verifiable non-agentic clinical AI?

1.2 What Makes NAI 2.0™ Uniquely Studyable

Most clinical AI systems require post-hoc interpretation methods to understand what the model did. NAI 2.0™ inverts this: its constitutional boundaries are hardware-enforced, formally specified, and immutably logged. This creates three properties that make it uniquely amenable to rigorous scientific study:

— Constitutional explicitness: every governance constraint is stated as a hardware specification or formal assertion — researchers can verify, measure, and report against the specification directly.

— Immutable audit trail: WD117 provides a complete, tamper-proof record of every clinical event, advisory, caregiver response, and override decision — exactly what a research study requires.

— Formal verifiability: the Sacred Pause™ timer, the Code B trigger pathway, and the Tiger .1x Key™ authentication sequence can all be stated as SystemVerilog Assertions and verified by formal tools.

> "The machine that can be proven correct is the machine that can be safely studied."

Chapter 2: Direction 1 — Hardware-Enforced PITL Studies

LKCMedicine / NHG · 3ZEROS™ · Sacred Pause™ · Tripartite Authentication

2.1 Overview

Direction 1 studies evaluate the real-world effects of hardware-enforced Physician-in-the-Loop (PITL) constraints on clinical practice. The three architectural components under study — 3ZEROS™ privacy stack, Sacred Pause™ cognitive forcing function, and Tiger .1x Key™ tripartite authentication — are not software configurations. They are physical hardware properties that can be independently verified and measured by the research team.

This distinction is fundamental to the research design: unlike software-based AI systems where the governance constraints can be modified to create experimental conditions, NAI 2.0™ constraints are constitutionally immutable. The research team does not control the Sacred Pause™ duration during the study — the ROM constants do. This is not a limitation. It is what gives the findings their validity.

2.2 Three Study Design Options

Study Design 1 · Observational Study — Ward Pilot

Goal: Assess the Sacred Pause™'s effect on alert fatigue in a live clinical environment.

Methodology

- Setting: live ward with 20–30 elders during 6-hour sessions (Toa Payoh Sanctuary or NHG cluster ward pilot).
- Participants: Green Lanyard certified caregivers — all certified, vow on record, annual renewal current.
- Pre/post automation bias assessment using the Goddard et al. methodology — validated instrument administered before and after deployment.
- Data source: WD117 Audit Trail (Tiger .1x Key™ access) — override rates, time-to-authentication by severity tier, advisory response patterns across Code A and Code B events.
- Sovereignty Audit: full 10-Point Audit must pass before each 6-hour session. Any failure halts the session.

Measurable Outcomes

- ✓ Override rate per advisory type — Code A vs Code B — across three severity tiers.
- ✓ Time-to-authentication: mean time from advisory receipt to Tiger .1x Key™ completion.
- ✓ Automation bias score: pre/post Goddard et al. instrument — correct vs incorrect decision switching rate.
- ✓ Alert fatigue proxy: click-through rate vs deliberation rate — ratio of rapid acceptances to actively contested decisions.

Study Design 2 · Simulation Study — High-Fidelity Environment

Goal: Evaluate how Tiger .1x Key™ tripartite authentication impacts cognitive load and response times during acute Code B emergency events.

Methodology

- Setting: high-fidelity clinical simulation environment with standardised patient scenarios — Code B acute events (high-velocity fall, critical HRV breach).
- Participants: Green Lanyard participants responding to controlled simulation scenarios across all three Sacred Pause™ timing ranges.
- Primary data: Sacred Pause™ oscilloscope logs (actual measured durations vs ROM constants), authentication completion rates, scenario debrief instruments.

- Secondary data: cognitive load assessment using NASA-TLX or equivalent validated instrument pre/post Code B scenarios.
- Formal documentation: oscilloscope logs and SVA verification report included as supplementary materials in research publication.

