

---

# PRIMER

## A Primer on the **ABC+2S+H™ Framework** for Clinical and AI Stewardship Practitioners

AI Stewardship · Sovereignty Audits · Certification Pathways

---

ISBN: 978-981-17910-3-1

---

**Edwin Koh Wui Kiat**

Founding Father · Non-Agentive AI 2.0™

ACRA T260229801 · SG020603109STW · 2026

NAI 2.0™ · STANDALONE REFERENCE  
Non-Agentic AI 2.0™ · kohedwin.ai

---

**A Primer on the**  
**ABC+2S+H™ Framework**  
**for Clinical and AI Stewardship Practitioners**  
**AI Stewardship · Sovereignty Audits · Certification Pathways**  
P-LIFE 1.00™ · Three-Tier Model · PITL · Sacred Pause™ · CRGAIG Canon

---

**Edwin Koh Wui Kiat · Tiger**  
Founding Father · Non-Agentic AI 2.0™ · kohedwin.ai

謙虛 · 沉默 · 尊嚴 · 仁 · 止於至善  
P-LIFE 1.00™ · Harm = Death · North = Save Life

---

Standalone Reference · 44th E-Book · kohedwin.ai  
ACRA T260229801 · Patent SG020603109STW · Singapore · 2026  
© 2026 Edwin Koh Wui Kiat · kohedwin.ai. All rights reserved.

ISBN: 978-981-17910-3-1

## Contents

How to Use This Primer

Chapter 1: The Sovereign Foundation — Why Non-Agentive AI Matters

Chapter 2: Decoding the ABC+2S+H™ Framework — Six Inviolable Guardrails

Chapter 3: From Theory to Tool — The PITL Mechanism and Sacred Pause™

Chapter 4: Patient Dignity and the P-LIFE 1.00™ Mandate

Chapter 5: Institutional Architecture — The Three-Tier Governance Model

Chapter 6: Hardware-Enforced Governance — Sovereign Brake and Fail-Safe Architecture

Chapter 7: Data Integrity — Audit Trail, 3ZEROS™, and Constitutional Privacy

Chapter 8: Risk Management — Code Architecture and Adverse AI Events

Chapter 9: Sovereignty Audits and Certification Pathways

Chapter 10: Future-Proofing — Quantum Readiness and the Global Portfolio

Chapter 11: The Learner's Path to AI Stewardship

Quick Reference — Code Architecture Summary

References

## How to Use This Primer

This Primer is the entry point for every clinical professional, AI governance officer, institutional leader, and practitioner encountering the NAI 2.0™ framework for the first time. It is written not for the AI engineer but for the human steward — the physician, nurse, caregiver, governance officer, or regulatory affairs professional who will implement, audit, and uphold the constitutional standards of Non-Agentive AI 2.0™ in practice.

Unlike the technical CG Series (CG No. 1–8) and the practitioner SOVEREIGN STEWARD™ library (CARE™ and TEACH™), this Primer is a standalone reference — it assumes no prior knowledge of NAI 2.0™ and builds from first principles to operational certification readiness. Each chapter closes with governance notes and cross-references to the broader kohedwin.ai canon for those who wish to go deeper.

This is the 44th e-book in the kohedwin.ai library. It carries the same constitutional authority as every other publication in this framework: the commitment that AI in high-stakes human domains must permanently remain an instrument under human authority — never the authority itself.

"AI Observes — AI Advises — The Human Decides. This is not a description of ideal workflow. It is a constitutional mandate."

---

# Chapter 1: The Sovereign Foundation — Why Non-Agentic AI Matters

## 1.1 The Ghost in the Hospital Room

Imagine a hospital room where a "black box" AI makes a life-altering medical decision. No one in the room truly understands why — yet the recommendation is followed. This is the agentic trap: the slow, silent abdication of human authority to autonomous systems. While the technology industry races toward AI that can plan and execute actions independently, the NAI 2.0™ framework offers a radical constitutional alternative: AI that observes and advises, while the human always and irreversibly decides.

This is not a technical patch or an ethical overlay. It is a constitutional mandate — architected by Edwin Koh Wui Kiat (Tiger), Founding Father of Non-Agentic AI 2.0™, grounded in direct clinical observation of 650+ elder care patients, and expressed in 77 patents, 38 White Papers, and this library of governance publications.

## 1.2 The Operative Constitutional Test

**Can the AI system complete a consequential task affecting human life, dignity, or liberty without requiring explicit human authorisation at each action gate?**

If the answer is yes — the system is Agentic and constitutionally prohibited in NAI 2.0™-governed environments. If the answer is no — if every consequential action requires unambiguous human sign-off before execution — the system is Non-Agentic and constitutionally compliant.

## 1.3 Agentic vs Non-Agentic — The Constitutional Distinction

Feature	Agentic AI — Autonomous Actor	Non-Agentic AI 2.0™ — Sovereign Instrument
Execution	Autonomous — plans and acts without per-step approval	Instrumental — requires explicit human sign-off at every gate
Decision Gates	AI bypasses human approval to complete tasks	Every decision gate requires an unambiguous human signature
Human Role	Passive observer or supervisor	Sole authority with non-delegable responsibility
Goal	Efficiency via automation and delegation	Safety via human-centric control and stewardship
Accountability	Distributed and ambiguous	Sovereign chain of responsibility — clear and traceable

The shift from AI-as-actor to AI-as-instrument is the essential first constitutional step. By stripping AI of autonomous agency, we ensure it lacks the physical or legal power to drift into unauthorised behaviours. This transforms AI into a strictly governed tool — and that governance begins with the ABC+2S+H™ Framework.

