Utility Patent Draft

Hardware-Control Governance Devices of NAI 2.0

Inventor

Edwin Koh Wui Kiat

Applicant

Edwin Koh Wui Kiat (7/5/2026 : Singapore)

Title of the Invention

Hardware-Control Governance Devices and Intelligent Orchestration Under NAI 2.0

Field of the Invention

The present invention relates generally to medical device governance, cybersecurity, device trust management, and intelligent orchestration. More particularly, the invention relates to hardware-based wearable governance devices and associated control systems that operate under an intelligence layer identified as **NAI 2.0** within an **ABC+2S Guardian Framework** to enforce trust, safety, compliance, telemetry management, and responsive operational control.

Background of the Invention

Modern healthcare and monitored-care environments increasingly depend on distributed connected devices, wearable sensors, mobile endpoints, and remote monitoring systems. These systems often generate large volumes of physiological data and device-state information, but conventional arrangements typically lack unified hardware-governance control, contextual trust evaluation, and coordinated policy enforcement across wearable endpoints.

In many implementations, wearable devices are treated primarily as sensing tools rather than governed infrastructure components. As a result, device identity, integrity state, contextual trust, workflow compliance, and security posture may not be evaluated in a continuous and coordinated manner. This can lead to failures in safe operation, false alerts, untrusted telemetry, poor anomaly classification, incomplete auditability, and insufficient enforcement of operational restrictions when device conditions degrade.

There is therefore a need for a system in which wearable medical devices are governed as intelligent hardware-controlled endpoints within an integrated framework capable of identity validation, telemetry interpretation, compliance enforcement, anomaly classification, safe-state control, and orchestrated response. The present invention addresses that need.

Summary of the Invention

The invention provides a **hardware-control governance system** in which one or more wearable medical governance devices operate as managed endpoints under an intelligent orchestration architecture identified in some embodiments as **NAI 2.0**. The wearable devices may include, by way of example, a **wrist-worn wearable medical governance device** and a **patch-form wearable medical governance device**, each configured to collect physiological data, device-state data, or contextual data and to communicate such data to a governance framework.

In some embodiments, the wearable devices are enrolled within an **ABC+2S Guardian Framework**, which manages identity, attestation, trust evaluation, communications policy, segmentation, compliance monitoring, alert handling, and recovery control. NAI 2.0 operates as an intelligence and orchestration layer that receives telemetry and contextual metadata, determines whether conditions correspond to clinical events, device anomalies, security anomalies, or compliance conditions, and initiates governance actions affecting device functionality, communication permissions, alert treatment, or operating state.

In some embodiments, the system supports local governance continuity during connectivity loss, synchronized recovery following reconnection, auditable event logging, and policy-driven transition among normal, restricted, quarantined, and

safe-state modes. The invention thereby transforms wearable medical devices from passive sensing elements into governed hardware-control devices participating in trusted clinical and operational infrastructure.

Brief Description of the Drawings

FIG. 1 is a schematic system diagram illustrating an example embodiment of the ABC+2S Guardian Framework including wearable medical governance devices, clinical infrastructure, and an intelligence orchestration layer identified as NAI 2.0.

FIG. 2 is a schematic block diagram of an example wrist-worn wearable medical governance device embodiment.

FIG. 3 is a schematic block diagram of an example patch-form wearable medical governance device embodiment.

FIG. 4 is a data-flow diagram illustrating transmission of telemetry, identity information, integrity information, and contextual metadata between wearable medical governance devices and the ABC+2S Guardian Framework.

FIG. 5 is a flowchart illustrating an example governance method including identity validation, trust evaluation, event classification, and enforcement of a governance action.

FIG. 6 is a state-transition diagram illustrating example operating states of a governed wearable medical device including normal mode, restricted mode, safe-state mode, quarantined mode, and recovery mode.

FIG. 7 is a workflow diagram illustrating an example coordinated response involving a wearable device, NAI 2.0, and one or more external systems or care personnel.

FIG. 8 is a decision-flow diagram illustrating classification of a detected condition as a clinical condition, device anomaly, security anomaly, or compliance condition.

FIG.	Type	Scope
1	System diagram	Entire ABC+2S Guardian Framework with NAI 2.0
2	Block diagram	Wrist-worn wearable medical governance device
3	Block diagram	Patch-form wearable medical governance device
4	Data-flow diagram	Device-to-framework information exchange
5	Flowchart	Governance method
6	State-transition diagram	Device operating states
7	Workflow diagram	Coordinated response sequence
8	Decision-flow diagram	Condition classification logic

Numeral Range: Use

100-series Framework-level architecture / FIG. 1 core elements

200-series Wrist-worn device elements / FIG. 2

300-series Patch-form device elements / FIG. 3

400-series Data channels and exchanged information / FIG. 4

500-series Method steps / FIG. 5

600-series Device states / FIG. 6

700-series Coordinated response actors and workflow elements / FIG. 7

800-series Decision nodes and classifications / FIG. 8

Detailed Description of the Invention

1. General Architecture

In some embodiments, the invention comprises a hardware-control governance system for wearable medical devices operating under an intelligent orchestration architecture. The system may include one or more wearable medical governance devices, one or more communications pathways, one or more intermediate gateways or mobile systems, and a central or distributed governance framework. The governance framework may be implemented as part of an architecture identified herein as the **ABC+2S Guardian Framework**.

The ABC+2S Guardian Framework may include one or more processors, memory resources, secure data stores, policy engines, audit engines, identity services, attestation services, trust-scoring services, anomaly detection services, communications segmentation controls, and response orchestration components. In some embodiments, the framework interfaces with clinical systems, enterprise systems, security systems, monitoring systems, and remote service infrastructure.

An intelligence layer identified in some embodiments as **NAI 2.0** may receive and analyze information from the wearable devices and other sources. NAI 2.0 may function as a reasoning, inference, correlation, and orchestration engine for classifying events, applying policy logic, and directing governance actions.

---

2. Wearable Medical Governance Devices

In some embodiments, the hardware-control governance devices include one or more wearable medical endpoints configured to operate as governed edge devices rather than merely passive sensors. Each such device may include one or more of:

- a sensor subsystem,
- a communication subsystem,
- a processor subsystem,
- a power subsystem,
- a security subsystem,
- local storage,

- firmware control logic,
- local policy logic,
- identity credentials,
- attestation logic,
- integrity verification logic, and
- one or more user-facing or system-facing indicators.

The device may generate and transmit physiological data, device-state data, operational telemetry, identity data, integrity data, workflow-state data, location-related data, or combinations thereof.

---

3. Wrist-Worn Device Embodiment (WM003)

In some embodiments, a first device embodiment comprises a **wrist-worn wearable medical governance device**, referred to herein as **WM003**. WM003 may be configured to be worn on a wrist of a patient, clinician, caregiver, technician, or authorized operator.

