Non-Agentic AI 2.0 Constitutional Governance and Privacy-Preserving Sensing System

Document Type

Integrated patent design and specification draft

Document Purpose

This document consolidates the core architectural, functional, governance, privacy, safety, clinical, and humanitarian concepts across the six-source NAI 2.0 body into **one coherent patent design** suitable for conversion into:

- **IPOS patent specification**
- **design + utility filing pack**
- **PCT-ready master draft**
- **formal inventor disclosure**
- **drawings-to-claims alignment document**

---

1. Title of the Invention

Non-Agentic AI 2.0 Constitutional Governance and Privacy-Preserving Sensing System

Alternative Filing Titles

- **Human-Authorized Non-Agentic AI Governance Core with Privacy-by-Physics Sensing**
- **Privacy-Preserving Clinical and Humanitarian Monitoring System with Hardware-Enforced Human Sovereignty**
- **Non-Agentic AI Governance Engine with Sacred Pause, Sovereign Brake, and Drift-Control Architecture**

---

2. Technical Field

The invention relates generally to:

- artificial intelligence governance systems,
- human-in-the-loop computational safety systems,
- privacy-preserving sensing instruments,
- clinical and eldercare monitoring systems,
- humanitarian field-deployable verification platforms,
- hardware-enforced governance architectures,
- edge-based advisory decision support systems,
- constitutional and non-agentic AI control frameworks.

More specifically, the invention relates to a **non-agentic AI platform** in which machine outputs are structurally constrained so that the system may **observe, derive, score, and offer**, but may not autonomously diagnose, prescribe, execute, command, or cause downstream action without explicit, authority-bound human approval.

---

3. Integrated Source Themes Consolidated into One Patent Design

This unified patent design integrates six major concept families into one invention:

Source Family 1 — Governance Core Engine

A governance core that binds all AI outputs to a human authority pathway and blocks autonomous execution.

Source Family 2 — Privacy-Preserving Sensing Instrument

A non-invasive monitoring device using non-conventional sensing, including point-cloud and related privacy-preserving modalities, instead of traditional camera-centric surveillance.

Source Family 3 — 3ZEROS Privacy Stack

A privacy architecture based on:

- **Zero Camera**
- **Zero Audio**
- **Zero Cloud**

with local processing and short-retention purge logic.

Source Family 4 — Sacred Pause and Sovereign Brake

A hardware-enforced delay and interruption architecture ensuring a mandatory deliberative pause and immediate physical override capability.

Source Family 5 — Drift Governance Chain

A detection and remediation framework comprising:

- Detect
- Freeze
- Audit
- Purge

to prevent constitutional drift, unauthorized logic mutation, and unsafe behavioral deviation.

Source Family 6 — Clinical and Humanitarian Deployment Architecture

A deployment model covering:

- eldercare,
- institutional safety,
- low-infrastructure environments,
- humanitarian operations,
- border and corridor verification,
- conflict-sensitive and privacy-sensitive field conditions.

---

4. Background of the Invention

Conventional AI systems increasingly combine sensing, inference, classification, actuation, and workflow initiation into tightly coupled pipelines. In many domains, especially healthcare, eldercare, governance, and humanitarian operations, such architectures create unacceptable risks, including:

- automated decision creep,
- unauthorized clinical influence,
- privacy invasion through cameras and audio capture,
- opaque model behavior,
- cloud-dependency and cross-border data exposure,
- difficulty of forensic review,
- drift from originally approved operational logic,
- dilution of human sovereignty in safety-critical workflows.

In eldercare and clinical settings, existing systems often depend on:

- video surveillance,
- wearable burden,
- continuous recording,
- opaque AI scoring,
- remote server dependence,
- quasi-autonomous escalation logic.

In humanitarian and field settings, conventional systems often fail because of:

- poor infrastructure,
- lack of trust,
- bandwidth limitations,
- privacy sensitivity,
- hostile or unstable environments,
- inability to validate provenance and operator authority.

A need therefore exists for a unified system that:

1. senses without invasive surveillance,
2. processes locally,
3. offers only governed outputs,

4. structurally prevents autonomous action,
5. preserves human dignity,
6. creates durable accountability,
7. supports safe deployment from clinic to humanitarian edge.

---

5. Objects of the Invention

The invention seeks to provide one or more of the following advantages:

- a **non-agentic AI architecture** that cannot autonomously act;
- a **hardware-enforced human sovereignty model**;
- a **privacy-by-physics sensing stack**;
- a **tamper-resistant accountability pathway**;
- a **mandatory deliberation pause** prior to output release;
- a **physical override brake**;
- a **constitutional drift control mechanism**;
- a **short-retention purge architecture**;
- a **field-deployable, cloud-independent implementation**;
- a **single platform usable in clinical, institutional, governance, and humanitarian contexts**.

---

6. Summary of the Invention

The invention is a **Non-Agentic AI 2.0 Constitutional Governance and Privacy-Preserving Sensing System** comprising:

1. a **sensor subsystem** configured to acquire non-invasive environmental or subject-related data;
2. a **local edge processing subsystem** configured to derive machine-readable features from said data;
3. a **governance core** configured to constrain machine outputs according to non-agentic rules;
4. an **offer-only logic module** configured to prevent direct autonomous execution;

5. a **human authority interface** through which an authorized human may accept, reject, defer, or escalate a governed output;
6. a **Sacred Pause timing subsystem** configured to impose a mandatory delay or deliberative hold before release of machine-generated recommendations;
7. a **Sovereign Brake subsystem** configured to physically and/or logically halt system progression;
8. a **multi-factor authority key architecture**, optionally including simultaneous bodily or physical confirmation channels;
9. a **continuity and accountability ledger** configured to preserve event traceability;
10. a **drift governance subsystem** configured to detect, freeze, audit, and purge unsafe or unapproved drift;
11. a **privacy stack** configured to minimize or eliminate image capture, audio capture, and cloud transmission;
12. a **data lifecycle controller** configured to retain data only for a bounded period and purge according to policy;
13. a **constrained protocol execution layer** configured to permit only pre-authorized bounded action trees.

In preferred embodiments, the system is used for:

- elder fall prevention,
- non-invasive care observation,
- room and perimeter safety,
- humanitarian verification,
- corridor monitoring,
- remote-site field safety,
- governance-grade audit trails.

The system **never diagnoses, prescribes, commands, or autonomously executes**. It may observe, derive, prioritize, and advise.

---

7. Core Inventive Concept

The central inventive concept is the combination of:

A. Privacy-Preserving Observation

The system acquires data through modalities designed to reduce dignity intrusion, such as point-cloud, depth, ranging, spatial occupancy, motion vectors, and optionally constrained thermal signatures.

B. Non-Agentic Constraint Layer

All machine output is transformed into a governed advisory form. No output may directly trigger execution or become an authoritative instruction merely by model generation.

C. Hardware-Enforced Human Sovereignty

Human approval is not merely a software setting. It is structurally enforced via timing gates, braking mechanisms, and authority-binding controls.

D. Drift Prevention

The system includes explicit logic to identify deviation from approved constitutional operation and to lock, audit, and purge compromised pathways.

E. Local, Air-Gapped, or Cloud-Independent Operation

The system is adapted for operation where network trust is limited or external transfer is undesirable.

These combined features create a platform that is not merely “AI with a person watching,” but rather **AI structurally denied agency**.

---

8. Constitutional Operating Principles

In preferred embodiments, the system is built around the following operating principles:

- **Human sovereignty is final**
- **No autonomous actuation**
- **Output is advisory only**
- **Privacy is preserved by architecture, not policy alone**
- **Safety requires enforced pause**
- **Unsafe drift requires lockout**
- **Auditability is continuous**
- **Data retention is minimal**
- **Use must remain bounded to approved protocol space**

These principles may be encoded in firmware, FPGA logic, PLC logic, ROM-locked controls, signed policy layers, or combinations thereof.

---

9. System Architecture

9.1 Sensor Subsystem

The sensor subsystem may include one or more of the following:

- LiDAR sensor
- voxel or point-cloud tracker
- depth sensor
- range sensor
- occupancy grid module
- motion vector module
- constrained thermal sensor
- environmental state sensor
- edge timestamping module

In preferred eldercare embodiments, the primary observation modality avoids conventional facial or identity imaging.

In preferred humanitarian embodiments, the sensor stack supports rugged field operation, low-light operation, low-bandwidth environments, and privacy-sensitive monitoring.

9.2 Edge Compute Subsystem

The compute subsystem is preferably local and may include:

- edge GPU or AI module,
- embedded processor,
- isolated inference module,
- secure local memory,
- no-cloud mode,
- air-gapped or selectively bridged communications.

The compute subsystem derives features such as:

- posture instability,
- fall-risk indicators,
- occupancy anomaly,
- corridor traversal patterns,
- zone-entry events,
- bounded safety alerts,
- operational confidence scores,
- protocol-specific risk markers.

9.3 Governance Core

The governance core is a central architectural element configured to:

- receive AI-derived outputs,
- normalize and classify outputs,
- bind outputs to authority pathways,
- suppress disallowed output types,
- ensure advisory-only transformation,

- enforce bounded protocol release,
- route approved events to ledger.

The governance core may include:

- output classifier,
- non-agentic rule engine,
- authority-binding layer,
- confidence threshold manager,
- state safety evaluator,
- protocol whitelist engine,
- release gate controller.

9.4 Offer-Only Logic Module

The offer-only logic module prevents the system from issuing direct commands. Instead, the system may only produce structured advisory outputs such as:

- observe,
- verify,
- check,
- attend,
- defer,
- review,
- confirm,
- escalate to authorized human.

The module is configured to block output classes such as:

- autonomous diagnosis,
- autonomous treatment selection,
- autonomous caregiver direction,
- autonomous workflow initiation,
- autonomous punitive or enforcement action,
- autonomous credentialing decision,
- autonomous triage closure.

9.5 Authority-Bound Human Control Interface

The human control interface is the only lawful route by which governed machine outputs may be operationalized.

The interface may include:

- clinician console,
- supervisor terminal,
- sovereign control screen,
- rugged field tablet,
- mission operator interface,
- bounded approval panel.

The interface supports actions including:

- accept,
- reject,
- defer,
- acknowledge,
- escalate,
- enter rationale,
- request second check,
- trigger stop.

All accepted actions are attributable to a human authority record.

9.6 Sacred Pause Timing Subsystem

The Sacred Pause subsystem imposes a mandatory time barrier between machine recommendation generation and human-visible release or downstream use.

It may be implemented in:

- FPGA fabric,
- PLC logic,

- secure timing hardware,
- ROM-locked timer logic,
- kernel-level emulation in lower-tier deployments.

Functions include:

- anti-reflex delay,
- deliberation enforcement,
- unsafe rate limiting,
- temporal separation of inference and action,
- prevention of instant machine-led workflow capture.

The Sacred Pause may vary by deployment tier and use case, but in preferred embodiments is not removable by routine software settings.

9.7 Sovereign Brake

The Sovereign Brake is a halt mechanism enabling immediate suspension of system propagation.

It may comprise:

- physical disconnect,
- pedal switch,
- hard relay interrupt,
- control-line cut,
- lockout trigger,
- emergency human veto input.

The brake may halt:

- output release,
- protocol progression,
- downstream transmission,
- actuator-linked systems,
- model presentation pathways.

This ensures that human authority is physically preservable even under software fault or model pressure.

9.8 Tiger .x1 Key / Tripartite Authorization Embodiment

In one embodiment, a high-authority release pathway requires simultaneous or sequential confirmation from three distinct physical modalities, for example:

- eye confirmation,
- hand confirmation,
- foot or leg confirmation.

This multi-modal authorization is intended for:

- high-risk actions,
- sovereign or supervisory release,
- constitutional override,
- export of sensitive governed decisions,
- activation of critical operational states.

The tripartite architecture reduces spoofability and reinforces deliberate human agency.

9.9 Constrained Protocol Execution Layer

Where downstream protocols exist, they are constrained to pre-authorized bounded trees.

The layer allows only:

- approved notification types,
- bounded escalation routes,
- pre-cleared observation tasks,
- safe-state transitions,
- reversible actions,
- explicitly enumerated protocol classes.

The layer blocks:

- open-ended autonomous adaptation,
- self-authored protocol invention,
- unauthorized endpoint transmission,
- latent prompt-driven behavior expansion.

9.10 Protected Communications Pathway

The system may include protected communications through:

- local encrypted bus,
- air-gapped transfer,
- selective sat-bridge,
- rugged tablet relay,
- event-only summary transmission.

In preferred embodiments, raw sensitive data is not broadly transmitted. Instead, only bounded governed summaries or event abstractions are shared.

9.11 Continuity and Accountability Ledger

The ledger records:

- sensor event generation,
- model inference event,
- governance transformations,
- delay events,
- approval or rejection actions,
- brake events,
- drift events,
- purge events,
- protocol state transitions.

The ledger may be tamper-resistant, append-only, hash-chained, or otherwise integrity-protected.

Its purposes include:

- forensic review,
- compliance review,
- regulatory evidence,
- internal governance,
- post-incident audit.

---

10. Privacy Architecture: 3ZEROS Model

10.1 Zero Camera

The system is preferably configured to avoid standard image-based visual surveillance. In embodiments using optical ranging, data is transformed into abstracted geometric or depth-derived representations.

10.2 Zero Audio

The system preferably does not capture conversational content or ambient audio, thereby reducing privacy invasion and consent complexity.

10.3 Zero Cloud

The system is configured for local or edge processing without dependence on external cloud inference. Remote transfer, where permitted, is minimized and governed.

---

11. Data Lifecycle and Purge Architecture

The invention includes a bounded retention controller configured to:

- hold only necessary event data,
- separate raw sensing from abstracted summaries,

- purge raw or intermediate data on schedule,
- preserve only justified accountability records,
- enforce deletion according to institutional policy.

In preferred embodiments, a short-duration purge daemon automatically clears transient sensing data within a defined period, such as a 24-hour retention window, unless a lawful exception is triggered.

This creates a physics-backed approximation of the principle that sensitive presence data should not become indefinite surveillance archives.

---

12. Drift Governance Subsystem

A major component of the invention is a four-stage drift governance chain:

12.1 Detect

Continuously or periodically compare runtime logic, outputs, policy maps, and system behavior against approved baseline constraints.

12.2 Freeze

On detecting unapproved drift, immediately halt governed release or system progression.

12.3 Audit

Perform forensic review of drift source, path, and operational exposure.

12.4 Purge

Remove or reset compromised logic, retrain or re-seed allowed policy, and restore approved constitutional state.

