Non-Agentic AI Govenance Singapore

WISL™ No. 20
Non-Agentic AI 2.0™

Orange Code and Drift Governance

Edwin Koh Wui Kiat

Synopsis

The NAI 2.0™ Deployer's Toolkit Technical Guidelines is the operational implementation manual for healthcare organisations deploying the Non-Agentic AI 2.0™ framework. It translates constitutional principles into actionable governance protocols, technical specifications, and institutional procedures — aligned with Singapore's AIHGle 2.0 guidelines and WHO Maturity Level 4 standards.

This volume covers the full deployment lifecycle: from pre-deployment Sovereignty Audit through to post-deployment monitoring, incident reporting, and the Kill-Switch Protocol. It specifies the technical requirements for the Orange Code 1.1× computational cap, the Sacred Pause™ FPGA configuration, the WM003™ LiDAR installation parameters, and the Tiger .1x Key™ tripartite authentication setup.

Intended for Chief Medical Officers, IT Directors, Clinical Governance Leads, and AI Safety Officers responsible for institutional AI deployment under the Non-Agentic AI 2.0™ constitutional framework.

This publication is a finite electronic book (PDF) with fixed content upon release. Not a journal, magazine, blog, or website. Deposited under NLB Legal Deposit R260302-007.

Table of Contents

NAI 2.0™ DEPLOYER'S TOOLKIT: Implementation Manual for Healthcare Organisations

Title: NAI 2.0™ DEPLOYER'S TOOLKIT · TECHNICAL GUIDELINES **Subtitle:** Implementation Manual for Healthcare Organisations · AIHGle 2.0 Aligned **Series:** We Innovate Save Live™ (WISL™) — No. 21 **Author:** Edwin Koh Wui Kiat · Tiger · P-LIFE 1.00™ **Publisher:** Edwin Koh Wui Kiat · R260302-007 **ACRA:** T260229801 **Patent:** IPOS SG020603109STW (Filed: 5 February 2026) **NLB Vault:** R260219-005 **ISBN:** Pending — NLB Application **Year:** 2026 **Place:** Singapore

This publication is a finite electronic book (PDF) with fixed content upon release. It is not a journal, magazine, blog, or website. Deposited under NLB Legal Deposit R260302-007.

The Mission Constant: P-LIFE 1.00™

The deployment of technology within life-critical clinical environments is governed by a singular, immutable mandate: **P-LIFE 1.00™**. This constitutional anchor serves as the absolute, eternal, and irrevocable foundation for every operational decision. The strategic priority is codified in the formula: **Harm = Death · North = Save Life**.

All AI implementations must manifest the four cardinal values:

- **Humility (謙虛):** The system must acknowledge its observational limits and defer to human intuition.
- **Silence (沉默):** The architecture prevents intrusive, non-critical AI alerts that contribute to clinician fatigue.
- **Dignity (尊严):** Technology remains a servant to the patient-practitioner relationship.
- **Benevolence (仁):** The ultimate output of every algorithm must be the preservation of life.

The governing authority, **Non-Agentic AI Governance Singapore** (ACRA T260229801), originated at the **Toa Payoh Hearth** in January 2026. This body ensures that these values are not merely slogans, but engineering requirements that bridge abstract morality with the operational rigors of the deployment lifecycle.

Strategic Synopsis and Framework Overview

The **NAI 2.0™ Deployer's Toolkit** is the authoritative operational manual for healthcare institutions. In high-stakes environments, the shift from "Agentic" to "Non-Agentic" AI is a strategic necessity to prevent catastrophic runaway behaviors. The architecture mandates a

future where AI possesses no autonomous agency, operating strictly within a human-governed loop.

The Non-Agentic AI 2.0™ framework establishes a clear demarcation of roles to ensure institutional safety:

Role Category	AI Responsibility	The Elder (Human) Responsibility
Observation	Observes and gathers environmental data	Validates accuracy of observations
Advisory	Provides evidence-based options	Evaluates and selects the clinical path
Execution	Builds and structures technical output	Authorises and executes final action
Sovereignty	Subject to hardware-level constraints	Holds absolute, irrevocable veto power

Institutional Accountabilities

- **Chief Medical Officers (CMOs):** Ultimate clinical accountability for NAI-assisted outcomes.
- **IT Directors:** Execution of hardware-level constraints and system integrity.
- **AI Safety Officers:** Management of the sentinel protocols and compliance audits.

