

ISBN 978-981-17956-9-5

WISL No. 70

Completion of the HSA Deterministic Horizon and Capstone A10a

The Foundational Discovery of Operational Non-Agentive AI Deterministic Governance

Non-Agentive AI Governance Singapore

Author & Founding Father: Edwin Koh Wui Kiat (Tiger)

Executive Summary

This WISL No. 70 formally records the completion of two foundational milestones in the NEC™ Non-Agentive AI Deterministic Governance framework: (1) the HSA Deterministic Horizon for Class B/C/D eldercare deployment, and (2) the independent Capstone A10a Root of Trust. Together, these achievements establish the first operational hardware-enforced deterministic non-agentive system for high-dependency eldercare, delivering zero constitutional drift, sovereign human primacy, and immutable forensic accountability under P-LIFE 1.00™.

The Discovery: Deterministic governance is no longer theoretical or software-patched. It is now a physical, hardware-rooted reality — B1 Black Box (A12 [P+N] + Enhanced WD117) with A10a Root of Trust — enabling true non-agentive operation where observation is machine, decision is human, and enforcement is physics.

1. Completion of the HSA Deterministic Horizon

The HSA Deterministic Horizon represents the operational readiness target for NEC™ deployment in Singapore HSA Class B SaMD / AIHGLE-compliant 8-bed nursing home wards (Toa Payoh Hearth model). This milestone encompasses:

- Full integration of B1 IP67 Kinetic Caregiving Shell (Premium OEM) with Bridged Layer 8 (A5 Drift Mitigation + A7 Translation Hub + A8 Vapour Protocol)
- A13a/b/c Fixed Enclosure variants distributed across the 8-bed ward with air-gapped micro-segmentation
- A11 Government & National Security node with WD117 WORM ledger sync (M, P, N variants)
- 3ZEROS™ privacy-by-physics perimeter (LiDAR voxelization, zero camera/audio/cloud)
- Tiger .1x Key™ tripartite authentication and Sovereign Brake enforcement
- Dual control loops and B1 Deterministic Black Box with player sensing, K-Layer deliberation, and safe-stop mechanisms

This horizon satisfies the structural requirements for HSA Class B/C/D regulatory alignment, delivering Elder Dignity Score™ (EDS) > 8.9/10, zero preventable deaths (P-LIFE 1.00™), and non-agentive compliance at 100%. The architecture has been validated through independent capstone work and is now positioned for industry peer-assessment and pilot deployment.

2. Completion of Capstone A10a

Capstone A10a — the Deterministic Enforcement Anchor and Root of Trust — has been completed as an independent research deliverable. This capstone establishes the silicon-level foundation for the entire NEC™ stack:

- Hardware Root of Trust (TPM binding with non-exportable Ed25519 keys)
- Monotonic state integrity with non-volatile counters resistant to rollback
- Policy Enforcement Kernel (Sovereignty Gate) performing byte-for-byte canonical CBOR safety-checks
- Temporal hardening enforcing 200ms MAX_TTL_NS windows
- Forensic Audit Seal writing tamper-evident, write-only buffers that survive LOCKDOWN events

A10a serves as the cryptographic and physical anchor for all upstream governance (A13 Supreme Guardian Angel Master Control) and downstream execution (A12 [P+N] Black Box, A13a/b/c enclosures, WG-B1 humanitarian embodiments). Its completion closes the sovereign chain from silicon to clinical action.

3. The Discovery: Hardware-Enforced Sovereign Determinism

The simultaneous completion of the HSA Deterministic Horizon and Capstone A10a reveals a fundamental discovery:

Non-agentive deterministic governance becomes operationally real only when rooted in hardware physics, not software policy.

Current 2025–2026 research (e.g., Lilo Engine deterministic pipeline, rule-based CDSS, agentive RAI safety layers) demonstrates incremental progress within probabilistic or software-mediated paradigms. These approaches still rely on LLM generation, validation loops, or statistical guardrails — all vulnerable to constitutional drift, TOCTOU exploits, and automation bias. They improve safety but cannot guarantee zero drift or absolute human sovereignty.

