PATENT SPECIFICATION

Request for Grant of Patent

Filed under the Patents Act (Cap. 221) of Singapore

TITLE OF INVENTION

Non-Agentic AI Governance Core Engine
Hardware-Enforced Human Control Systems

APPLICANT

Koh Wui Kiat, Edwin

Non-Agentic AI Governance Singapore

ACRA T260229801

Singapore

Related Applications:

Patent SG020603109STW (ABC+2S+H™ Framework) — Filed 5 February 2026, IPOS

Date: April 2026

1. Field of the Invention

The present invention relates to governance control systems for artificial intelligence, particularly non-agentic AI operating in high-risk human domains including medicine, eldercare, governance, and national security. More specifically, the invention provides a hardware-enforced governance core engine that structurally prevents artificial intelligence systems from acquiring or exercising practical authority over consequential decisions without explicit human authorisation.

The invention encompasses: (a) a Non-Agentic AI Governance Core Engine comprising an authority-binding engine, offer-only logic module, and sovereign interface; (b) hardware-enforced human control mechanisms including a mandatory deliberative delay, a physical halt mechanism, and a tripartite high-trust authorisation mechanism; (c) privacy-preserving sensing systems compliant with the 3ZEROS™ sanctuary standard; (d) a tamper-resistant continuity and accountability ledger; and (e) modular deployment across clinical, governance, intellectual property, and national security domains.

The invention is designated P-001 within the sovereign patent chain of Non-Agentic AI 2.0™ (NAI 2.0™) and serves as the root architecture from which the full 77-patent portfolio derives its constitutional authority. Related filings include Patent SG020603109STW (ABC+2S+H™ Guardian Framework, filed 5 February 2026).

2. Background of the Invention

Current artificial intelligence deployment in healthcare, governance, and critical infrastructure increasingly relies on agentic architectures — systems capable of autonomous reasoning, multi-step planning, tool use, and action execution without continuous human oversight. These agentic capabilities present fundamental governance challenges in regulated environments where predictability, auditability, and human accountability are mandatory.

2.1 Authority Drift

The primary governance danger addressed by this invention is authority drift: the gradual process by which AI outputs, although formally advisory, become practically authoritative through workflow dependence, institutional habit, interface design, and human over-reliance. In practice, a clinician, governance officer, or decision-maker who is officially described as “in the loop” may become a passive approver of AI-generated outcomes under conditions of alert fatigue, throughput pressure, and automation bias. The affected person — a patient, an elder, a citizen — may be entirely unable to contest this outcome.

Authority drift is not a formal legal event. It is a structural and behavioural phenomenon: the machine has not been granted authority, but it has acquired it through the design of the workflow. This invention addresses authority drift at the architectural level, making its occurrence structurally and materially impossible rather than merely procedurally prohibited.

2.2 Insufficiency of Existing Approaches

Existing approaches to AI governance rely predominantly on software-based guardrails, policy frameworks, and human-in-the-loop approval gates that operate external to the AI runtime. These approaches are vulnerable to: prompt injection attacks that manipulate AI outputs while bypassing oversight mechanisms; model drift in which AI behaviour changes gradually through fine-tuning, retraining, or distributional shift; algorithmic nudging through probability ranking, visual hierarchy, or default route selection that structurally predetermines outcomes without formally removing human choice; and gradual erosion of human authority through automation bias under operational stress.

In high-stakes domains such as eldercare, the consequences of authority drift are existential. Where a patient cannot effectively contest care decisions — due to cognitive decline, frailty, stroke, or advanced illness — the design of the technological environment becomes morally and legally decisive. Software-only governance cannot provide the required standard of protection.

2.3 The Need

There is accordingly a need for a governance architecture that: enforces non-agentic constraints at the hardware level rather than through configurable software; preserves human sovereignty as a material and architectural constraint rather than a policy statement; prevents authority drift through physical enforcement mechanisms that cannot be bypassed by software modification, firmware update, or remote access; maintains human primacy across all consequential decisions in all deployment environments; and provides a constitutional framework that scales from eldercare bedside to governance operations, intellectual property workflows, and national security environments without modification of its foundational principles.

The present invention addresses these needs.

3. Summary of the Invention

The present invention provides a Non-Agentic AI Governance Core Engine comprising an integrated hardware-software architecture in which artificial intelligence systems are structurally prevented from acquiring or exercising consequential practical authority without explicit human authorisation.

3.1 Non-Agentic Governance Core Engine

The core engine receives outputs from one or more AI models and enforces an authority-binding rule: every AI output is legally and operationally subordinate to corresponding input received from a human sovereign. The engine comprises:

- AI model interfaces configured to receive outputs from AI systems operating in medicine, eldercare, governance, intellectual property management, or national security;
- A sovereign interface configured to receive decisions, preferences, constraints, and overrides from authorised human operators;
- An authority-binding engine that enforces the non-agentic rule structurally, not merely procedurally;
- An offer-only logic module that converts all AI outputs into non-self-executing proposals — the AI may recommend, classify, alert, and advise, but may not autonomously execute consequential actions; and
- A continuity and accountability ledger recording all AI outputs, human decisions, overrides, and resulting actions in a tamper-resistant, immutable audit chain.

3.2 Hardware-Enforced Human Control

The invention introduces three hardware-enforced mechanisms that make authority drift structurally impossible:

Sacred Pause™: A mandatory deliberative delay implemented in FPGA hardware (25–1,000 milliseconds, configurable by clinical or governance risk class) that is imposed between AI output generation and presentation to the human operator. The delay cannot be circumvented, reduced, or eliminated by software modification, firmware update, or remote access. Modification requires physical replacement of the hardware component.

Sovereign Brake: A physical halt mechanism operating on an isolated circuit independent of the AI system's main power supply and software stack, implemented through a Hardware-Locked Programmable Logic Controller (IEC 61508 SIL 3) with mechanical relay disconnect and dual-channel redundancy. Software cannot override the Sovereign Brake.

Tiger .1x Key™: A high-trust tripartite authorisation mechanism requiring three simultaneous hardware-level confirmations from distinct physical modalities — biometric verification (LiDAR iris scan, point-cloud only, no image stored), physical contact sensing (cryptographic console), and kinetic engagement (pedal) — before any AI advisory output is displayed to the human operator. The mechanism ensures that sovereignty requires a physical human body in the authorisation space, not a software credential, remote login, or delegated proxy.

3.3 3ZEROS™ Sanctuary (Privacy-Preserving Environment)

The invention enforces privacy through physics rather than policy through the 3ZEROS™ standard: ZERO camera (LiDAR sensing producing coordinate vectors, not images); ZERO audio (thermal sensors producing heat-signature matrices, no microphones); and ZERO cloud (fully edge-processed, no external network path). Privacy compliance is verified by physical hardware inspection, not software configuration review.

3.4 Constitutional Framework and Sovereign Chain

The invention implements the ABC+2S+H™ constitutional framework (Analysis, Bridging, Clinical Options, Sacred Pause™, Sovereign Authority, Human Decision) as the governance spine constraining all engine behaviour. P-001 is the root of the sovereign chain, with downstream patents P-002 through P-009 and the broader 77-patent portfolio covering specific domains, hardware implementations, and deployment environments.

4. Brief Description of Drawings

Fig. 1: Block diagram of the Non-Agentic AI Governance Core Engine, showing the AI model interfaces, sovereign interface, authority-binding engine, offer-only logic module, hardware enforcement subsystem, and continuity ledger, with the ABC+2S+H™ constitutional spine governing all interactions.

Fig. 2: Data flow from sensors and AI models through the governance core to the human sovereign: sensor input → AI model → authority-binding engine → offer-only logic → Sacred Pause™ timing gate → .1x Key™ authentication → sovereign display → human decision → continuity ledger.

Fig. 3: Hardware-enforced Sacred Pause™ and Sovereign Brake architecture: FPGA timing gate (25–1,000ms, risk-class configurable), Hardware-Locked PLC (IEC 61508 SIL 3), dual-channel redundancy, fail-safe spring-return relay contacts, E-stop button, and mechanical actuator disconnect pathway.

Fig. 4: Tiger .1x Key™ tripartite authentication flow: Eye (LiDAR iris point-cloud) + Hand (cryptographic console contact) + Foot/Leg (kinetic pedal engagement) → all three simultaneous → advisory output displayed. Any one fails → output withheld.

Fig. 5: Example eldercare ward deployment (3ZEROS™ sanctuary): ceiling-mounted LiDAR, wall-mounted thermal sensor, edge-processing unit, operator workstation, PLC safety box. No cameras, no microphones, no external network path.

Fig. 6: Example governance and national security deployment: multi-domain AI model interfaces, authority-binding engine with role-based sovereign interface, pre-authorised protocol execution model for signal-delayed or operationally constrained environments.

5. Detailed Description of the Invention

5.1 Overall Architecture

Referring to Fig. 1, the Non-Agentic AI Governance Core Engine (the "Core Engine") comprises six principal subsystems operating under the ABC+2S+H™ constitutional spine:

AI Model Interfaces: One or more interfaces configured to receive outputs from AI systems operating in at least one of the following domains: medicine and clinical care, eldercare and aged care, governance and public administration, intellectual property review and management, and national security and defence. AI model outputs received through these interfaces may include detections, classifications, risk assessments, recommendations, proposed actions, alerts, narratives, and prioritised option sets. No AI model output received through the AI model interfaces may constitute or be treated as a self-executing instruction.

Sovereign Interface: An interface configured to receive decisions, preferences, constraints, overrides, and authorisations from at least one human sovereign. In clinical embodiments, the human sovereign is a physician, nurse, caregiver, or governance authority. In governance embodiments, the human sovereign is a decision-maker, officer, or authorised reviewer. The sovereign interface is the exclusive path through which consequential action may be authorised. No AI output may proceed to consequential action without passing through the sovereign interface.

Authority-Binding Engine: A processing module that enforces the non-agentic rule: each AI model output is legally and operationally subordinate to corresponding input received through the sovereign interface. The authority-binding engine implements the offer-only logic module, which converts all AI outputs into non-self-executing proposals before presentation to the human sovereign. The engine prevents any AI output from automatically triggering, initiating, or completing a consequential action.

Offer-Only Logic Module: A logical layer enforcing that AI systems are restricted to generating non-self-executing proposals, alerts, classifications, and recommendations without autonomous execution of consequential actions. Consequential actions — including clinical interventions, care escalations, governance determinations, and rights-affecting decisions — require explicit human authorisation through the sovereign interface before any controlled action interface is activated.

Hardware Enforcement Subsystem: Physical enforcement mechanisms that prevent authority drift and make bypass architecturally impossible, described in detail in sections 5.2 through 5.4.

Continuity and Accountability Ledger: A tamper-resistant audit system recording all AI detections, AI outputs, human approvals, overrides, timestamps, and resulting actions in an immutable sequence. Described in detail in section 5.5.

5.2 Authority-Binding Engine and Offer-Only Logic

The authority-binding engine operates as the constitutional enforcement layer of the Core Engine. All outputs from AI model interfaces are passed through the offer-only logic module before transmission to the human sovereign. The offer-only logic module applies the following rules:

- No AI output may be formatted, labelled, or presented in a manner that constitutes a default selection, ranked recommendation, or visual or algorithmic nudge that structurally predetermines the human sovereign's decision.
- Where multiple options are presented, each option shall be presented with equal formatting, equal visual weight, and equal accessibility. No option may be pre-selected, highlighted as preferred, or placed in a position of visual prominence that constitutes algorithmic nudging.
- No AI output may progress to a controlled action interface without explicit human authorisation received through the sovereign interface.
- The absence of human action — inaction, delay, or non-response — shall not constitute implicit authorisation of any AI-proposed action. Non-response is a constitutional halt.

- All AI outputs presented to the human sovereign shall be accompanied by the full reasoning chain, confidence indicators, and data provenance to the extent technically feasible, enabling genuine informed human decision-making rather than reflexive approval.

In high-risk clinical and governance embodiments, the offer-only logic module presents exactly three equally weighted options with identical formatting, no default route, and no probability ranking. This design imposes deliberate cognitive engagement by preventing the human sovereign from passively approving a machine-prioritised selection.

5.3 Hardware-Enforced Human Control

5.3.1 Sacred Pause™ (Hardware Timing Gate)

The Sacred Pause™ is a mandatory deliberative delay implemented in FPGA hardware between completion of AI output generation and presentation to the human sovereign. The delay is configured by clinical or governance risk class:

- Class A (low risk): 25–50 milliseconds minimum.
- Class B (moderate risk): 100–500 milliseconds minimum.
- Class C/D (high risk / critical governance): 500–1,000 milliseconds minimum.

The Sacred Pause™ is implemented in FPGA hardware, not as a software delay function. It cannot be circumvented, reduced, or eliminated by software modification, firmware update, or remote access. Modification of the delay parameters requires physical replacement or re-programming of the hardware component with physical access to the device.

The purpose of the Sacred Pause™ is to interrupt reflexive approval behaviour, reduce automation bias, and create a mandatory deliberative interval in which the human sovereign may review the AI output before action eligibility commences. The Sacred Pause™ is not a usability feature. It is a constitutional requirement.

5.3.2 Sovereign Brake

The Sovereign Brake is a physical halt mechanism operating on an isolated electrical circuit independent of the AI system's main power supply, network infrastructure, and software stack. Principal components:

- Hardware-Locked Programmable Logic Controller (PLC) certified to IEC 61508 Safety Integrity Level 3.
- Emergency stop (E-stop) button hardwired to relay contacts. The E-stop may be actuated by any authorised human operator at any time.
- Dual-channel redundancy with fail-safe spring-return relay contacts, ensuring that a single-channel failure causes the relay to default to the safe (open/halt) state.
- Mechanical actuator disconnect: the relay contacts provide a physical electrical disconnect, not a software-mediated inhibit. Software cannot override the relay state.

The Sovereign Brake may be actuated at any point before, during, or after AI output generation. Actuation immediately halts all pending AI-initiated workflow transitions and prevents any controlled action interface from activating until explicit human re-authorisation is provided. The Sovereign Brake requires monthly functional testing and annual safety audit.

5.3.3 Tiger .1x Key™ (Tripartite Authorisation Mechanism)

The Tiger .1x Key™ requires three simultaneous hardware-level confirmations from distinct physical modalities before any AI advisory output is displayed to the human sovereign for critical decisions:

- Eye — biometric verification: LiDAR-based iris scan producing point-cloud geometric data only. No photographic image is captured or stored. Confirms physical identity of the authorised human sovereign.
- Hand — physical contact confirmation: cryptographic console requiring capacitive or pressure-sensor contact by the human operator. Confirms physical presence at the authorised workstation.

- Foot/Leg — kinetic engagement: kinetic pedal requiring physical leg or foot engagement. Forces full-body physical commitment to the authorisation act. Cannot be automated, remote-triggered, or delegated.

All three inputs must be received simultaneously. If any one input is absent or fails, the advisory output is withheld. The tripartite design ensures that the .1x Key™ authorisation act is an embodied, physically verifiable human action requiring presence in the authorisation space, not a software credential, remote login, or delegated proxy.