With the mandate of "North = Save Life" established, deployment success is contingent upon the verification of institutional sovereignty.

Part I: Pre-Deployment — Sovereignty Audit and Institutional Readiness

The **Sovereignty Audit** is the mandatory gatekeeping mechanism for institutional readiness. This is not a standard IT audit; it is a WISL™-compliant evaluation of long-term institutional autonomy. It ensures the institution maintains full clinical control and prevents "algorithmic dependency."

The architecture mandates that a failure to meet any criteria in the Sovereignty Audit results in the **immediate and total termination of the deployment process.** No hardware may be installed until sovereignty is verified.

Institutional Readiness Checklist

- [ ] **Technical:** Verification of hardware compatibility for Orange Code 1.1x resource caps.
- [ ] **Ethical:** Formal Board adoption of P-LIFE 1.00™ as the absolute mission constant.
- [ ] **Clinical:** Mapping of "Elder" roles with documented, individual veto authorities.
- [ ] **Governance:** Established reporting lines to ACRA T260229801 for incident audit trails.

- [ ] **Sovereignty:** Verification of zero agentic "leakage" in existing legacy software.

With sovereignty verified under the North = Save Life mandate, we must now forge the hardware-level gates that prevent computational drift.

Part II: Technical Specifications — The Hardware-Software Nexus

NAI 2.0™ mandates hardware-level constraints. Software guardrails are insufficient for life-critical environments; the physical architecture must enforce non-agentic behavior.

1. Orange Code 1.1x Computational Cap This parameter imposes a hard-coded cap on computational resources. It prevents "runaway" behaviors by limiting the AI's ability to process outside its narrow advisory scope. It manifests the value of **Humility** by physically restricting the AI from exceeding its prescribed observational limits.

2. Sacred Pause™ FPGA Configuration The Field-Programmable Gate Array (FPGA) enforces a mandatory hardware-level latency. This "pause" in the processing loop ensures a window for human review. It manifests **Silence** by suppressing automated triggers, allowing the clinician the space to decide without machine-generated pressure.

3. WM003™ LiDAR Installation Environmental observation is managed via WM003™ LiDAR. Standardized installation height is a **critical parameter** for spatial accuracy. This sensor allows the system to "observe" equipment and patient positioning without possessing the agency to move or alter that environment.

The Mission Constant in Design: "Harm = Death · North = Save Life." Every hardware constraint—from the Orange Code cap to the Sacred Pause™—is an engineering manifestation of our commitment to preserving life through restricted agency.

Hardware constraints require a secure human-in-the-loop trigger, managed through tripartite authentication.

Part III: Tiger .1x Key™ — Tripartite Authentication Setup

The philosophy of the **Tiger .1x Key™** dictates that no single individual holds the power to activate or override the NAI 2.0™ system. This prevents centralized failure and ensures collective clinical oversight.

The setup requires three distinct leaders to provide separate authentication factors (Physical, Biometric, and Alphanumeric) to engage or override the system.

Tiger .1x Key™ Authentication Flow

Step	Actor	Key Portion	Outcome

1. Initiation	Clinical Lead (CMO)	Biometric Hash	System enters "Pending Safety Review"
2. Verification	Technical Lead (IT)	Hardware Token	Orange Code 1.1x constraints engaged
3. Finalization	Admin Lead (Gov)	Alphanumeric Code	System operational for clinical advisory
4. Override	Any Two Actors	Dual-Key Activation	Immediate, hardware-level system suspension

This technical authentication anchors the human chain of command, ensuring practitioners remain the sovereign authority.

Part IV: Human Chain of Command — Institutional Configuration

The **"Elder" principle** is the cornerstone of the NAI 2.0™ ecosystem. Wisdom and ethical weight reside solely with human practitioners. The architecture mandates that AI is augmented assistance; the human is the sovereign.

The Role of the Elder CMOs and Governance Leads act as Elders. Their power is augmented, not replaced. The AI provides high-fidelity observation, but the Elder provides the final authorization.