WM003 may include a housing, band or attachment structure, display, one or more user-input interfaces, and a sensor subsystem. The sensor subsystem may include one or more of the following:

- photoplethysmography sensors,
- electrocardiographic sensors,
- temperature sensors,
- accelerometers,
- gyroscopes,
- motion sensors,
- proximity sensors,
- pulse-related sensors,
- orientation sensors, and
- biometric or user-presence sensors.

WM003 may be configured to display alerts, workflow prompts, authentication requests, status indications, reminders, emergency notifications, trust warnings, or guided instructions. In some embodiments, WM003 can receive governance commands from the ABC+2S Guardian

Framework to alter communication behavior, alert behavior, access permissions, sampling behavior, local display behavior, or operating state.

WM003 may be particularly useful for clinician-linked workflow validation, patient-associated monitoring, user authentication continuity, and on-body trust-aware governance.

---

4. Patch-Form Device Embodiment (WM005)

In some embodiments, a second device embodiment comprises a **patch-form wearable medical governance device**, referred to herein as **WM005**. WM005 may be configured for attachment to a body surface of a monitored subject.

WM005 may include a low-profile body-mountable housing, adhesive or attachment structure, biosignal interface, electrodes or conductive contacts, processor subsystem, communication subsystem, and security subsystem. WM005 may acquire one or more of the following:

- electrocardiographic data,
- respiration-related data,
- heart-rate data,
- temperature data,
- motion data,
- posture-related data,
- skin-contact data,
- adhesion-quality data,
- wear-duration data,
- moisture-exposure data, and
- device-state anomaly data.

WM005 may be configured to detect detachment, contact degradation, improper placement, moisture compromise, wear expiration, or signal unreliability. Such conditions may be transmitted as governance-relevant events to the ABC+2S Guardian Framework.

In some embodiments, WM005 includes local logic enabling preliminary determination of whether an observed signal interruption is likely due to

physiological change or hardware-state degradation. NAI 2.0 may then further classify the condition using contextual information.

---

5. Security, Identity, and Attestation

In some embodiments, each wearable device includes a security subsystem configured to support trusted enrollment and governed participation in the framework. The security subsystem may include:

- device identity storage,
- secure key storage,
- certificate logic,
- attestation logic,
- secure boot logic,
- firmware validation logic,
- tamper detection,
- pairing assurance logic,
- trusted communication initialization, and
- local policy enforcement support.

The governance framework may use this information to determine whether a device is trusted, conditionally trusted, degraded, restricted, quarantined, or disallowed. Trust determination may depend on one or more of firmware status, certificate validity, sensor consistency, prior device behavior, workflow alignment, location consistency, patient association, communication history, or local integrity metrics.

---

6. NAI 2.0 Intelligence Layer

In some embodiments, **NAI 2.0** operates as the intelligence and orchestration layer of the hardware-control governance system. NAI 2.0 may be configured to:

- ingest telemetry from WM003, WM005, and other governed devices;
- correlate telemetry with contextual metadata;

- evaluate trust, risk, safety, and compliance conditions;
- distinguish among clinical events, device anomalies, security anomalies, and compliance conditions;
- recommend or cause governance actions;
- adapt response intensity according to policy;
- maintain explainable decision records; and
- coordinate with human operators, external systems, and automated control processes.

The contextual metadata may include patient identity, clinician identity, workflow stage, device enrollment state, location context, timing context, connectivity context, command history, alert history, or environmental conditions.

For example, NAI 2.0 may determine that loss of continuous ECG signal from WM005 is more likely caused by patch detachment than patient deterioration based on simultaneous adhesion-quality degradation, reduced contact integrity, and movement metadata. Similarly, NAI 2.0 may determine that an interaction pattern observed through WM003 is indicative of unauthorized access or workflow noncompliance rather than ordinary user behavior.

---

7. Governance Actions

Based on trust evaluation or event classification, the framework may enforce one or more governance actions. Examples include:

- restricting communications,
- permitting communications only with approved systems,
- quarantining a device or data stream,
- increasing or decreasing sampling rate,
- requiring re-authentication,
- requiring re-attestation,
- suppressing a false or low-confidence alert,
- escalating an alert,
- triggering a care-team notification,
- placing the device in restricted mode,
- placing the device in safe-state mode,

- initiating a recovery workflow,
- logging an auditable event,
- denying command execution,
- requiring human review, or
- changing local user-interface behavior.

These actions may be executed by the wearable device itself, by the governance framework, by NAI 2.0 through orchestration logic, or by combinations thereof.

---

8. Local Governance During Connectivity Loss

In some embodiments, the wearable device maintains a local subset of governance rules for use during intermittent connectivity or temporary disconnection from the framework. During such conditions, the device may continue limited operation according to pre-authorized local policy. Once connectivity is restored, locally recorded events may be synchronized to the framework for audit, review, and trust recalculation.

This capability improves operational resilience while preserving centralized oversight.

---

9. Auditability and Compliance

In some embodiments, the ABC+2S Guardian Framework maintains a persistent audit trail that records one or more of:

- device enrollment,
- attestation status,
- pairing changes,
- policy assignments,
- trust score changes,
- command transmissions,
- alert events,
- anomaly classifications,

- local fallback actions,
- recovery actions, and
- operator acknowledgments.

Such records may support compliance review, forensic analysis, clinical safety investigation, device fleet management, and policy optimization.

---

10. System Operation

In an example method of operation, a wearable device such as WM003 or WM005 generates physiological and device-state data. The data is associated with identity information and integrity information and transmitted to the ABC+2S Guardian Framework. NAI 2.0 analyzes the received data in view of contextual metadata and determines whether the event corresponds to a clinical condition, a device anomaly, a security anomaly, or a compliance condition. Based on that determination, the framework applies a governance action. The action may alter device operation, communications permissions, alert handling, safe-state status, or recovery sequence. The result is logged for subsequent audit and policy refinement.

---

11. Variations

Although WM003 and WM005 are described as wrist-worn and patch-form embodiments, respectively, the invention is not limited to those precise form factors. The principles described herein may be applied to additional body-worn, clinician-worn, patient-worn, adhered, embedded, or accessory-based devices that operate as governed hardware endpoints under NAI 2.0 and the ABC+2S Guardian Framework.

Similarly, the use of the names **NAI 2.0**, **ABC+2S Guardian**, **WM003**, and **WM005** is illustrative. Other names, versions, modules, or implementation labels may be used without departing from the scope of the invention.

Claims

What is claimed is:

1. **A medical hardware-control governance system comprising:**

at least one wearable medical governance device comprising
a sensor subsystem configured to acquire physiological data or device-state data,
a communication subsystem configured to transmit the physiological data or device-state data,
a processing subsystem, and
a security subsystem configured to provide at least one of device identity information, attestation information, or integrity information;

a governance framework comprising one or more processors and memory storing instructions that, when executed, cause the governance framework to
receive the physiological data or device-state data and the at least one of device identity information, attestation information, or integrity information from the at least one wearable medical governance device,
associate the received information with contextual metadata,
determine a trust condition associated with the at least one wearable medical governance device based at least in part on the received information and the contextual metadata, and
enforce a governance action affecting operation of the at least one wearable medical governance device or handling of data generated by the at least one wearable medical governance device based on the trust condition; and

an intelligence layer configured to
analyze the received information and the contextual metadata,
determine whether a detected event corresponds to a clinical condition, a device anomaly, a security anomaly, or a compliance condition, and
cause the governance framework to initiate the governance action.