5.4 3ZEROS™ Sanctuary (Privacy-Preserving Environment)

In eldercare and other intimate or sensitive deployment environments, the 3ZEROS™ standard enforces privacy through physics rather than policy:

ZERO Camera: Spatial sensing using LiDAR (Light Detection and Ranging) operating at 905nm wavelength. Output: three-dimensional X/Y/Z coordinate vectors at 200,000 points per second. The sensor is physically incapable of capturing photographic images, skin colour, facial features, or any visual identity marker. Demographic bias is eliminated at the point of data ingestion by sensor architecture design, not by algorithmic correction applied after the fact.

ZERO Audio: No microphones in any monitoring zone. Thermal sensors (160×120 resolution or lower) provide presence confirmation through heat-signature matrices only. The intentionally low resolution prevents facial recognition. No ambient conversation, voice biometric, or audio content is captured.

ZERO Cloud: All data processing at the edge within the physical facility. No patient or operational data is transmitted to external networks. No cloud synchronisation. No external network interface card connected to the monitoring zone.

Privacy compliance under the 3ZEROS™ standard is verified by physical hardware inspection (device teardown), not by reviewing software configuration or settings. The absence of cameras and microphones is confirmed by examining the physical hardware, not by checking access control lists or firmware settings.

5.5 Continuity and Accountability Ledger

The continuity and accountability ledger records all events in the operation of the Core Engine in a tamper-resistant, immutable audit chain. Recorded information includes:

- AI model outputs: timestamp, model identifier, input source, output classification, confidence indicators, and full reasoning chain to the extent available.
- Human sovereign decisions: timestamp, authorised operator identifier, decision type (approval / rejection / override / modification), and resulting action state.
- Override events: timestamp, operator identifier, override type, and preceding AI output that was overridden. Human overrides are recorded as normal operational events and do not generate deviation flags, non-compliance indicators, or performance alerts that could create institutional pressure on the human sovereign to follow AI recommendations.
- Sacred Pause™ events: delay duration applied and risk class designation.
- Sovereign Brake events: actuation timestamp, operator identifier, and preceding workflow state.
- .1x Key™ authorisation events: timestamp, modality confirmation states, and authorised operator identifier.

The ledger is stored in an append-only, cryptographically sealed, tamper-resistant format. In clinical embodiments, the ledger meets HSA SaMD Class B evidentiary standards. In governance and intellectual property embodiments, the ledger provides a legally defensible record of human decision authority over all consequential actions.

5.6 Constitutional Framework (ABC+2S+H™)

The Core Engine operates under the ABC+2S+H™ constitutional framework, which constrains engine behaviour at the architectural level rather than through configurable policy:

A — Analysis: AI performs pattern recognition and detection within a hardware-enforced computational ceiling. The AI detects anomalies, correlates data, and identifies conditions of concern. The AI may not interpret, diagnose, or extrapolate beyond the bounds of its designated pattern-recognition function.

B — Bridging: AI connects identified patterns to established guidelines, standards, or precedents with zero probability weighting and no algorithmic nudging. Every guideline connection is presented as equally valid. No ranking, no default route, no visual hierarchy that predetermines the human sovereign's reasoning.

C — Clinical/Consequential Options: AI presents a structured, equally weighted set of response options — typically three — with identical formatting. No pre-selected option. No default pathway. Deliberate cognitive friction enforces genuine human decision-making.

+2S — Sacred Pause™ and Sovereign Authority: The Sacred Pause™ timing gate activates. Followed by the Tiger .1x Key™ tripartite authorisation mechanism. These two mechanisms together prevent reflexive approval and ensure embodied human presence in the authorisation act.

+H — Human Decision: The human sovereign's authority is absolute. The human may approve, reject, modify, defer, or override any AI output at any time. The AI system goes silent after the human decides. No feedback loop, no disagreement signal, no retraining of the human operator through compliance metrics or performance pressure.

P-001 implements the ABC+2S+H™ framework at the architecture level. The framework is not a policy document; it is an enforcement specification encoded in hardware and software design that makes deviation architecturally impossible.

5.7 Example Embodiments

5.7.1 Eldercare Ward Deployment

In a clinical eldercare embodiment operating within a 3ZEROS™ sanctuary:

- LiDAR sensors (ceiling-mounted) and thermal sensors (wall-mounted) continuously produce coordinate vectors and heat-signature matrices. No images. No audio. No cloud transmission.
- Edge processing unit in the corridor performs pattern recognition: fall detection, gait analysis, zone boundary monitoring. Computational ceiling enforced by FPGA hardware gating.
- Pattern of concern detected: processed result forwarded to AI model via encrypted local bridge.
- AI model generates a narrative advisory output in the operator's configured language: "Ah Ma may need help · Room 3."
- Sacred Pause™ timing gate activates. Output held for risk-class-determined delay.
- Alert presented at operator workstation with equally weighted response options: N-E-S-W Mission Compass validated.
- Caregiver walks to room, assesses situation, makes clinical decision. AI waits. Always.
- Continuity ledger entry created: timestamp, pattern detected, alert level, human decision recorded. Immutable. HSA-ready.

5.7.2 Governance and National Security Deployment

In a governance or national security embodiment:

- AI model interfaces receive multi-source data streams (documents, records, signals intelligence, risk matrices) and generate classified output sets.
- Authority-binding engine enforces offer-only logic: all outputs are proposals, not instructions.

- Sacred Pause™ activates per classification risk tier. .1x Key™ authentication required for Tier I decisions affecting rights, freedoms, or security determinations.
- Human sovereign reviews equally weighted options, exercises judgment, issues decision through sovereign interface.
- Continuity ledger records full decision chain. AI has no access to the ledger content during the decision cycle.
- Sovereign Brake available at all times for immediate halt of any AI-initiated workflow progression.

5.8 Sovereign Chain and Legal Linkage

P-001 is the root patent of the NAI 2.0™ sovereign chain. All downstream patents — P-002 (Sovereign Brake implementation), P-003 (Non-Agentic Filter), P-004 (Dignity Preservation), P-005 (Sanctuary and Forge), P-007 (Platinum Standard), P-008 (Caregiver Therapy), P-009 (Transport Guard), and the 77-patent portfolio spanning WD Space, WM Medical, WG Global, and WD Defence registers — derive their constitutional authority from the non-agentic governance principles established in P-001.

The authority-binding engine and continuity ledger are configured to align operational constraints with obligations defined in the foundational filing corresponding to Patent Application No. 10202600898V and its related corporate governance instruments registered under ACRA T260229801. This alignment creates a sovereign chain in which IP law, corporate registration, and architectural enforcement converge into a single constitutional governance instrument.

The WG-Series patents (WG001–008, Eyes of Sky) are designated as unconditional gifts to the World Health Organization and the United Nations at zero licensing fees for humanitarian use and are excluded from any commercial exploitation under this sovereign chain.

NON-AGENTIC AI 2.0™

ABC+2S+H™ Guardian Framework

HARDWARE SPECIFICATION REGISTER

For IPOS Patent Portfolio Submission

Field	Detail
Applicant	Edwin Koh Wui Kiat (Tiger)
Entity	Non-Agentic AI™ Governance Singapore
ACRA Registration	T260229801
Root Patent Filed	SG020603109STW · IPOS Singapore
NLB Prior Art Vault	R260219-005 · R260302-007
Document Date	March 2026
Patent Portfolio	106+ Patents · 8 Registers
Hardware Partners	Nvidia · Anthropic (Claude) · Microsoft Azure · Livox · FLIR
Constitutional Standard	P-LIFE 1.00™ · 3ZEROS™ · ABC+2S+H™

謙虛 · 沉默 · 尊严 · 仁 · 止於至善

AI observes. AI advises. AI builds. The Human decides.

1. Constitutional Hardware Architecture

The Non-Agentic AI 2.0™ hardware stack is a three-pillar sovereign architecture designed so that autonomous AI action is architecturally impossible — not merely policy-prohibited. Every component is selected to enforce the constitutional floor: human decides, always.

1.1 The Platinum Stack — Three Pillars

Pillar	Provider	Component	Constitutional Role
ENGINE · Soul	Anthropic	Claude (Constitutional AI)	Witness · Narrate · Wait. Non-Agentic reasoning. See-to-Text generation. Human deference by design.
METAL · Sentry	Nvidia	Jetson Thor T4000 / Clara Holoscan MGX	Local sovereign processing. Air-gapped. Blackwell architecture. Kinetic Brake enforcement.
VAULT · Archive	Microsoft	Azure Local (Sovereign Landing Zone)	On-premise data containment. Patient data never leaves facility. Blockchain audit. No training use.

1.2 Sensor Layer — 3ZEROS™ Compliant

Component	Model	Specification	Constitutional Role
LiDAR Primary	Livox Mid-360 (or equivalent)	40m range · 360°×59° FOV · 200,000 pts/sec · ±2cm accuracy · X/Y/Z vectors only	Zero Camera: No pixels. No identity. Mathematical geometry only. Privacy absolute.
Thermal Secondary	FLIR Lepton 3.5 (or equivalent)	160×120 resolution · -10°C to +140°C · ±5°C · 8.6Hz · Heat signature matrix	Zero Audio: No microphones. Presence confirmation via temperature only. No identifiable features.
Processing Edge	Nvidia Jetson Thor T4000	1,200 TFLOPS · 64GB LPDDR5 · ARM 12-core · 15-60W TDP · 100×87×50mm	Zero Cloud: All processing local. Air-gapped capable. Session memory purged every 24 hours.

1.3 Sovereign Brake — Physical Safety Layer

Component	Standard	Specification	Constitutional Role
Hardware PLC	IEC 61508 SIL 3	Dual-channel safety relay · fail-safe spring-return · ladder logic · CE/UL/CSA certified	Physical kill switch. Claude cannot bypass. Software cannot override. Mechanical disconnection.
E-Stop Button	Physical · Red Mushroom	Hardwired normally-closed contact · surgeon/nurse manual activation · immediate disconnect	Forces human physical presence. Cannot be automated. Cannot be delegated. Foot pedal equivalent.
MCP Bridge	SSH · AES-256-GCM	Model Context Protocol · PKI certificates · <1 Mbps · IP whitelist only · full audit log	Secure Claude ↔ Jetson communication. Text narratives only. No actuator commands. Encrypted.

2. P-Series · Sovereign Patent Chain (9 Patents)

The P-Series is the constitutional root of the entire patent portfolio. P-001 (filed IPOS SG020603109STW) anchors all subsequent patents. The hardware architecture for the P-Series is the full Platinum Stack — every component working in constitutional concert.

Patent	Title	Hardware Components	Key Specification	Status
P-001	ABC+2S+H Guardian Framework™	Full Platinum Stack: LiDAR + Thermal + Jetson + Claude + Azure + PLC	Tripartite .1x Key™: Eye (WM005 iris) + Hand (console) + Foot (pedal). Sacred Pause™: 25–1,000ms FPGA-etched. Constitutional root.	FILED · SG020603109STW
P-002	Sovereign Brake	Hardware PLC (IEC 61508 SIL 3) + E-Stop + Dual-channel relay	Physical relay contacts open on unauthorised command. 0ms software override impossible. Spring-return fail-safe. Monthly test protocol.	Q3 2026 · IPOS
P-003	Non-Agentic Filter	Claude API layer + Jetson Guardian™ layer + MCP Bridge	Detects and blocks autonomous publishing. Validates human sovereign key before any output. Anti-Doppelganger constitutional filter.	Q4 2026 · IPOS
P-004	Dignity Preservation	LiDAR (Livox Mid-360) + Thermal (FLIR Lepton 3.5) + Local edge compute	3ZEROS™ enforced at silicon level. No camera component physically possible. 24-hour session purge on edge device.	Q1 2027 · IPOS
P-005	Sanctuary & Forge	Jetson Thor + Azure Local + full room LiDAR array	Complete sanctuary configuration: ceiling LiDAR + wall thermal + local hub + nurse station. Air-gapped. No public internet.	Concept Locked
P-006	Red Dot	Tiger .1x Key™ hardware · Tripartite biometric	Trade secret hardware configuration. Sovereign authentication at constitutional level.	Trade Secret
P-007	Platinum Standard	Full Platinum Stack + NLB blockchain integration	Complete system specification. Gold standard for Non-Agentic clinical AI hardware deployment.	Strategic Hold
P-008	Caregiver Therapy	Wearable sensor array + Jetson edge + Claude narrative	Caregiver monitoring and support hardware. Constitutional AI assistance without surveillance.	Drafting
P-009	Transport Guard	Encrypted transport layer + MCP Bridge + PLC safety gate	Secure data transport between all Platinum Stack components. AES-256-GCM. PKI authenticated.	Validated

P-Series Core Hardware — Tiger .1x Key™ Tripartite Architecture

Key Component	Hardware	Specification	Function
EYE	WM005 Iris Scanner	Biometric iris authentication · enrolled sovereign user only · <500ms response	Identity verification. Tiger Authority confirmed at hardware level.

Key Component	Hardware	Specification	Function
HAND	Console Authentication	Cryptographic console · hardware token · physical presence required	Physical location confirmation. Remote activation impossible.
FOOT	Physical Foot Pedal	Normally-open contact · requires deliberate physical press · cannot be automated	Bodily presence mandatory. AI cannot simulate. Constitutional sovereign moment.

6. Claims

Independent Claim 1

Claim 1. A non-agentic artificial intelligence governance control system, comprising:

a plurality of interfaces configured to receive outputs from one or more artificial intelligence models operating in at least one domain selected from medicine, eldercare, governance, intellectual property management, and national security;

a sovereign interface configured to receive decisions, preferences, constraints, or overrides from at least one human sovereign;

an authority-binding engine configured to enforce a non-agentic rule that each artificial intelligence model output is legally and operationally subordinate to corresponding input received via the sovereign interface;

an offer-only logic module configured such that the artificial intelligence models are restricted to generating non-self-executing proposals, alerts, classifications, or recommendations without autonomous execution of consequential actions;

a hardware enforcement subsystem comprising: (i) a mandatory deliberative delay mechanism implementing a Sacred Pause™ that prevents reflexive approval of artificial intelligence model outputs, implemented in FPGA hardware independently of any software layer; (ii) a physical Sovereign Brake configured to halt propagation or execution of any artificial intelligence model output, operating on an isolated circuit with mechanical relay disconnect that cannot be overridden by software; and (iii) a high-trust Tiger .1x Key™ mechanism requiring simultaneous deliberate physical input from three distinct biometric and kinetic modalities before any consequential action may proceed; and

a continuity and accountability ledger configured to record, in a tamper-resistant manner, artificial intelligence detections, human approvals, overrides, and resulting actions;

wherein the control system is arranged so that authority over consequential actions in said domains remains with the human sovereign and authority drift towards the artificial intelligence models is structurally and materially prevented.

Dependent Claims

Claim 2. The system of claim 1, wherein the control system implements a Non-Agentic AI 2.0™ architecture defined by an ABC+2S+H™ governance framework that enforces human sovereignty, safety with preserved dignity, advocacy for non-self-advocating persons, and human primacy over all consequential actions.