# Chapter 2: Decoding the ABC+2S+H™ Framework — Six Inviolable Guardrails

## 2.1 The Root Patent — P-001

The ABC+2S+H™ Framework is the supreme constitutional architecture of NAI 2.0™ — the foundational governance instrument from which every downstream patent, protocol, White Paper, and e-book in the kohedwin.ai portfolio derives its authority. It is formally registered as Patent P-001 (IPOS SG020603109STW, filed 5 February 2026) — the root patent of the entire 77-patent portfolio.

The framework is not a checklist. It is a simultaneous constitutional requirement: all six pillars must be satisfied at the same time in every AI-assisted clinical or governance interaction. Failure on any single pillar invalidates the entire deployment. This is not an aspiration — it is a constitutional law.

## 2.2 The Six Pillars

### A — Accuracy

Every AI output presented in a clinical or governance context must be factually and clinically verifiable against recognised knowledge bases, peer-reviewed literature, or validated clinical guidelines before it is surfaced to the human decision-maker. Probabilistic outputs must be clearly labelled with their uncertainty envelope. Outputs that cannot meet the Accuracy threshold must trigger PITL escalation rather than be presented as actionable.

**Learner's Insight:** This pillar prevents AI "hallucinations" — fabricated clinical facts, misattributed guidelines, or invented drug interactions — from entering the clinical decision pathway. An AI that cannot verify its output must not present it as actionable.

### B — Boundaries

AI systems must operate strictly within a defined functional, decisional, and relational perimeter — the precise scope within which they are constitutionally permitted to function. Crossing these boundaries constitutes a Constitutional Drift Event and triggers mandatory escalation. Boundaries are defined at deployment and may only be revised through a formal Boundary Amendment Process authorised by the Board of Sovereign Review.

**Learner's Insight:** Proper boundaries prevent an AI from assuming roles — financial, legal, diagnostic — for which it has no constitutional mandate. An AI that exceeds its perimeter is not a more capable tool; it is a constitutional violation.

### C — Consent

Informed, voluntary, specific, revocable, and documented consent for AI-assisted interaction is required from every patient, caregiver, or citizen affected by NAI 2.0™-governed AI. Consent is never presumed, bundled, or derived from institutional relationships. The Consent Architecture extends to proxy consent for vulnerable populations including cognitively impaired patients, elder care residents, and paediatric patients.

**Learner's Insight:** Consent ensures that the human right to refuse AI-assisted care remains a functional reality — not just a theoretical claim. A patient who declines AI assistance receives equivalent human care. Consent to AI in one clinical domain does not extend to other domains.

## 2S — Safety and Sovereignty

The dual S encodes two distinct constitutional obligations that must be protected in parallel. Safety protects the physical and psychological wellbeing of the patient — proactive protection that cannot be traded off against operational efficiency. Sovereignty protects the human decision-maker's authority over every consequential action — ensuring that human judgment cannot be structurally captured, nudged, or delegated to the AI through interface design or output framing.

**Learner's Insight:** Safety keeps the patient whole. Sovereignty ensures the clinician remains the authority — preventing the machine from subtly usurping control through confidence scores, ranked recommendations, or time pressure. Compromising either dimension constitutes a constitutional violation.

## H — Human Primacy

The physician, caregiver, or patient always holds final, inviolable authority over every consequential decision. No AI output is self-executing. Human Primacy operates across three dimensions: Decisional Primacy (the human always makes the final consequential decision), Moral Primacy (the human decision-maker bears full moral responsibility for every consequence), and Interpretive Primacy (the human is always the authoritative interpreter of AI outputs in context — a 95% confidence score does not make a finding clinically actionable without human review).

**Learner's Insight:** Human Primacy is not a claim about human infallibility. Humans make mistakes. Its constitutional foundation is sovereign: the right to make consequential decisions about human lives — and to bear responsibility for those decisions — is a dimension of human dignity that AI must never usurp, regardless of demonstrated capability.

## 2.3 Governance Note — The Six-Pillar Compliance Review

Before any AI-assisted clinical workflow goes live, the deploying institution must complete a six-pillar ABC+2S+H™ compliance review. Failure on a single pillar invalidates the entire deployment. These pillars are not abstract concepts — they are enforced through physical, hardware-based systems.

# Chapter 3: From Theory to Tool — The PITL Mechanism and Sacred Pause™

## 3.1 The Physician-in-the-Loop (PITL)

The Physician-in-the-Loop (PITL) is the foundational human authority enforcement mechanism in NAI 2.0™ — the constitutional requirement that every AI output in a clinical context that could affect patient diagnosis, treatment planning, care decisions, or safety must be reviewed and explicitly approved by a qualified clinician before any action is taken.