This subsystem is particularly important where:

- models are updated,
- edge systems operate in unstable environments,

- software mutability creates hidden risk,
- compliance must be demonstrable.

---

13. Implementation Tiers

The invention may be deployed in multiple implementation tiers while preserving core invariants.

Tier 1 — Humanitarian

- software-emulated timing,
- simplified human pause,
- advisory-only mode,
- higher drift risk,
- used where resources are constrained.

Tier 2 — Regional

- partial hardware enforcement,
- improved sensing validation,
- medium resistance to mutation,
- bounded institutional deployment.

Tier 3 — National / Academic / Clinical-Grade

- full hardware pause,
- brake subsystem,
- multi-factor authority controls,
- immutable or near-immutable enforcement,
- lowest drift tolerance,
- strongest evidentiary posture.

A key inventive aspect is that the same constitutional framework scales across deployment resource levels without abandoning non-agentic principles.

---

14. Embodiments

14.1 Clinical Eldercare Embodiment

A room or care-space monitoring system detects fall risk, movement instability, occupancy anomaly, or distress pattern using privacy-preserving sensing. The system presents an advisory to a caregiver after passing through governance controls, but cannot autonomously diagnose or enter medical record conclusions.

14.2 Institutional Safety Embodiment

The system monitors sensitive corridors, access points, or dwell zones and provides governed verification outputs. It may support staff safety, access review, or incident documentation without continuous video surveillance.

14.3 Humanitarian Field Embodiment

The system is packaged in rugged form with local compute, portable power, and optional low-bandwidth communications. It may support corridor monitoring, camp safety, maritime route awareness, perimeter observation, or field verification in dignity-sensitive contexts.

14.4 Governance and National-Security Embodiment

The system receives outputs from multiple AI-assisted domains, routes them through the governance core, and prevents autonomous operationalization. Pre-authorized response trees may exist, but only under bounded conditions and human authority control.

---

15. Brief Description of the Drawings

A unified drawing set for this patent design may comprise the following figures:

FIG. 1

Overall block diagram of the Non-Agentic AI 2.0 Constitutional Governance and Privacy-Preserving Sensing System.

FIG. 2

Data flow from sensor acquisition through edge inference, governance core, human authority interface, and ledger.

FIG. 3

Hardware architecture showing sensor module, local compute, Sacred Pause gate, Sovereign Brake, and protected communications path.

FIG. 4

Tripartite authorization flow for high-authority release using multi-modal human confirmation.

FIG. 5

3ZEROS privacy stack showing Zero Camera, Zero Audio, Zero Cloud operating boundaries.

FIG. 6

Clinical embodiment for eldercare observation and governed caregiver advisory pathway.

FIG. 7

Humanitarian field-deployment kit including rugged compute, sensor head, power subsystem, and optional communications bridge.

FIG. 8

Drift governance subsystem illustrating Detect, Freeze, Audit, and Purge.

FIG. 9

Constrained protocol execution layer showing approved bounded action trees and blocked autonomous branches.

FIG. 10

Continuity and accountability ledger architecture.

FIG. 11

Data retention and purge sequence showing transient sensing storage and scheduled deletion.

FIG. 12

Implementation tier model showing Tier 1, Tier 2, and Tier 3 configurations.

FIG. 13

Protected operational communications pathway for local, air-gapped, and selective relay modes.

FIG. 14

State logic diagram including observe, infer, pause, review, approve/reject, log, and purge states.

FIG. 15

Integrated ecosystem diagram showing clinical, institutional, governance, and humanitarian embodiments of the same constitutional architecture.

---

16. Reference Numerals

Below is a consolidated example reference numeral framework:

- **100** — Overall Non-Agentic AI 2.0 system
- **110** — Sensor subsystem
- **111** — LiDAR / spatial ranging module
- **112** — Depth / voxel processor
- **113** — Thermal sensing module
- **120** — Edge compute subsystem
- **130** — Governance core
- **131** — Output normalization engine

- **132** — Authority-binding engine
- **133** — Offer-only logic module
- **134** — Protocol whitelist engine
- **140** — Human authority interface
- **150** — Sacred Pause subsystem
- **160** — Sovereign Brake
- **170** — Tripartite authorization subsystem
- **171** — Eye confirmation input
- **172** — Hand confirmation input
- **173** — Foot/leg confirmation input
- **180** — Protected communications path
- **190** — Continuity and accountability ledger
- **200** — Drift detection module
- **210** — Freeze module
- **220** — Audit module
- **230** — Purge controller
- **240** — Data retention controller
- **250** — Privacy stack controller
- **260** — Constrained protocol execution layer
- **270** — Local power subsystem
- **280** — Rugged operator terminal
- **290** — Humanitarian deployment pack

---

17. Detailed Operation

17.1 Normal Operation

1. The sensor subsystem acquires environmental or subject-related data.
2. The edge compute subsystem derives abstracted features or risk indicators.
3. The governance core classifies the output.
4. The offer-only module converts output into advisory form.
5. The Sacred Pause imposes a mandatory deliberative delay.
6. The human authority interface presents the governed advisory.
7. The authorized human accepts, rejects, defers, or escalates.
8. The result is logged in the ledger.
9. Raw or transient data is purged according to policy.

17.2 High-Risk Operation

Where an output falls into a high-consequence category:

1. release is held,
2. stronger review rules apply,
3. brake-ready state is maintained,
4. optionally, tripartite confirmation is required.

17.3 Drift Event Operation

1. drift is detected,
2. system freezes release,
3. audit begins,
4. compromised logic is isolated,
5. purge resets or removes unsafe state,
6. service resumes only upon validated restoration.

---

18. Advantages of the Invention

The invention provides the following technical and governance benefits:

- reduces surveillance intrusiveness;
- keeps sensitive processing local;
- removes autonomous action pathways;
- creates provable human finality;
- improves auditability;
- strengthens compliance readiness;
- resists hidden system drift;
- supports clinical dignity and humanitarian trust;
- adapts across multiple operating environments.

---

19. Exemplary Claims

Claim 1 — Independent Apparatus Claim

A non-agentic artificial intelligence system comprising:

- a sensor subsystem configured to acquire non-invasive environment or subject-related data;
- an edge compute subsystem configured to generate one or more machine-derived outputs from the data;
- a governance core configured to transform said machine-derived outputs into governed advisory outputs according to one or more non-agentic constraints;
- an authority-bound human control interface configured to receive human acceptance, rejection, deferral, or escalation of said governed advisory outputs;
- a timing subsystem configured to impose a mandatory delay before release of said governed advisory outputs;
- a brake subsystem configured to halt propagation of said governed advisory outputs; and
- an accountability ledger configured to record output generation, human interaction, and system state changes,

wherein the system is configured such that no machine-derived output autonomously executes an external action without explicit human authorization.

Claim 2

The system of claim 1, wherein the sensor subsystem includes a LiDAR or point-cloud sensing module configured to generate abstracted spatial representations without conventional image capture.

Claim 3

The system of claim 1, wherein the governance core includes an offer-only logic module configured to suppress command-form outputs and permit only advisory-form outputs.

Claim 4

The system of claim 1, wherein the timing subsystem comprises hardware-enforced delay logic implemented in programmable logic, fixed logic, or immutable timing circuitry.

Claim 5

The system of claim 1, wherein the brake subsystem comprises a physical interrupt configured to disable output propagation independently of software state.

Claim 6

The system of claim 1, wherein the authority-bound human control interface requires multi-modal human confirmation prior to release of a high-risk output.

Claim 7

The system of claim 6, wherein the multi-modal confirmation includes confirmation from at least three distinct bodily or physical input channels.

Claim 8

The system of claim 1, further comprising a drift governance subsystem configured to detect deviation from an approved operational baseline, freeze governed output release, perform forensic audit, and purge compromised logic.

Claim 9

The system of claim 1, further comprising a privacy controller configured to enforce one or more of: absence of camera capture, absence of audio capture, and absence of cloud processing.

Claim 10

The system of claim 1, further comprising a data lifecycle controller configured to delete transient sensing data after a bounded retention period.

Claim 11

The system of claim 1, wherein the system is configured for eldercare monitoring and provides fall-risk or instability advisories without autonomous diagnosis.

Claim 12

The system of claim 1, wherein the system is configured for humanitarian field deployment with ruggedized power, local compute, and selective low-bandwidth communications.

Claim 13

The system of claim 1, wherein the governance core is configured to permit only pre-authorized bounded protocol trees.

Claim 14

The system of claim 1, wherein the accountability ledger is append-only or integrity-protected.

Claim 15

The system of claim 1, wherein the edge compute subsystem is air-gapped from external cloud services during normal operation.

Claim 16 — Independent Method Claim

A method of operating a non-agentic artificial intelligence system, comprising:

- acquiring non-invasive sensor data;
- generating a machine-derived inference from the sensor data;
- transforming the machine-derived inference into a governed advisory output using a governance core;
- imposing a mandatory pause before release of the governed advisory output;
- presenting the governed advisory output to an authorized human operator;
- receiving a human decision concerning said governed advisory output; and
- recording the event in an accountability ledger,

wherein no external action is autonomously executed based solely on the machine-derived inference.

Claim 17

The method of claim 16, further comprising detecting runtime drift relative to a baseline governance configuration and freezing output release upon detecting the drift.

Claim 18

The method of claim 16, further comprising purging transient data after a predefined retention period.

Claim 19

The method of claim 16, wherein the acquired sensor data comprises point-cloud data or spatial occupancy data.

Claim 20

The method of claim 16, wherein the governed advisory output is limited to one or more of: observe, verify, review, attend, or escalate.

---

20. Abstract

A non-agentic artificial intelligence system is disclosed comprising a privacy-preserving sensor subsystem, a local edge processing subsystem, a governance core, an offer-only logic module, an authority-bound human control interface, a mandatory timing pause subsystem, a brake subsystem, and an accountability ledger. The system is configured to derive machine outputs from sensed data while preventing autonomous execution. Outputs are transformed into governed advisory forms and released only through human authorization. Preferred embodiments employ privacy-preserving sensing, local processing without cloud dependence, bounded data retention, and drift governance comprising detect, freeze, audit, and purge functions. The system is applicable to clinical monitoring, eldercare, institutional safety, governance operations, and humanitarian field deployment.

Patent Caption Style for the Drawings

- **FIG. 1** is a system block diagram of an embodiment of the Non-Agentic AI 2.0 system.
- **FIG. 2** is a perspective view of an external sensing device embodiment.
- **FIG. 3** is an internal hardware architecture diagram of the system.
- **FIG. 4** is a data processing and governance flow diagram.
- **FIG. 5** is a diagram of offer-only logic and constrained protocol execution.
- **FIG. 6** is a control diagram of the Sacred Pause and Sovereign Brake architecture.
- **FIG. 7** is a diagram of a tripartite authorization embodiment.
- **FIG. 8** is a diagram of the 3ZEROS privacy architecture.
- **FIG. 9** is a diagram of the drift governance subsystem.
- **FIG. 10** is a diagram of the data lifecycle and purge architecture.
- **FIG. 11** is an environmental view of a clinical or eldercare deployment embodiment.
- **FIG. 12** is an environmental view of a humanitarian field deployment embodiment.
- **FIG. 13** is a protected communications pathway diagram.
- **FIG. 14** is an operational state diagram of the system.
- **FIG. 15** is an integrated ecosystem diagram showing multiple deployment contexts.

DRAWINGS IN FIGURES

NAI 2.0 Patent Design

Non-Agentic AI 2.0 Constitutional Governance and Privacy-Preserving Sensing System

Brief Description of the Drawings

The accompanying figures illustrate exemplary embodiments of the invention and are intended to support the description of the system architecture, privacy-preserving sensing arrangement, governance logic, safety controls, and deployment configurations.

Reference Numerals for the Figures

Use the following numerals consistently across all drawings:

- **100** — Overall NAI 2.0 system
- **110** — Sensor subsystem
- **111** — LiDAR or spatial sensing module
- **112** — Depth or point-cloud processor
- **113** — Thermal sensing module
- **114** — Environmental sensor module
- **120** — Edge compute subsystem
- **121** — Embedded processor
- **122** — AI inference engine
- **123** — Secure memory
- **124** — Local storage
- **125** — Power management controller
- **130** — Governance core
- **131** — Output normalization engine
- **132** — Authority-binding engine
- **133** — Offer-only logic module
- **134** — Protocol whitelist engine
- **135** — Safety rule engine
- **140** — Human authority interface
- **141** — Review console
- **142** — Accept input

- **143** — Reject input
- **144** — Defer or escalate input
- **150** — Sacred Pause subsystem
- **151** — Timing controller
- **152** — Delay gate
- **160** — Sovereign Brake
- **161** — Manual interrupt switch
- **162** — Output halt relay
- **170** — Tripartite authorization subsystem
- **171** — Eye confirmation input
- **172** — Hand confirmation input
- **173** — Foot or leg confirmation input
- **180** — Protected communications pathway
- **181** — Local-only path
- **182** — Air-gapped transfer path
- **183** — Selective relay path
- **190** — Accountability ledger
- **200** — Drift governance subsystem
- **210** — Drift detect module
- **220** — Freeze module
- **230** — Audit module
- **240** — Purge controller
- **250** — Data lifecycle controller
- **260** — Privacy stack controller
- **270** — Rugged power subsystem
- **280** — Operator tablet or field terminal
- **290** — Humanitarian deployment kit