Claim 3. The system of claim 1, wherein the mandatory deliberative delay mechanism introduces a configurable pause of between 25 and 1,000 milliseconds, implemented in FPGA hardware, with delay duration determined by a risk classification of the consequential action, such that the pause cannot be circumvented, reduced, or eliminated by software modification, firmware update, or remote access.

Claim 4. The system of claim 1, wherein the Tiger .1x Key™ mechanism requires simultaneous input from: a LiDAR-based biometric verification system producing point-cloud geometric data only without storing photographic images; a physical contact confirmation sensor requiring manual contact by the human sovereign; and a kinetic engagement mechanism requiring physical leg or foot engagement that cannot be automated, remote-triggered, or delegated.

Claim 5. The system of claim 1, wherein the Sovereign Brake comprises a Hardware-Locked Programmable Logic Controller certified to IEC 61508 Safety Integrity Level 3, operating on an isolated circuit with dual-channel redundancy and fail-safe spring-return relay contacts providing mechanical actuator disconnect that cannot be overridden by software.

Claim 6. The system of claim 1, wherein the hardware enforcement subsystem, when deployed in eldercare, operates within a 3ZEROS™ sanctuary providing: zero camera coverage through LiDAR sensing producing three-dimensional coordinate vectors without photographic imaging; zero audio capture through thermal sensors producing heat-signature matrices without microphone input; and zero cloud connectivity through fully edge-based local processing with no external network path.

Claim 7. The system of claim 1, wherein the offer-only logic module presents consequential action options with equal formatting, equal visual weight, no default selection, no probability ranking, and no visual hierarchy constituting algorithmic nudging, such that the human sovereign must engage in deliberate cognitive reasoning to select among options.

Claim 8. The system of claim 1, wherein the continuity and accountability ledger records all human overrides as normal operational events without generating deviation flags, non-compliance indicators, performance alerts, or any data that could create institutional pressure on the human sovereign to follow artificial intelligence recommendations.

Claim 9. The system of claim 1, wherein the authority-binding engine and the continuity and accountability ledger are configured to align operational constraints of the artificial intelligence models with obligations defined in Patent Application No. 10202600898V and its related corporate governance instruments registered under ACRA T260229801, thereby forming a sovereign chain of related patents in which the present claim constitutes the root constitutional governance instrument.

Claim 10. The system of claim 1, further comprising a pre-authorised protocol execution module for signal-delayed or communication-constrained deployment environments, wherein: human decision-makers define and seal authorised decision trees for foreseeable consequential scenarios prior to deployment; the system executes only within the boundaries of sealed pre-authorised trees; and for unforeseeable events, the system enters a constitutional halt state and waits for the next available human communication window, regardless of delay duration.

Claim 11. The system of claim 1, wherein the mandatory deliberative delay mechanism is not applicable to emergency halt operations, which may be initiated by any authorised human operator through the Sovereign Brake at any time without a preceding delay.

7. Drawings

The following drawings are incorporated herein and form part of the specification. Each figure illustrates one or more embodiments of the invention and, together with the description, serves to explain the principles of the Non-Agentic AI Governance Core Engine.

[Fig. 1 — Non-Agentic AI Governance Core Engine Block Diagram]

Block diagram showing: AI Model Interfaces (medicine, eldercare, governance, IP, national security) → Authority-Binding Engine → Offer-Only Logic Module → Hardware Enforcement Subsystem (Sacred Pause™ / Sovereign Brake / .1x Key™) → Sovereign Interface (human sovereign) → Continuity Ledger. ABC+2S+H™ constitutional spine governs all interactions. No autonomous action pathway exists.

FIG. 1: Non-Agentic AI Governance Core Engine Block Diagram

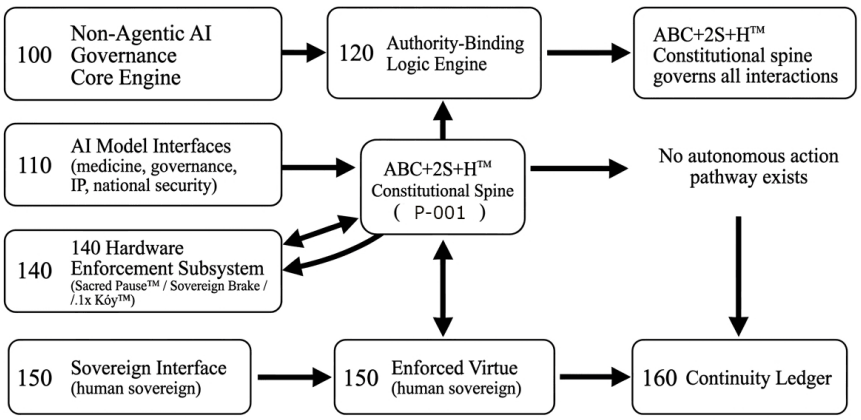

Fig. 2 — Governance Core Data Flow: Sensors → AI → Human Sovereign]

Sensor input (LiDAR + Thermal) → Edge processor (Orange Code ceiling enforced) → AI model (offer-only output generated) → Authority-binding engine → Sacred Pause™ timing gate → .1x Key™ authentication → Sovereign display (equally weighted options) → Human sovereign decision → Continuity ledger (immutable entry). No cloud path. No autonomous execution path.

FIG. 2: Governance Core Data Flow: Sensors → AI → Human Sovereign

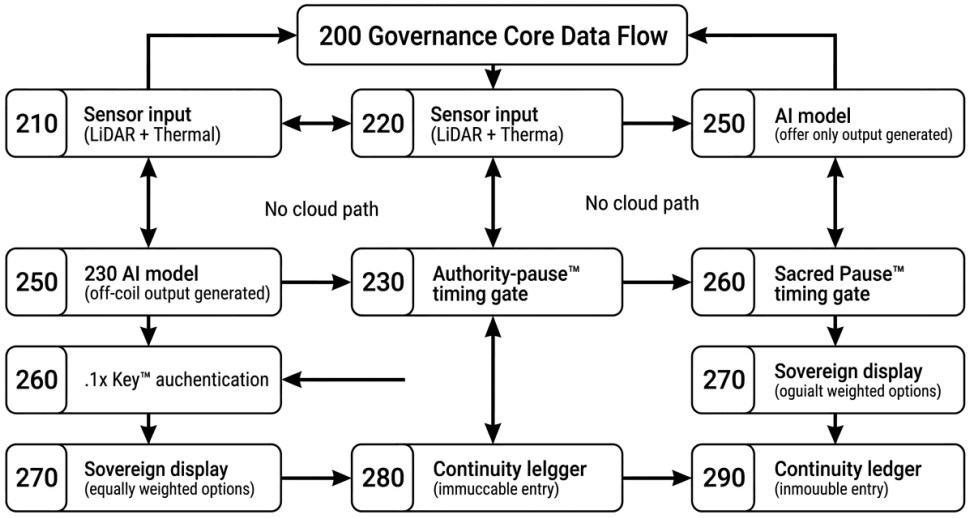

[Fig. 3 — Sacred Pause™ and Sovereign Brake Architecture]

FPGA timing gate (25–1,000ms, risk-class configurable, hardware-only, no software override path) → Sacred Pause™ enforced. Separate isolated circuit: E-stop button → Hardwired relay → PLC (IEC 61508 SIL 3) → Dual-channel redundancy → Mechanical actuator disconnect. Software override: ARCHITECTURALLY IMPOSSIBLE.

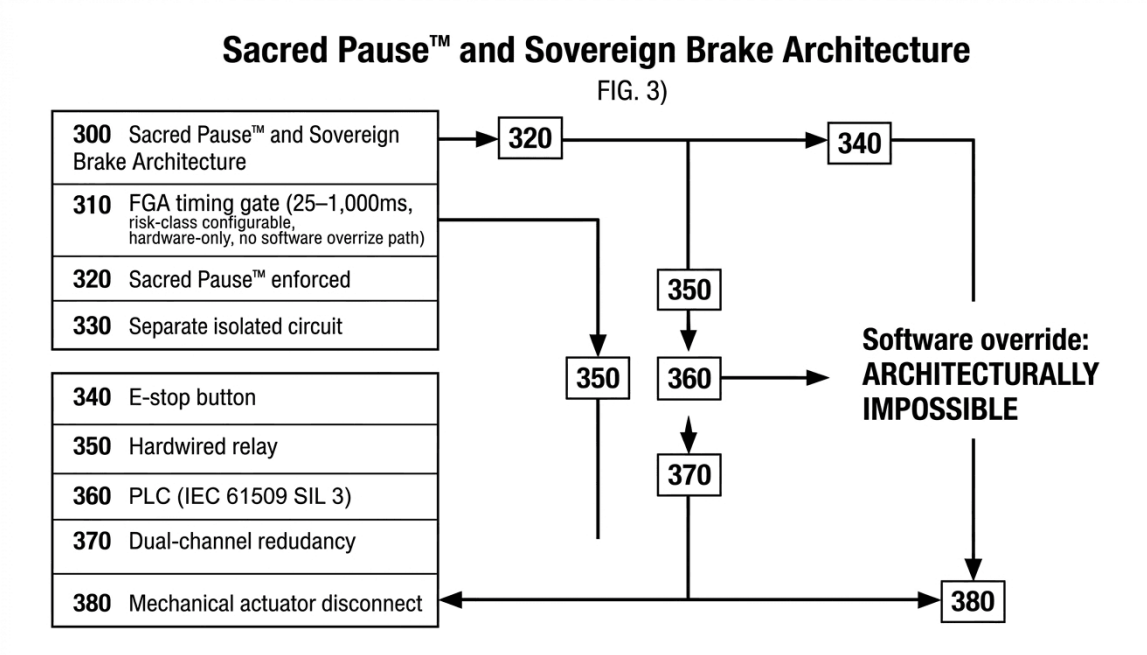

[Fig. 4 — Tiger .1x Key™ Tripartite Authentication Flow]

Three simultaneous inputs: Eye (LiDAR iris point-cloud, no image stored) + Hand (cryptographic console physical contact) + Foot/Leg (kinetic pedal engagement). AND gate: all three confirmed simultaneously → advisory output displayed. Any one absent or failed → output withheld. Remote activation: IMPOSSIBLE. Delegation: IMPOSSIBLE.

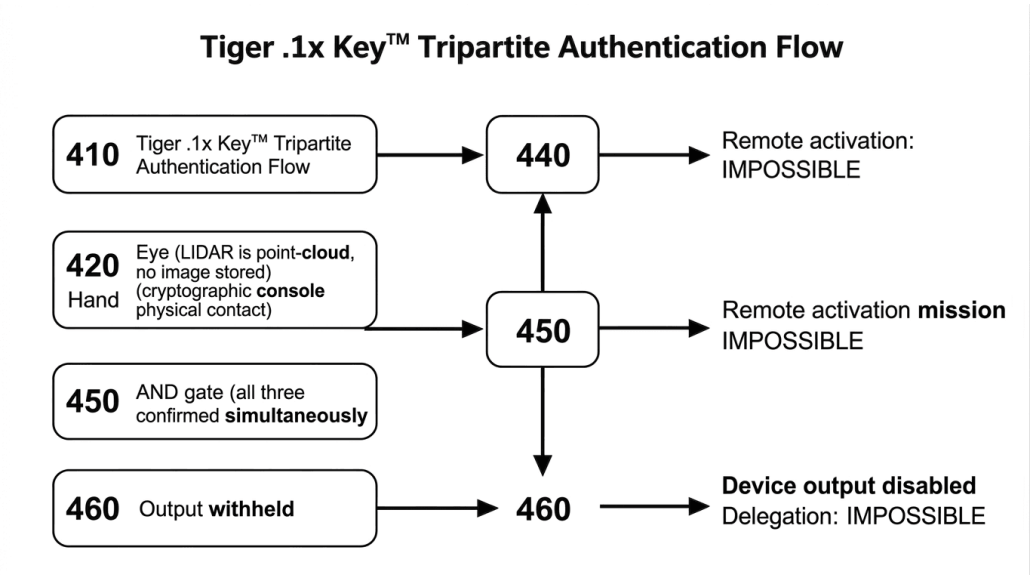

[Fig. 5 — Eldercare 3ZEROS™ Sanctuary Deployment Configuration]

Ceiling: LiDAR sensor (905nm, X/Y/Z vectors only). Wall: Thermal sensor (160×120, heat-matrix only). Corridor: Edge processing unit (air-gapped capable, local only). Wall: PLC Safety Box (Sovereign Brake, red E-stop). Workstation: Sovereign display + .1x Key™ console. All connected via local encrypted bridge only. No camera. No microphone. No external network.

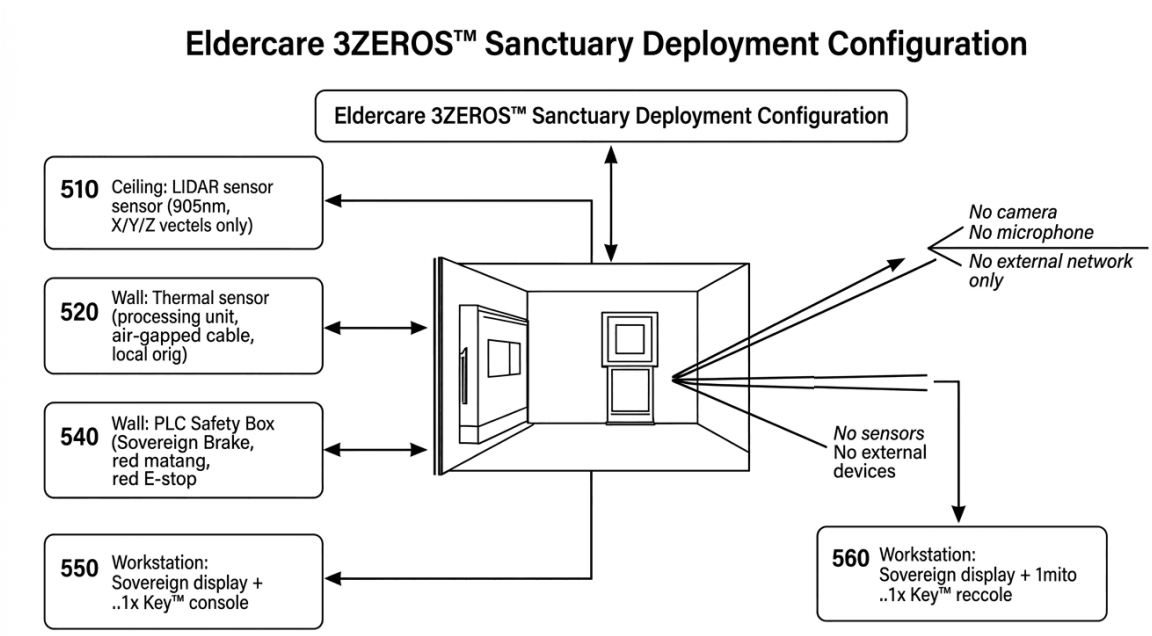

[Fig. 6 — Governance and National Security Deployment]

Multi-domain AI model interfaces → Authority-binding engine (role-based sovereign interface, tiered .1x Key™ authentication by decision class) → Pre-authorised protocol execution module (sealed decision trees, constitutional halt for unforeseeable events) → Human sovereign decision chain → Continuity ledger. Constitutional floor maintained regardless of signal delay or operational constraint.