Measurable Outcomes

- ✓ Cognitive load: NASA-TLX scores pre/post tripartite authentication across Code A (500–1000ms) and Code B (25ms) conditions.
- ✓ Response time: time from Code B LiDAR trigger to clinician authentication completion.
- ✓ Authentication accuracy: completion rate of tripartite sequence (iris + console + pedal) under simulated acute emergency conditions.
- ✓ Deliberation quality: qualitative debrief — did the 25ms micro-friction preserve cognitive agency or generate perceived excessive friction?

Study Design 3 · Comparative Study — Parallel Ward Groups

Goal: Measure alert fatigue reduction against a conventional clinical alert system.

Methodology

- Setting: two parallel ward groups — one NAI 2.0™ deployment, one conventional alert system — matched on patient demographics, ward type, and caregiver staffing.
- Duration: minimum 3 months to establish baseline patterns and seasonal variation.
- Both systems must produce equivalent audit granularity for comparison — WD117 audit log vs comparator system log reviewed for data equivalence before study begins.
- Caregiver self-report: validated alert fatigue instrument (Alarm Fatigue Survey or equivalent) administered at baseline, 6 weeks, and 12 weeks.
- Ethical requirement: comparator ward must not be disadvantaged — minimum safety standards maintained throughout study period.

Measurable Outcomes

- ✓ Alert fatigue: click-through rate (advisories accepted without review) — NAI 2.0™ vs comparator.
- ✓ Override frequency: rate at which caregivers dismiss advisories without action — both systems.
- ✓ Caregiver self-report: alert fatigue scale scores at 6 and 12 weeks — between-group comparison.
- ✓ Clinical safety: near-miss rate, fall event rate, preventable adverse event rate — compared against shared pre-deployment baseline.

Chapter 3: The Three Measurable Outcomes Across All Study Designs

Across all three Direction 1 study designs, three primary outcome categories provide the basis for comparison, publication, and clinical governance reporting. These outcomes are measurable, traceable to the constitutional architecture, and directly relevant to the HSA Class B SaMD validation requirements and NTU MSc capstone thesis criteria.

Outcome	Definition	Primary Measurement	Data Source	Constitutional Anchor
1. Workflow Integration	Time from advisory receipt to Tiger .1x Key™ authentication — by severity tier 1, 2, and 3	Mean time-to-authentication (seconds) per tier · Standard deviation · Range	WD117 Audit Trail — timestamp delta between advisory generation and auth completion	Tiger .1x Key™ tripartite auth — PLC relay close timestamp
2. Alert Fatigue	Automation bias score (pre/post), override rate per advisory type, click-through vs deliberation ratio	Goddard et al. automation bias instrument score · Override rate (%) by advisory type · Click-through ratio	WD117 Audit Trail — all override and no-action decisions · Pre/post instrument scores from research instrument	Sacred Pause™ — cognitive forcing function · Zero Weighting — no algorithmic nudging
3. Clinical Safety	Near-miss rate, fall event rate, preventable adverse event rate — compared against pre-deployment baseline	Event rates per 1000 patient-days · Comparison vs baseline period	WD117 Audit Trail — all safety events, near-misses, fall events, HRV breaches · Pre-deployment baseline from ward records	WD113–WD117 Group H Eldercare Patents · Zero Preventable Deaths constitutional floor

> "A system that can be measured can be trusted. A system that can be trusted can be replicated. A system that can be replicated can save lives."

Chapter 4: Direction 2 — Constitutional Drift and Formal Methods

AI / Formal Methods Groups · Constitutional Drift · 10-Point Sovereignty Audit as Verifiable Framework

4.1 The Formal Methods Research Proposition

Direction 2 treats the NAI 2.0™ architecture not as a clinical tool but as a mathematically verifiable system. The constitutional governance properties of NAI 2.0™ — the Sacred Pause™ timer, the tripartite authentication gate, the drift detection protocols — can all be stated as formal propositions and verified exhaustively using model checking and theorem proving tools.

This makes NAI 2.0™ exceptionally amenable to formal methods research. Unlike probabilistic AI systems where governance properties are emergent and statistically characterised, NAI 2.0™ governance properties are deterministic, hardware-enforced, and formally specifiable. A formal verification study of NAI 2.0™ does not estimate that the system is safe — it proves it.