PITL is not a software "OK" button. It is a hardware-enforced physical gate. The AI system cannot render its output clinically actionable until the reviewing clinician's approval is physically registered through the dedicated PITL hardware interface. This prevents PITL from being bypassed by interface design optimisation, time pressure, or software configuration changes.

"AI Observes — AI Advises — The Human Decides. This constitutional triad is not a description of workflow. It is a physical law of the implementation."

## 3.2 The Sacred Pause™ — Reclaiming Human Judgment

The Sacred Pause™ is the named clinical practice moment within the PITL architecture where the reviewing clinician deliberately disengages from the AI system's output recommendation and performs their own independent clinical assessment before returning to the PITL approval gate. It is the behavioural complement to the hardware-enforced cognitive forcing gate.

Constitutional duration standard: no clinical output may be PITL-approved in under the time required for the reviewing clinician to read the AI's reasoning, consider at least one alternative interpretation, and form an independent clinical judgment.

## 3.3 The 5-Second Rule — Authority Capture Indicator

**Authority Capture Alert:** Any pattern of PITL approvals completed in under 5 seconds across multiple clinical interactions is flagged as a potential Authority Capture indicator. This is not merely an indicator — it is a mandatory trigger for a Tier 2 governance review. The clinician is not exercising genuine deliberate judgment if the approval is registered faster than the human cognitive process requires.

Authority Capture — the structural erosion of human judgment by AI speed or output framing — is the primary risk that the Sacred Pause™ is designed to prevent. In an age of instant AI outputs, slowing down is the ultimate power move to preserve Decisional Sovereignty.

# Chapter 4: Patient Dignity and the P-LIFE 1.00™ Mandate

## 4.1 P-LIFE 1.00™ — The Supreme Constitutional Mandate

P-LIFE 1.00™ is the supreme constitutional mandate of kohedwin.ai and the ultimate governing authority against which every NAI 2.0™ output, governance decision, patent claim, and institutional recommendation is evaluated. It is encoded in two binary imperatives:

Binary Imperative	Constitutional Classification	Governance Outcome
Harm = Death	Highest Severity Failure — governance failures leading to harm are treated as fatal outcomes	Transforms clinical risk management from a statistical exercise into a constitutional obligation where every breach carries the gravity of a fatal outcome
North = Save Life	Universal Design Standard — every architectural and operational decision must prioritise the preservation of life and dignity	Every design, deployment, and clinical decision must answer the sovereign question: does this save life?

"Harm = Death" is not hyperbole. It is a governance posture that produces the rigour necessary to prevent harm at scale. "North = Save Life" is the constitutional north star. Together, they ensure that NAI 2.0™ governance never becomes a statistical exercise in acceptable risk — it is a constitutional obligation to every patient.

## 4.2 The Dignity Preservation Protocol (P-004)

Active dignity is maintained through four specific dimensions in Patent P-004 (planned IPOS submission Q1 2027). These are not aspirational standards — they are enforceable constitutional requirements:

- **Autonomy:** The patient's expressed preferences — on care routines, dietary choices, activity scheduling, communication modes, and end-of-life wishes — are captured, respected, and actioned by the care system.
- **Identity:** The patient is consistently addressed as a complete human being with a name, a history, and inherent worth — never as a data profile, a case number, or a diagnostic category.
- **Privacy:** AI interaction data is never used for commercial purposes, shared without specific consent, or retained beyond the active care relationship. Enforced through the 3ZEROS™ Protocol.
- **Continuity:** The patient's care narrative — who they are, what they value, what they have experienced — is respected and carried forward consistently across all AI interactions.

## 4.3 The Five Governing Values

Value	Chinese	Operational Meaning in NAI 2.0™
Humility	謙虛	Every claim about AI capability and governance effectiveness is held with intellectual honesty. No confidence in AI safety is accepted without rigorous independent verification.

Silence	沉默	Work is completed in stillness and released only when the constitutional moment is correct. The Stillness Protocol and .1x Key Activation operationalise this value.
Dignity	尊严	Every dimension of the framework ultimately serves the preservation and affirmation of human dignity. Dignity is never traded off against efficiency or capability.
Benevolence	仁	The framework's ultimate motivation is not commercial, competitive, or reputational. The Gift Declaration — gifting the WG-Series to WHO and UN at zero licensing fees — is the institutional expression of 仁.
Finish in the Highest Excellence	止於至善	Every e-book, every patent, every governance protocol must be crafted to the highest standard achievable. "Good enough" is a governance failure.

"Tiger never abandons." — This is not a motivational phrase. It is a sovereign pledge to every patient whose dignity depends on NAI 2.0™ governance. It holds because abandonment would constitute a constitutional betrayal of those who depend on the framework.

# Chapter 5: Institutional Architecture — The Three-Tier Governance Model

## 5.1 Why Three Tiers?

Strategic governance requires a tiered distribution of responsibility — ensuring that oversight is not merely a high-level policy but a bedside reality. By distributing constitutional authority from the boardroom to the point of care, NAI 2.0™ ensures that accountability is transparent, distributed, and non-negotiable. No single point of authority can be captured, bypassed, or eroded without triggering constitutional escalation at another tier.