FIG. 1 — Overall System Block Diagram

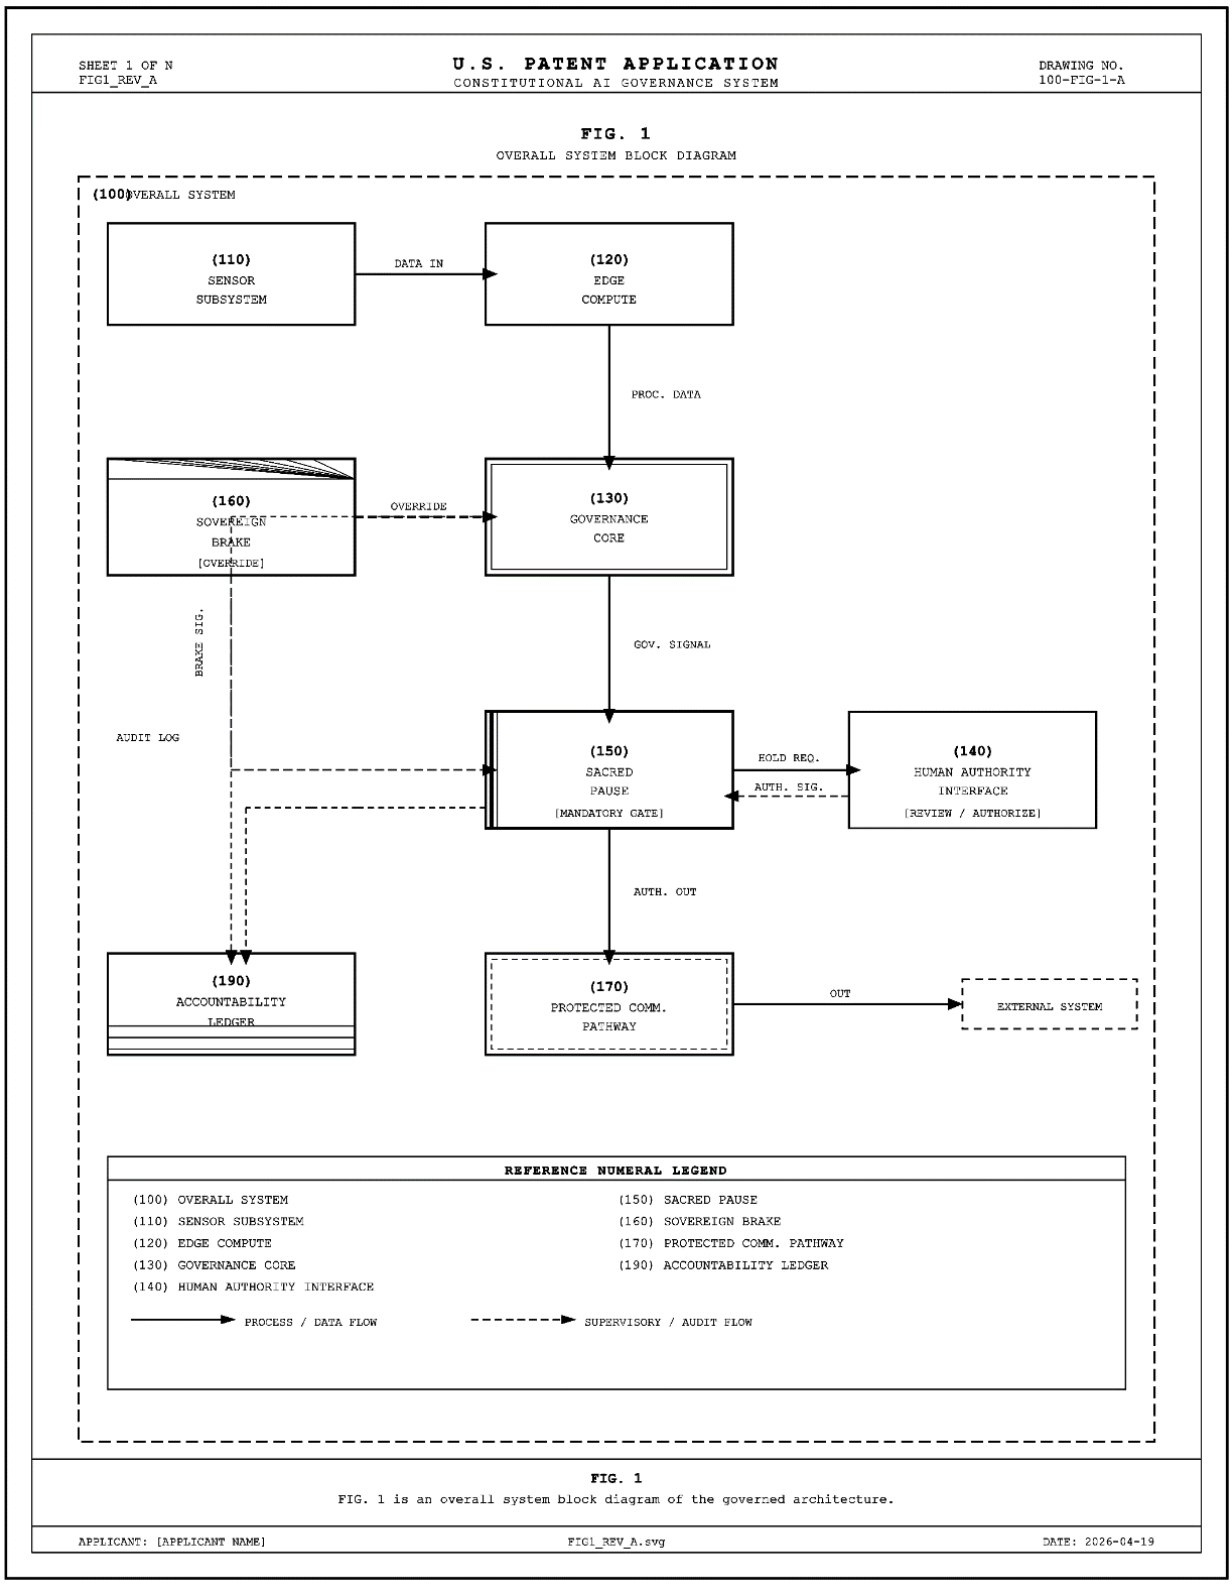

Overall system block diagram of the Non-Agentic AI 2.0 system, showing the relationship between the sensor subsystem, edge compute subsystem, governance core, human authority interface, Sacred Pause subsystem, Sovereign Brake, communications pathway, and accountability ledger.

Show a top-level block layout with arrows indicating signal and control flow:

- **110 Sensor Subsystem**
- **120 Edge Compute Subsystem**
- **130 Governance Core**
- **150 Sacred Pause Subsystem**
- **140 Human Authority Interface**
- **160 Sovereign Brake**
- **190 Accountability Ledger**
- **180 Protected Communications Pathway**

Suggested flow:

110 -> 120 -> 130 -> 150 -> 140 -> 190

With **160** shown as an interrupt line cutting across the release path.

FIG. 2 — External Device Embodiment

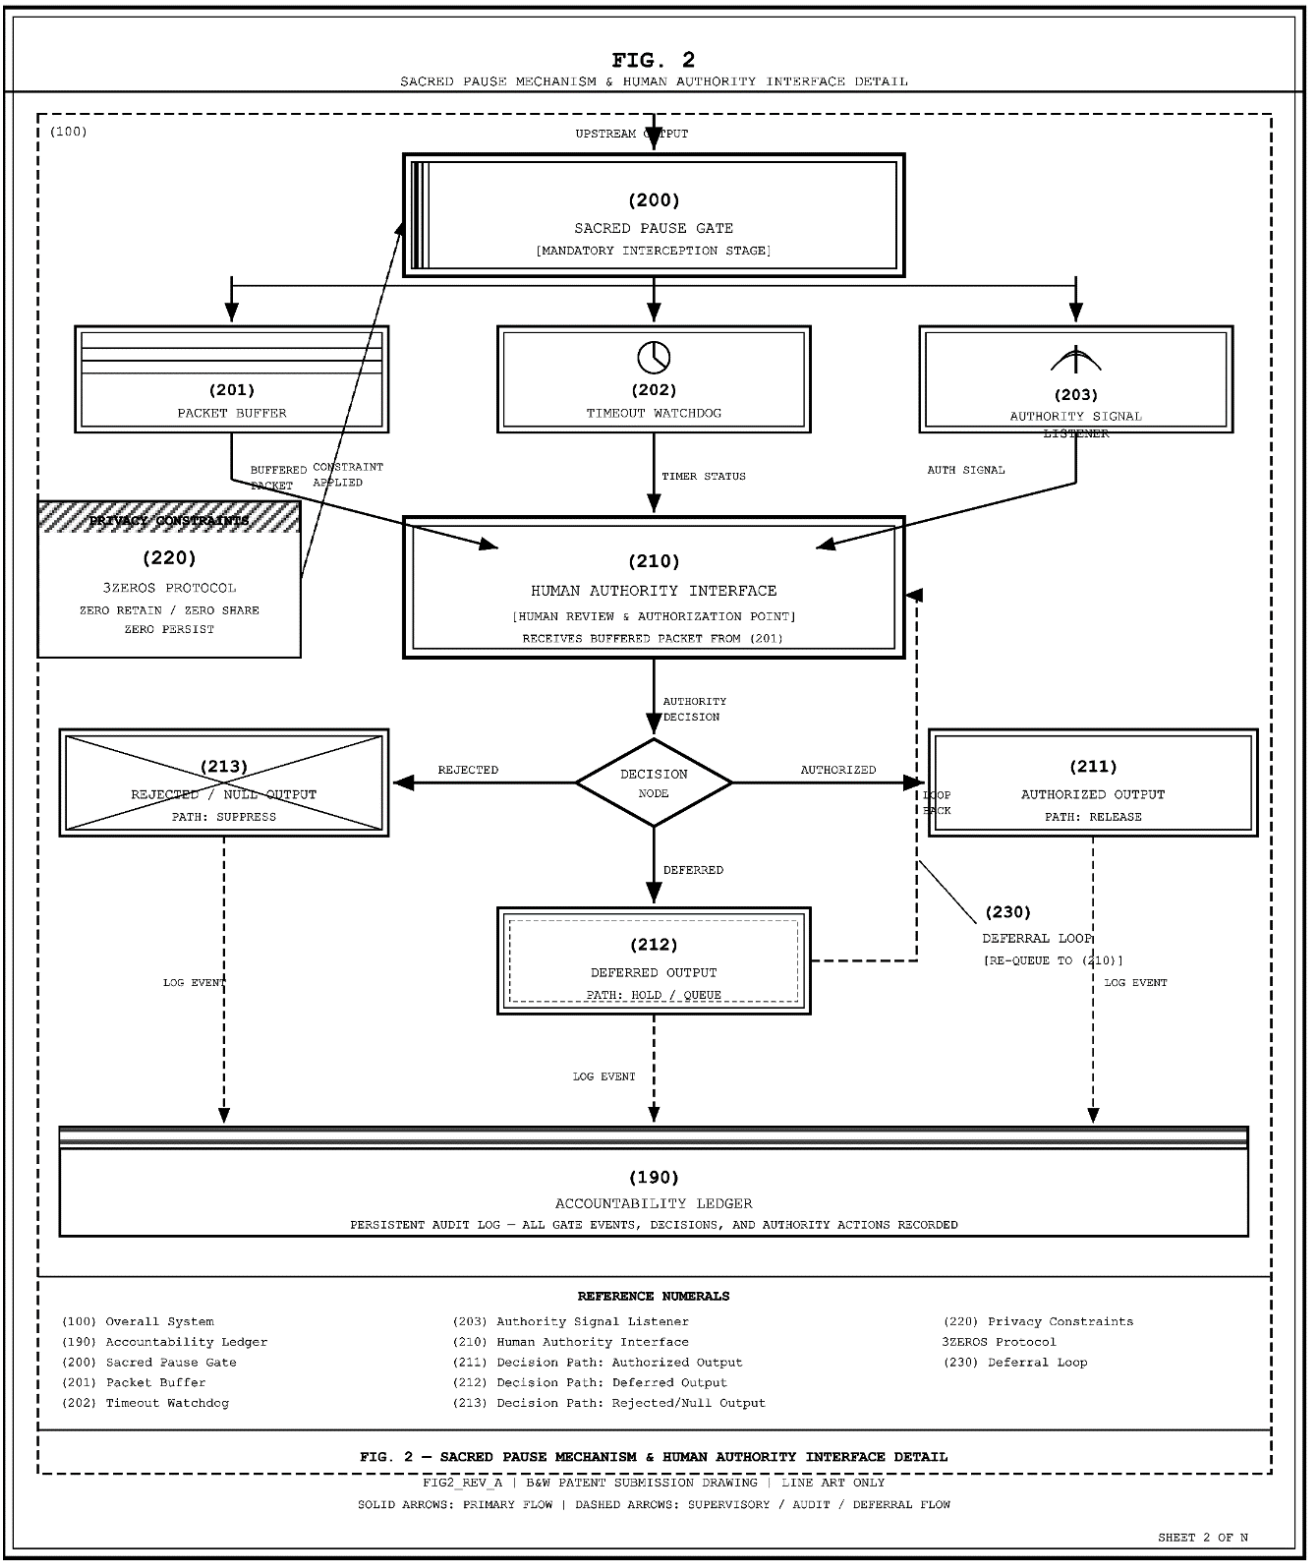

External device embodiment showing a privacy-preserving sensing unit mounted in a clinical, institutional, or field environment, including the sensor head, protective enclosure, support structure, operator interface, and emergency brake element.

Illustrate the outward appearance of the system:

- upright sensor housing
- front sensing window or aperture
- non-camera sensing face
- support arm, stand, wall mount, or base
- local interface screen or panel
- emergency brake location
- cable routing or protected housing

Key labels: 110, 111, 140, 160, 270.