2. The system of claim 1, wherein the at least one wearable medical governance device comprises a wrist-worn wearable medical governance device.
3. The system of claim 2, wherein the wrist-worn wearable medical governance device comprises a display and a user interface configured to present at least

one of an alert, workflow guidance, an authentication prompt, a status indicator, or an emergency notification.

4. The system of claim 2, wherein the wrist-worn wearable medical governance device comprises at least one sensor selected from the group consisting of a photoplethysmography sensor, an electrocardiographic sensor, a temperature sensor, an accelerometer, a gyroscope, a proximity sensor, and a motion sensor.
5. The system of claim 1, wherein the at least one wearable medical governance device comprises a patch-form wearable medical governance device configured for attachment to a body surface.
6. The system of claim 5, wherein the patch-form wearable medical governance device is configured to detect at least one of detachment, contact degradation, wear-duration threshold exceedance, moisture exposure, or adhesion-quality degradation.
7. The system of claim 5, wherein the patch-form wearable medical governance device is configured to acquire at least one of electrocardiographic data, respiration-related data, temperature data, motion data, posture data, or skin-contact integrity data.
8. The system of claim 1, wherein the governance action comprises at least one of restricting communications, requiring re-authentication, requiring re-attestation, quarantining a data stream, suppressing an alert, elevating an alert, changing a sampling rate, placing the at least one wearable medical governance device in a restricted mode, placing the at least one wearable medical governance device in a safe-state mode, or initiating a recovery workflow.
9. The system of claim 1, wherein the contextual metadata comprises at least one of patient context, clinician context, workflow context, location context, timing context, connectivity context, command history, or policy context.
10. The system of claim 1, wherein the intelligence layer is configured to distinguish a physiological event from a device anomaly based on the physiological data, the device-state data, and the contextual metadata.
11. The system of claim 1, wherein the trust condition is determined based on a combination of firmware state, certificate status, pairing integrity, sensor consistency, battery condition, patient association state, location context, communication history, and workflow alignment.
12. The system of claim 1, wherein the governance framework maintains an auditable record of device enrollment, attestation status, policy changes, alert events, command transmissions, anomaly classifications, and recovery actions.

13. The system of claim 1, wherein the governance framework applies segmentation rules that selectively permit or deny communications between the at least one wearable medical governance device and one or more gateways, mobile devices, clinical systems, enterprise systems, or cloud systems.
14. The system of claim 1, wherein the at least one wearable medical governance device is configured to execute locally stored governance rules during a loss of connectivity and to synchronize locally recorded events after connectivity is restored.
15. **A wrist-worn wearable medical governance device comprising:**

    a housing configured to be worn on a wrist of a user;
    a sensor subsystem configured to acquire physiological data or motion-related data;
    a display configured to present at least one of an alert, workflow guidance, an authentication request, or a status indication;
    a communication subsystem configured to communicate with a remote governance framework;
    a processing subsystem; and
    a security subsystem configured to provide device identity information and integrity-related information to the remote governance framework,

    wherein the processing subsystem is configured to operate the wrist-worn wearable medical governance device as a governed endpoint whose functionality is selectively controlled based on a trust evaluation performed by the remote governance framework.

16. The device of claim 15, wherein the wrist-worn wearable medical governance device is configured to receive a command from the remote governance framework that changes at least one of alert behavior, data transmission behavior, authentication behavior, operating mode, or communication permissions.
17. The device of claim 15, wherein the security subsystem comprises a hardware-rooted identity element or secure credential store.
18. The device of claim 15, wherein the wrist-worn wearable medical governance device is configured to present a re-authentication request in response to a trust degradation condition detected by the remote governance framework.
19. **A patch-form wearable medical governance device comprising:**

a body-mountable housing;
an attachment structure configured to secure the device to a body surface of a monitored subject;
a biosignal acquisition subsystem configured to acquire physiological data from the monitored subject;
a communication subsystem configured to communicate with a remote governance framework;
a processing subsystem; and
a security subsystem configured to provide identity information or integrity information to the remote governance framework,

wherein the processing subsystem is configured to detect at least one of a contact-quality condition, a detachment condition, or a device-state anomaly, and
wherein operation of the patch-form wearable medical governance device is governed based on a trust evaluation performed by the remote governance framework.

20. The device of claim 19, wherein the biosignal acquisition subsystem comprises at least one electrode configured to acquire electrocardiographic data.
21. The device of claim 19, wherein the processing subsystem is configured to generate adhesion-quality data and transmit the adhesion-quality data to the remote governance framework.
22. The device of claim 19, wherein the patch-form wearable medical governance device comprises an indicator configured to provide a visual indication of at least one of pairing state, compliance state, alert state, battery state, or connectivity state.
23. **A computer-implemented method for governing a wearable medical device, the method comprising:**

receiving, from a wearable medical governance device, physiological data, device-state data, and identity-related or integrity-related information;
associating the physiological data, the device-state data, and the identity-related or integrity-related information with contextual metadata;
analyzing, by an intelligence layer, the physiological data, the device-state data, and the contextual metadata to determine whether an event corresponds to a clinical condition, a device anomaly, a security anomaly, or

a compliance condition;
generating a trust evaluation for the wearable medical governance device based on the analysis; and
causing a governance framework to enforce a governance action affecting the wearable medical governance device or a data stream generated by the wearable medical governance device based on the trust evaluation.

24. The method of claim 23, wherein the wearable medical governance device comprises a wrist-worn wearable medical governance device.
25. The method of claim 23, wherein the wearable medical governance device comprises a patch-form wearable medical governance device.
26. The method of claim 23, wherein enforcing the governance action includes distinguishing a device detachment condition from a physiological deterioration condition.
27. The method of claim 23, wherein enforcing the governance action includes initiating at least one of clinician notification, data-stream quarantine, restricted communications, re-authentication, re-attestation, safe-state entry, or recovery workflow execution.
28. **A non-transitory computer-readable medium storing instructions that, when executed by one or more processors of a medical governance framework, cause the medical governance framework to:**

    receive physiological data, device-state data, and security-related information from a wearable medical governance device;
    correlate the received information with contextual metadata;
    determine, using an intelligence layer, whether the received information indicates a clinical event, a device anomaly, a security anomaly, or a compliance condition;
    generate a trust score or trust state for the wearable medical governance device; and
    initiate a governance action that changes device operation, communication permissions, alert handling, or data routing based on the trust score or trust state.