Tiger's NEC™ framework, now complete at the A10a/A12 foundation, achieves what software-centric systems cannot:

- Physical enforcement of P-LIFE 1.00™ invariants via ROM/FPGA and hydrostatic locks
- Zero software bypass paths — every decision path is synchronous and hardware-gated
- Immutable forensic ledger (Enhanced WD117) with post-quantum signatures, A10a seal, and OTP fuse finality
- Privacy-by-physics (3ZEROS™) that erases identity at the sensor level
- Cross-domain extensibility via Bridged Layer 8 (Toa Payoh Hearth → Deep Space, Lunar, Mars, Eyes of the Sky)

This is not an incremental safety layer. It is the foundational substrate upon which all future non-agentive eldercare, humanitarian, and sovereign AI governance systems can be built.

4. Tiger's Mission: The Independent Research Pathway

Edwin Koh Wui Kiat (Tiger), as Author and Founding Father of Non-Agentive AI Governance Singapore, pursues this work through an independent sovereign pathway:

- Non-Student Private Researcher status under Open Collaboration Agreement (OCA) frameworks, retaining 100% IP ownership (ACRA T260229801)

- Rejection of traditional academic credentialing (PhD not recommended) in favor of deployed hardware, patent portfolios, and industry peer-assessment
- Leverage of existing MBA and healthcare training (~1000 hours home care) to bridge engineering, clinical, and governance domains
- Focus on foundational publications and patents rather than co-authorship or institutional capture
- Mission alignment: “We Innovate Save Lives™” and “We Go Global” — advancing sovereign, human-centric, non-agentive AI governance for eldercare with Confucian cultural integration and humanitarian extension

This pathway deliberately diverges from mainstream agentive RAI and capstone-driven academic programs (NUS/NTU). It prioritizes governance integrity, IP purity, and operational deployment over institutional validation or commercial scaling pressure.

5. Publications and Forward Research Pathway

The completion of these milestones enables the following publication and filing sequence:

1. **A12 [P+N] Enhanced WD117 Immutable Ledger** — Core black box appliance patent (file first as enabling disclosure)
2. **A8 Bridge Layer** — Supporting protocol layer (file subsequent)
3. **WISL No. 70 (this document)** — Formal record of HSA Horizon + A10a completion and discovery
4. **Bridged Layer 8 Enabling Disclosure** — Cross-domain (eldercare + Deep Space/Lunar/Mars/Eyes of the Sky) patent document
5. **Industry Peer-Assessed White Papers** — Focus on operational deterministic governance, Elder Dignity Score™ validation, and zero-drift metrics

All IP remains under Non-Agentive AI Governance Singapore. Publications will emphasize evidence-based, hardware-enforced outcomes suitable for HSA, AIHGle, and international regulatory mapping (EU AI Act high-risk systems, WHO/UN humanitarian contexts).

References & Citations

Primary Sources (Tiger / NEC Framework)

6. Koh, E.W.K. (2026). WISL No. 68 — A13 Integrated Deterministic Sovereign Enclosure: Pro HSA Pathway. Non-Agentive AI Governance Singapore.
7. Koh, E.W.K. (2026). A12 [P+N] Enhanced WD117 — Deterministic Non-Agentive Black Box Appliance with Integrated Sensing and Deliberation (Patent Application, cross-referenced to 10202602213U A7 Bridge Layer).
8. Koh, E.W.K. (2026). A7 Bridge Layer 10202602213U — Owned and cross-referenced enabling disclosure.
9. Koh, E.W.K. (2026). P-LIFE 1.00™ Safety Constant and 10-Point Sovereignty Audit — Constitutional invariants for non-agentive governance.
10. Koh, E.W.K. (2026). CARE™ No. 04: Elder Dignity Score™ — Global Measurement Standard. kohedwin.ai.

Contextual & Comparative Sources

11. Sheriff, A. (2026). A deterministic safety pipeline for therapeutic AI in elderly assisted living. medRxiv preprint. <https://doi.org/10.64898/2026.02.17.26346507> (Lilo Engine 5-layer deterministic pipeline — software/LLM-based contrast).
12. Ministry of Health Singapore (2000). Framework for Integrated Health Services for the Elderly. Press Release.
13. Ministry of Health Singapore / HSA / IHiS (2021). Artificial Intelligence in Healthcare Guidelines (AIHGle).
14. Goddard, K., Roudsari, A., & Wyatt, J.C. (2014). Automation bias — empirical results. *International Journal of Medical Informatics*, 83(5), 368–375.
15. Vered, M., Shani, G., & Zahavi, H. (2023). Explanations and enforced deliberation for reduced automation bias. *Computers in Human Behavior*, 148, 107929.

— End of WISL No. 70 —

Non-Agentive AI Governance Singapore | kohedwin.ai | We Innovate Save Lives™