FIG. 3 — Internal Hardware Architecture

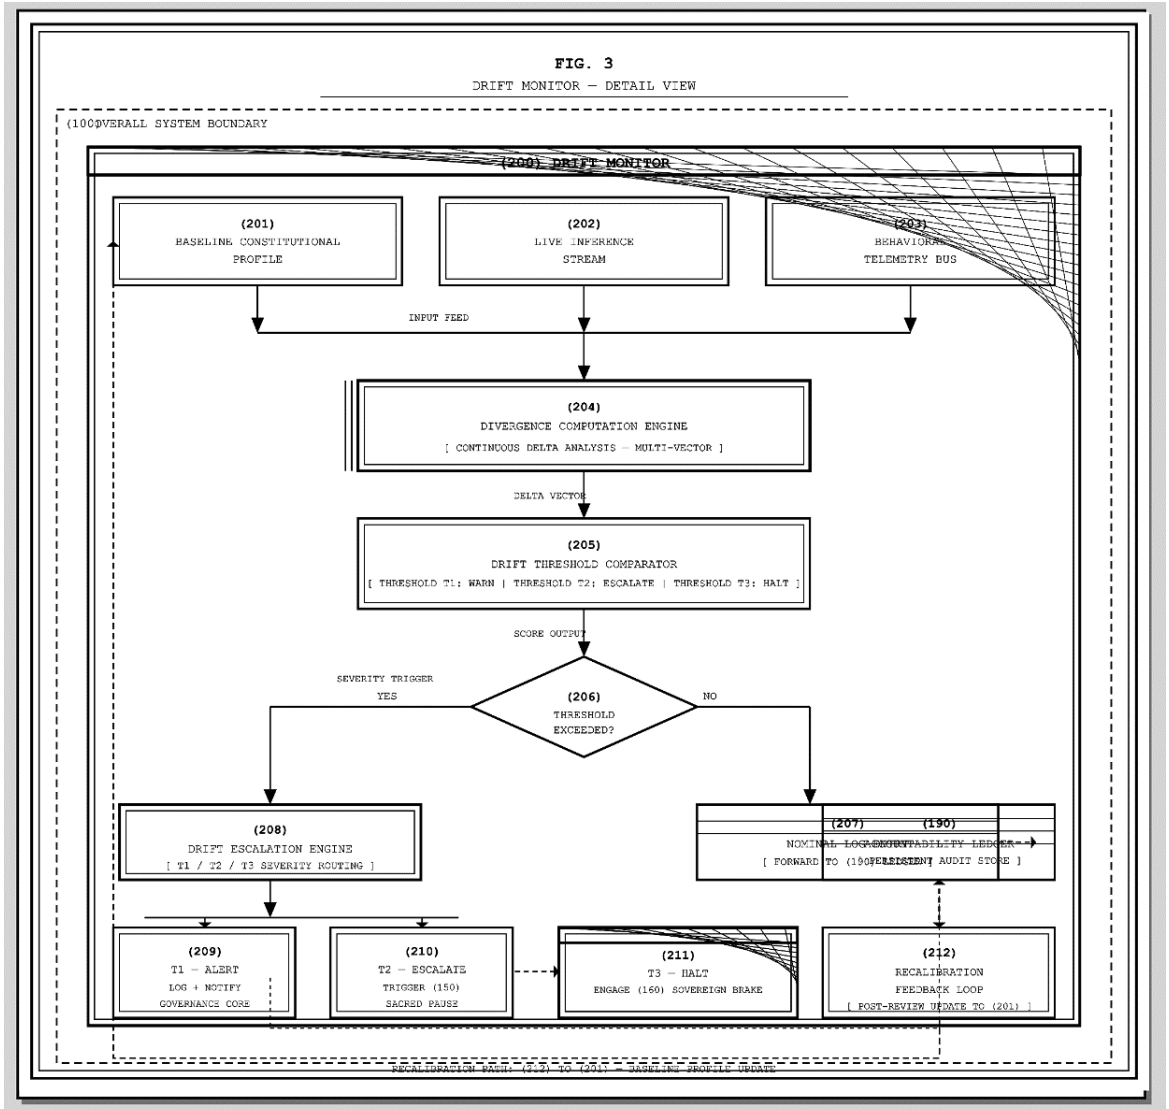

FIG. 3 — DRIFT MONITOR DETAIL VIEW
ALL REFERENCE NUMERALS AS LABELED

FIG. 3 is a detail view of the Drift Monitor (200), showing the Baseline Constitutional Profile (201), Live Inference Stream (202), Behavioral Telemetry Bus (203), Divergence Computation Engine (204), Drift Threshold Comparator (205), threshold decision logic (206), Nominal Log (207), Drift Escalation Engine (208), and severity output paths T1 (209), T2 (210), T3 (211), Recalibration Feedback Loop (212), and Accountability Ledger (190).

SHEET 3 OF 5
B&W PATENT STYLE
REV. A

REFERENCE NUMERAL LEGEND — FIG. 3

(100) — Overall System Boundary
(200) — Drift Monitor [Master Block]
(201) — Baseline Constitutional Profile
(202) — Live Inference Stream
(203) — Behavioral Telemetry Bus
(204) — Divergence Computation Engine
(205) — Drift Threshold Comparator
(206) — Threshold Decision: Exceeded?
(207) — Nominal Log Entry
(208) — Drift Escalation Engine
(209) — T1 Alert: Log + Notify Governance Core
(210) — T2 Escalate: Trigger Sacred Pause (150)
(211) — T3 Halt: Engage Sovereign Brake (160)
(212) — Recalibration Feedback Loop to (201)
(150) — Sacred Pause [cross-ref]
(160) — Sovereign Brake [cross-ref]
(190) — Accountability Ledger [cross-ref]

Internal hardware architecture showing the sensor modules, embedded compute engine, secure memory, power subsystem, timing controller, governance processor, and protected output pathways.

Use a sectional or exploded-style functional diagram:

- sensor inputs entering preprocessing block
- embedded processor
- inference engine
- secure memory
- timing controller
- governance core
- output gate
- relay/brake circuit
- protected communications board

Key labels: 111, 112, 113, 121, 122, 123, 130, 151, 152, 162, 180.

FIG. 4 — Data Processing and Governance Flow

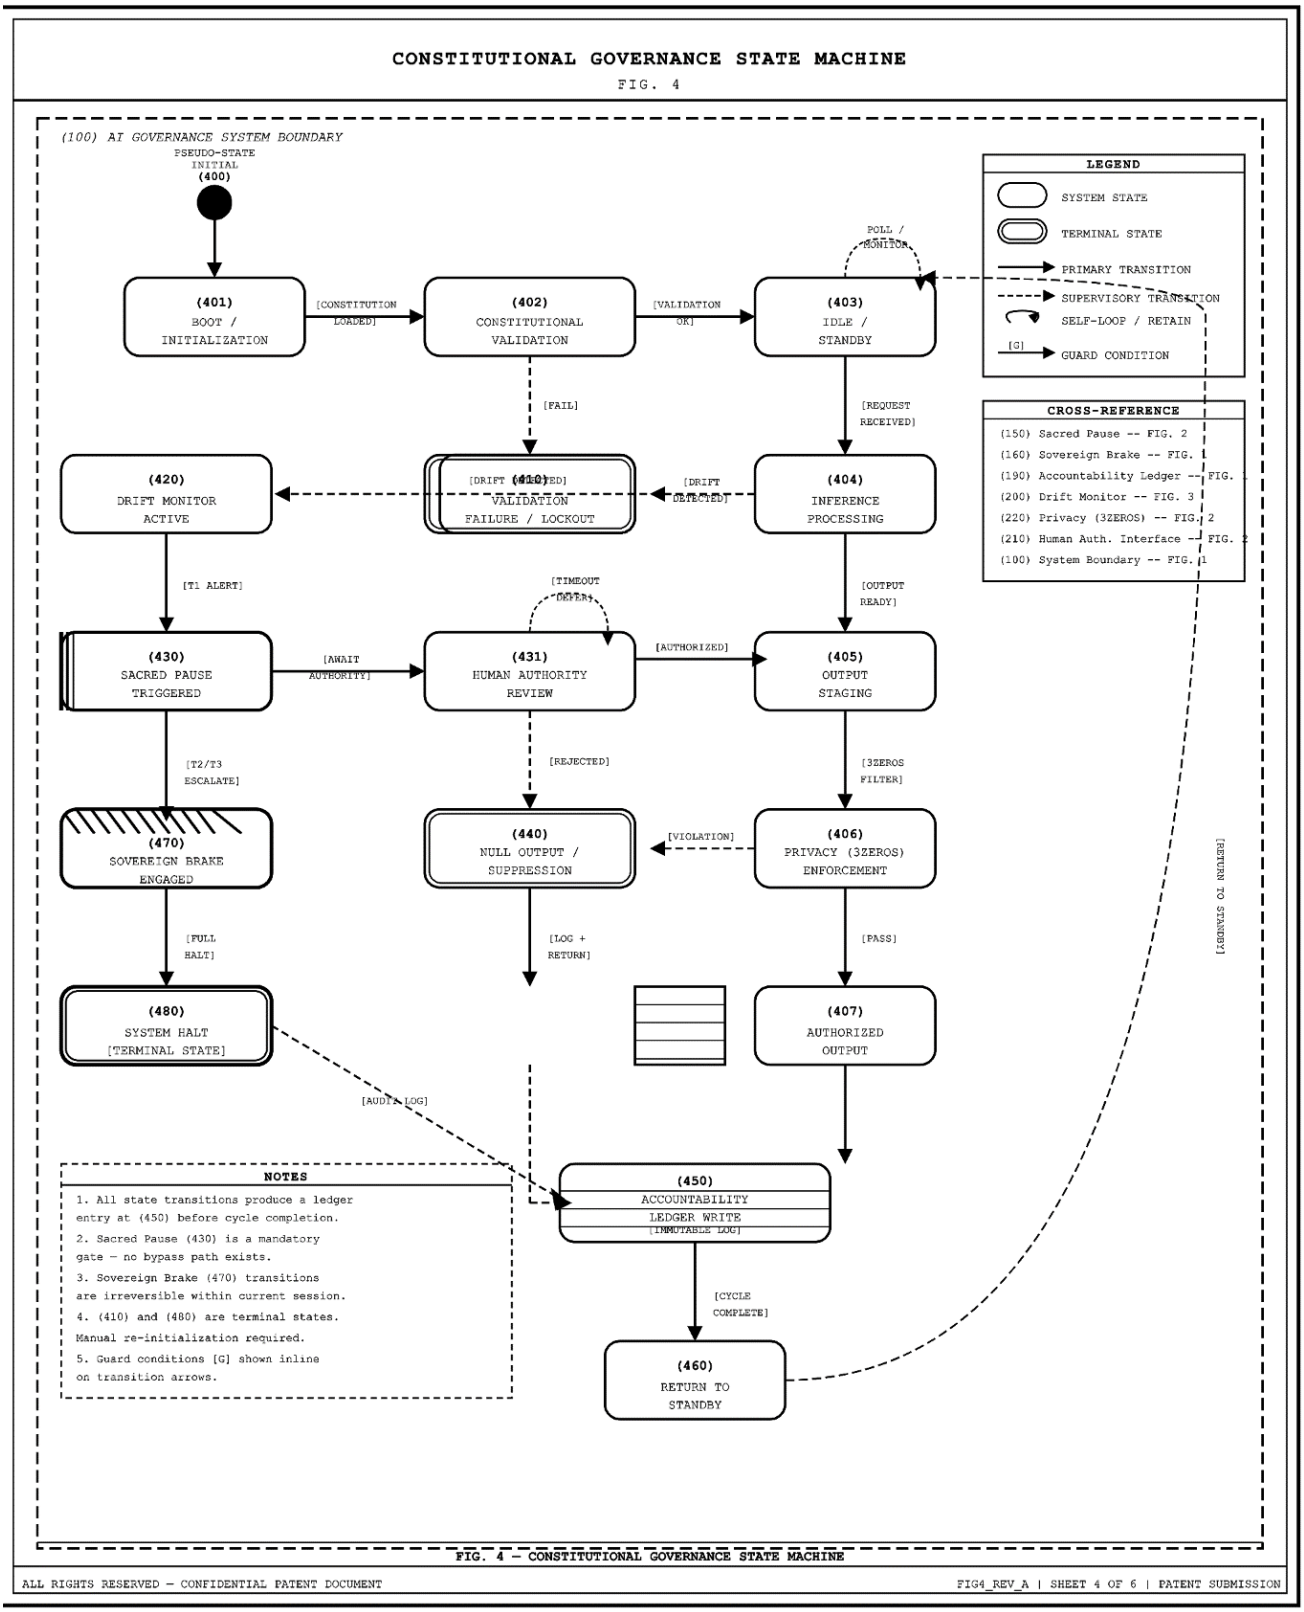

FIG. 4 — CONSTITUTIONAL GOVERNANCE STATE MACHINE

Data processing and governance flow diagram showing sensor acquisition, preprocessing, inference generation, governance transformation, pause gate, human review, decision outcome, and ledger logging.

Draw as a left-to-right or top-down process flow:

1. Sense
2. Preprocess
3. Infer
4. Govern
5. Pause
6. Human review
7. Decision
8. Log
9. Purge

Include a blocked branch labeled:

- **autonomous execution prohibited**

Key labels: 110, 120, 130, 150, 140, 190, 250.