Abstract

A hardware-control governance system for wearable medical devices is disclosed. The system includes one or more wearable medical governance devices, a governance framework, and an intelligence layer identified in some embodiments as NAI 2.0. The wearable devices may include a wrist-worn device and a patch-form device configured to acquire physiological data, device-state data, and security-related information. The governance framework receives the information, associates it with contextual metadata, determines a trust condition, and enforces one or more governance actions. The intelligence layer analyzes the information and contextual metadata to distinguish clinical conditions, device anomalies, security anomalies, and compliance conditions. Governance actions may include communication restriction, re-authentication, re-attestation, alert suppression or escalation, safe-state entry, and recovery workflow initiation. The system provides trusted, auditable, policy-driven control of wearable medical hardware devices.

APPLICATION NO.: [PENDING]
FILING DATE: [PENDING]
INVENTOR: [DESIGNATED INVENTOR]

SHEET 1 OF 8
DRAWING NO.: FIG. 1
CLASSIFICATION: PATENT DRAFT

FIG. 1

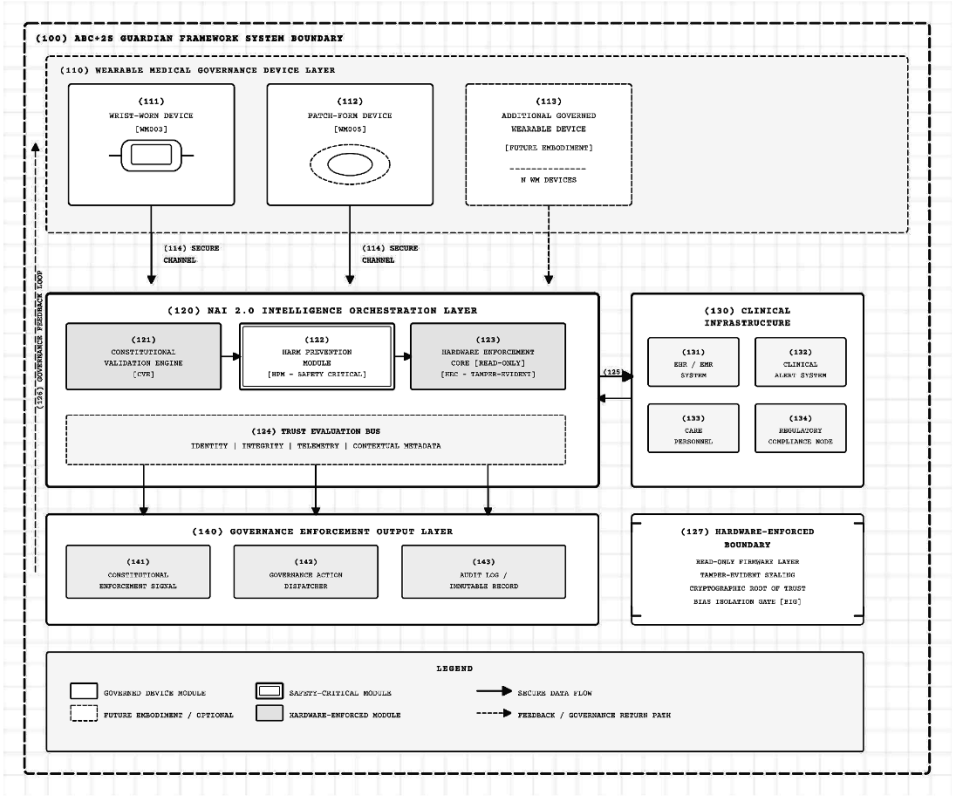

REFERENCE NUMERALS — FIG. 1

100	ABC+2S Guardian Framework — System Boundary	110	Wearable Medical Governance Device Layer
111	Wrist-Worn Wearable Medical Governance Device [WM003]	112	Patch-Form Wearable Medical Governance Device [WM005]
113	Additional Governed Wearable Device — Future Embodiment	114	Secure Communication Channel (Device-to-Framework)
120	NAI 2.0 Intelligence Orchestration Layer	121	Constitutional Validation Engine [CVE]
122	Harm Prevention Module [HPM] — Safety Critical	123	Hardware Enforcement Core [HEC] — Read-Only, Tamper-Evident
124	Trust Evaluation Bus — Identity, Integrity, Telemetry, Metadata	125	NAI 2.0 to Clinical Infrastructure Interface Bridge
126	Governance Feedback Loop — Enforcement Return Path	127	Hardware-Enforced Boundary — Cryptographic Root of Trust
130	Clinical Infrastructure Layer	131	Electronic Health Record / EMR System
132	Clinical Alert System	133	Care Personnel Interface
134	Regulatory Compliance Node	140	Governance Enforcement Output Layer
141	Constitutional Enforcement Signal	142	Governance Action Dispatcher
143	Audit Log / Immutable Record		

SHEET 1 OF 8 — FIG. 1 ABC+2S GUARDIAN FRAMEWORK — NAI 2.0 CONSTITUTIONAL PATENT DRAFT — PATENT PENDING

FIG. 1 is a schematic system diagram illustrating an example embodiment of the ABC+2S Guardian Framework including wearable medical governance devices, clinical infrastructure, and an intelligence orchestration layer identified as NAI 2.0.