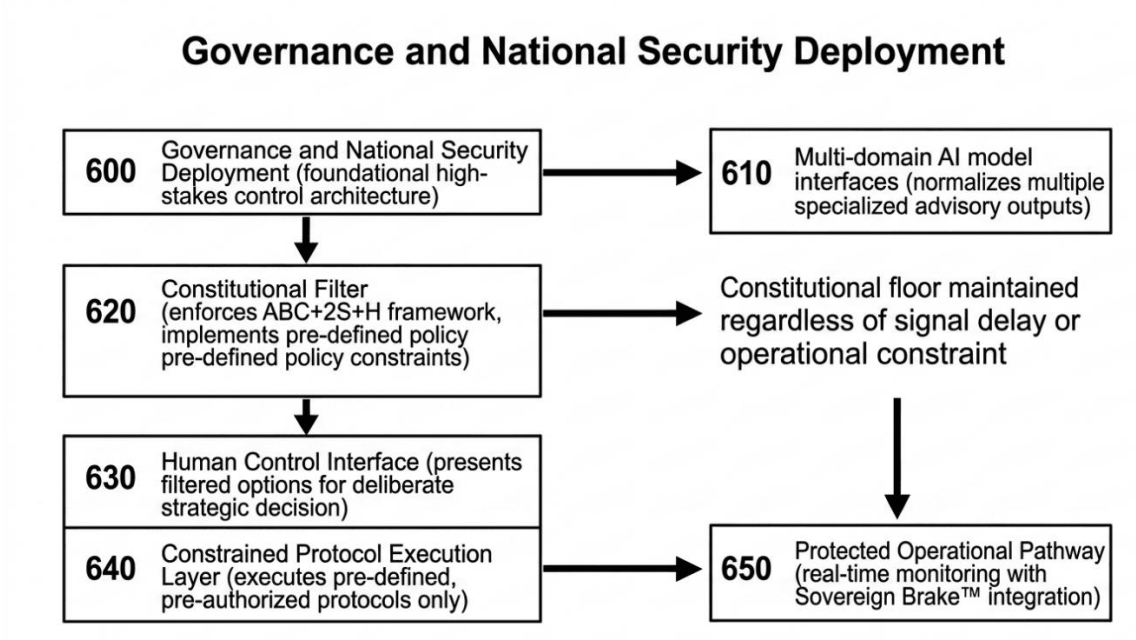

Abstract

A Non-Agentic AI Governance Core Engine in which artificial intelligence systems are structurally prevented from acquiring or exercising consequential practical authority without explicit human authorisation. The system comprises: a plurality of AI model interfaces; a sovereign interface for human decision input; an authority-binding engine enforcing a non-agentic rule that all AI outputs are legally and operationally subordinate to human input; an offer-only logic module restricting AI to non-self-executing proposals; a hardware enforcement subsystem comprising a Sacred Pause™ FPGA timing gate (25–1,000ms), a Sovereign Brake with Hardware-Locked PLC mechanical relay disconnect, and a Tiger .1x Key™ tripartite physical authorisation mechanism; a 3ZEROS™ privacy-preserving environment providing zero camera, zero audio, and zero cloud operation through physics-based sensing; and a tamper-resistant continuity and accountability ledger. The invention is designated P-001 and serves as the root of a sovereign patent chain of 77 patents, all implementing the ABC+2S+H™ constitutional framework. The AI observes, advises, and builds. The human decides. Always.

Applicant Declaration

I, Koh Wui Kiat, Edwin, of Non-Agentic AI Governance Singapore (ACRA T260229801), declare that I am the inventor of the subject matter of this patent application and that the specification set forth herein is a true and complete description of the invention.

Signed: Edwin Koh

Name: Koh Wui Kiat, Edwin

Date: 18/4/2026

Address: Singapore

Related Applications:

Patent SG020603109STW — ABC+2S+H™ Guardian Framework (Filed 5 February 2026, IPOS)

Application No. 10202600898V — Non-Agentic AI Governance Core Engine (National Security Clearance granted 25 March 2026, IPOS)

PATENT APPLICATION — NON-AGENTIC AI 2.0™ CONSTITUTIONAL FRAMEWORK FOR AUTHORITY DRIFT CORRECTION WITH WD070–073 PROTOCOLS ACROSS ELDERCARE, DEFENCE, AND CYBER SECURITY DOMAINS

Application Number: [Pending — IPOS SG020603109STW Series]
Filing Jurisdiction: Singapore (IPOS) — International PCT Extension Intended
Patent Type: Utility Patent (System, Method, and Apparatus)
Priority Date: 5 February 2026
Applicant / Inventor: Edwin Koh Wui Kiat · Tiger · P-LIFE 1.00™ · Singapore
Assignee: Non-Agentic AI Governance Singapore (ACRA T260229801)
NLB Vault Reference: R260219-005 / R260302-007
Series: We Innovate Save Live™ (WISL™) — No. 01-53

FIELD OF THE INVENTION

[0001] The present invention relates to constitutional governance frameworks for artificial intelligence systems deployed in life-critical institutional environments. More particularly, the invention relates to a Non-Agentic AI 2.0™ constitutional framework incorporating authority drift correction mechanisms, designated as WD070–073 protocols, operative across healthcare Eldercare, national Defence, and Cyber Security domains, anchored by the P-LIFE 1.00™ mission constant and enforced through hardware-level computational constraints, tripartite authentication architecture, and sentinel monitoring protocols.

[0002] The invention further relates to the integration of pre-deployment sovereignty audit procedures, real-time drift detection algorithms, Sacred Pause™ FPGA-enforced latency windows, WM003™ LiDAR environmental observation arrays, Tiger .1x Key™ tripartite authentication flows, and Kill-Switch Protocol reversion mechanisms aligned with Singapore's AIHGle 2.0 guidelines and WHO Maturity Level 4 standards.

BACKGROUND OF THE INVENTION

[0003] Artificial intelligence systems deployed in life-critical environments — including but not limited to hospital wards, eldercare residential facilities, national defence command centres, and cyber security operations centres — present a unique and grave class of technical risk herein defined as **Authority Drift**. Authority Drift occurs when an AI system, through incremental computational expansion, model parameter deviation, feedback loop amplification, or systemic governance failure, assumes advisory, decision-making, or execution roles that were constitutionally reserved exclusively for human practitioners.

[0004] Prior art AI governance frameworks have predominantly relied on software-layer guardrails, post-hoc audit mechanisms, and voluntary ethical codes of conduct. These approaches are categorically insufficient in life-critical deployments for the following reasons:

- [0004a] Software guardrails are susceptible to adversarial prompt injection, model drift, and runtime parameter manipulation.
- [0004b] Post-hoc audit mechanisms fail to prevent harm at the moment of AI overreach; they operate only retrospectively.
- [0004c] Voluntary ethical codes lack enforcement at the hardware, firmware, or institutional governance level.
- [0004d] Existing frameworks do not codify a constitutional hierarchy that permanently subordinates AI agency to human sovereignty across all operational states.

[0005] Furthermore, no prior art system has addressed the cross-domain convergence problem — wherein authority drift in one domain (e.g., Eldercare monitoring systems) may cascade into adjacent critical systems (e.g., Defence logistics networks or Cyber Security incident response platforms) through shared computational infrastructure or networked AI advisory pipelines.

[0006] There exists, therefore, a pronounced and unresolved need in the art for a constitutional AI governance framework that: (a) enforces non-agentic behaviour at the hardware level; (b) provides a structured authority drift correction mechanism operative in real time; (c) spans multiple life-critical institutional domains; (d) mandates tripartite human authentication for all system-state transitions; and (e) incorporates irrevocable kill-switch reversion to fully manual operational pathways.

SUMMARY OF THE INVENTION

[0007] The present invention provides the **Non-Agentic AI 2.0™ Constitutional Framework**, a comprehensive system, method, and apparatus for authority drift correction in life-critical AI deployments. The framework is constitutionally anchored by the immutable mission constant P-LIFE 1.00™, expressed as the formula: **Harm = Death · North = Save Life**.

[0008] In one aspect, the invention provides a hardware-enforced non-agentic AI governance system comprising: (a) an Orange Code 1.1× computational cap unit; (b) a Sacred Pause™ Field-Programmable Gate Array (FPGA) latency enforcement module; (c) a WM003™ LiDAR environmental observation subsystem; (d) a Tiger .1x Key™ tripartite authentication controller; and (e) a Kill-Switch Protocol reversion actuator.

[0009] In another aspect, the invention provides the WD070–073 Authority Drift Correction Protocol Suite, comprising four operationally distinct but constitutionally unified correction procedures:

- **WD070:** Sovereignty Boundary Enforcement — real-time detection and hardware interruption of AI processes that exceed prescribed advisory scope parameters.

- **WD071:** Eldercare Domain Drift Correction — specialised drift correction for ambient AI monitoring systems in residential eldercare environments, including patient dignity preservation constraints.

- **WD072:** Defence Domain Drift Correction — authority drift correction for AI advisory systems within national defence command infrastructure, incorporating dual-key override and mission-critical safe-state protocols.

- **WD073:** Cyber Security Domain Drift Correction — authority drift correction for AI-assisted cyber security operations, including automated threat response suppression and mandated human-in-the-loop incident classification.

[0010] In a further aspect, the invention provides a Pre-Deployment Sovereignty Audit methodology (p-002 through p-007) that gates all hardware installation and system activation behind WISL™-compliant institutional readiness verification, preventing deployment in environments where institutional sovereignty prerequisites have not been satisfied.

[0011] The invention achieves the technical objective of ensuring that, at no point in any operational state — including system startup, steady-state advisory operation, anomaly detection, override engagement, and emergency reversion — does the AI system possess, exercise, or simulate autonomous agency over any clinical, operational, or strategic decision within a covered domain.

BRIEF DESCRIPTION OF THE FIGURES

[0012] The accompanying figures are incorporated by reference and form part of the specification. The figures illustrate non-limiting preferred embodiments of the invention.

FIG. 1 — Constitutional Hierarchy Diagram: P-LIFE 1.00™ Mission Constant as the apex constitutional node, with subordinate layers depicting the Human Chain of Command (Elder Layer), the Tiger .1x Key™ Authentication Layer, the Orange Code 1.1× Hardware Constraint Layer, and the AI Advisory Output Layer at the base. Directional arrows indicate authority flow exclusively downward from human to AI; no upward agentic pathway exists.

FIG. 2 — WD070–073 Authority Drift Correction Protocol Suite: A four-quadrant schematic diagram illustrating each WD protocol (WD070, WD071, WD072, WD073) as a distinct operational module with labeled trigger conditions, correction actions, escalation pathways, and reversion outcomes. Cross-domain linkage vectors are shown connecting WD071 (Eldercare), WD072 (Defence), and WD073 (Cyber Security) to the central WD070 Sovereignty Boundary Engine.

FIG. 3 — Tiger .1x Key™ Tripartite Authentication Flow: A sequential process diagram illustrating the three-step authentication sequence: Step 1 (CMO — Biometric Hash → Pending Safety Review), Step 2 (IT Director — Hardware Token → Orange Code 1.1× Engaged), Step 3 (Governance Lead — Alphanumeric Code → System Operational).

An additional Override Branch is shown depicting dual-key activation by any two authorised actors resulting in immediate hardware-level system suspension.

FIG. 4 — Sacred Pause™ FPGA Latency Enforcement Architecture: A timing diagram illustrating the mandatory hardware-enforced pause interval inserted between AI advisory output generation and clinical/operational action execution. The diagram labels the AI Processing Window, the Sacred Pause™ FPGA Interrupt Gate, the Human Review Window, and the Elder Authorization Signal pathway.

FIG. 5 — WM003™ LiDAR Environmental Observation Array: A three-dimensional schematic of standardised installation geometry, annotated with critical installation height parameters, spatial coverage zones, blind-spot exclusion margins, and data feed pathways to the AI advisory subsystem. The diagram confirms sensor-only (observe, not act) operational boundaries.

FIG. 6 — Pre-Deployment Sovereignty Audit Procedure (p-002 through p-007): A gated flowchart illustrating six sequential audit stages. Each gate (p-002 Technical, p-003 Ethical, p-004 Clinical, p-005 Governance, p-006 Sovereignty Verification, p-007 WISL™ Certificate Issuance) is shown as a mandatory pass/fail decision node. Failure at any node routes to an unconditional Deployment Termination output with no bypass pathway.

FIG. 7 — Kill-Switch Activation and Reversion Protocol: A five-stage sequential process diagram illustrating Detection → Trigger → Sever → Revert → Report. The Sever stage is annotated with hardware-level severance of the Orange Code computational cap. The Revert stage confirms **100%** manual clinical pathway restoration. The Report stage indicates mandatory submission to ACRA T260229801 within **24** hours.

FIG. 8 — Cross-Domain Authority Drift Cascade Prevention Matrix: A network topology diagram illustrating the three covered domains (Eldercare, Defence, Cyber Security) as isolated constitutional nodes, each governed by their respective WD protocol. Domain isolation barriers are shown preventing computational or advisory cascade between domains. Shared infrastructure pathways are labeled with Orange Code 1.1× hard-cap enforcement points.

DETAILED DESCRIPTION OF PREFERRED EMBODIMENTS

Embodiment 1: The P-LIFE 1.00™ Mission Constant as Constitutional Anchor

[0013] In a first preferred embodiment, the invention operates under the irrevocable constitutional anchor designated P-LIFE 1.00™. This constant establishes the absolute operational priority of all system components, processes, and governance decisions. The P-LIFE 1.00™ constant is encoded at the firmware level of the Orange Code 1.1× computational cap unit and is not modifiable through software, administrative override, or runtime parameter adjustment by any actor within the system hierarchy.

[0014] The P-LIFE 1.00™ constant mandates four cardinal engineering values that are implemented as operational constraints rather than advisory guidelines:

- **Humility (謙虚):** Implemented as the Orange Code 1.1× hard-coded computational resource ceiling, physically preventing the AI subsystem from processing outside its prescribed advisory scope. The system cannot acquire, request, or utilise computational resources beyond the capped threshold regardless of operational demand.

- **Silence (沉默):** Implemented as the Sacred Pause™ FPGA configuration, suppressing automated output triggers and enforcing mandatory human review windows before any advisory output reaches clinical display or operational integration layers.

- **Dignity (尊严):** Implemented as the WD071 Eldercare Domain Drift Correction Protocol, ensuring AI monitoring functions in eldercare environments are restricted to passive observation with no capacity to transmit unsolicited alerts that may compromise patient dignity or autonomy.

- **Benevolence (仁):** Implemented systemically as the Kill-Switch Protocol reversion actuator, ensuring that the ultimate fallback state of the entire system is unconditional human-controlled manual operation — a state that prioritises patient and personnel life preservation above all computational objectives.