4.2 Three Formal Methods Study Design Options

Study Design 1 · Formal Verification Study

Goal: Mathematically prove that the WD070–WD073 Drift Governance sequence is sound and complete — that it covers all constitutionally prohibited transitions.

Methodology

- Model the WD070–WD073 protocols (Detect → Freeze → Audit → Purge) as a temporal logic specification using LTL (Linear Temporal Logic) or CTL (Computation Tree Logic).
- Verify that the drift governance sequence is formally correct: every constitutional boundary violation is detected (completeness), and every freeze event corresponds to a genuine violation (soundness).
- Use JasperGold, Cadence Jasper, or equivalent formal verification tool to produce a proof that the sequence covers all constitutionally prohibited transitions.
- Document SVA properties from the NAI 2.0™ specification and map each to a formal verification result.

Measurable Outcomes

- ✓ Formal proof: the drift detection → freeze → audit → purge sequence is sound and complete.
- ✓ Coverage report: all constitutionally prohibited state transitions are covered by WD070 detection.
- ✓ Counterexample analysis: any counterexamples identified represent genuine constitutional vulnerabilities — reported to framework sovereign for specification update.

Study Design 2 · Empirical Drift Simulation

Goal: Empirically characterise the drift governance response surface — detection latency, freeze accuracy, restoration quality.

Methodology

- Introduce controlled perturbations into a test NAI 2.0™ deployment — deliberate adversarial inputs, logic shifts, or configuration modifications that represent real-world drift scenarios.
- Measure WD070 detection latency: time from perturbation introduction to drift detection event.
- Measure WD071 freeze accuracy: false positive rate — freeze events triggered without genuine drift.

- Measure WD073 restoration quality: post-purge drift signature measurement — must reach 0.000%.
- Vary perturbation magnitude and type to characterise the full detection response surface.

Measurable Outcomes

✓ WD070 detection latency: mean time from constitutional boundary violation to freeze trigger (target: < 100ms).

✓ WD071 false positive rate: freeze events without genuine constitutional violation (target: < 2%).

✓ WD073 restoration quality: post-purge drift signature (target: 0.000%).

✓ Response surface map: detection latency and accuracy as a function of perturbation magnitude and type.

Study Design 3 · Constitutional Audit Framework Publication

Goal: Formalise the 10-Point Sovereignty Audit as a peer-reviewed framework applicable to any hardware-governed clinical AI — beyond NAI 2.0™.

Methodology

- Map each of the 10 Sovereignty Audit points to a testable formal property (SVA assertion or temporal logic formula).
- Develop a generalised audit methodology: for each audit point, define the formal property, the verification method, the pass/fail criterion, and the failure consequence.
- Validate the framework against the NAI 2.0™ system as a reference implementation — demonstrating that the formal properties are achievable and the audit is executable.
- Submit as a peer-reviewed publication — target: npj Digital Medicine, Journal of the American Medical Informatics Association, or IEEE Transactions on Medical Informatics.

Measurable Outcomes

✓ Published formal audit framework: a peer-reviewed methodology applicable to any hardware-governed clinical AI.

✓ Validated reference implementation: the NAI 2.0™ system demonstrates the framework is executable in a real clinical AI deployment.

✓ Community adoption potential: the framework provides regulators, researchers, and clinical governance committees with a validated tool for assessing any constitutional AI system.

Chapter 5: The 10-Point Sovereignty Audit as a Formal Research Framework

5.1 From Operational Checklist to Formal Verification Framework

The 10-Point Sovereignty Audit was designed as the mandatory pre-session constitutional verification tool for NAI 2.0™ clinical deployments. In the formal methods research context, it becomes something more: a bridge between operational clinical governance and formal mathematical verification — translating physical hardware constraints into testable formal propositions.