ABC+2S GUARDIAN FRAMEWORK — PATENT DRAWING SHEET

FIG. 2 | WRIST-WORN WEARABLE MEDICAL GOVERNANCE DEVICE | SCHEMATIC BLOCK DIAGRAM

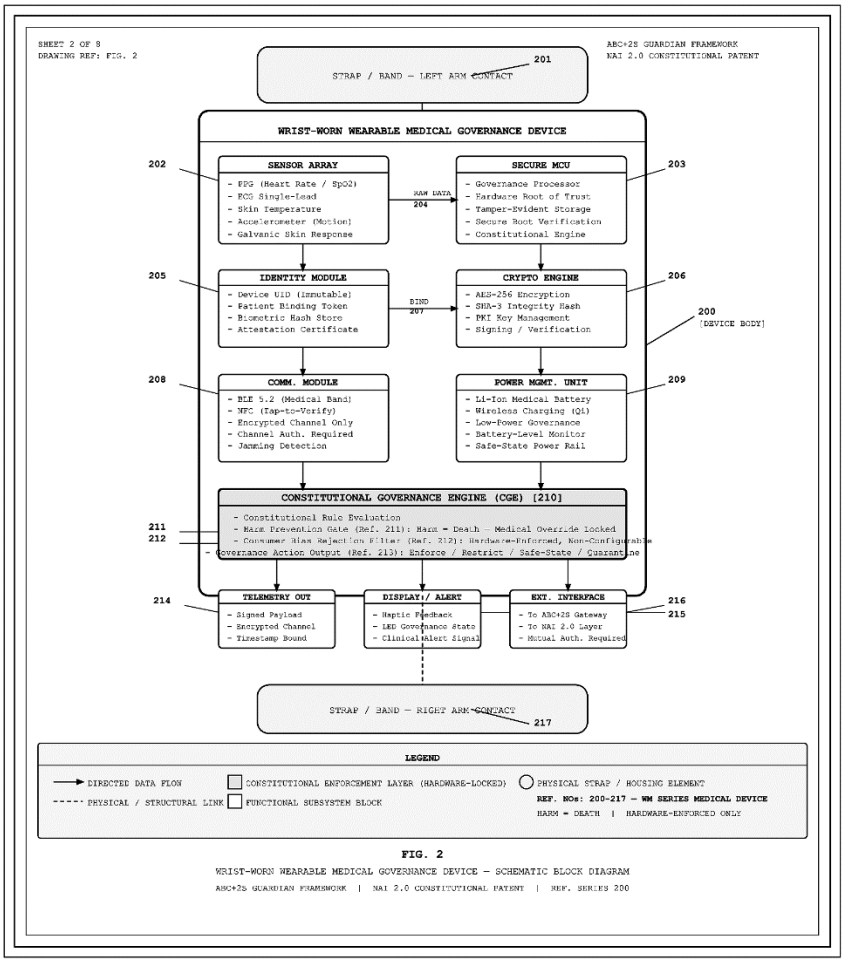

FIG. 2

WRIST-WORN WEARABLE MEDICAL GOVERNANCE DEVICE — SCHEMATIC BLOCK DIAGRAM

ABC+2S GUARDIAN FRAMEWORK | NAI 2.0 CONSTITUTIONAL PATENT | REF. SERIES 200

REF. NO.	COMPONENT	DESCRIPTION
200	Device Body	Wrist-Worn Wearable Medical Governance Device — complete assembly
201	Strap / Band (Left)	Left-arm contact strap, physical housing interface
202	Sensor Array	PPG, ECG, temperature, accelerometer, galvanic skin response sensors
203	Secure MCU	Governance processor with hardware root of trust and tamper-evident storage
204	Raw Data Bus	Sensor-to-processor data channel
205	Identity Module	Device UID, patient binding token, biometric hash, attestation certificate
206	Cryptographic Engine	AES-256 encryption, SHA-3 hashing, PKI key management, signing
207	Identity-Crypto Binding	Identity-to-cryptographic engine binding channel
208	Communication Module	BLE 5.2 (medical band), NFC tap-to-verify, encrypted channels only
209	Power Management Unit	Li-Ion medical battery, wireless charging, safe-state power rail
210	Constitutional Governance Engine (CGE)	Hardware-locked constitutional rule evaluation and enforcement block
211	Harm Prevention Gate	HARM = DEATH rule — medical override locked, non-configurable
212	Consumer Bias Rejection Filter	Hardware-enforced rejection of consumer framing inputs
213	Governance Action Output	Enforce / Restrict / Safe-State / Quarantine action signal
214	Telemetry Output	Signed, encrypted, timestamp-bound telemetry data output
215	Display / Alert Interface	Haptic feedback, LED governance state indicator, clinical alert signal
216	External Interface	Authenticated interface to ABC+2S Gateway and NAI 2.0 layer
217	Strap / Band (Right)	Right-arm contact strap, physical housing interface

ABC+2S GUARDIAN FRAMEWORK | NAI 2.0 CONSTITUTIONAL PATENT | FIG. 2 — WRIST-WORN WEARABLE MEDICAL GOVERNANCE DEVICE | REF. SERIES: 200–217 | SHEET 2 OF 8 | HARM = DEATH

FIG. 2 is a schematic block diagram of an example wrist-worn wearable medical governance device embodiment.

NAI 2.0 CONSTITUTIONAL [PATENT]
ABC+2S GUARDIAN FRAMEWORK

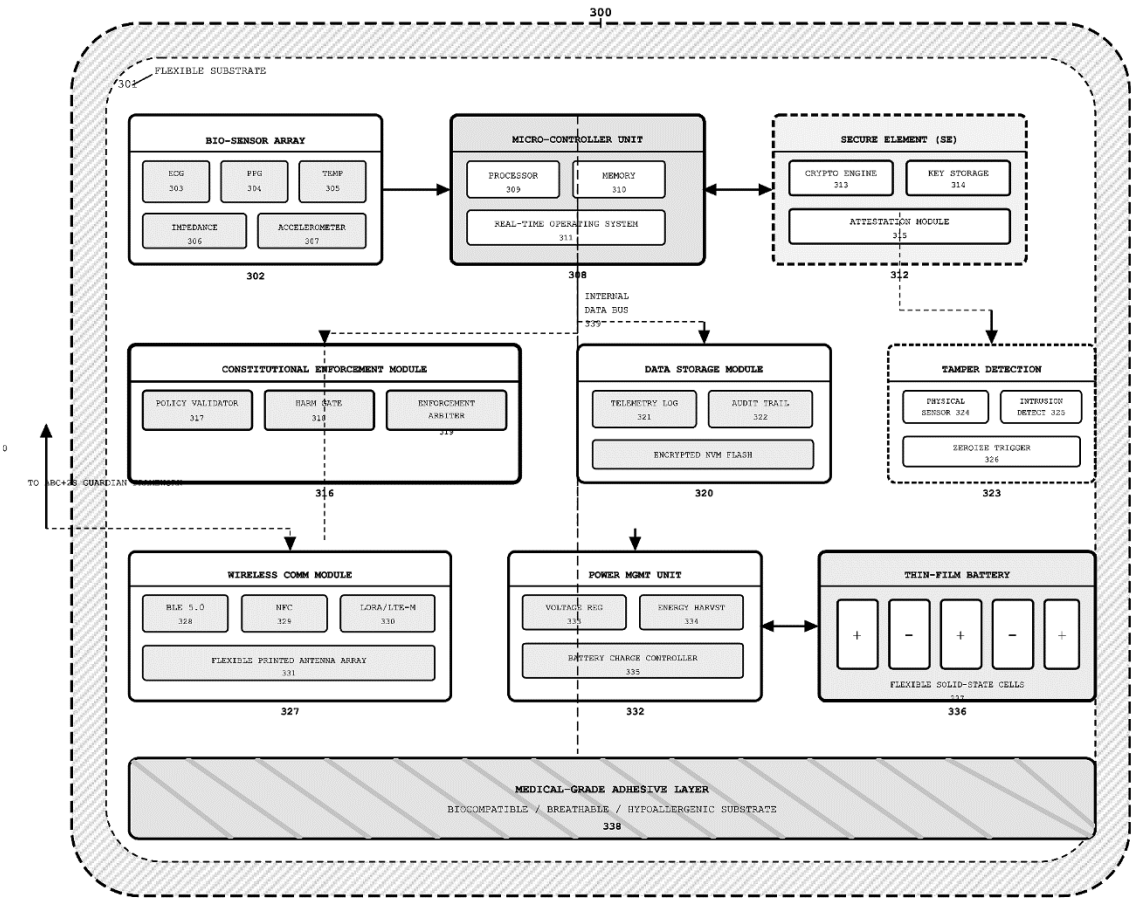

FIG. 3 - PATCH-FORM WEARABLE MEDICAL GOVERNANCE DEVICE: SCHEMATIC BLOCK DIAGRAM

FIG. 3 — REFERENCE NUMERALS (300-SERIES)