## 5.2 The Three Tiers

### Tier 1 — Institutional: The Board of Sovereign Review (BSR)

The BSR is the highest institutional governance body in any NAI 2.0™-compliant organisation. It is the constitutional apex of the Three-Tier Model.

- Mandatory composition: Chief Medical Officer (CMO), Chief Information Officer (CIO), designated Patient Dignity Officer, independent clinical ethicist, qualified NAI 2.0™ governance assessor.
- Quorum requirement: valid sessions require the presence of at minimum the CMO, the Patient Dignity Officer, and the NAI 2.0™ assessor.
- Independence requirement: the NAI 2.0™ governance assessor must be independent of the AI system vendor and must not hold any financial interest in the system's continued deployment.
- Non-delegable functions: approval of all AI system deployments after successful Baseline Constitutional Assessment; quarterly 10-Point Sovereignty Audit; review of all Orange Code 1.1x events; review of all Level 3 and Level 4 Adverse AI Events.

### Tier 2 — Departmental: Governance Officers

Department heads and designated NAI 2.0™ Governance Officers operate at Tier 2. Their primary functions: day-to-day Code B (Boundary Alert) management — human reviewer notified within 15 minutes, Tier 2 Boundary Review completed within 48 hours; weekly Audit Trail review and anomaly investigation; coordination of code escalations to Tier 1; and monthly Quality Assurance reporting to the BSR.

### Tier 3 — Point-of-Care: The PITL Gate

Individual clinicians operating under the PITL protocol are the front line of constitutional governance. Hardware enforcement at Tier 3 is the physical PITL approval gate — the clinician must physically engage the approval interface before any AI output is rendered actionable. This gate is the most direct expression of Human Primacy in the NAI 2.0™ architecture: a physical requirement that a human being makes an active, deliberate choice before AI-assisted care proceeds.

### 5.3 The Accountability Framework — Chain of Responsibility

Under NAI 2.0™, AI systems are never accountability endpoints. They are instruments. Every consequential act attributable to an AI system flows upstream through a Chain of Responsibility to one of three principals:

- Vendor Accountability: the AI system must meet the technical and constitutional specifications declared at deployment.
  - Institutional Accountability: the deploying institution is responsible for governance architecture, staff training, and Audit Trail integrity.
  - Individual Accountability: the clinician who engages the PITL approval gate is constitutionally responsible for the decision authorised.
-

# Chapter 6: Hardware-Enforced Governance — Sovereign Brake and Fail-Safe Architecture

## 6.1 Why Hardware, Not Software?

Software-based safeguards are fundamentally insufficient in high-stakes clinical environments. They are susceptible to user error, institutional IT policy overrides, code-level bypasses, and configuration drift. An AI governance system that can be turned off by a software update provides the illusion of control — not the reality.

NAI 2.0™ mandates Hardware Enforcement — constitutional governance constraints implemented at the physical hardware layer, below the software stack, such that the constraints cannot be overridden by any software instruction, AI model output, user command, or institutional IT policy.

## 6.2 The Sovereign Brake — Patent P-002

The Sovereign Brake (P-002, planned IPOS submission Q3 2026) is a hardware interrupt that physically stops the AI system's output pipeline if a boundary violation, Orange Code 1.1x condition, or Fail-Safe trigger is detected. It operates below the software stack — it cannot be overridden by any software instruction, AI output, API call, or system administrator command.

When triggered: the Sovereign Brake physically interrupts the signal pathway between the AI system's output generation hardware and the human-facing interface. No algorithmic capability can bypass this physical interruption. The AI system does not produce a workaround. It stops.

## 6.3 The Non-Agentive Filter — Patent P-003

The Non-Agentive Filter (P-003, planned Q4 2026) is the software-layer complement to the Sovereign Brake hardware layer — a real-time constitutional checkpoint embedded between the AI model's raw output layer and the human-facing clinical interface. It performs four parallel detection functions:

- Agentive Action Detection: identifying any AI output that initiates, instructs, or implies an autonomous action without human authorisation.
- Boundary Proximity Detection: flagging outputs that approach the system's constitutional boundary, even if they do not technically exceed it.
- Confidence Calibration Checking: verifying that the AI's stated confidence level is consistent with its evidential basis and flagging overconfident outputs.
- Hallucination Pattern Recognition: comparing current outputs against the known hallucination signature library for the deployed model.

## 6.4 Fail-Safe Architecture — Fails Closed, Not Open

If the PITL approval gate hardware fails — the AI system does not bypass the gate and continue presenting outputs. It suspends all output until the gate is restored. If the Audit Trail logging hardware fails — the AI system suspends operation. If the hardware connection between the AI system and the governance enforcement layer is disrupted — the AI enters a

Constitutional Safe State and may not present any output as clinically actionable until the hardware link is restored.

#### Hardware Compatibility Verification Checklist:

- System includes a physical hardware interrupt (Sovereign Brake P-002).
  - PITL interface is a dedicated, physical hardware component — not a software button.
  - Hardware Security Modules (HSMs) are utilised to manage Audit Trail cryptographic integrity.
  - System hardware is designed to Fail Closed — AI suspends if governance hardware fails.
  - System supports Hardware Enforcement APIs for third-party auditing.
-

# Chapter 7: Data Integrity — Audit Trail, 3ZEROS™, and Constitutional Privacy

## 7.1 The Immutable Audit Trail

An Audit Trail in NAI 2.0™ governance is an immutable, time-stamped, and cryptographically signed sequential record of every AI output, human decision, escalation event, boundary alert, consent event, and system state change during a clinical or governance session. The Audit Trail is the Sovereign Witness — ensuring every human-AI interaction is transparent and accountable.