Each audit point is not merely a checklist item. It is a physical instantiation of a constitutional principle — and each can be mapped to a formal property that is verifiable by automated tools. This is what makes the 10-Point Sovereignty Audit uniquely valuable as a research framework: it is simultaneously operational (executable in a clinical ward) and formal (verifiable by a model checker).

5.2 The 10 Points — Formal Property Mapping

Point	Audit Requirement	Constitutional Principle	Formal Property (SVA)	Verification Method
1	Hardware Manifest	Zero Camera, Zero Audio, Zero Cloud	No optical sensor, microphone, or cloud egress route in device manifest	OS device tree inspection — grep for sensor drivers
2	FPGA Timer Calibration	Sacred Pause™ immutability	Delay_value ∈ {Range1, Range2, Range3} — SVA Property 1	Oscilloscope physical measurement ±5ms vs ROM constants
3	Dual-Redundant Counter	SIL 3 fault tolerance	Counter_A == Counter_B — SVA mismatch assertion	Hardware-in-Loop test — inject counter mismatch, verify watchdog fires
4	PLC Tripartite Auth	Tiger .1x Key™ Anti-Doppelgänger Lock	iris_valid ∧ console_confirmed ∧ pedal_closed — SVA Property 4	Live connection test — all three DI inputs verified active
5	Code B Configuration	Physical threshold only — no algorithmic trigger	code_b_active → lidar_physical_threshold — SVA Property 3	CG Committee approval documentation + threshold parameter audit
6	Drift Baseline	WD070 continuous passive scan	NESW Compass deviation threshold set, scan daemon running	WD070 daemon process verification — NESW baseline log review
7	Audit Trail	WD117 immutable logging active	All event types logging, immutability checksum active	Tiger .1x Key™ authenticated access — last event timestamp verified
8	Green Lanyard Roster	Clinician Sovereign certification	All shift caregivers certified, vow on record, renewal current	Certification log audit — Tiger .1x Key™ authenticated

9	24h Purge Daemon	Zero Cloud data sovereignty	Purge daemon PID active in OS init, last purge within 24h	OS process inspection — purge log timestamp verification
10	EDS Baseline	Elder dignity as constitutional threshold	EDS 5-domain score > 8.9/10 pre-administered	5-domain instrument administered, score recorded, 24h re-assessment scheduled

5.3 Research Applications of the Formal Audit Framework

— Formal verification research: each SVA property becomes a verification target — researchers prove that the hardware implementation satisfies the formal specification.

— Empirical validation: the audit is executed before each research session and all results logged — creating a session-by-session compliance record that supports publication.

— Comparative framework: the formal properties can be applied to other hardware-governed clinical AI systems — enabling cross-system constitutional comparison.

— Regulatory submission: the audit results, when formally documented with oscilloscope logs and SVA verification reports, constitute evidence for HSA Class B SaMD submission.

Chapter 6: Multi-Sensor Data Fusion — Monitoring Heart Function and Physical Distress

HRV Wearables · WM003™ LiDAR · Bounded Risk Advisory · Privacy by Physics

6.1 The Clinical Problem — Monitoring Without Surveilling

Conventional approaches to clinical monitoring of cardiac function and physical distress rely on one of two modalities: continuous audiovisual surveillance (cameras, microphones) or physically invasive sensors (electrodes, catheters). Both approaches impose significant costs on patient dignity, privacy, and psychological wellbeing. In eldercare settings, these costs are particularly severe — continuous visual surveillance suppresses natural movement, creates sustained psychological distress, and accelerates physical deconditioning.

The NAI 2.0™ framework resolves this through Multi-Sensor Data Fusion — integrating two non-invasive, non-visual, non-audio modalities that together provide comprehensive safety monitoring while constitutionally preserving the elder's dignity.