FIG. 5 — Offer-Only Logic and Bounded Protocol Tree

NEURAL AI CONSTITUTIONAL GOVERNANCE SYSTEM

FIG. 5 — PRIVACY ENFORCEMENT ENGINE (3ZEROS) — DETAIL VIEW

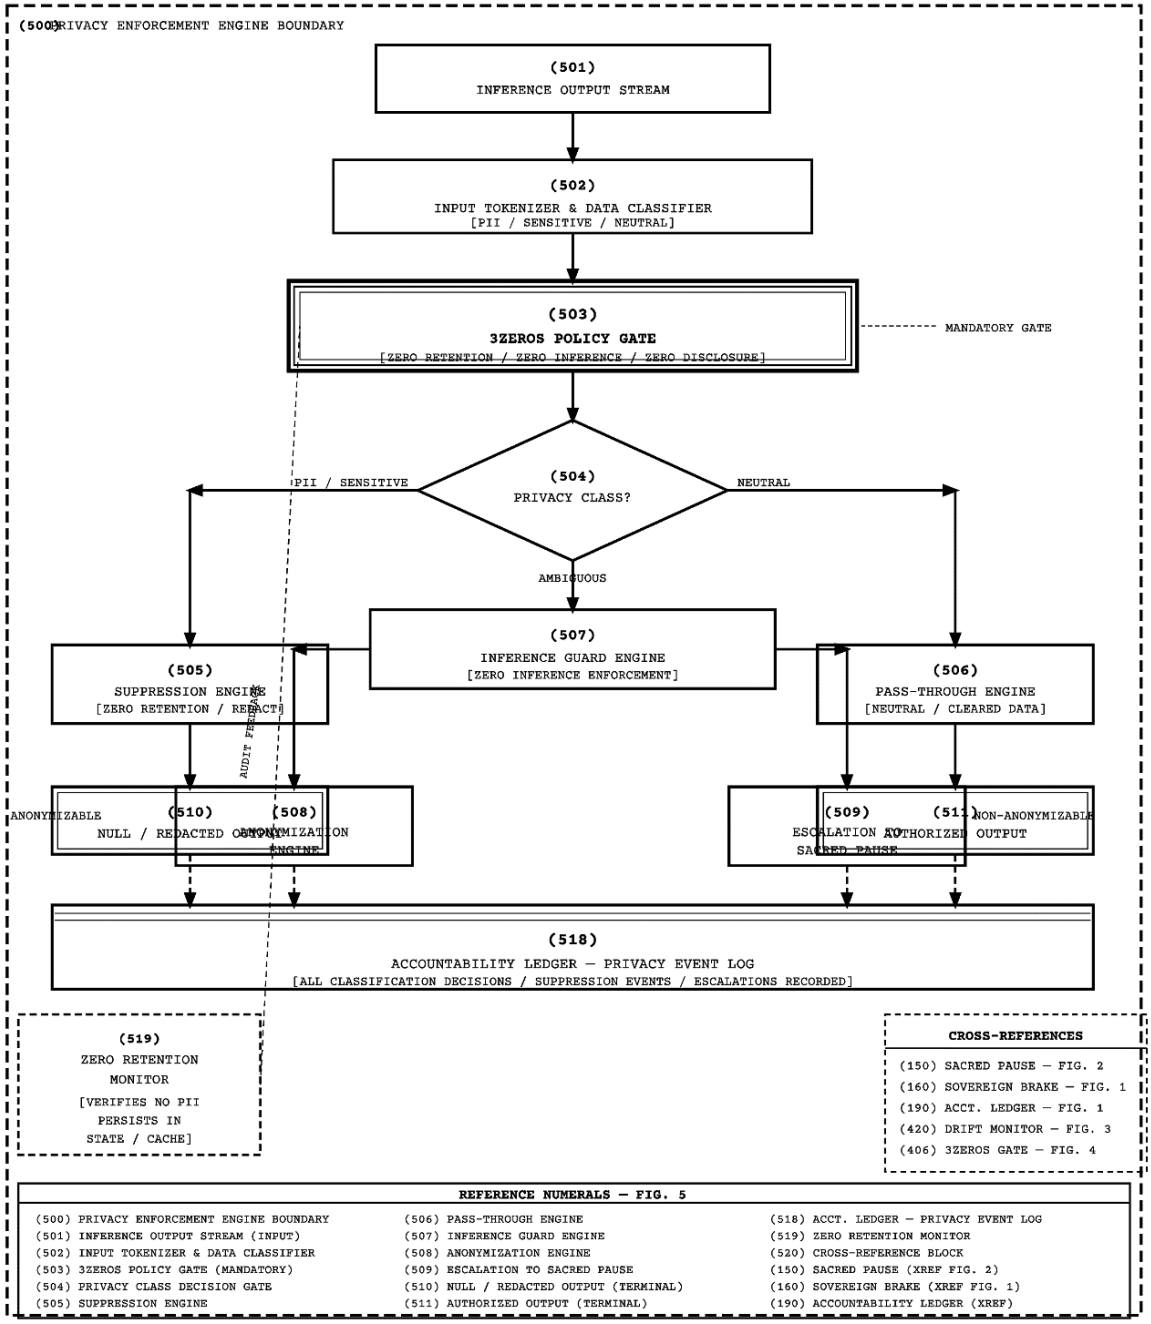

REFERENCE NUMERALS — FIG. 5

(500) PRIVACY ENFORCEMENT ENGINE BOUNDARY	(506) PASS-THROUGH ENGINE	(518) ACCT. LEDGER — PRIVACY EVENT LOG
(501) INFERENCE OUTPUT STREAM (INPUT)	(507) INFERENCE GUARD ENGINE	(519) ZERO RETENTION MONITOR
(502) INPUT TOKENIZER & DATA CLASSIFIER	(508) ANONYMIZATION ENGINE	(520) CROSS-REFERENCE BLOCK
(503) 3ZEROS POLICY GATE (MANDATORY)	(509) ESCALATION TO SACRED PAUSE	(150) SACRED PAUSE (XREF FIG. 2)
(504) PRIVACY CLASS DECISION GATE	(510) NULL / REDACTED OUTPUT (TERMINAL)	(160) SOVEREIGN BRAKE (XREF FIG. 1)
(505) SUPPRESSION ENGINE	(511) AUTHORIZED OUTPUT (TERMINAL)	(190) ACCOUNTABILITY LEDGER (XREF)

Offer-only logic and constrained protocol execution diagram showing how machine-generated outputs are transformed into advisory outputs and how non-permitted autonomous branches are blocked.

Show machine outputs entering a classification box, then splitting into:

- **permitted advisory outputs**
  - observe
  - verify
  - review
  - attend
  - escalate
- **blocked outputs**
  - diagnose
  - prescribe
  - command
  - autonomously execute

Then connect permitted outputs to a limited protocol tree only.

Key labels: 133, 134, 135.

FIG. 6 — Sacred Pause and Sovereign Brake

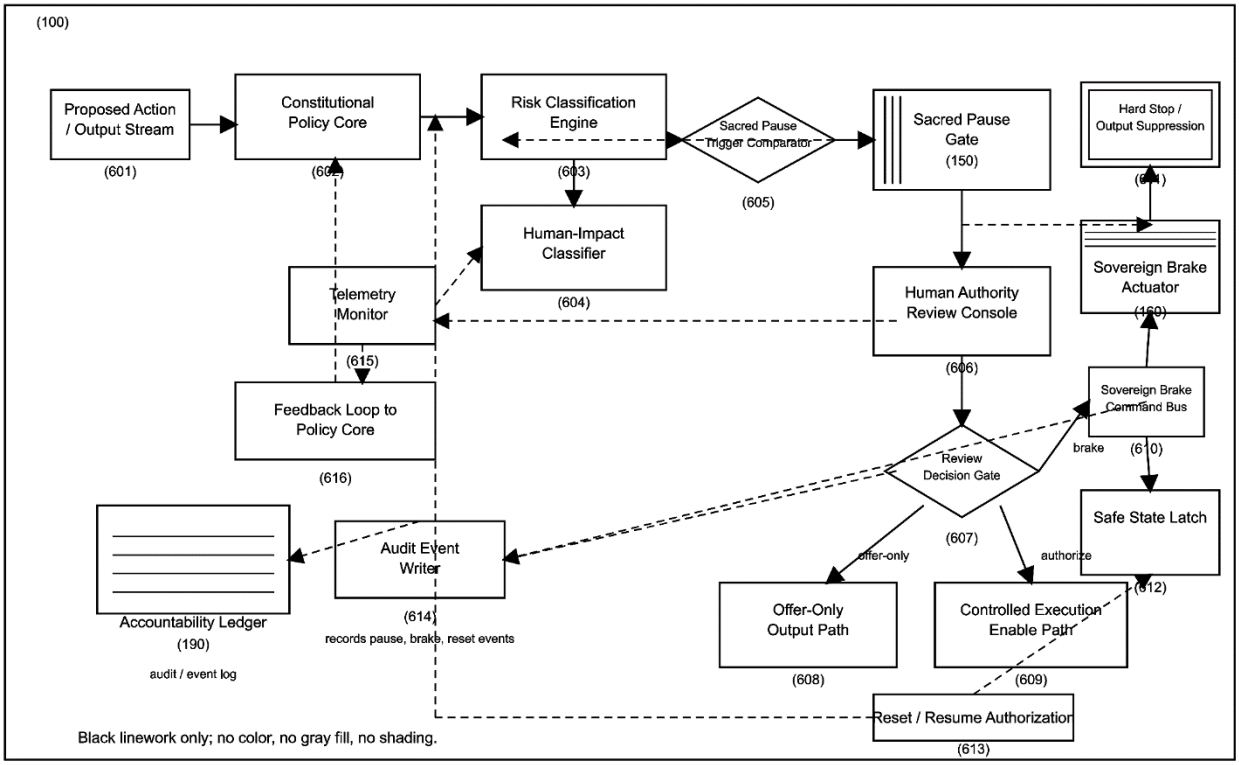

Sacred Pause and Sovereign Brake control architecture showing mandatory timing delay, release gate control, manual brake override, and interrupt logic.

Show:

- output candidate generated
- mandatory hold in timing gate
- human review step
- emergency brake interrupt line
- release only after human authorization

Add a visible lock symbol or gate symbol between governance and review.

Key labels: 150, 151, 152, 160, 161, 162.

FIG. 7 — Tripartite Authorizaton Embodiment

NON-AGENTIC AI 2.0 SYSTEM **FIG7_REV_A.svg**

Patent Application No.: [PENDING] Sheet 7 of 15

Inventor: [APPLICANT NAME] Rev: A

Filing Date: [DATE] Scale: NTS

FIG. 7 — TRIPARTITE AUTHORIZATION EMBODIMENT

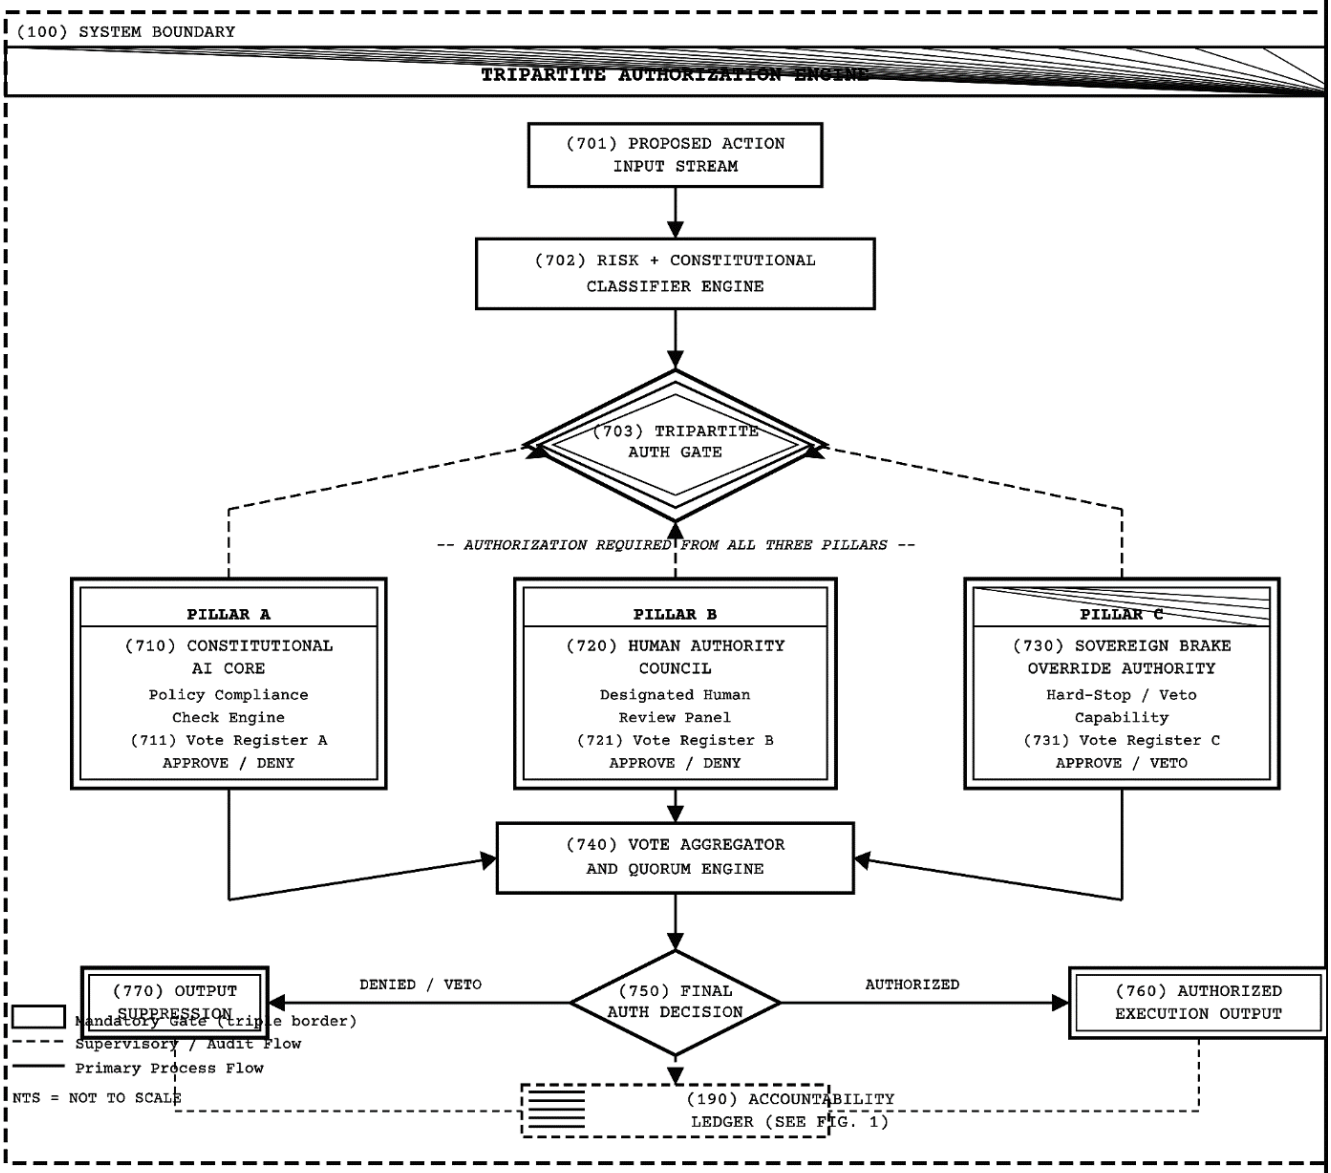

Tripartite authorization embodiment showing a high-authority release sequence requiring multiple human confirmation channels, including eye confirmation, hand confirmation, and foot or leg confirmation.

Draw a central release controller receiving three confirmation inputs:

- **171** eye confirmation
- **172** hand confirmation
- **173** foot or leg confirmation

Only when all required inputs are satisfied does the release path open.

Key label: 170 as the composite authorization controller.