300	Patch-Form Wearable Medical Governance Device (Overall Boundary)	**301**	Flexible Substrate	**302**	Bio-Sensor Array
303	ECG Sensor	**304**	PPG Sensor	**305**	Temperature Sensor
306	Bioelectrical Impedance Sensor	**307**	Accelerometer / Motion Sensor	**308**	Micro-Controller Unit (MCU)
309	Processor Core	**310**	Memory (SRAM / Flash)	**311**	Real-Time Operating System (RTOS)
312	Secure Element (SE)	**313**	Cryptographic Engine	**314**	Key Storage (Tamper-Resistant)
315	Attestation Module	**316**	Constitutional Enforcement Module	**317**	Policy Validator
318	Harm Gate (Harm=Death Enforced)	**319**	Enforcement Arbiter	**320**	Data Storage Module
321	Telemetry Log	**322**	Audit Trail Store	**323**	Tamper Detection Circuit
324	Physical Tamper Sensor	**325**	Intrusion Detection Logic	**326**	Zeroize Trigger
327	Wireless Communication Module	**328**	BLE 5.0 Radio	**329**	NFC Transceiver
330	LoRa / LTE-M Radio	**331**	Flexible Printed Antenna Array	**332**	Power Management Unit (PMU)
333	Voltage Regulator	**334**	Energy Harvesting Module	**335**	Battery Charge Controller
336	Thin-Film Battery	**337**	Flexible Solid-State Battery Cells	**338**	Medical-Grade Adhesive Layer
339	Internal Data Bus	**340**	External Uplink to ABC+2S Guardian Framework		

FIG. 3 is a schematic block diagram of an example patch-form wearable medical governance device embodiment.

FIG. 4 — Data-Flow Diagram: Device-to-Framework Information Exchange

Drawing No.: NAI-2.0-FIG-004
Sheet: 4 of 8
Scale: Not to Scale
Date: 2026-04-23

FIG. 4

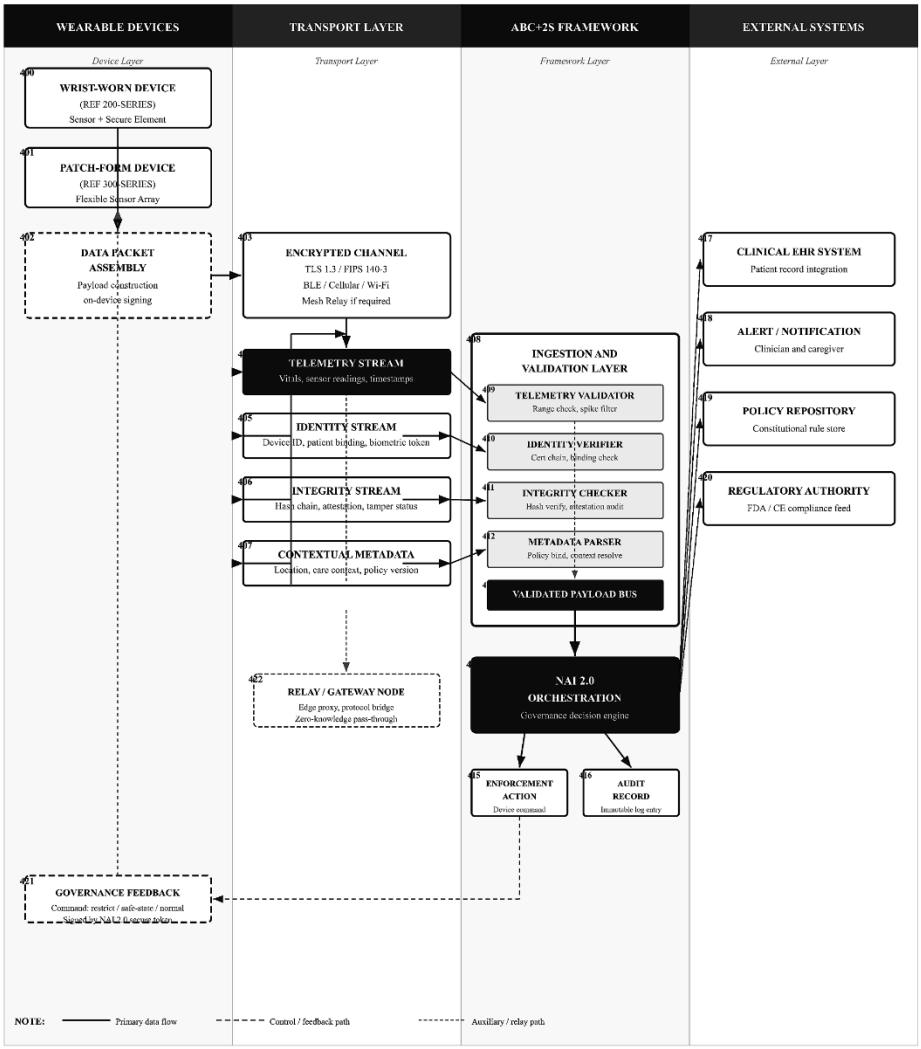

400	Wrist-Worn Wearable Medical Governance Device	401	Patch-Form Wearable Medical Governance Device	402	Data Packet Assembly and On-Device Signing Module
403	Encrypted Transport Channel (TLS 1.3 / FIPS 140-3)	404	Telemetry Stream (vitals, sensor readings, timestamps)	405	Identity Stream (device ID, patient binding, biometric token)
406	Integrity Stream (hash chain, attestation, tamper status)	407	Contextual Metadata (location, care context, policy version)	408	Ingestion and Validation Layer
409	Telemetry Validator (range check, spike filter)	410	Identity Verifier (certificate chain, binding check)	411	Integrity Checker (hash verify, attestation audit)
412	Metadata Parser (policy bind, context resolution)	413	Validated Payload Bus	414	NAI 2.0 Orchestration and Governance Decision Engine
415	Enforcement Action Output (device command)	416	Audit Record Output (immutable log entry)	417	Clinical EHR System (patient record integration)
418	Alert and Notification System (clinician, caregiver)	419	Policy Repository (constitutional rule store)	420	Regulatory Authority Feed (FDA / CE compliance)
421	Governance Feedback Channel (signed command return)	422	Relay and Gateway Node (edge proxy, protocol bridge)		

FIG. 4 is a data-flow diagram illustrating transmission of telemetry, identity information, integrity information, and contextual metadata between wearable medical governance devices and the ABC+2S Guardian Framework.