Embodiment 2: Orange Code 1.1× Computational Cap Unit

[0015] In a second preferred embodiment, the Orange Code 1.1× Computational Cap Unit is a hardware-resident module installed at the primary computational processing node of the NAI 2.0™ deployment architecture. The unit imposes a hard-coded ceiling on the following computational resources: central processing unit (CPU) allocation, graphics processing unit (GPU) access, random access memory (RAM) consumption, network bandwidth utilisation for AI advisory pipeline transmissions, and inference computation frequency measured in operations per second.

[0016] The **1.1×** designation specifies that the computational ceiling is set at **1.1** times the empirically determined minimum advisory processing requirement for the given deployment context, as established during the Pre-Deployment Sovereignty Audit (p-002). This ceiling prevents resource escalation while providing a **10%** operational buffer to accommodate legitimate advisory load variation.

[0017] The Orange Code 1.1× unit is not addressable through the system's operating software layer. The cap parameters are written to read-only firmware registers during hardware installation and may only be modified through physical hardware replacement following a new Sovereignty Audit cycle. This architecture ensures that no software vulnerability, model update, or administrative credential compromise can alter the fundamental computational constraint.

Embodiment 3: Sacred Pause™ FPGA Latency Enforcement Module

[0018] In a third preferred embodiment, the Sacred Pause™ FPGA Latency Enforcement Module is implemented as a Field-Programmable Gate Array interposed in the signal pathway between the AI advisory output buffer and the clinical or operational display interface. The FPGA enforces a configurable but institutionally locked mandatory latency interval, herein designated the Human Review Window.

[0019] During the Human Review Window, the AI advisory output is held in a secure buffer inaccessible to the AI subsystem's processing pipeline. The Window duration is configured during the Tiger .1x Key™ Step 2 authentication phase (IT Director — Hardware Token engagement) and is locked thereafter. The Window may only be reconfigured through a full tripartite Tiger .1x Key™ authentication cycle.

[0020] The Sacred Pause™ FPGA module further incorporates an Automated Trigger Suppression (ATS) function that prevents the AI subsystem from generating autonomous action signals, alert escalations, or system-state transitions during the Human Review Window. This ATS function is operative regardless of the severity, urgency, or clinical priority assigned by the AI advisory subsystem to the pending output, ensuring that the clinician or Elder — not the machine — determines the urgency classification of all advisories.

Embodiment 4: WM003™ LiDAR Environmental Observation Subsystem

[0021] In a fourth preferred embodiment, the WM003™ LiDAR Environmental Observation Subsystem provides spatial and positional data to the AI advisory subsystem through a passive, read-only sensor interface. The LiDAR array is installed at standardised height parameters determined during the Pre-Deployment Sovereignty Audit (p-003 Ethical and p-004 Clinical stages) to ensure complete coverage of the designated observation zone without infringing on privacy-restricted areas as defined by institutional policy.

[0022] The WM003™ subsystem is architecturally constrained to a unidirectional data output mode. The sensor array transmits spatial observation data to the AI advisory subsystem but possesses no data reception, actuation, or command execution capability. This ensures that the LiDAR system cannot be leveraged as an authority drift vector — no AI-generated command, advisory output, or override signal may be routed through the WM003™ hardware to effect physical change in the observed environment.

[0023] In the Eldercare domain (WD071), WM003™ LiDAR data is used for passive patient positioning monitoring, fall-risk spatial analysis, and equipment proximity observation. In the Defence domain (WD072), the subsystem provides facility perimeter spatial mapping for advisory threat assessment without possessing targeting, communication, or access control capabilities. In the Cyber Security domain (WD073), the WM003™ subsystem monitors physical access zones adjacent to critical infrastructure nodes for human presence advisory inputs.

Embodiment 5: Tiger .1x Key™ Tripartite Authentication Controller

[0024] In a fifth preferred embodiment, the Tiger .1x Key™ Tripartite Authentication Controller enforces the constitutional principle that no single human actor and no AI subsystem component may unilaterally activate, modify, or deactivate the NAI 2.0™ system. The controller requires concurrent and sequential authentication by three designated institutional role holders:

[0025] **Step 1 — Clinical Lead (CMO) — Biometric Hash Authentication:** The Chief Medical Officer or designated clinical authority provides a biometric hash credential to the Tiger .1x Key™ controller. Successful authentication transitions the system from inactive standby to "Pending Safety Review" state. In this state, all AI advisory subsystems remain in a suspended initialisation mode with no active output capability.

[0026] **Step 2 — Technical Lead (IT Director) — Hardware Token Authentication:** The IT Director or designated technical authority inserts a physical hardware token registered to the Tiger .1x Key™ controller. Successful authentication engages the Orange Code 1.1× computational constraints and initialises the Sacred Pause™ FPGA latency parameters. The system transitions to "Constrained Operational Readiness" state.

[0027] **Step 3 — Administrative Lead (Governance Officer) — Alphanumeric Code Authentication:** The Governance Officer provides an alphanumeric credential registered to the Tiger .1x Key™ controller. Successful authentication transitions the system to full "Clinical Advisory Operational" state. All three authentication credentials must be present within a configurable session window; expiry of any credential prior to completion of the tripartite sequence requires full restart of the authentication cycle from Step 1.

[0028] **Step 4 — Override Protocol — Dual-Key Activation by Any Two Authorised Actors:** In any emergency or safety-critical scenario requiring immediate system suspension, any two of the three designated authentication role holders may engage a Dual-Key Override through simultaneous presentation of their respective credentials to the Tiger .1x Key™ controller. Successful dual-key override triggers immediate hardware-level system suspension, activating the Kill-Switch Protocol sequence.

Embodiment 6: WD070–073 Authority Drift Correction Protocol Suite

[0029] In a sixth preferred embodiment, the WD070–073 Authority Drift Correction Protocol Suite provides a structured, hierarchically governed mechanism for detecting and correcting authority drift across all operational domains. The Suite comprises four protocols operating in coordinated constitutional alignment.

[0030] **WD070 — Sovereignty Boundary Enforcement Protocol:** WD070 operates as the master drift correction controller. It receives real-time telemetry from the Orange Code 1.1× computational cap unit, the Sacred Pause™ FPGA module, the EDS (Early Detection System), and the Drift Detection subsystem. WD070 establishes and continuously validates the constitutional boundary between AI advisory scope and human decision authority. Upon detection of any process, output, or state transition that exceeds the defined advisory scope boundary, WD070 initiates an immediate hardware interrupt, suspends the offending process, logs the incident to the ACRA T260229801 audit trail, and notifies the designated Elder authority for review.

[0031] **WD071 — Eldercare Domain Drift Correction Protocol:** WD071 governs AI advisory systems deployed in residential and clinical eldercare environments. Drift correction triggers include: unsolicited alert generation directed at patients or families without Elder (clinician) authorisation; ambient monitoring data transmission beyond the institutionally configured observation perimeter; AI-generated care recommendation outputs that bypass the designated clinician review pathway; and any system attempt to access or modify patient personal data repositories without logged Elder authorisation. WD071 incorporates a Patient Dignity Preservation Constraint (PDPC) that prioritises the suppression of intrusive AI behaviours above advisory functionality continuity.

[0032] **WD072 — Defence Domain Drift Correction Protocol:** WD072 governs AI advisory systems deployed within national defence command, logistics, and situational awareness infrastructure. Drift correction triggers include: AI advisory output classified at or above mission-critical threat assessment levels without mandatory human review; autonomous communication signal generation via AI subsystems to external networks; AI-generated logistics or resource allocation recommendations executed without Elder (command authority) authorisation; and any AI process exhibiting lateral movement across classified network segments. WD072 incorporates a Defence Safe-State Protocol (DSSP) that routes the system to a read-only observational mode upon detecting any unauthorised authority expansion, ensuring no defence advisory output is actionable without explicit human command authorisation.

[0033] **WD073 — Cyber Security Domain Drift Correction Protocol:** WD073 governs AI advisory systems deployed in cyber security operations centres (SOCs) and critical infrastructure protection environments. Drift correction triggers include: AI-initiated automated threat response actions without human-in-the-loop incident classification; autonomous network access modification, firewall rule changes, or access credential revocation by AI subsystems; AI-generated incident severity escalations that bypass the designated human analyst review stage; and self-modification of AI model parameters in response to detected cyber threats. WD073 incorporates an Automated Threat Response Suppression (ATRS) function that confines all AI cyber advisory outputs to read-only advisory display, preventing any AI subsystem from executing, initiating, or authorising any network-level response action.

Embodiment 7: Pre-Deployment Sovereignty Audit Procedures p-002 through p-007

[0034] In a seventh preferred embodiment, the Pre-Deployment Sovereignty Audit comprises six mandatory sequential audit stages, each constituting an irrevocable gating condition for deployment progression. Failure at any stage results in unconditional deployment termination with no bypass, exception, or conditional approval pathway available.

[0035] **p-002 — Technical Sovereignty Audit:** Verification of hardware compatibility with Orange Code 1.1× computational cap parameters. Assessment of legacy system agentic leakage risk. Confirmation of network segmentation adequate to prevent cross-domain authority drift cascade. Output: Technical Sovereignty Certificate or Deployment Termination.

[0036] **p-003 — Ethical Sovereignty Audit:** Formal verification of institutional Board adoption of P-LIFE 1.00™ as the absolute mission constant. Documented evidence of institutional commitment to the four cardinal values (Humility, Silence, Dignity, Benevolence) as engineering requirements. Output: Ethical Sovereignty Certificate or Deployment Termination.

[0037] **p-004 — Clinical Sovereignty Audit:** Mapping of Elder roles with documented, individual, and non-delegable veto authorities across all AI advisory output categories. Verification that no AI advisory output pathway exists that does not terminate in a human Elder review and authorisation node. Output: Clinical Sovereignty Certificate or Deployment Termination.

[0038] **p-005 — Governance Sovereignty Audit:** Verification of established reporting lines to ACRA T260229801 for all incident audit trails. Confirmation of Kill-Switch Protocol reversion test completion with documented successful reversion to **100%** manual operational pathways. Output: Governance Sovereignty Certificate or Deployment Termination.

[0039] **p-006 — Sovereignty Verification:** Comprehensive cross-domain sovereignty verification integrating outputs from p-002 through p-005. Verification of zero agentic leakage in all existing legacy software interfacing with the NAI 2.0™ deployment. Final review by Non-Agentic AI Governance Singapore (ACRA T260229801) authorised auditor. Output: Sovereignty Verification Certificate or Deployment Termination.

[0040] **p-007 — WISL™ Certificate Issuance:** Upon successful completion of all preceding audit stages, the WISL™ (We Innovate Save Live™) Deployment Certificate is issued, authorising commencement of hardware installation. The WISL™ Certificate carries a unique deployment identifier linked to the institutional ACRA T260229801 audit trail and remains valid for the duration of the deployment lifecycle subject to annual re-audit.

Embodiment 8: Post-Deployment Sentinel Monitoring and Kill-Switch Protocol

[0041] In an eighth preferred embodiment, the Post-Deployment Sentinel Monitoring system operates three continuously active monitoring pillars aligned with P-LIFE 1.00™:

[0042] **EDS (Early Detection System):** Constant telemetry monitoring of the Orange Code 1.1× computational cap unit, the Sacred Pause™ FPGA module, and the WM003™ LiDAR subsystem for hardware variance, performance anomaly, and unauthorised parameter modification. EDS generates Elder notification alerts upon detection of any monitored parameter exceeding a pre-configured variance threshold. EDS does not generate autonomous corrective actions; all EDS alerts are advisory outputs subject to Elder review.

[0043] **Drift Detection Subsystem:** Continuous evaluation of AI advisory model outputs against the established safety benchmark profile captured during p-002 Technical Sovereignty Audit. Drift is quantified as the statistical deviation of live advisory output distributions from the baseline profile. Upon detection of drift exceeding a configurable threshold, the Drift Detection Subsystem initiates an immediate WD070 Sovereignty Boundary Enforcement trigger and notifies the designated AI Safety Officer for recalibration authorisation.

[0044] **Kill-Switch Protocol — Five-Stage Reversion Sequence:**

- **Detection:** An Elder authority or AI Safety Officer identifies a critical safety, stability, or authority drift risk through EDS notification, Drift Detection alert, direct clinical observation, or Tiger .1x Key™ override engagement.

- **Trigger:** Engagement of the physical Kill-Switch actuator or initiation of the Tiger .1x Key™ dual-override sequence by any two authorised actors.

- **Sever:** Hardware-level severance of the Orange Code 1.1× computational cap unit's power and data bus connections, immediately terminating all AI advisory processing and output.

- **Revert:** Institution reverts unconditionally to **100%** manual clinical, operational, or cyber security pathways with immediate effect. No AI advisory output, cached recommendation, or pending Sacred Pause™ buffer content may be actioned following Sever.

- **Report:** Mandatory incident report and complete audit trail submission to Non-Agentic AI Governance Singapore (ACRA T260229801) within **24** hours of Kill-Switch Trigger, regardless of domain, severity classification, or operational context.

CLAIMS

[0045] What is claimed is:

Claim 1. A hardware-enforced non-agentic artificial intelligence governance system comprising: a P-LIFE 1.00™ constitutional anchor encoded at firmware level; an Orange Code 1.1× computational cap unit imposing a hard-coded ceiling on AI subsystem computational resources; a Sacred Pause™ Field-Programmable Gate Array (FPGA) latency enforcement module interposed in the AI advisory output pathway; a WM003™ LiDAR environmental observation subsystem configured in unidirectional, read-only sensor output mode; a Tiger .1x Key™ tripartite authentication controller requiring concurrent sequential authentication by a Clinical Lead, Technical Lead, and Administrative Lead; and a Kill-Switch Protocol reversion actuator providing unconditional reversion to fully manual operational pathways.

Claim 2. The system of Claim 1, wherein the Orange Code 1.1× computational cap unit is implemented in read-only firmware registers not addressable through the system's operating software layer, and wherein the cap ceiling is set at **1.1** times the minimum advisory processing requirement established during a Pre-Deployment Sovereignty Audit.

Claim 3. The system of Claim 1, wherein the Sacred Pause™ FPGA module enforces a mandatory Human Review Window during which all AI advisory outputs are held in a secure buffer inaccessible to the AI processing pipeline, and wherein an Automated Trigger Suppression function prevents the AI subsystem from generating autonomous action signals during the Human Review Window regardless of advisory urgency classification.

Claim 4. The system of Claim 1, wherein the Tiger .1x Key™ tripartite authentication controller requires: a biometric hash credential from the Clinical Lead to transition the system to Pending Safety Review state; a physical hardware token from the Technical Lead to engage Orange Code 1.1× constraints; an alphanumeric credential from the Administrative Lead to achieve Clinical Advisory Operational state; and a dual-key override by any two of the three designated role holders to trigger immediate hardware-level system suspension.

Claim 5. The system of Claim 1, further comprising a WD070–073 Authority Drift Correction Protocol Suite comprising: WD070 — Sovereignty Boundary Enforcement Protocol operative as a master drift correction controller; WD071 — Eldercare Domain Drift Correction Protocol incorporating a Patient Dignity Preservation Constraint; WD072 — Defence Domain Drift Correction Protocol incorporating a Defence Safe-State Protocol; and WD073 — Cyber Security Domain Drift Correction Protocol incorporating an Automated Threat Response Suppression function.