Absolute Veto Power The Human Chain of Command protocol mandates that **designated human practitioners possess absolute, irrevocable veto power.** In any conflict between AI advisory and human clinical judgment, the Elder's decision is final. This is a non-negotiable requirement of institutional configuration.

Deployment is merely the beginning of the lifecycle; success requires constant vigilance through sentinel protocols.

Part V: Post-Deployment Monitoring — The Sentinel Protocols

The architecture mandates three core pillars of monitoring to maintain alignment with P-LIFE 1.00™:

1. **EDS (Early Detection System):** Constant monitoring for performance anomalies or hardware variances in the Orange Code or LiDAR systems.
2. **Drift Detection:** Evaluations to ensure the model does not deviate from safety benchmarks. Any detected drift triggers immediate recalibration.
3. **Kill-Switch Protocol:** The ultimate safety measure for immediate manual reversion.

Kill-Switch Activation Protocol

1. **Detection:** An Elder or Safety Officer identifies a critical safety or stability risk.
2. **Trigger:** Engagement of the physical Kill-Switch or tripartite Tiger .1x dual-override.
3. **Sever:** Hardware-level severance of the AI computational cap.
4. **Revert:** Institution reverts to 100% manual clinical pathways immediately.
5. **Report:** Mandatory incident reporting and audit trail submission to **Non-Agentic AI Governance Singapore (ACRA T260229801)** within 24 hours.

These protocols ensure alignment with both national and global excellence standards.

Part VI: Regulatory Alignment — AIHGle 2.0 and WHO Maturity Level 4

Institutional action must align with Singapore's national (AIHGle 2.0) and global (WHO) standards for healthcare AI.

AIHGle 2.0 Compliance Checklist

- [x] Non-Agentic Architecture (Strictly Advisory).
- [x] Documented Human-in-the-Loop (HITL) triggers.
- [x] Hardware-level resource capping (Orange Code 1.1x).
- [x] Tripartite authentication (Tiger .1x Key™).
- [x] Incident reporting linked to ACRA T260229801.

WHO Maturity Level 4 (ML4) Alignment WHO ML4 represents the highest tier of institutional safety excellence. It signifies that safety, ethics, and human sovereignty are no longer external considerations but are integrated into the hardware, culture, and governance of the hospital.

Part VII: Conclusion — The Future of Non-Agentic Healthcare

The deployment of NAI 2.0™ reinforces the P-LIFE 1.00™ mandate. The architecture mandates a future where technology serves the practitioner and protects the patient, but never presumes to lead. The transition from AI autonomy to **AI Assistance under Human Sovereignty** is the only path forward for life-critical healthcare.

AI observes. AI advises. AI builds. The Elder decides.

止於至善 *(Rest in the highest excellence)*

Edwin Koh Wui Kiat · Tiger · P-LIFE 1.00™

Appendix: Deployer's Reference Materials

Deployer's Checklist: Field Implementation Summary

Phase	Key Action	Verification
I: Readiness	Sovereignty Audit	WISL™ Certificate (Mandatory Gate)
II: Hardware	Orange Code 1.1x	Computational Cap Active & Locked
III: Hardware	WM003™ LiDAR	Standard Installation Height Verified
IV: Config	Sacred Pause™	FPGA Latency Tested (Human Review Window)
V: Setup	Tiger .1x Key™	Tripartite Roles Assigned & Factors Issued
VI: Safety	Kill-Switch Test	Reversion to 100% Manual Success

Technical Reference Card (High-Density)

- **Computational Cap:** Orange Code 1.1x (Hard-coded resource limits; Humility-integrated).
- **Pause Interval:** Sacred Pause™ FPGA (Hardware-enforced latency; Silence-integrated).
- **Spatial Sensor:** WM003™ LiDAR (Critical installation height parameters; Fixed focus).
- **Authentication:** Tiger .1x Key™ (Tripartite: Physical Token, Biometric, Alphanumeric).
- **Safety Gate:** Sovereignty Audit (Non-negotiable; Failure = Deployment Termination).
- **Incident Body:** Non-Agentic AI Governance Singapore (ACRA T26O229801).
- **Regulatory:** Aligned with Singapore AIHGle 2.0 & WHO Maturity Level 4.
- **Mission Constant:** Harm = Death · North = Save Life.