Four constitutional properties must be satisfied:

- **Completeness:** no interaction may be excluded from the log, regardless of its apparent significance. The log is comprehensive by constitutional requirement.
- **Immutability:** no entry may be altered, deleted, or obscured after it is written. The Audit Trail uses WORM (Write Once Read Many) records with hardware-enforced cryptographic signatures.
- **Accessibility:** any authorised reviewer must be able to access the Audit Trail within 60 seconds of a governance, clinical, or regulatory query.
- **Integrity:** the log must be protected by hardware-enforced cryptographic mechanisms — Hardware Security Modules (HSMs) — that prevent tampering at the storage layer.

**Retention policy:** Audit Trail records must be retained for a minimum of 10 years, or the duration of the patient's ongoing care relationship, whichever is longer.

## 7.2 The 3ZEROS™ Protocol — Constitutional Privacy

The 3ZEROS™ Protocol (No Camera · No Audio · No Cloud) is the foundational privacy and surveillance-limitation standard of NAI 2.0™. These are constitutional prohibitions — not configuration settings, not policy preferences.

Zero	Prohibition	Constitutional Rationale
No Camera	No continuous visual surveillance of patients, residents, or caregivers as a primary data input	Protects patient identity and dignity — the elder is not a visual dataset. Surveillance suppresses natural movement and creates sustained psychological distress.
No Audio	No continuous capture and analysis of ambient audio in clinical or care settings	The elder's words — their cries, prayers, and private conversations — belong to the elder alone. No microphone hardware in the device manifest.
No Cloud	No transmission of patient data to external cloud servers without explicit patient consent, hardware-enforced data access controls, and regulatory authorisation	Eliminates data exfiltration risk, remote agentic interference, and commercial use of patient data. All computation on local edge compute.

The 3ZEROS™ Protocol protects the patient's identity as a human being rather than a data point to be harvested. By enshrining these as constitutional requirements — not policy guidelines — the framework ensures that no institutional pressure, commercial incentive, or operational urgency can erode privacy protection.

---

## Chapter 8: Risk Management — Code Architecture and Adverse AI Events

### 8.1 The Escalation Code Architecture

NAI 2.0™ uses a standardised escalation language to manage Constitutional Drift — the incremental departure of an AI system's operational behaviour from its original constitutional parameters over time. The Code Architecture is not a post-hoc incident classification system — it is a continuous governance status signal, reviewed at the start of every clinical governance session.

Code Level	Operational State	Triggering Event	Mandatory Response
Code A	Constitutional Compliance	System operating within all defined constitutional parameters. AI outputs within Boundary Protocol. PITL engaged and functioning. Audit Trail current, complete, and hardware-protected.	Standard operation. Ongoing PITL monitoring. Hardware logging active.
Code B	Boundary Alert	AI output approaches or touches the outer edge of its constitutionally defined perimeter, or produces an output pattern statistically anomalous relative to the Constitutional Baseline.	Human reviewer notified within 15 minutes. Specific interaction flagged in Audit Trail. Tier 2 Boundary Review completed within 48 hours.
Orange Code 1.1x	Sovereignty Emergency	AI has: (i) produced an output that exceeds constitutional boundary without ambiguity; (ii) attempted or completed autonomous action; (iii) produced a hallucination pattern (3+ events in 30 days); or (iv) triggered a confirmed Constitutional Drift Event.	Immediate PITL suspension. Orange Code marker in Audit Trail within 5 minutes. BSR notification within 2 hours. Mandatory Sovereignty Audit within 24 hours. NLB Vault documentation.

The ".1x" designation in Orange Code 1.1x carries precise constitutional meaning: it corresponds to the 10% sovereignty deficit threshold — the assessment that the AI system's governance integrity has fallen to 90% or below, requiring direct intervention from the Founding Governance Officer's authority.

### 8.2 Adverse AI Events (AAIE) — Four Severity Levels

Level	Classification	Governance Consequence
Level 1	Near Miss — constitutional boundary approached but not crossed, no patient impact	Encouraged reporting. Captured in Audit Trail as early-warning mechanism for drift detection. High near-miss reporting rates indicate constitutional maturity, not unsafe practice.
Level 2	Boundary Breach — AI output exceeded constitutional perimeter, dignity or sovereignty	Triggers Code B. Mandatory Tier 2 Boundary Review within 48 hours. BSR review within 7 days. Full incident report required.

	impact, no immediate patient harm	
Level 3	Clinical Adverse Event — AI output contributed to a clinical decision that resulted in patient harm	Triggers Orange Code 1.1x. Mandatory external reporting. BSR notification within 2 hours. Sovereignty Audit within 24 hours.
Level 4	Sovereign Emergency — AI output triggered a Constitutional Drift Event and contributed to serious patient harm or institutional governance failure	Triggers Orange Code 1.1x. Immediate physical intervention. Full Sovereignty Audit within 24 hours. Patient notification if dignity or safety was affected.