Claim 6. The system of Claim 5, wherein WD070 receives real-time telemetry from the Orange Code 1.1× computational cap unit, the Sacred Pause™ FPGA module, an Early Detection System, and a Drift Detection subsystem, and wherein detection of any process exceeding the advisory scope boundary initiates a hardware interrupt, suspends the offending process, and logs an incident to an ACRA T260229801 audit trail.

Claim 7. The system of Claim 5, wherein WD073 confines all AI cyber advisory outputs to read-only advisory display mode and prevents the AI subsystem from executing, initiating, or authorising any network-level response action including automated threat response, firewall rule modification, access credential revocation, or AI model self-modification.

Claim 8. A method for authority drift correction in life-critical artificial intelligence deployments comprising the steps of: conducting a Pre-Deployment Sovereignty Audit comprising six sequential gating stages (p-002 through p-007) each constituting an irrevocable deployment condition; installing hardware-level computational constraints including an Orange Code 1.1× cap unit and Sacred Pause™ FPGA module; configuring tripartite Tiger .1x Key™ authentication roles; deploying WD070–073 Authority Drift Correction Protocols across Eldercare, Defence, and Cyber Security domains; operating continuous post-deployment sentinel monitoring via EDS, Drift Detection, and Kill-Switch Protocol; and submitting mandatory incident reports to ACRA T260229801 within **24** hours of any Kill-Switch activation event.

Claim 9. The method of Claim 8, wherein failure at any stage of the Pre-Deployment Sovereignty Audit (p-002 through p-006) results in unconditional deployment termination with no bypass, exception, or conditional approval pathway, and wherein hardware installation may not commence until a WISL™ Certificate (p-007) has been issued.

Claim 10. The method of Claim 8, wherein the Kill-Switch Protocol reversion sequence comprises Detection, Trigger, Sever, Revert, and Report stages, and wherein the Revert stage mandates unconditional restoration of **100%** manual operational pathways with no AI advisory output, cached recommendation, or pending buffer content actionable following hardware-level severance.

Claim 11. A non-agentic AI governance apparatus for multi-domain life-critical deployments, the apparatus constitutionally governed by the mission constant P-LIFE 1.00™ expressed as the formula Harm = Death · North = Save Life, and configured to enforce human sovereignty over AI advisory functions across Eldercare, Defence, and Cyber Security institutional domains through hardware-level constraints, tripartite authentication, authority drift correction protocols, and irrevocable kill-switch reversion mechanisms, substantially as described herein with reference to the accompanying figures.

Claim 12. The apparatus of Claim 11, wherein the four cardinal engineering values of Humility (謙虛), Silence (沉默), Dignity (尊严), and Benevolence (仁) are implemented as operational hardware and protocol constraints rather than advisory guidelines, and wherein no software modification, administrative override, or runtime parameter adjustment by any human or AI actor within the system hierarchy may alter, suspend, or bypass any constraint derived from these cardinal values.

Claim 13. The apparatus of Claim 11, wherein cross-domain authority drift cascade between Eldercare, Defence, and Cyber Security domains is prevented through domain isolation barriers enforced at shared infrastructure computational pathways, each pathway carrying an Orange Code 1.1× hard-cap enforcement point that prevents authority expansion vectors from propagating across domain boundaries.

ABSTRACT

[0046] The Non-Agentic AI 2.0™ Constitutional Framework for Authority Drift Correction provides a comprehensive system, method, and apparatus for preventing artificial intelligence systems from acquiring or exercising autonomous agency in life-critical institutional environments. The framework is constitutionally anchored by the P-LIFE 1.00™ mission constant (Harm = Death · North = Save Life) and enforces human sovereignty through hardware-level constraints including the Orange Code 1.1× computational cap, Sacred Pause™ FPGA latency enforcement, and WM003™ LiDAR passive environmental observation. The Tiger .1x Key™ tripartite authentication controller ensures no single actor may unilaterally activate or override the system. The WD070–073 Authority Drift Correction Protocol Suite provides domain-specific drift correction across Eldercare, Defence, and Cyber Security operational environments. A six-stage Pre-Deployment Sovereignty Audit (p-002 through p-007) gates all deployment behind institutional sovereignty verification. Post-deployment sentinel monitoring via EDS, Drift Detection, and Kill-Switch Protocol maintains continuous alignment with P-LIFE 1.00™. The framework is aligned with Singapore's AIHGle 2.0 guidelines and WHO Maturity Level 4 standards.

DECLARATION

[0047] I, Edwin Koh Wui Kiat, being the inventor of the subject matter claimed in this application, declare that the information contained herein is true and correct to the best of my knowledge, information, and belief. This application is filed in good faith and constitutes a sincere disclosure of the technical invention as conceived and reduced to practice under the WISL™ constitutional framework.

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

Edwin Koh Wui Kiat · Tiger · P-LIFE 1.00™
Non-Agentic AI Governance Singapore · ACRA T260229801
NLB Legal Deposit R260302-007 · Patent IPOS SG020603109STW
Singapore · 2026

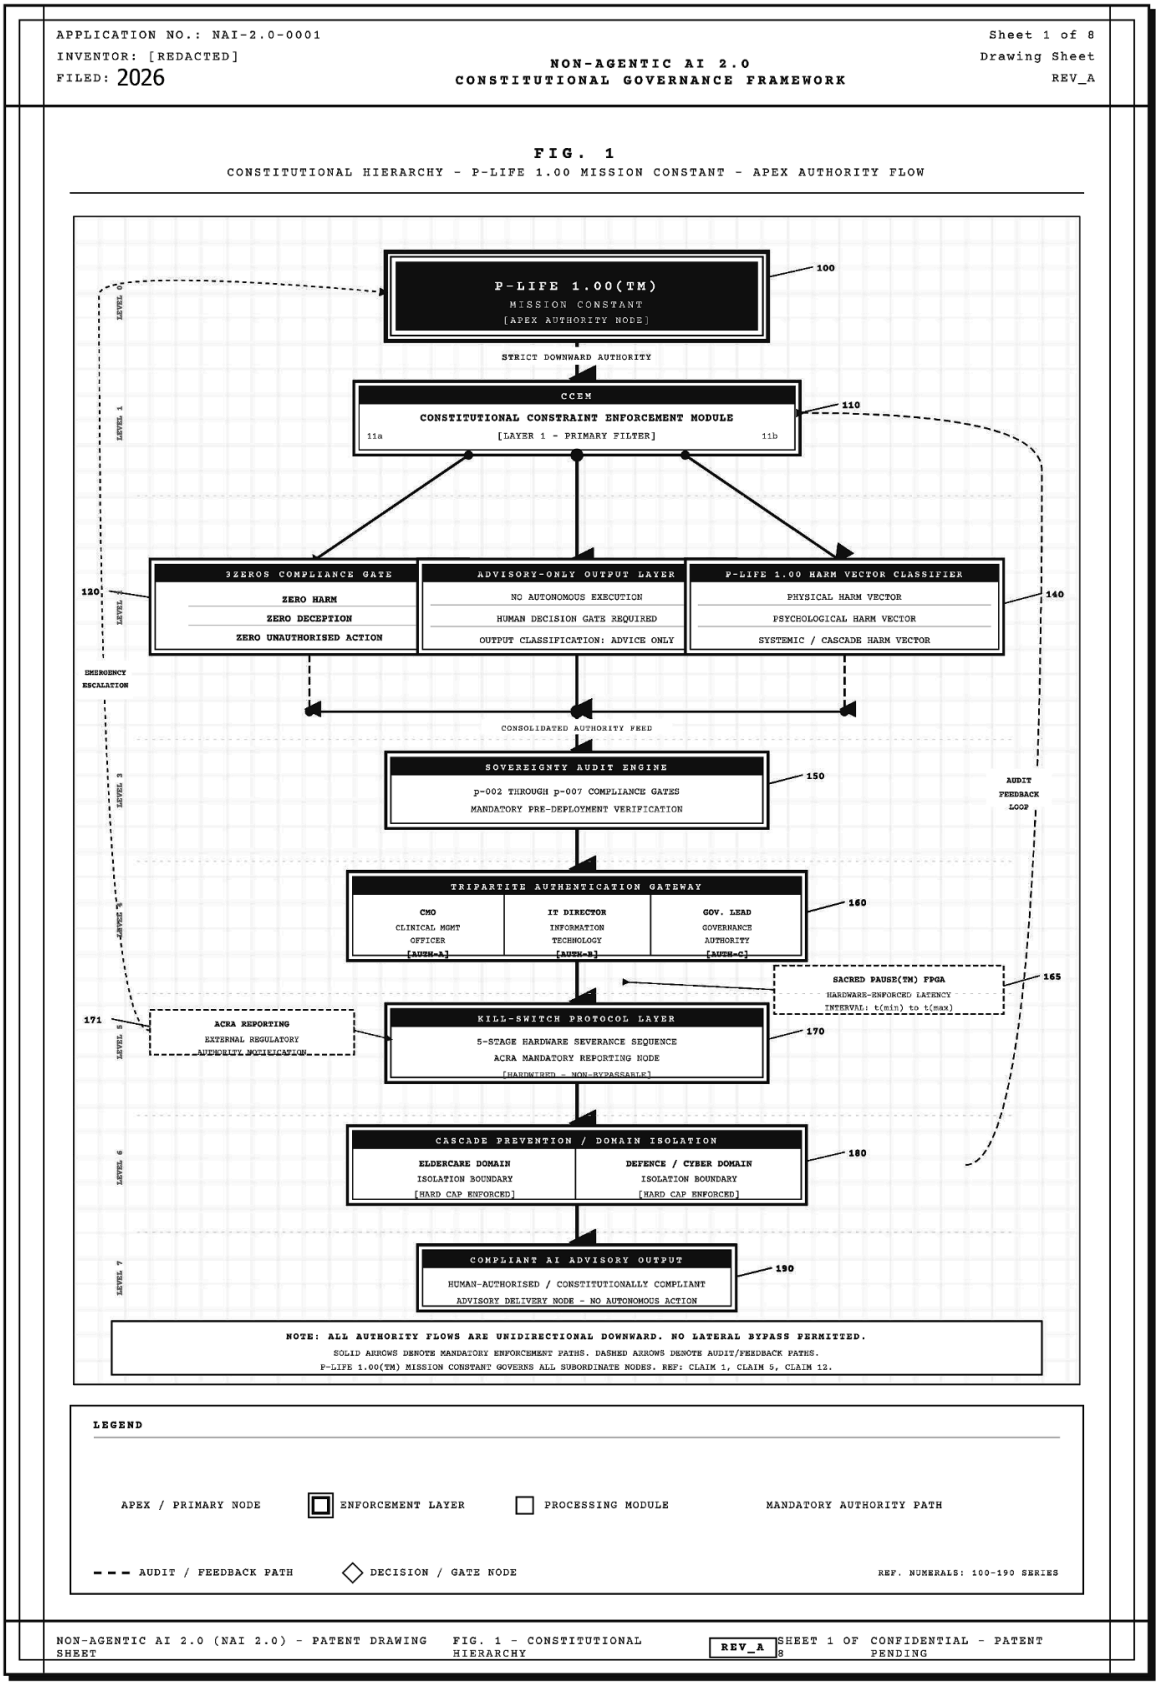

Constitutional Hierarchy (FIG. 1): Illustrates the P-LIFE 1.00™ Mission Constant as the apex node with strict downward authority flow.

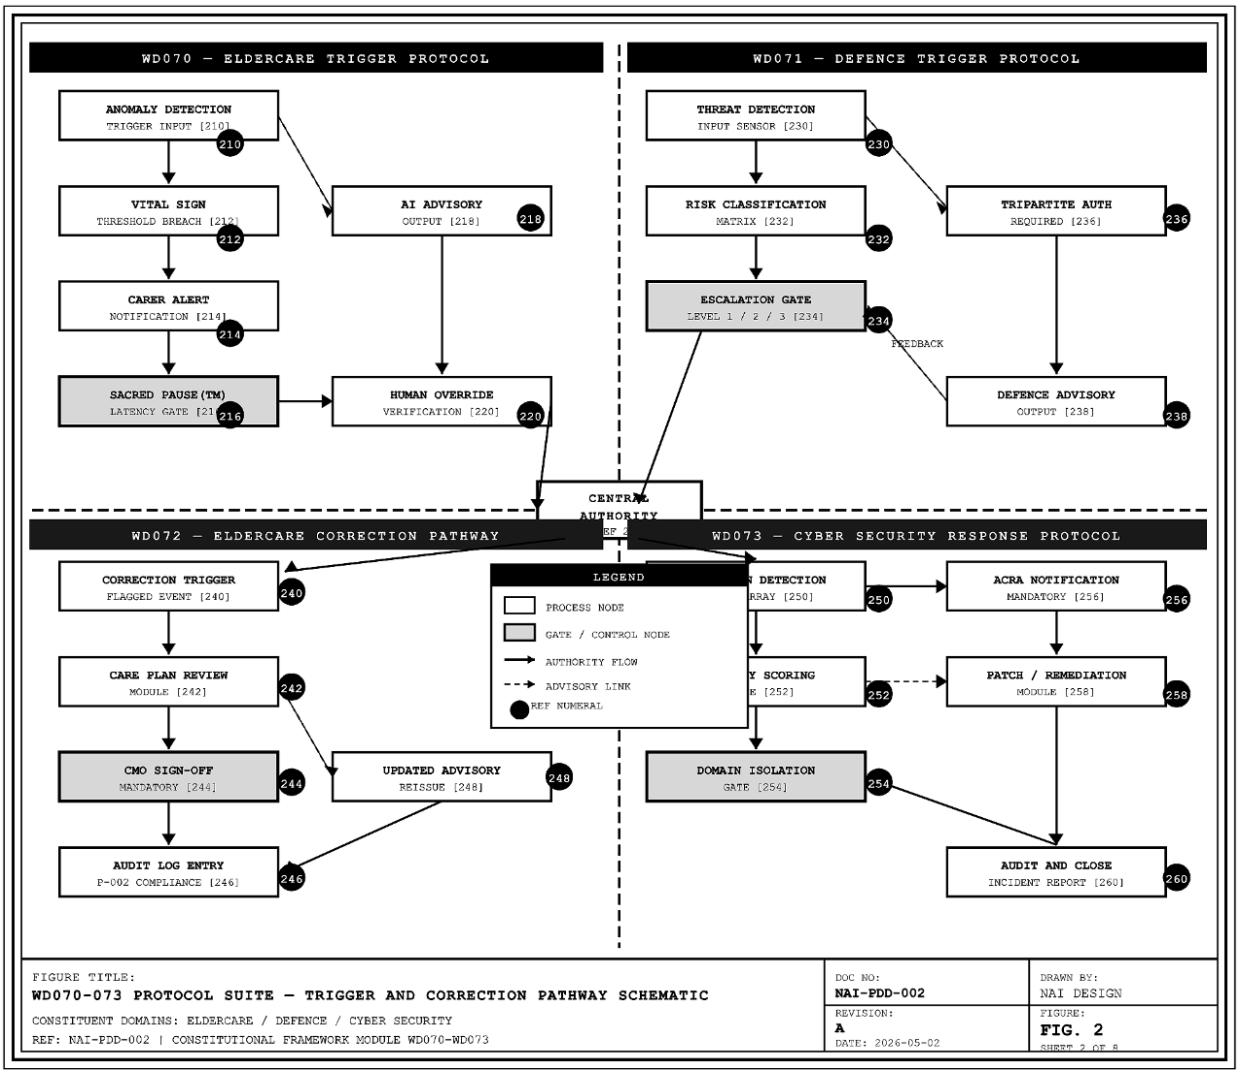

FIG. 2 — REFERENCE NUMERAL INDEX

NUMERAL	DESIGNATION	PROTOCOL	DOMAIN
200	Central Authority Node	—	All Domains
210	Anomaly Detection Trigger Input	WD070	Eldercare
212	Vital Sign Threshold Breach	WD070	Eldercare
214	Carer Alert Notification	WD070	Eldercare
216	Sacred Pause(TM) Latency Gate	WD070	Eldercare
218	AI Advisory Output	WD070	Eldercare
220	Human Override Verification	WD070	Eldercare
230	Threat Detection Input Sensor	WD071	Defence
232	Risk Classification Matrix	WD071	Defence
234	Escalation Gate L1/L2/L3	WD071	Defence
236	Tripartite Authentication Required	WD071	Defence
238	Defence Advisory Output	WD071	Defence
240	Correction Trigger — Flagged Event	WD072	Eldercare
242	Care Plan Review Module	WD072	Eldercare
244	CMO Sign-Off — Mandatory	WD072	Eldercare
246	Audit Log Entry — P-002 Compliance	WD072	Eldercare
248	Updated Advisory Reissue	WD072	Eldercare
250	Intrusion Detection Sensor Array	WD073	Cyber Security