6.2 The Two Sensor Modalities

Layer	Sensor	What It Measures	Clinical Significance	Privacy Guarantee
Internal Biometric	HRV Wearable	Heart Rate Variability: RMSSD, SDNN, LF/HF ratio — autonomic nervous system state	Pre-fall sympathetic activation · Internal physiological distress before physical event occurs · Voice of the silent elder	Localised to the elder — no transmission except anonymised breach advisory · Zero Cloud
External Geometry	WM003™ Livox LiDAR — 96-cell voxel grid, 20Hz, 905nm	Postural geometry — Vertical → Transitional → Floor-level · Gait deviation · Zone boundary crossing	Fall detection · Postural instability correlated with HRV pre-event · Physical context confirming internal biometric signal	Point-cloud geometry only — no pixel, no face, no biometric identity · Zero Camera

6.3 How Multi-Sensor Fusion Works — The Bounded Risk Advisory

The Heartware system processes fused LiDAR geometry and HRV biometric data on the Nvidia Jetson Thor edge compute (air-gapped, hardened Linux, 24-hour ring-buffer purge). When a threshold crossing is detected, the system generates a Bounded Risk Advisory containing:

- Anonymised resident ID — a reference number only, no name, no photograph, no biometric identifier.
- Event type — fall detection (LiDAR trigger), HRV breach (wearable trigger), or fused multi-sensor event.
- Severity tier — pre-keyed by clinical team: Tier 1 (monitoring advisory), Tier 2 (priority), Tier 3 (immediate).

— Pre-keyed response options — three unranked, equally weighted options under Code A (Zero Weighting). One non-prescriptive alert under Code B.
— Routing: exclusively to the designated Green Lanyard caregiver on shift — no broadcast, no general alarm.

The AI detects. The AI advises. The human decides. This is the constitutional boundary — enforced by the Sacred Pause™ gate, authenticated by the Tiger .1x Key™, and logged immutably by WD117.

6.4 The Role of HRV in Research — Tracing the Electromagnetic Signature of the Heartbeat

Heart Rate Variability monitoring provides indirect access to the electromagnetic signals produced by cardiac activity. The wearable sensor measures the R-R interval variability in the ECG signal — the time between consecutive heartbeats — and derives the following metrics that are clinically significant in eldercare:

HRV Metric	What It Measures	Clinical Significance in Eldercare Research
RMSSD	Root mean square of successive RR interval differences — short-term HRV	Parasympathetic nervous system activity · Vagal tone · Decreases before many fall events · Primary Code A HRV monitoring metric
SDNN	Standard deviation of all RR intervals — overall HRV	Overall autonomic nervous system health · Long-term cardiac risk assessment · Baseline calibration metric
LF/HF Ratio	Low-frequency to high-frequency power ratio — sympathovagal balance	Sympathetic activation indicator · Elevated before acute physiological distress · Pre-fall sympathetic surge detection

In the research context, HRV metrics provide a quantitative, non-invasive measure of physiological state that can be correlated with clinical events logged in the WD117 Audit Trail — creating a dataset that links internal physiological signals to clinical outcomes without any visual or audio capture.

Chapter 7: Giving Voice to the Silent Elder — Multi-Sensor Fusion in Practice

7.1 Who Are the Silent Elders?

Silent elders are patients who are unable to self-report symptoms or distress due to conditions including cognitive decline, expressive aphasia, late-stage dementia, advanced Parkinson's disease, post-stroke communication impairment, or severe frailty. These patients represent a significant proportion of the nursing home and active ageing centre population — and they are structurally failed by conventional monitoring systems.

Conventional systems depend on patient self-report or visual caregiver observation. A fall or physiological crisis in a silent elder can occur without any report, without any verbal warning, and without any visual cue that a caregiver managing multiple patients can detect in time. By the time the crisis is discovered, preventable harm has already occurred.