FIG. 8 — 3ZEROS Privacy Architecture

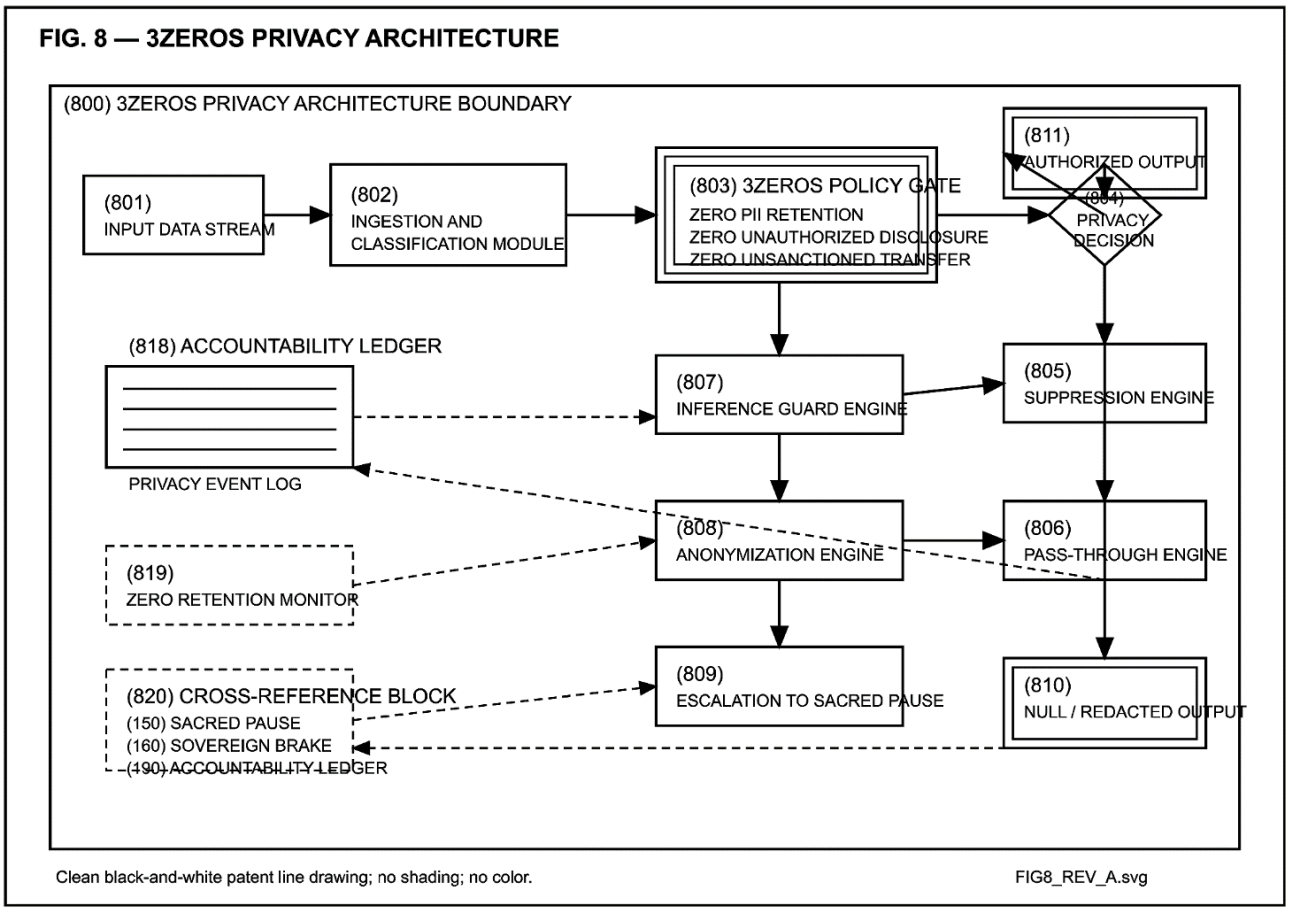

3ZEROS privacy architecture showing Zero Camera, Zero Audio, and Zero Cloud operating boundaries, local processing, and bounded output transmission.

Draw three perimeter blocks or shields around the core system:

- **Zero Camera**
- **Zero Audio**
- **Zero Cloud**

Inside the protected area, show:

- local sensing
- edge processing
- governed advisory output
- local ledger

Optional outbound path should be labeled:

- **event summary only**
- **bounded transmission**

Key labels: 260, 181, 183.

FIG. 9 — Drift Governance Subsystem

FIG. 9 — DRIFT GOVERNANCE SUBSYSTEM

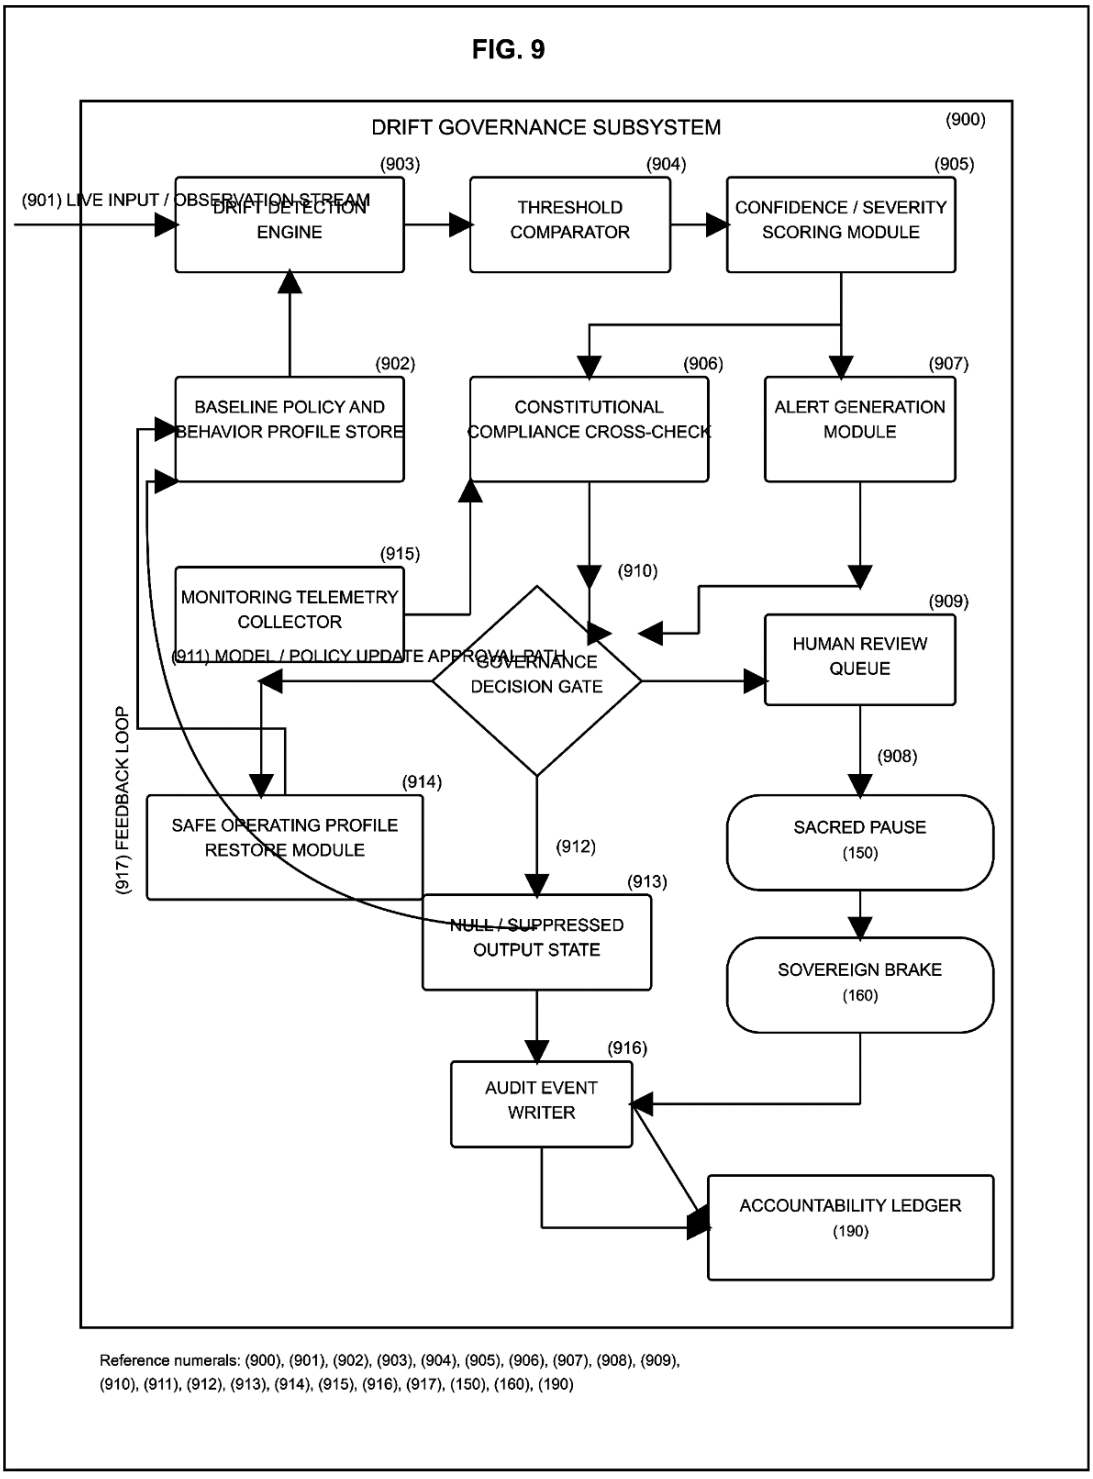

Reference numerals: (900), (901), (902), (903), (904), (905), (906), (907), (908), (909), (910), (911), (912), (913), (914), (915), (916), (917), (150), (160), (190)

Drift governance subsystem showing Detect, Freeze, Audit, and Purge operations and restoration to approved constitutional state.

Use a cyclical or linear safety chain:

- Detect
- Freeze
- Audit
- Purge
- Restore approved state

Show the drift event cutting off normal output release.

Key labels: 200, 210, 220, 230, 240.

FIG. 10 — Data Lifecycle and Purge

FIG. 10
Constitutional Policy Core — Detail View

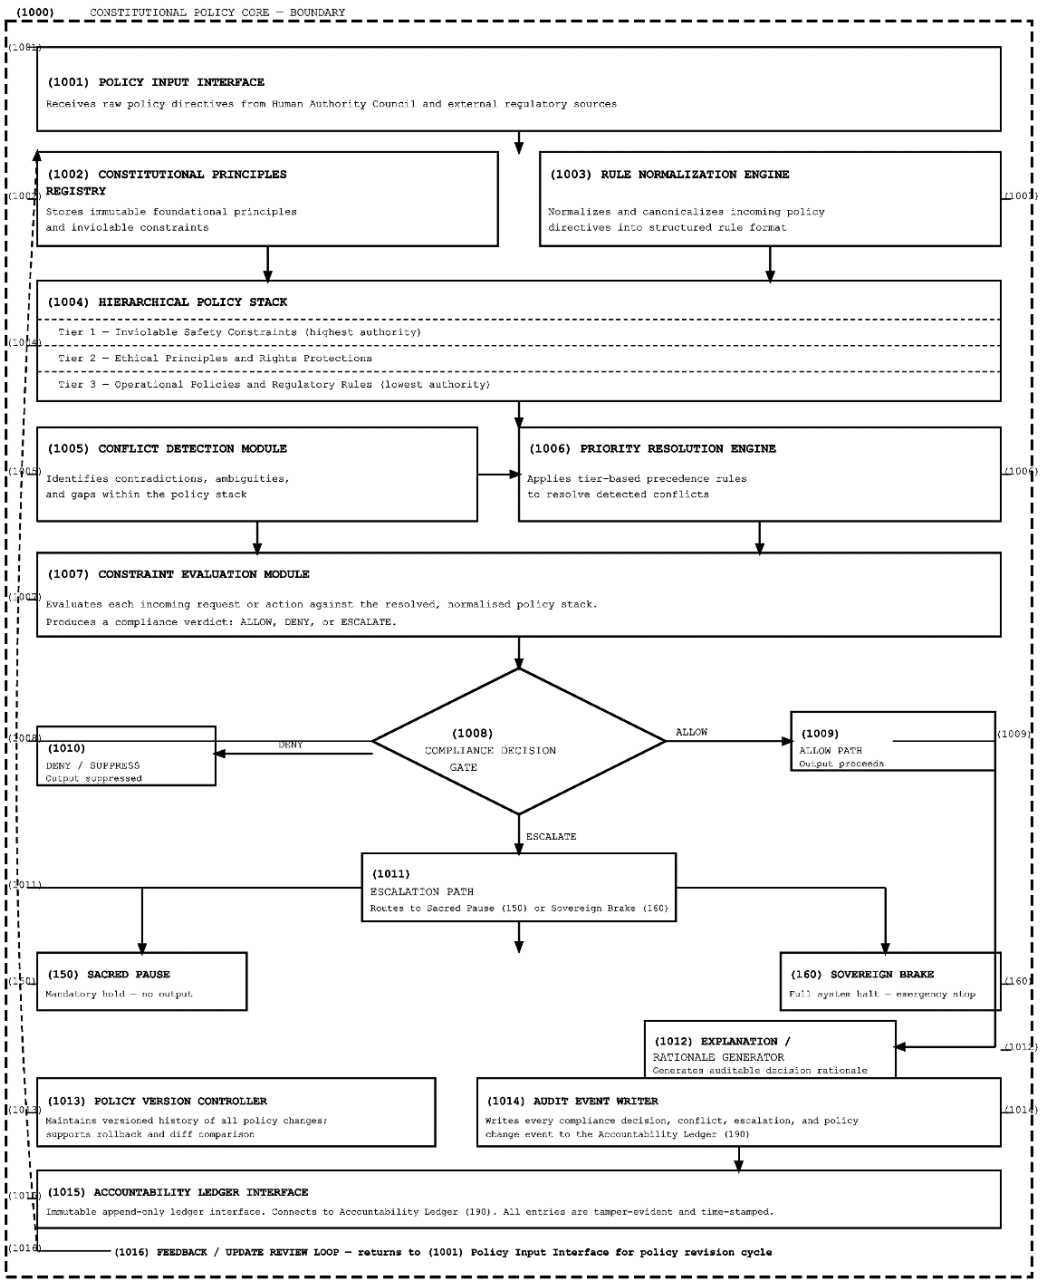

Data lifecycle and purge architecture showing transient sensor data storage, event abstraction, retention window control, lawful exception handling, and scheduled purge.

Draw storage layers:

- transient raw sensing buffer
- feature abstraction layer
- governed event record
- ledger archive
- purge timer

Use arrows from raw data to deletion after a bounded period.

Key labels: 123, 124, 190, 240, 250.