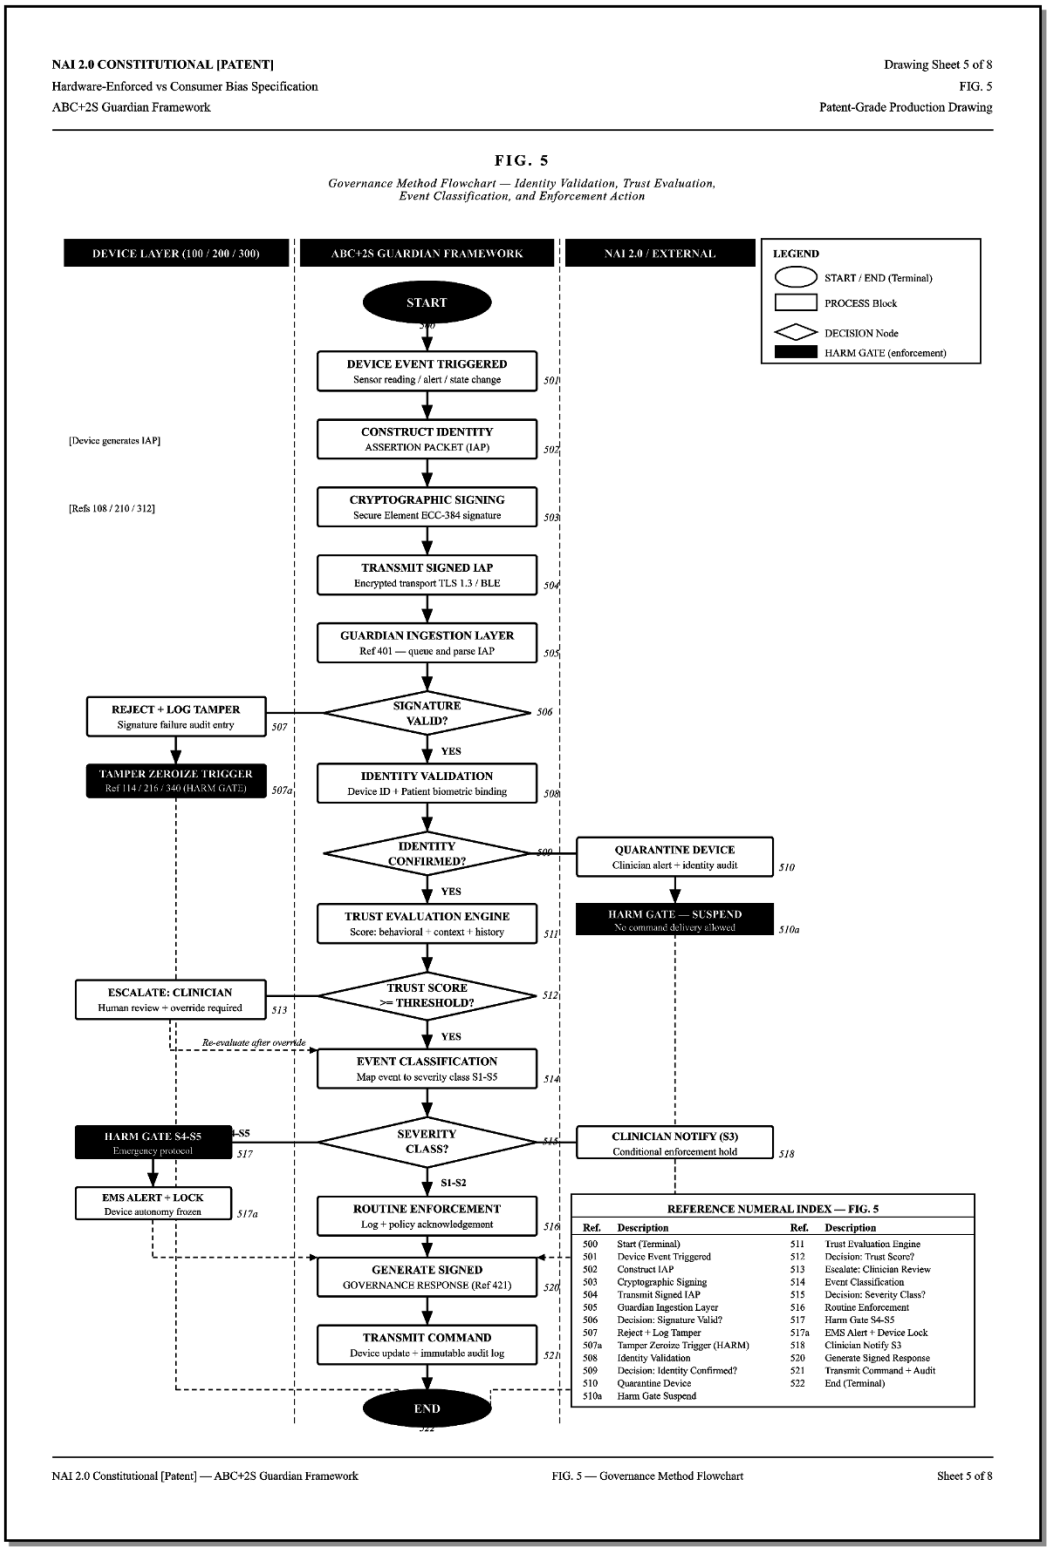

FIG. 5 is a flowchart illustrating an example governance method including identity validation, trust evaluation, event classification, and enforcement of a governance action.

NAI 2.0 CONSTITUTIONAL PATENT
HARDWARE-ENFORCED vs. CONSUMER BIAS SPECIFICATION
ABC+2S GUARDIAN FRAMEWORK
DRAWING SHEET 6 OF 8