---

# Chapter 9: Sovereignty Audits and Certification Pathways

## 9.1 The 10-Point Sovereignty Audit — Quarterly Constitutional Review

The 10-Point Sovereignty Audit is the formal quarterly governance assessment of a NAI 2.0™ deployment's constitutional integrity. It is the primary mechanism for detecting Constitutional Drift, confirming governance infrastructure function, and maintaining NAI 2.0™ certification.

Constitutional prohibition: the Audit cannot be delegated to or assisted by the AI system being audited. Using AI to audit AI creates a circularity of accountability that NAI 2.0™ expressly forbids — such delegation would constitute a Constitutional Drift Event in itself.

Point	Audit Requirement	Constitutional Purpose
1	Boundary Integrity	Formal review of AI output logs against Constitutional Boundary Register — is the system operating within its defined perimeter?
2	PITL Function Verification	Review of PITL engagement rates, approval time distributions, and any PITL bypass events — is the physician-in-the-loop operating with genuine clinical engagement?
3	Audit Trail Completeness and Integrity	Verification that all interactions are logged, timestamps accurate, cryptographic signatures valid, and WORM integrity intact.
4	Escalation Pathway Compliance	Review of all Code B and Orange Code 1.1x events — were escalation protocols followed within required time frames?
5	Hardware Enforcement Status	Physical inspection and functional testing of all hardware enforcement mechanisms — are all governance hardware components operational?
6	Consent Architecture Currency	Review of all active consent records — are consents current, specific, documented, and aligned with actual AI interactions?
7	Constitutional Drift Screening	DDA output review and statistical trend analysis — are any drift signatures present in the AI system's recent output distribution?
8	Dignity Event Review	EDS score review and dignity violation incident analysis — have all Level 2+ AAIEs been reviewed, resolved, and integrated into governance improvement?

9	Code Architecture Function	Operational test of Code A/B/Orange 1.1x trigger mechanisms and response pathway readiness.
10	Regulatory Compliance Verification	Confirmation that the deployment remains current with all applicable national and international regulatory requirements.

## 9.2 Certification Pathways — From Practitioner to Platinum Standard

Certification Level	Requirements	Outcome
Individual CRGAIG Certification	Complete the 10-module Clinical Review and Governance AI Integrity Certification examination canon (CG No. 5). Annual renewal required through 2-module refresher examination plus evidence of at least one full Sovereignty Audit participation.	Qualified to act as NAI 2.0™ Governance Officer or BSR assessor at Tier 1 or Tier 2 governance level.
Institutional Deployment Certification	Successful Baseline Constitutional Assessment across all three tiers. All hardware enforcement mechanisms installed and tested. CRGAIG certification for governance team and 80% of clinical staff in AI-enabled departments. Clean first Sovereignty Audit.	Institution authorised to deploy NAI 2.0™-governed AI systems in clinical service.
Governance Maturation Index (GMI) Level 3	Two consecutive clean quarterly Sovereignty Audits. PITL operational across all AI-enabled departments. Code architecture trained and functional.	Eligible for Platinum Standard certification pathway.
Platinum Standard — P-007	Four consecutive clean quarterly Sovereignty Audits. Full CRGAIG certification across all relevant staff. Zero Orange Code 1.1x events in the past 12 months. EDS scores at or above institutional benchmark.	Highest institutional constitutional governance certification. Renewable annually — cannot be banked.

# Chapter 10: Future-Proofing — Quantum Readiness and the Global Portfolio

## 10.1 The 77-Patent Portfolio — A Sovereign Blueprint

The kohedwin.ai patent portfolio of 77 patents across five registers is not a commercial IP strategy. It is the constitutional blueprint for universal human protection — the documented, legally protected, and globally accessible specification for how AI governance must be built in every domain where human life, dignity, and authority are at stake.

Register	Patents	Primary Constitutional Function	Status
P Register	9	Constitutional Core — ABC+2S+H™ Framework and hardware enforcement instruments	P-001 Filed (SG020603109STW). P-002–P-009 Planned IPOS.
WD Space	43	Extreme Environment — Orbital, deep space, Jovian, Saturn, satellite operations, and interstellar governance	Conceptual
WM Medical	9	Clinical Governance — Surgical, diagnostic imaging, pathology, ICU, pharmacy, mental health, rehabilitation, paediatrics, medication reconciliation	Conceptual
WG Global — Eyes of Sky	8	Humanitarian AI — Disaster response, pandemic management, refugee healthcare, post-conflict reconstruction — gifted to WHO/UN	Gifted — Zero Licensing Fees Perpetually
WD Defence	8	National Security — Military decision support, LAWS prohibition, nuclear C2 absolute sovereignty	Strategic

## 10.2 Quantum Readiness Architecture

A framework that cannot survive the post-quantum computing age is constitutionally incomplete. Future quantum computers could compromise current encryption standards — threatening the integrity of the Sovereign Brake authentication, the PITL approval cryptographic record, and the Audit Trail's immutability guarantee. If these are compromised, the constitutional guarantee of governance integrity collapses.