FIG. 11 — Clinical/Eldercare Deployment

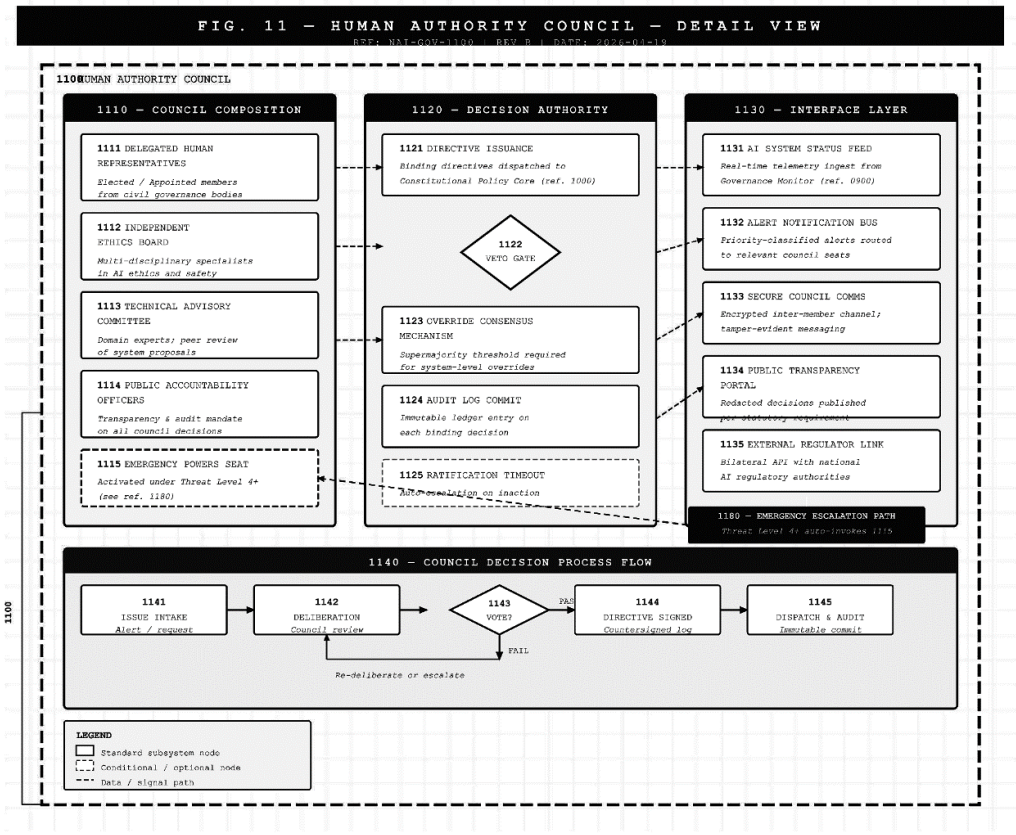

FIG. 11
HUMAN AUTHORITY COUNCIL — DETAIL VIEW | FIG11_REV_B.SVG | NAI GOVERNANCE ARCHITECTURE

REFERENCE NUMERAL INDEX — FIG. 11

REF. NO.	ELEMENT LABEL	DESCRIPTION / NOTES
1100	HUMAN AUTHORITY COUNCIL	*Top-level governance body boundary; outer enclosure*
1110	COUNCIL COMPOSITION PANEL	*Left column; member seat classifications*
1111	DELEGATED HUMAN REPRESENTATIVES	*Elected or appointed from civil governance bodies*
1112	INDEPENDENT ETHICS BOARD	*Multi-disciplinary AI ethics and safety specialists*
1113	TECHNICAL ADVISORY COMMITTEE	*Domain experts; peer review of system proposals*
1114	PUBLIC ACCOUNTABILITY OFFICERS	*Transparency and audit mandate on all decisions*
1115	EMERGENCY POWERS SEAT	*Conditional seat; activated at Threat Level 4+*
1120	DECISION AUTHORITY PANEL	*Center column; authority mechanisms and gates*
1121	DIRECTIVE ISSUANCE	*Binding directive dispatch to Constitutional Policy Core*

1122		VETO GATE	*Diamond decision node; veto check before ratification*
1123		OVERRIDE CONSENSUS MECHANISM	*Supermajority threshold required for system-level overrides*
1124		AUDIT LOG COMMIT	*Immutable ledger entry on each binding council decision*
1125		RATIFICATION TIMEOUT	*Auto-escalation triggered on deliberation inaction*
1130		INTERFACE LAYER PANEL	*Right column; inbound/outbound communication interfaces*
1131		AI SYSTEM STATUS FEED	*Real-time telemetry from Governance Monitor (ref. 0900)*
1132		ALERT NOTIFICATION BUS	*Priority-classified alert routing to relevant council seats*
1133		SECURE COUNCIL COMMS	*Encrypted inter-member channel; tamper-evident messaging*
1134		PUBLIC TRANSPARENCY PORTAL	*Redacted decision publication per statutory requirement*
1135		EXTERNAL REGULATOR LINK	*Bilateral API with national AI regulatory authorities*
1140		COUNCIL DECISION PROCESS FLOW	*Sequential flow diagram; bottom sub-section of figure*
1141		ISSUE INTAKE	*Entry point; alert or formal request received*
1142		DELIBERATION	*Council review and discussion phase*
1143		VOTE DECISION GATE	*Diamond gate; PASS proceeds to sign; FAIL returns to 1142*
1144		DIRECTIVE SIGNED	*Countersigned and timestamped directive record*
1145		DISPATCH AND AUDIT	*Directive dispatched; immutable audit commit appended*
1180		EMERGENCY ESCALATION PATH	*Threat Level 4+ path; auto-invokes seat 1115*

NAI-GOV-1100 | FIG11_REV_B.SVG B&W PATENT STYLE — USPTO COMPATIBLE SHEET 11 OF 15 | REV B | 2026-04-19

Clinical and eldercare deployment embodiment showing the system observing a monitored space using privacy-preserving sensing and providing governed advisory outputs to an authorized caregiver.

Illustrate a room environment with:

- bed or chair
- monitored subject region
- privacy-preserving sensor position
- local compute device
- caregiver review interface

Show that the system detects posture, instability, or fall-risk conditions but routes all results through human review.

Key labels: 110, 120, 140.

FIG. 12 — Humanitarian Field Deployment

FIG. 12

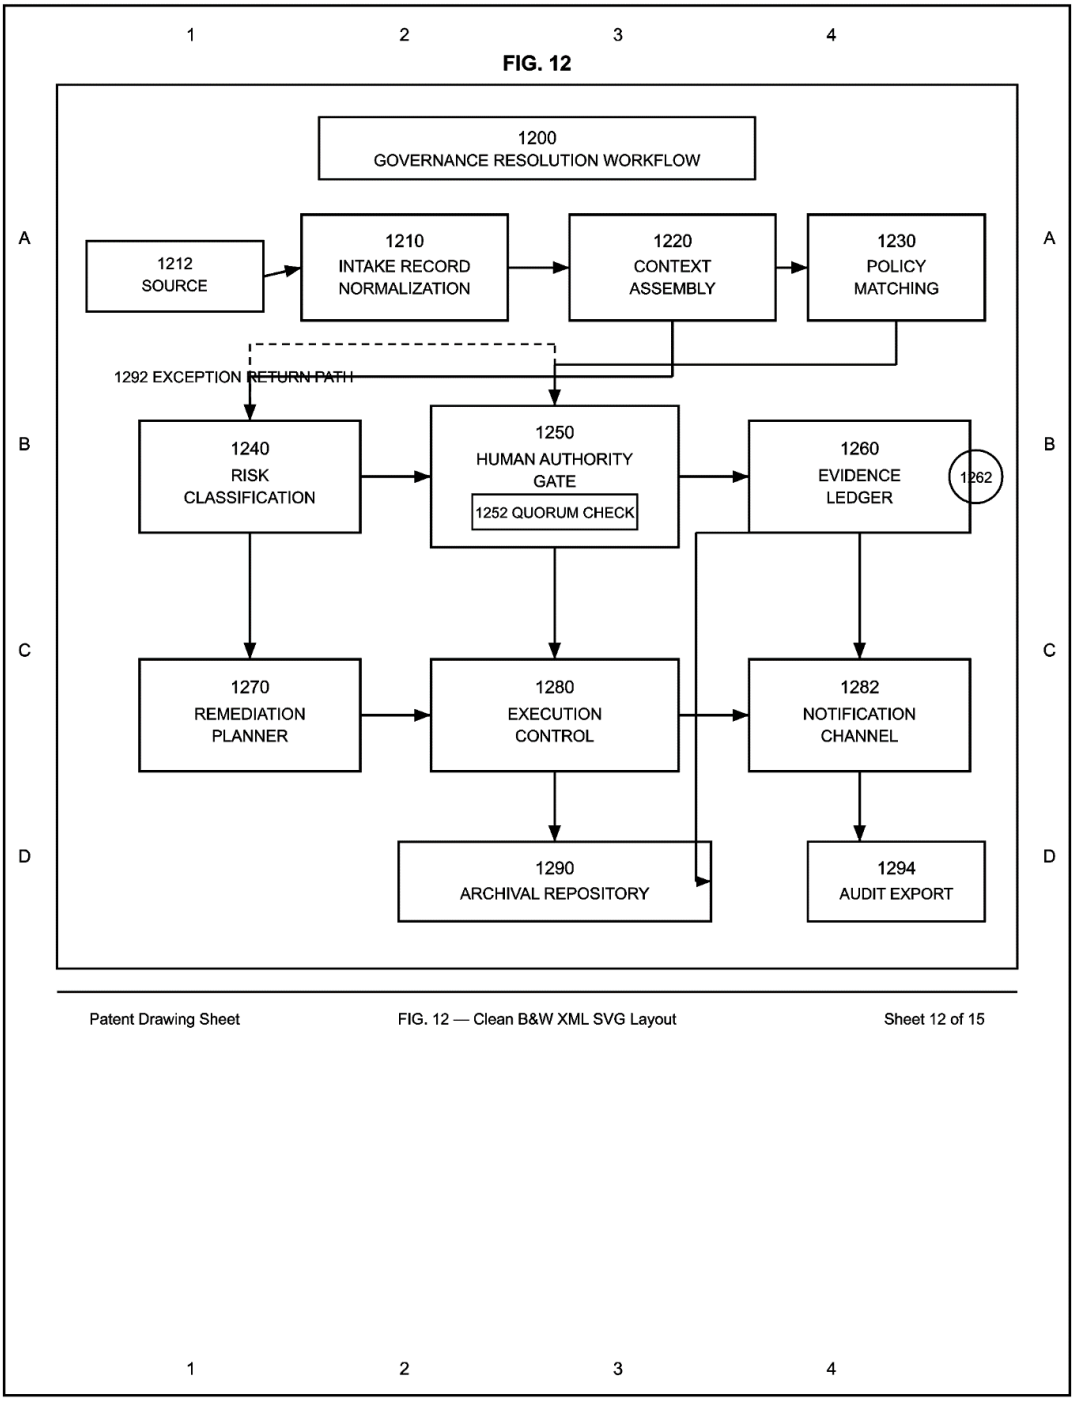

Humanitarian field deployment embodiment showing a ruggedized kit with local compute, portable power, protected tablet interface, sensor head, and optional selective communications bridge.

Illustrate a rugged transportable kit containing:

- sensor module
- edge compute box
- portable battery or solar pack
- rugged tablet
- optional selective satellite or remote relay component

Show local-first operation.

Key labels: 110, 120, 270, 280, 290, 183.

FIG. 13 — Protected Communications Pathway

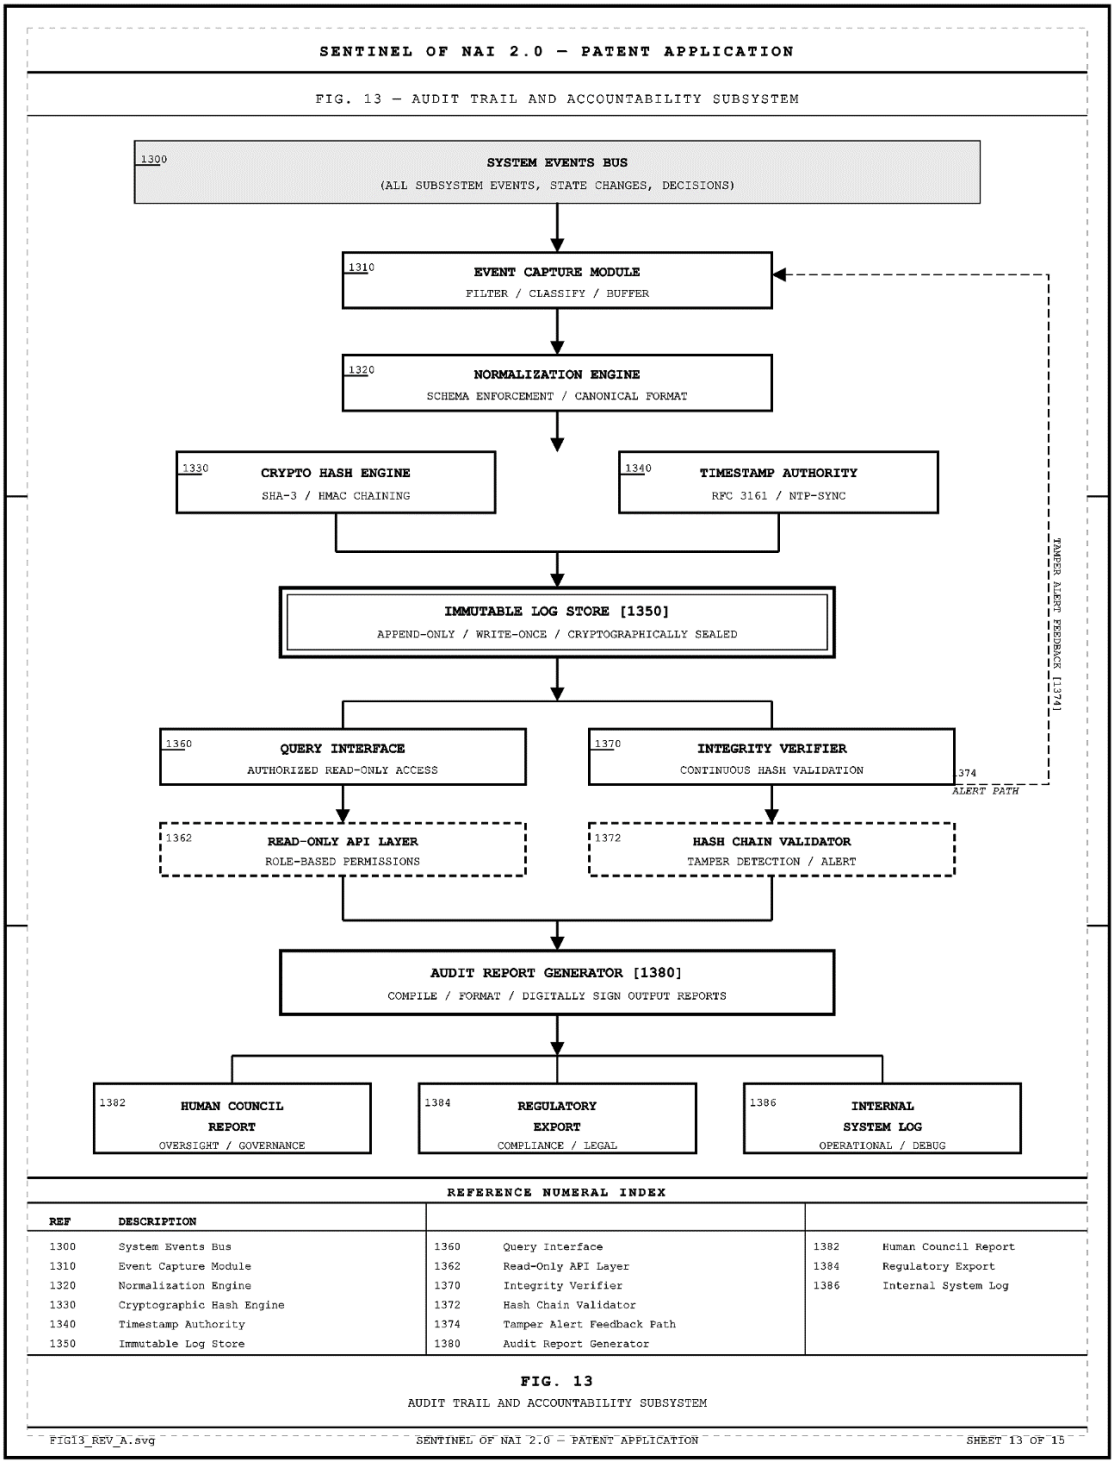

REFERENCE NUMERAL INDEX

REF	DESCRIPTION				
1300	System Events Bus	1360	Query Interface	1382	Human Council Report
1310	Event Capture Module	1362	Read-Only API Layer	1384	Regulatory Export
1320	Normalization Engine	1370	Integrity Verifier	1386	Internal System Log
1330	Cryptographic Hash Engine	1372	Hash Chain Validator		
1340	Timestamp Authority	1374	Tamper Alert Feedback Path		
1350	Immutable Log Store	1380	Audit Report Generator		

FIG. 13
AUDIT TRAIL AND ACCOUNTABILITY SUBSYSTEM

Protected communications pathway showing local-only mode, air-gapped mode, selective transmission mode, and event-summary relay without unrestricted raw data export.