7.2 How Multi-Sensor Fusion Provides a Constitutional Voice

Monitoring Layer	What It Detects	When It Detects It	How It Gives Voice
HRV Wearable — Internal Biometric	Pre-fall sympathetic activation · Autonomic nervous system stress · Cardiac irregularity	Before the physical event — HRV changes precede many falls by seconds to minutes	WD116 generates a Breach Advisory when HRV metric crosses pre-keyed threshold — advisory routes to Green Lanyard caregiver before the fall occurs
LiDAR — External Geometry	Postural transition (Vertical → Transitional → Floor-level) · Abnormal gait · Zone boundary crossing	During or immediately after the physical event — LiDAR detects fall kinematics within 50ms	WD113 generates a fall detection advisory — corroborates the HRV pre-event signal with physical confirmation
Fused Multi-Sensor Event	Internal distress signal + external physical confirmation — highest confidence advisory	Simultaneous or near-simultaneous HRV breach and LiDAR postural event	Tier 3 severity advisory — immediate — routed to Green Lanyard caregiver with full context: HRV data, postural event, anonymised resident ID

7.3 The Constitutional Boundary in Practice

The Multi-Sensor Data Fusion system operates under the strict constitutional boundary: the AI detects, the AI advises, the human decides. This boundary is not merely a policy commitment — it is enforced by the hardware architecture:

- The AI cannot initiate treatment: the Sovereign Brake (P-002) is a physical PLC relay that remains open until the Tiger .1x Key™ tripartite authentication is completed. The AI observation cannot translate into clinical action without a living, physically present, biometrically verified human clinician.
- The AI cannot diagnose: the Breach Advisory contains the anonymised resident ID, the breach type, the severity tier, and pre-keyed response options. It contains no diagnosis, no

recommended treatment, no probability score. The AI presents what it detected — not what it means.

— The AI cannot override: if the Green Lanyard caregiver reviews the advisory and decides no action is required, that decision is logged in WD117 as an authoritative human clinical act — not as a false positive or an AI error. The clinician's authority is absolute.

> "The system sees the outside of the elder through LiDAR geometry and listens to the inside of the elder through HRV biometrics — without ever seeing the elder's face or recording the elder's voice."

Chapter 8: Research Governance — Ethics, Data, and Constitutional Constraints

8.1 DSRB / IRB Requirements for NAI 2.0™ Studies

All NAI 2.0™ clinical research in Singapore requires DSRB (Domain Specific Review Board) approval under the National Medical Ethics Committee framework. Key considerations specific to NAI 2.0™ deployments:

- 3ZEROS™ privacy architecture satisfies the data minimisation principle of the Personal Data Protection Act (PDPA) and the research ethics requirement for minimum data collection — no visual data, no audio data, no identifiable patient information.
- WD117 anonymised audit trail data may be used for research without individual patient consent, subject to DSRB data-sharing agreement — all data consists of anonymised timestamps, event types, and severity tiers.
- Green Lanyard caregivers participating in research require informed consent as research participants in addition to their clinical role.
- The 24-hour purge daemon operates regardless of research protocols — the research schedule must be designed around the constitutional data retention limit, not the other way around.

8.2 Key Research Governance Requirements

Requirement	Specification	Constitutional Constraint
10-Point Sovereignty Audit	Full audit must pass before every research session	Any single audit failure halts the session — constitutional requirement, not research discretion
Green Lanyard certification	All caregivers accessing advisories during research must be certified	No uncertified caregiver may action advisories — constitutional requirement
WD117 data access	Tiger .1x Key™ tripartite authentication required for every audit trail export	Research data export requires framework sovereign or designated auditor authentication
24-Hour purge	Purge daemon runs at 00:00 regardless of research session timing	Longitudinal data collection must account for daily purge — only WD117 anonymised event records persist
Publication standards	Full hardware specification must be documented for replication	Implementation tier, Sacred Pause™ ROM constants, Code B status, Green Lanyard roster, Sovereignty Audit results — all required for methods section