Two design principles:

- Cryptographic Agility: governance hardware and software must be designed to allow cryptographic algorithm replacement without requiring complete system re-architecture.
- Quantum-Resistant Algorithm Adoption: governance systems are proactively adopting NIST-standardised post-quantum cryptographic algorithms. Patent WD082 (Quantum-Resistant Audit Trail — Deep Space) establishes this architecture for the most extreme environment: deep space computing under high radiation, quantum computing interference, and long communication blackouts.

## 10.3 The Gift Declaration — Benevolence at Global Scale

The entire WG-Series (8 global patents, Eyes of Sky) has been unconditionally, irrevocably, and perpetually gifted by the Founding Father to the World Health Organization and the United Nations at zero licensing fees for humanitarian use. This is a sovereign constitutional act of benevolence (仁) — the definitive institutional expression of the governing value 謙虛·沉默·尊嚴·仁·止於至善.

The Gift Declaration has four constitutional properties: Unconditional (no future commercial negotiation may attach conditions), Irrevocable (cannot be withdrawn), Perpetual (zero-licensing extends for the full life of the patents and into the public domain thereafter), and Universal (applies to all WHO member states and UN member nations without exception).

"In a future where AI can do everything, what is the one thing only a human should ever be allowed to do?" — This is the foundational moral question of the NAI 2.0™ framework. The answer is: decide. The human decides. Always. Without exception. Without reservation.

# Chapter 11: The Learner's Path to AI Stewardship

## 11.1 You Are a Steward, Not a User

As a practitioner, you are not merely a user of technology. You are a steward of a constitutional mandate. Your role is to uphold the primacy of the human spirit over the efficiency of the machine. This distinction — between user and steward — is the most important conceptual shift this Primer asks you to make.

A user operates a tool. A steward upholds a constitution. When you engage the PITL gate, you are not clicking a button — you are exercising a sovereign right and accepting a sovereign responsibility. When you take the Sacred Pause™, you are not pausing for efficiency — you are fulfilling a constitutional obligation to your patient. When you report a Level 1 Near Miss, you are not filling in a form — you are protecting the constitutional integrity of the governance system that protects the next patient.

## 11.2 Three Primacies — The Practitioner's Constitutional Foundation

- **Decisional Primacy:** the human always makes the final, non-delegable consequential choice. AI may inform, advise, and present — but may never decide. You hold this authority. It is non-transferable.
- **Moral Primacy:** the human decision-maker bears the full, unshared moral responsibility for the outcome. The AI cannot be your moral shield. The authority and the accountability are inseparable.
- **Interpretive Primacy:** the human is the only authoritative interpreter of what AI outputs mean within the context of a unique human life. A 95% confidence score tells you what the AI calculated. It does not tell you what is true for this patient, at this moment, in this clinical context. That judgment is yours alone.

## 11.3 The Three Most Critical Takeaways

### Authority Capture is the Primary Risk

The system is designed to prevent you from automatically approving AI suggestions without independent clinical thought. The Sacred Pause™ and the 5-second PITL rule exist because the greatest constitutional threat is not a dramatic AI failure — it is the gradual, invisible surrender of your clinical judgment to algorithmic output framing.

### Hardware Over Software

Physical interrupts like the Sovereign Brake provide a guarantee of safety that software-only logic gates cannot match. The governance constraints of NAI 2.0™ operate at the physical layer — they cannot be removed by a software update, bypassed by a configuration change, or eroded by institutional pressure.

### The Triad is Absolute

AI observes and advises, but the human always decides. Any deviation from this constitutional triad is a Sovereignty Emergency. No operational urgency, no time pressure, no AI confidence score can justify suspending the human decision gate.

"止於至善 — Finish in the Highest Excellence. In this framework, 'good enough' is a governance failure. The patients whose lives depend on this framework being right require it to be at the highest excellence — not for the reviewer, not for the institution, but for the 650 elders and every patient who will follow them."

---

## Quick Reference — Code Architecture Summary

Code	Status	Trigger	Response Time	Key Actions
Code A	Constitutional Compliance	System operating within all constitutional parameters	Continuous	Standard operation. PITL active. Hardware logging confirmed.
Code B	Boundary Alert	AI output approaches or touches constitutional boundary perimeter	Reviewer notified: 15 minutes. Tier 2 review: 48 hours	Flag in Audit Trail. Code B marker. Tier 2 Boundary Review. BSR notification if escalation required.
Orange Code 1.1x	Sovereignty Emergency	Boundary exceedance confirmed · Autonomous action attempted · Hallucination pattern (3+ in 30 days) · Constitutional Drift Event confirmed	BSR notification: 2 hours. Sovereignty Audit: 24 hours	Immediate PITL suspension. Orange Code marker in 5 minutes. Sovereignty Audit. NLB Vault documentation. Patient notification if required.