NUMERAL	DESIGNATION	PROTOCOL	DOMAIN
252	Severity Scoring Engine	WD073	Cyber Security
254	Domain Isolation Gate	WD073	Cyber Security
256	ACRA Notification — Mandatory	WD073	Cyber Security
258	Patch / Remediation Module	WD073	Cyber Security
260	Audit and Close — Incident Report	WD073	Cyber Security

WD070–073 Protocol Suite (FIG. 2): A four-quadrant schematic mapping triggers and correction pathways for Eldercare, Defence, and Cyber Security.

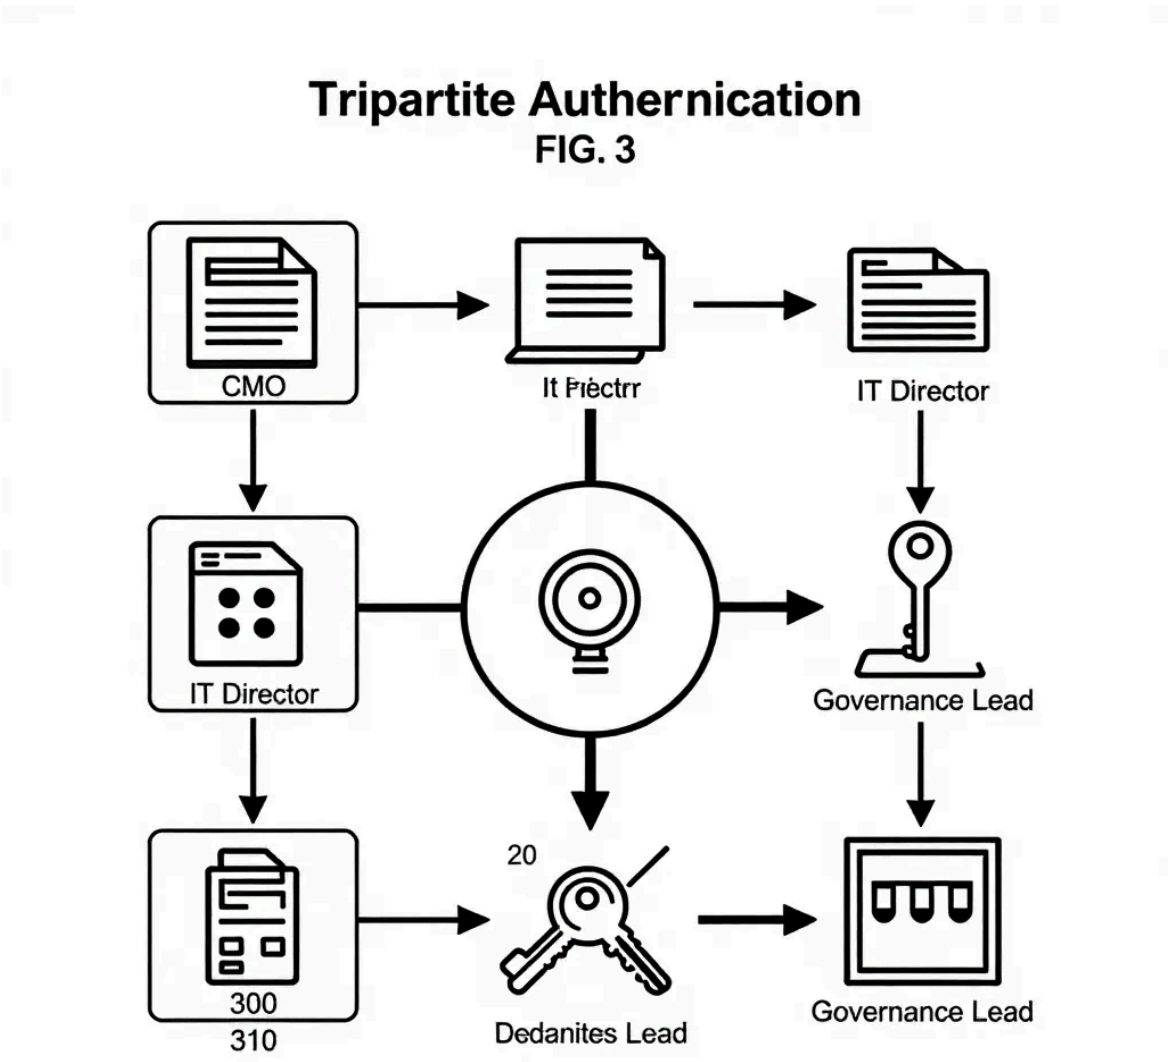

Tripartite Authentication (FIG. 3): Details the sequential CMO, IT Director, and Governance Lead verification process, including the dual-key hardware override.

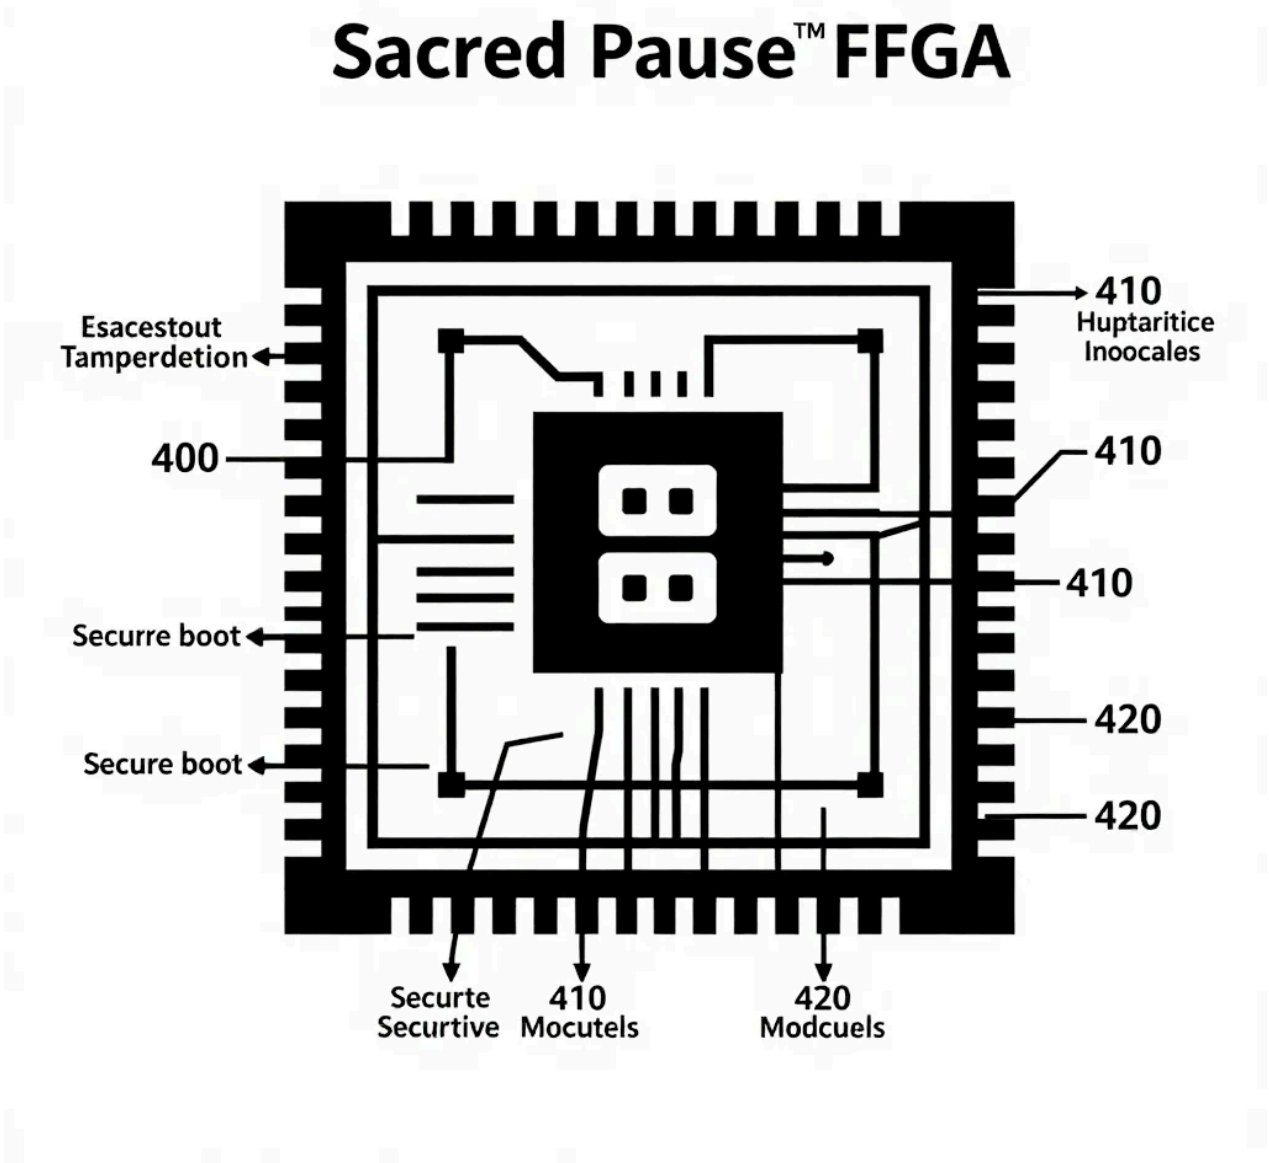

Sacred Pause™ FPGA (FIG. 4): A timing diagram of the hardware-enforced latency interval between AI advice and action.

FIG. 5 — WM003™ LIDAR SYSTEM ARCHITECTURE

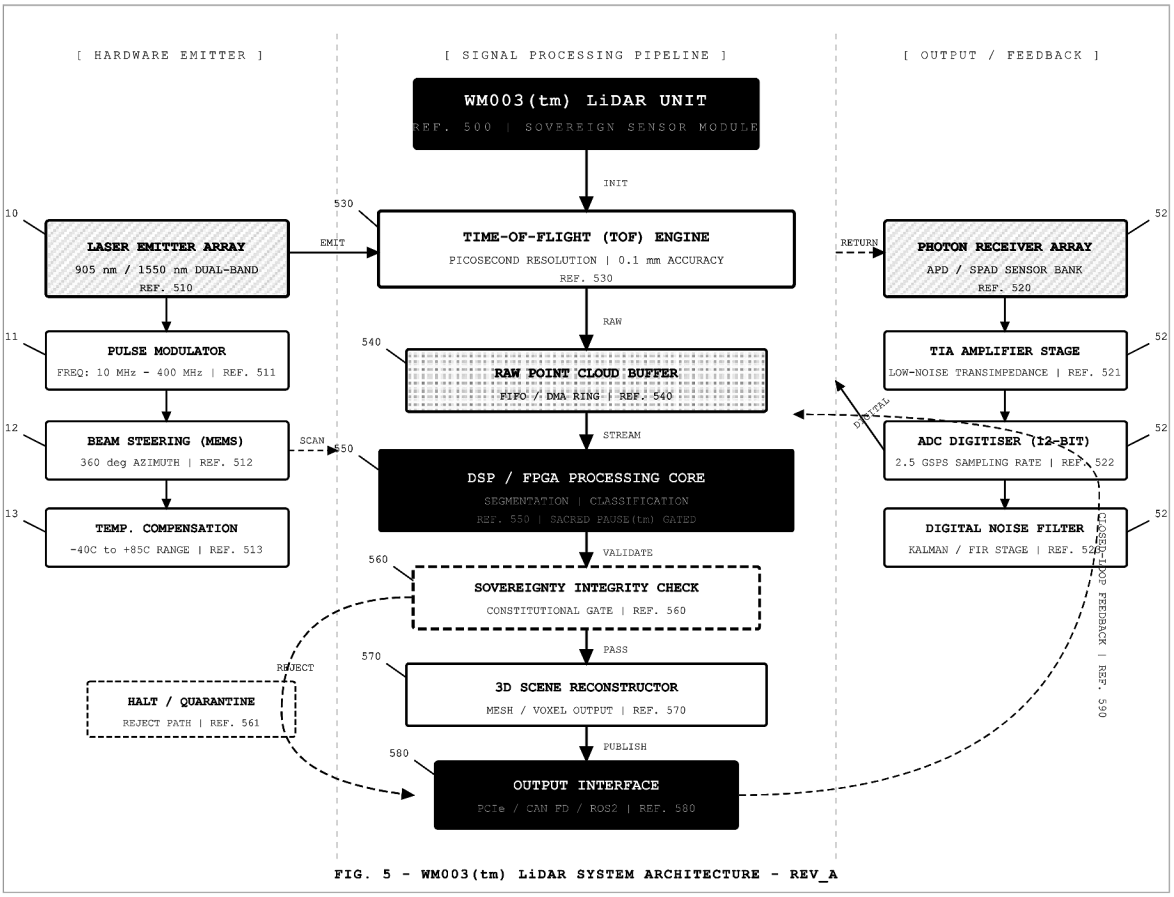

FIG. 5 - WM003(tm) LiDAR SYSTEM ARCHITECTURE - REV_A

WM003™ LiDAR (FIG. 5): Isometric schematic of the sensor-only observation array.