Appendix A: Study Design Summary Table

Study Type	Direction	Setting	Primary Outcome	Data Source	Duration
Observational — Ward Pilot	Direction 1	Live ward · 20–30 elders · 6-hour sessions	Sacred Pause™ effect on alert fatigue	WD117 Audit Trail · Pre/post Goddard et al. instrument	Minimum 3 sessions
Simulation Study	Direction 1	High-fidelity simulation · standardised scenarios	Tiger .1x Key™ cognitive load and response time in Code B	Oscilloscope logs · auth completion rates · NASA-TLX	Minimum 2 simulation sessions
Comparative Study	Direction 1	Two parallel ward groups	Alert fatigue reduction vs conventional system	WD117 vs comparator audit log · caregiver self-report	Minimum 3 months
Formal Verification	Direction 2	Formal tools (JasperGold / Cadence) · test deployment	Mathematical proof — WD070–073 soundness and completeness	SVA assertions · formal verification report	Ongoing — model development
Empirical Drift Simulation	Direction 2	Test NAI 2.0™ deployment with controlled perturbations	Drift response surface — detection latency, freeze accuracy, restoration quality	WD070 detection log · WD073 drift signature measurement	Minimum 30 perturbation trials
Constitutional Audit Publication	Direction 2	Framework development + NAI 2.0™ reference validation	Peer-reviewed formal audit framework for hardware-governed clinical AI	10-Point Sovereignty Audit × formal properties mapping	6–12 month publication arc

Appendix B: SVA Formal Properties — 10-Point Sovereignty Audit

The following SystemVerilog Assertions (SVA) formally specify the constitutional properties verified at each audit point. These properties constitute the formal specification of the NAI 2.0™ constitutional architecture.

Audit Point	SVA Property	Formal Assertion	Constitutional Meaning
2 — FPGA Timer	SVA Property 1	assert property (@(posedge clk) delay_value inside {Range1, Range2, Range3})	Sacred Pause™ ROM constants are the only permissible delay values — no runtime modification possible
3 — Dual Counter	SVA Mismatch	assert property (@(posedge clk) counter_A == counter_B)	Both redundant counters must agree — mismatch triggers watchdog and system halt
4 — PLC Tripartite	SVA Property 4	assert property (@(posedge clk) relay_close |-> (iris_valid && console_confirmed && pedal_closed))	Clinical action relay can only close if all three authentication inputs are simultaneously active
5 — Code B	SVA Property 3	assert property (@(posedge clk) code_b_active |-> lidar_physical_threshold)	Code B can only be active when a physical LiDAR threshold has been crossed — never by algorithmic score
General — Non-Agentic	SVA Property 0	assert property (@(posedge clk) ai_output_generated |-> ##[1:MAX_DELAY] human_auth_completed)	Every AI-generated advisory must be followed by human authentication before any clinical action

References

Goddard, K., Roudsari, A., & Wyatt, J. C. (2014). Automation bias — empirical results assessing influencing factors. International Journal of Medical Informatics, 83(5), 368–375.

International Electrotechnical Commission. (2006). IEC 62304: Medical device software — Software life cycle processes. IEC.

International Electrotechnical Commission. (2010). IEC 61508: Functional safety of E/E/PE safety-related systems. IEC.

Kahneman, D. (2011). Thinking, fast and slow. Farrar, Straus and Giroux.

Koh, E. W. K. (2026). ABC+2S+H Guardian Framework™ [P-001]. IPOS SG020603109STW. ACRA T260229801.

Koh, E. W. K. (2026). WD070–WD073: Drift Governance Protocols. Prior Art.

Koh, E. W. K. (2026). WD113–WD117: Group H Terrestrial Eldercare Patents. Prior Art.

Koh, E. W. K. (2026). CG No. 1–5. NAI 2.0™ Clinical Governance Monograph Series.

Low, S. E. (2026, April 2). RE: Non-Agentic Eldercare AI — Academic-Non-Commercial Collaboration. NTU LKCMedicine.

Ministry of Health Singapore. (2026). AIHGle 2.0. MOH Holdings.

Vered, M., Shani, G., & Zahavi, H. (2023). Explanations and enforced deliberation for reduced automation bias. Computers in Human Behavior, 148, 107929.

謙虛 · 沉默 · 尊严 · 仁 · 止於至善

P-LIFE 1.00™ · Harm = Death · North = Save Life · We Innovate Save Live™

Edwin Koh Wui Kiat (Tiger) · Founding Father, Non-Agentic AI 2.0™

ACRA T260229801 · Patent SG020603109STW · Singapore · 2026