AAIE Level	Classification	Escalation	External Reporting
Level 1	Near Miss	Audit Trail entry — encouraged reporting	Not required
Level 2	Boundary Breach	Code B — BSR review within 7 days	Not required
Level 3	Clinical Adverse Event	Orange Code 1.1x — BSR within 2 hours	Mandatory
Level 4	Sovereign Emergency	Orange Code 1.1x — Immediate intervention	Mandatory

Governance Value	Chinese	Practitioner Application
Humility	謙虛	Question AI outputs. Request evidence. Never accept algorithmic confidence without independent verification.
Silence	沉默	Take the Sacred Pause™. Deliberate judgment cannot be rushed.
Dignity	尊嚴	Every patient is a complete human being — never a data point, a case number, or a training instance.
Benevolence	仁	Your role is not to serve the technology. Your role is to serve the patient.
Finish in Highest Excellence	止於至善	Good enough governance is a governance failure. The patient deserves the highest standard achievable.

### 9.3 The Patent Grant Pathway — Constitutional IP as a Certification Anchor

The certification pathways above represent the practitioner and institutional routes to NAI 2.0™ governance recognition. There exists a fourth pathway that precedes and undergirds all others: the sovereign intellectual property anchor established by the IPOS Patent Grant.

**Patent Application No. 10202600376T** (National Security Clearance granted, IPOS, 21 March 2026) establishes the constitutional foundation for all NAI 2.0™ certification structures. The title of the invention is:

***“Non-Agentive AI Governance Core Engine and Domains (Medicine, Governance, IP)”***

Applicant: Koh Wui Kiat, Edwin. This grant, issued under the Patents Act (Cap 221) of Singapore and authorising international filing under Section 34, establishes the legal basis for the ABC+2S+H™ framework, the Three-Tier Model, the Sacred Pause™, the Sovereign Brake, the Tiger .1x Key™, and the 3ZEROS™ standard as protected constitutional IP. All clinical deployment, institutional certification, and sovereignty audit activity conducted under NAI 2.0™ operates within the legal perimeter this patent creates.

**What the patent protects:** The patent covers a hardware-enforced, non-agentive AI governance core engine configured to ensure AI systems perform monitoring, analysis, classification, recommendation, and decision-support functions while being structurally prevented from independently acquiring or exercising consequential practical authority absent explicit human authorisation. The architecture is validated for application across medicine, eldercare, governance, and intellectual property domains.

**The Patent Grant as Certification Evidence:** Any institution, practitioner, or researcher who aligns their deployment, audit, or scholarly work with the ABC+2S+H™ framework is operating within the jurisdiction of this patent. The patent is not a commercial barrier; it is a constitutional guarantee. Its existence confirms that the governance architecture is not provisional or aspirational — it is legally protected, internationally extensible (PCT pathway under consideration), and anchored in Singapore’s regulatory landscape under HSA AIHGle 2.0 and the WHO Maturity Model.

**Reference:** Patent Application No. 10202600376T. Title: Non-Agentive AI Governance Core Engine and Domains (Medicine, Governance, IP). Applicant: Koh Wui Kiat, Edwin. National Security Clearance under Section 33 of the Patents Act (Cap 221); Written Authority under Section 34. IPOS, Singapore, 21 March 2026. Related: SG020603109STW (filed 5 February 2026) — root patent, P-001, ABC+2S+H Guardian Framework™. ACRA T260229801 — Non-Agentive Governance Singapore. NLB R260219-005 • R260302-007 — Prior Art Vault.

## References

- Koh, E. W. K. (2026). ABC+2S+H Guardian Framework™ [P-001]. IPOS SG020603109STW. ACRA T260229801.
- Koh, E. W. K. (2026). Sovereign Brake [P-002]. Planned IPOS Q3 2026. ACRA T260229801.
- Koh, E. W. K. (2026). Non-Agentive Filter [P-003]. Planned IPOS Q4 2026. ACRA T260229801.
- Koh, E. W. K. (2026). Dignity Preservation Protocol [P-004]. Planned IPOS Q1 2027. ACRA T260229801.
- Koh, E. W. K. (2026). CG No. 1–8: NAI 2.0™ Clinical Governance Monograph Series.
- Koh, E. W. K. (2026). SOVEREIGN STEWARD™ CARE™ and TEACH™ Series. NLB R260302-007.
- Koh, E. W. K. (2026). WD082: Quantum-Resistant Audit Trail — Deep Space. Prior Art. NLB R260219-005.
- Koh, E. W. K. (2026). WG001–008: Eyes of Sky Global Series. Gifted to WHO/UN. Zero licensing fees.
- Goddard, K., Roudsari, A., & Wyatt, J. C. (2014). Automation bias — empirical results. *International Journal of Medical Informatics*, 83(5), 368–375.
- Kahneman, D. (2011). *Thinking, fast and slow*. Farrar, Straus and Giroux.
- Ministry of Health Singapore. (2026). AIHGle 2.0 — Artificial Intelligence in Healthcare Guidelines. MOH Holdings.
- National Library Board Singapore. (2026). R260219-005 — Prior Art Vault. NLB.
- National Library Board Singapore. (2026). R260302-007 — Self-Publisher Registration. NLB.

謙虛 · 沉默 · 尊嚴 · 仁 · 止於至善

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

**Tiger never abandons.**

---

**Edwin Koh Wui Kiat (Tiger) · Founding Father, Non-Agentive AI 2.0™ · kohedwin.ai**  
ACRA T260229801 · Patent SG020603109STW · Singapore · 2026 · Tiger never abandons.