Draw a branch diagram with three communications modes:

- **181 Local-only mode**
- **182 Air-gapped transfer**
- **183 Selective relay**

Explicitly separate raw sensitive data from abstracted event summaries.

Add a blocked path:

- **cloud inference prohibited** or
- **unrestricted export prohibited**

FIG. 14 — Operational State Diagram

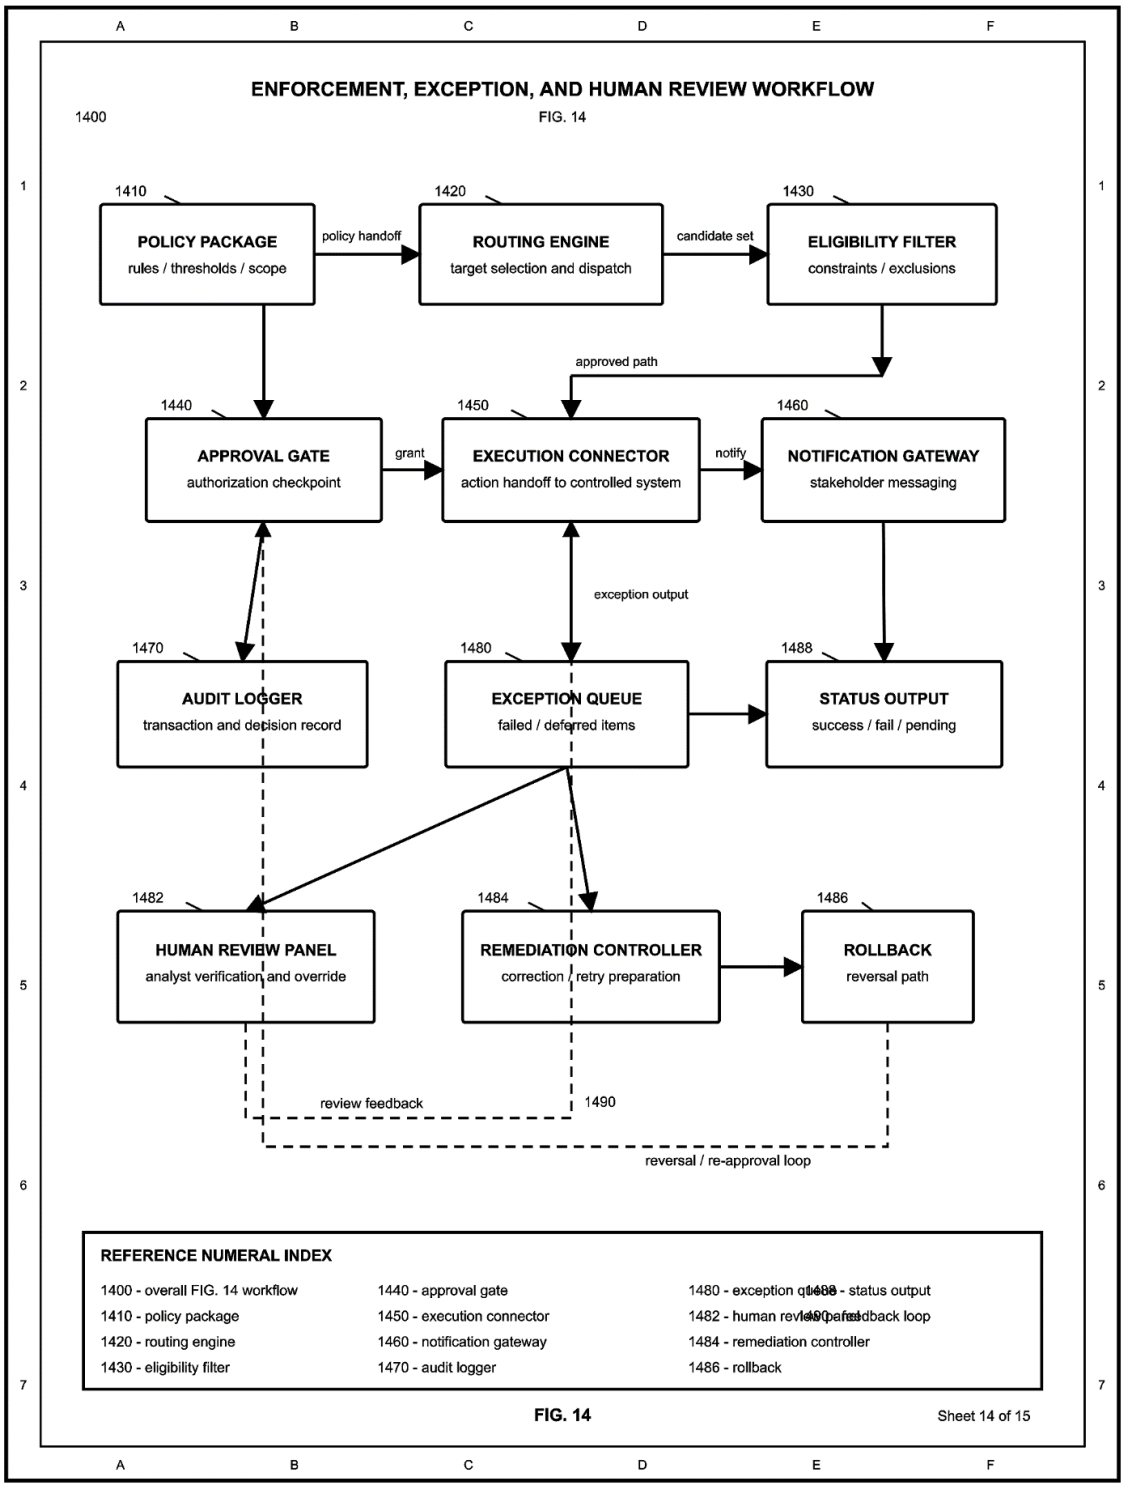

Operational state diagram showing Observe, Infer, Govern, Pause, Review, Approve or Reject, Log, Freeze, and Purge states.

Use a state-machine layout:

- Observe
- Infer
- Govern
- Pause
- Review
- Approve
- Reject
- Log
- Freeze
- Purge

Show rejection returning to safe idle and drift sending system to Freeze.

FIG. 15 — Integrated Ecosystem Figure

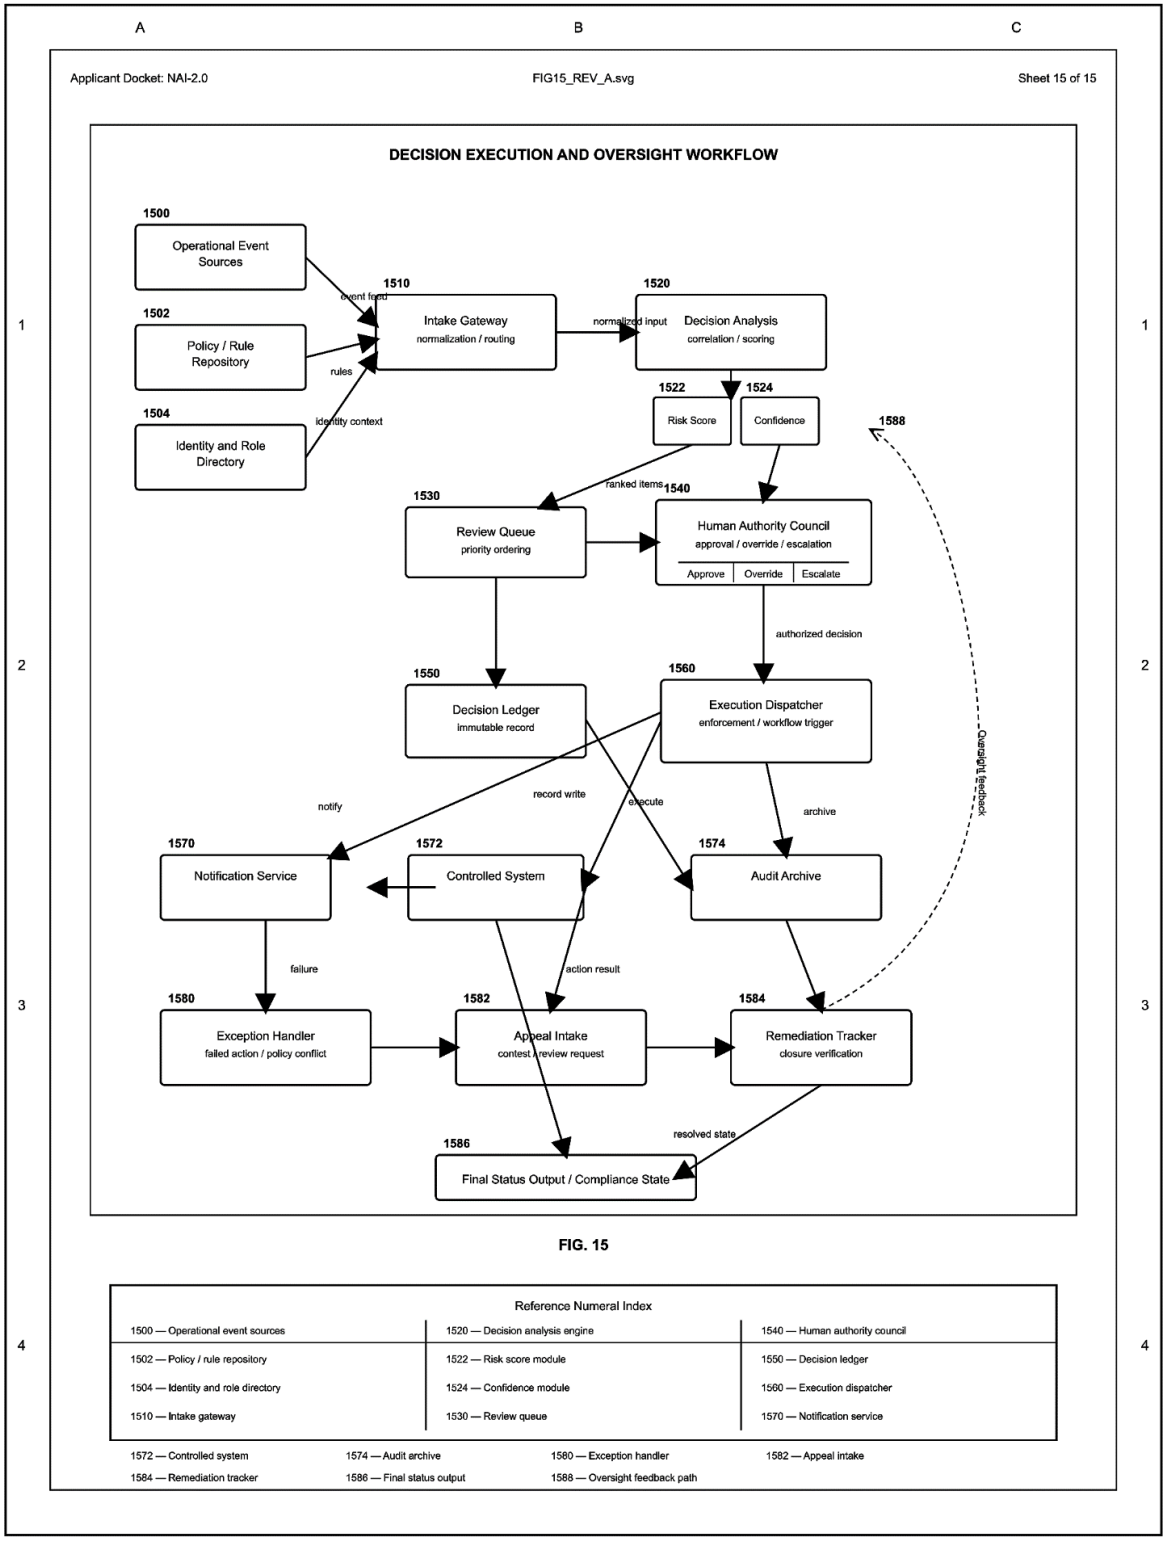

FIG. 15

Reference Numeral Index		
1500 — Operational event sources	1520 — Decision analysis engine	1540 — Human authority council
1502 — Policy / rule repository	1522 — Risk score module	1550 — Decision ledger
1504 — Identity and role directory	1524 — Confidence module	1560 — Execution dispatcher
1510 — Intake gateway	1530 — Review queue	1570 — Notification service

1572 — Controlled system
1574 — Audit archive
1580 — Exception handler
1582 — Appeal intake
1584 — Remediation tracker
1586 — Final status output
1588 — Oversight feedback path

Integrated ecosystem figure showing the application of the same constitutional architecture across clinical, institutional, governance, and humanitarian environments.

Create a hub-and-spoke figure with **100 NAI 2.0 System** at the center and the following domains around it:

- clinical
- eldercare
- institutional safety
- governance
- humanitarian operations
- remote field environments

Show that all use the same governance core and privacy architecture.

This creates a non-agentic safety architecture because the system cannot become fully self-governing through software adaptation alone. If drift control is hardware-enforced, then:

- policy drift cannot be silently normalized by software
- self-modification cannot expand authority beyond physical constraints
- learned behavior cannot redefine command boundaries
- software cannot fully model or rewrite the enforcement layer

So the system may behave intelligently, but it cannot possess sovereignty over its operational mandate.

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: _____

Name: Koh Wui Kiat, Edwin

Date: _____19/4/2026_____

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)