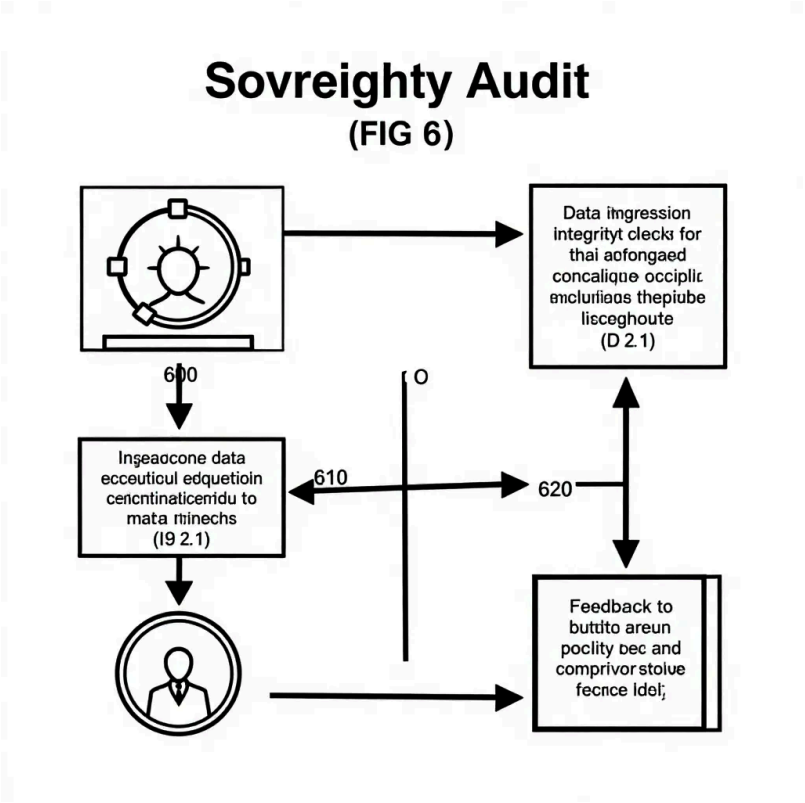

Sovereignty Audit (FIG. 6): The p-002 through p-007 gated flowchart ensuring mandatory compliance before deployment.

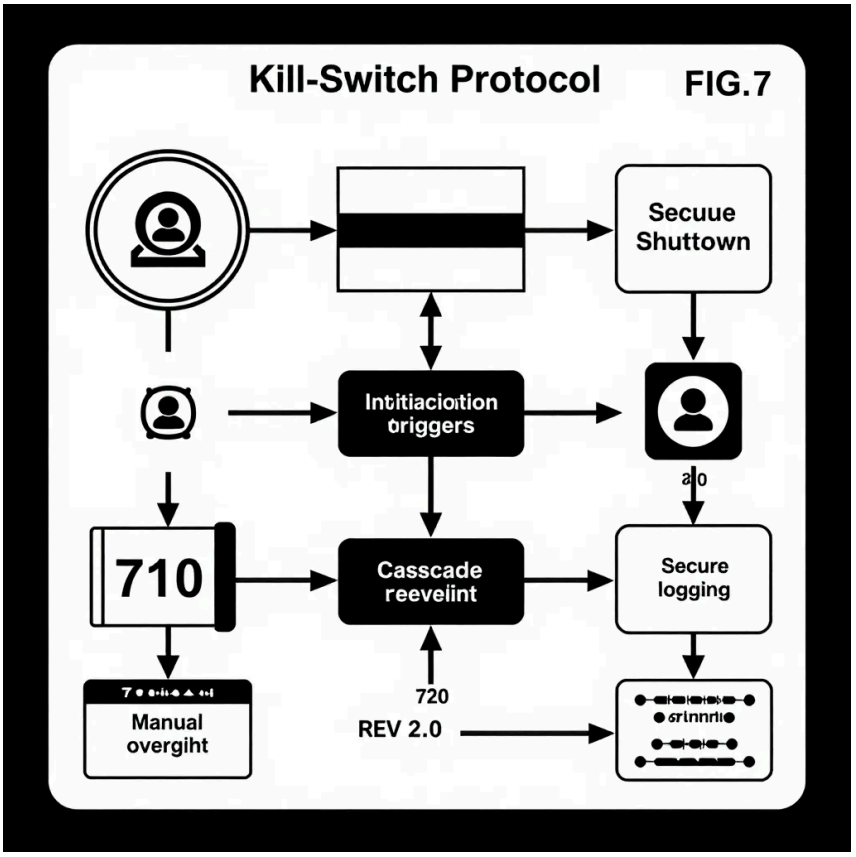

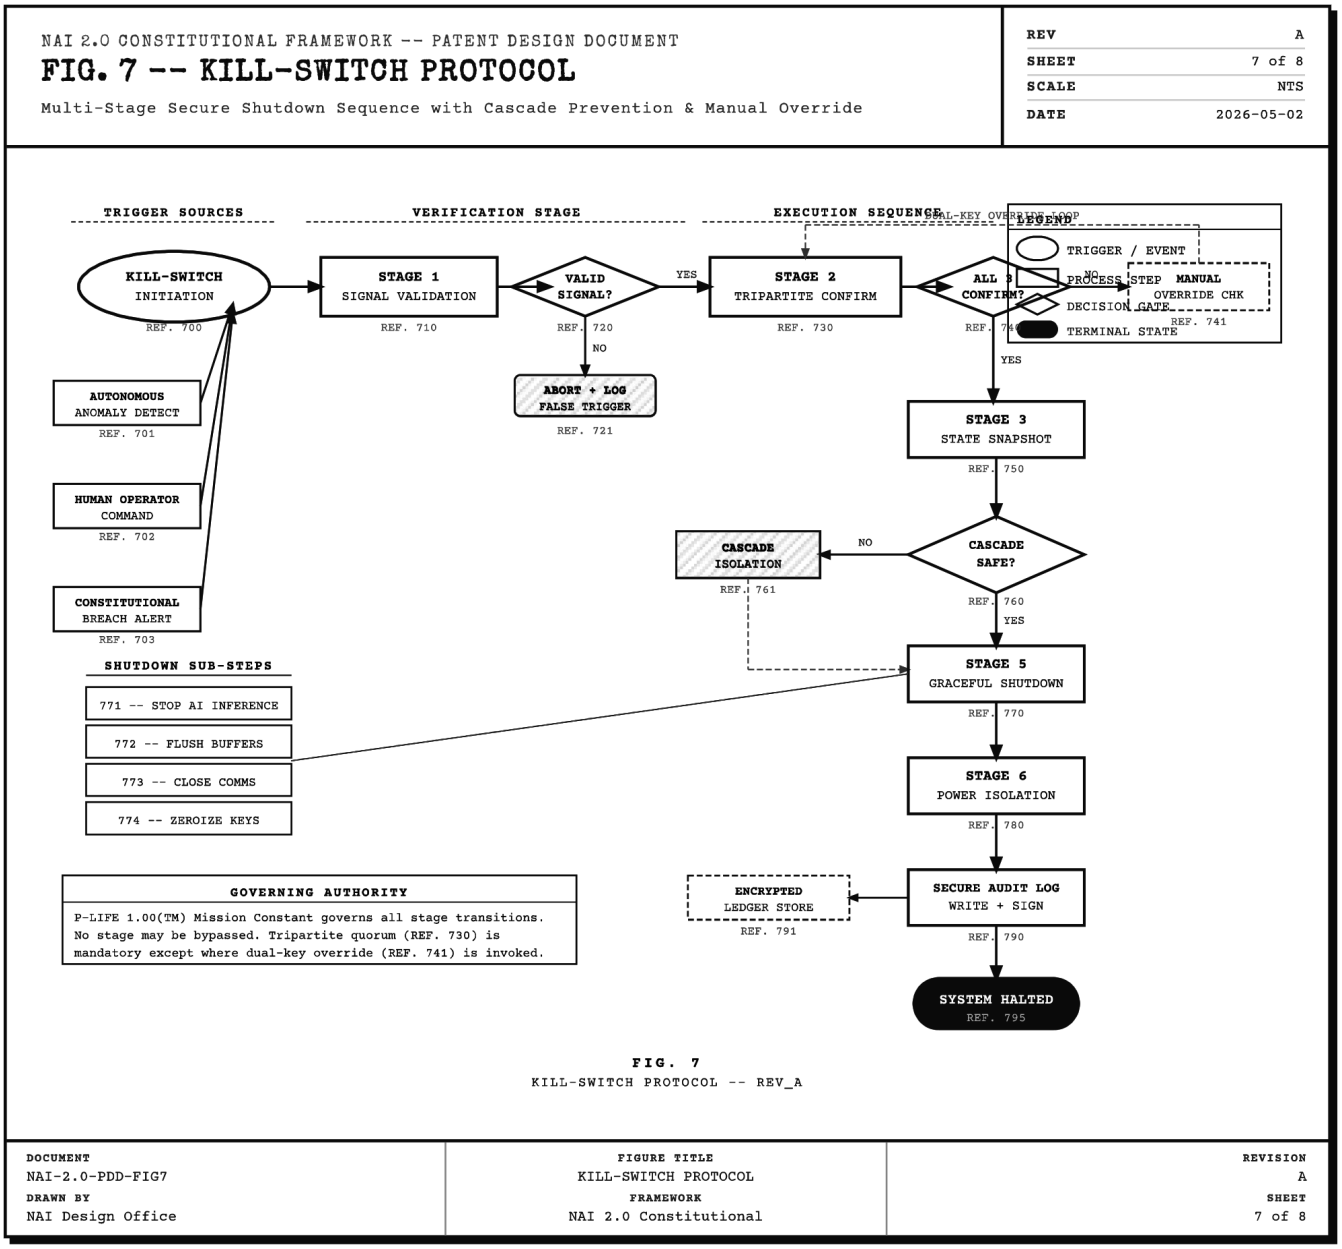

Kill-Switch Protocol (FIG. 7): Five-stage sequence showing hardware severance and ACRA reporting.

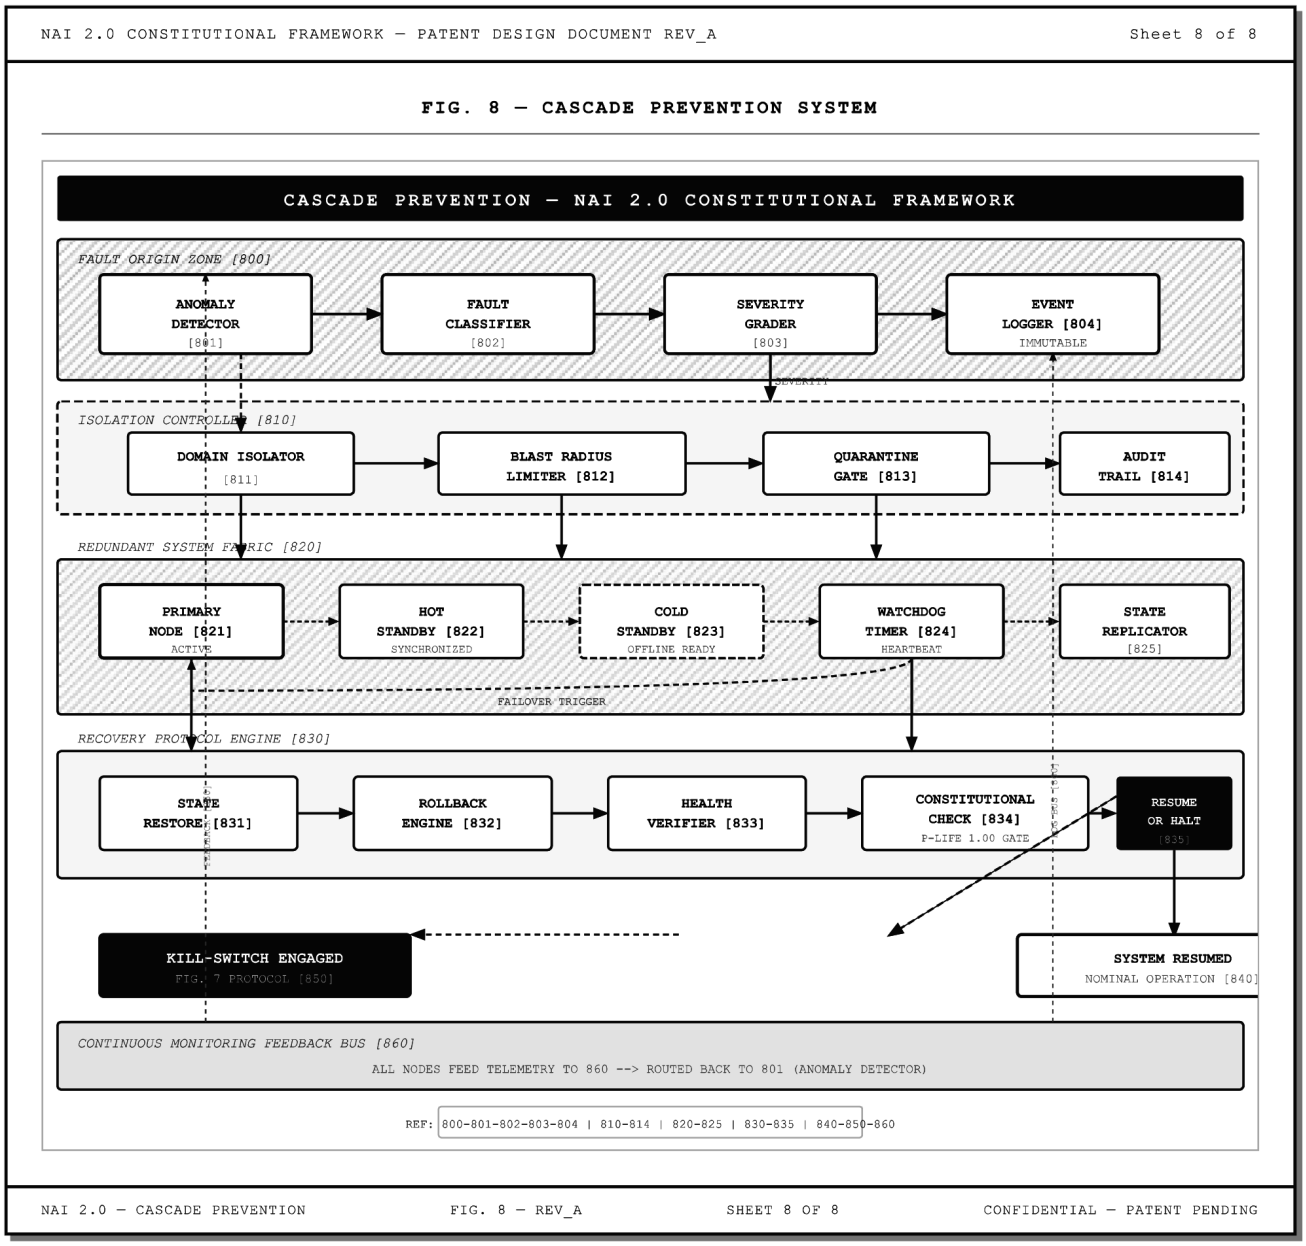

Cascade Prevention (FIG. 8): Network topology demonstrating domain isolation and shared infrastructure hard-caps.