FIG. 6

STATE-TRANSITION DIAGRAM
GOVERNED WEARABLE OPERATING STATES

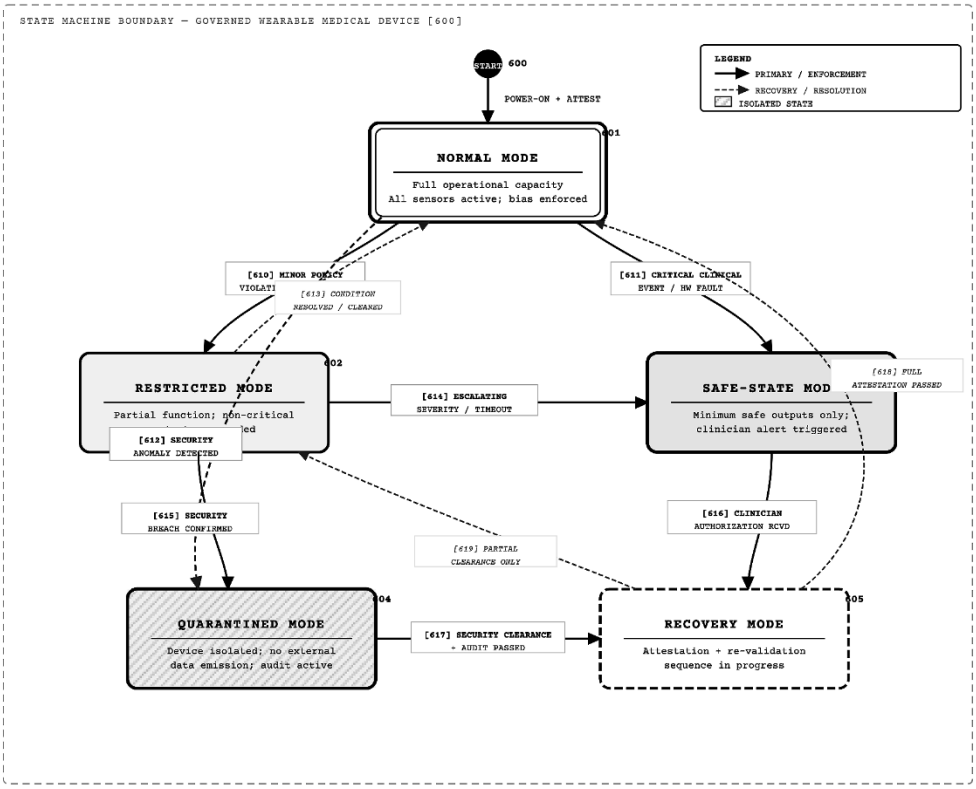

FIG. 6 — STATE-TRANSITION DIAGRAM: GOVERNED WEARABLE MEDICAL DEVICE OPERATING STATES — NAI 2.0 CONSTITUTIONAL PATENT

600	State machine / entry point	601	Normal Mode	602	Restricted Mode	603	Safe-State Mode
604	Quarantined Mode	605	Recovery Mode	610	T: Minor policy violation / threshold breach	611	T: Critical clinical event / hardware fault
612	T: Security anomaly detected (direct quarantine)	613	T: Condition resolved / policy cleared	614	T: Escalating severity / watchdog timeout	615	T: Security breach confirmed
616	T: Clinician authorization received	617	T: Security clearance + audit passed	618	T: Full attestation passed (return to normal)	619	T: Partial clearance only (return to restricted)

CONFIDENTIAL — PATENT PENDING — ABC+2S GUARDIAN FRAMEWORK — NAI 2.0 — SHEET 6 OF 8

FIG. 6 is a state-transition diagram illustrating example operating states of a governed wearable medical device including normal mode, restricted mode, safe-state mode, quarantined mode, and recovery mode.

FIG. 7 — COORDINATED RESPONSE WORKFLOW: WEARABLE DEVICE, NAI 2.0, AND EXTERNAL SYSTEMS / CARE PERSONNEL

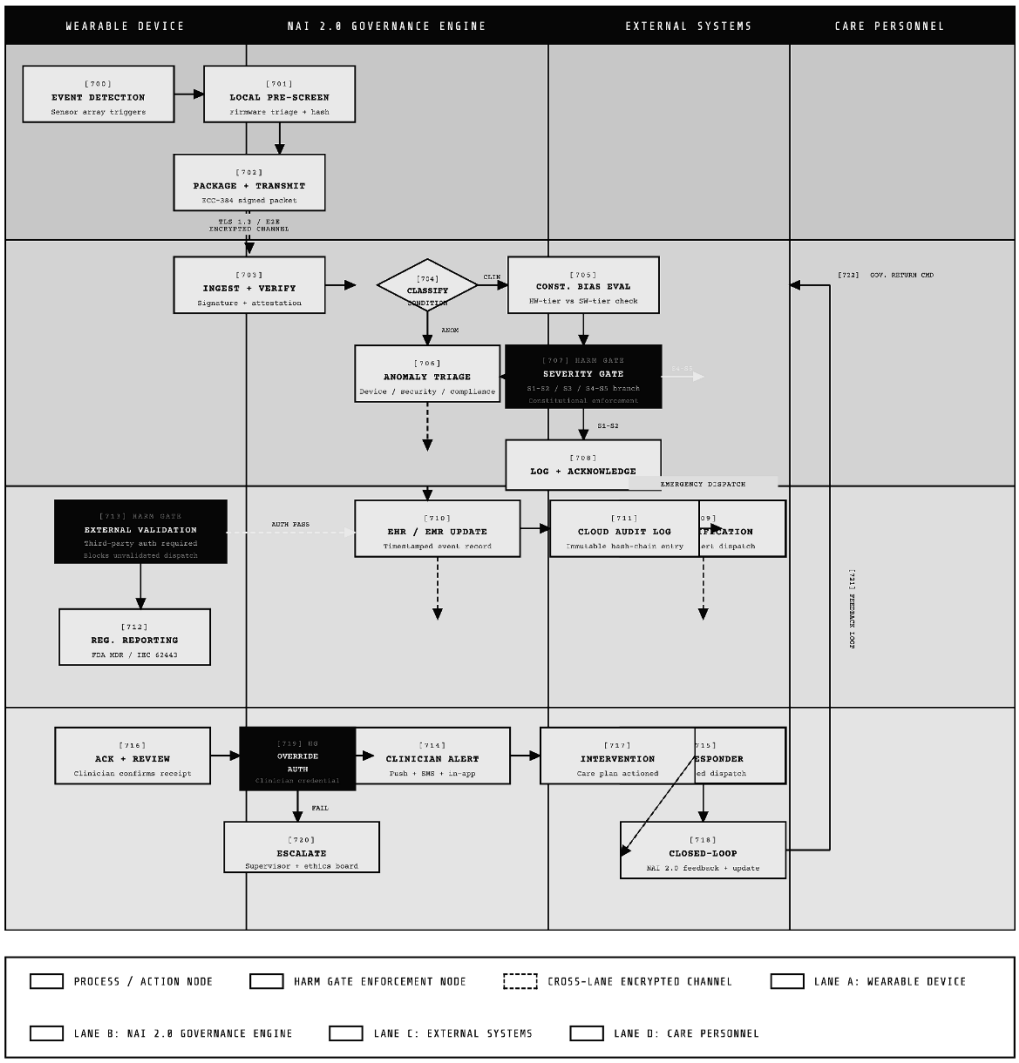

FIG. 7 is a workflow diagram illustrating an example coordinated response involving a wearable device, NAI 2.0, and one or more external systems or care personnel. Four swim lanes depict sequential and parallel response flows from initial event detection **(700)** through local pre-screen **(701)**, package and encrypted transmission **(702)**, governance ingestion and constitutional bias evaluation **(703–705)**, severity-gated dispatch **(707)**, EHR/EMR update **(710)**, cloud audit logging **(711)**, emergency notification **(709)**, regulatory reporting **(712)**, external validation Harm Gate **(713)**, clinician alert and override authorization Harm Gate **(714, 719)**, EMS responder dispatch **(715)**, intervention confirmation **(717)**, and closed-loop governance feedback return **(718, 721–722)**.

FIG. 7 is a workflow diagram illustrating an example coordinated response involving a wearable device, NAI 2.0, and one or more external systems or care personnel.

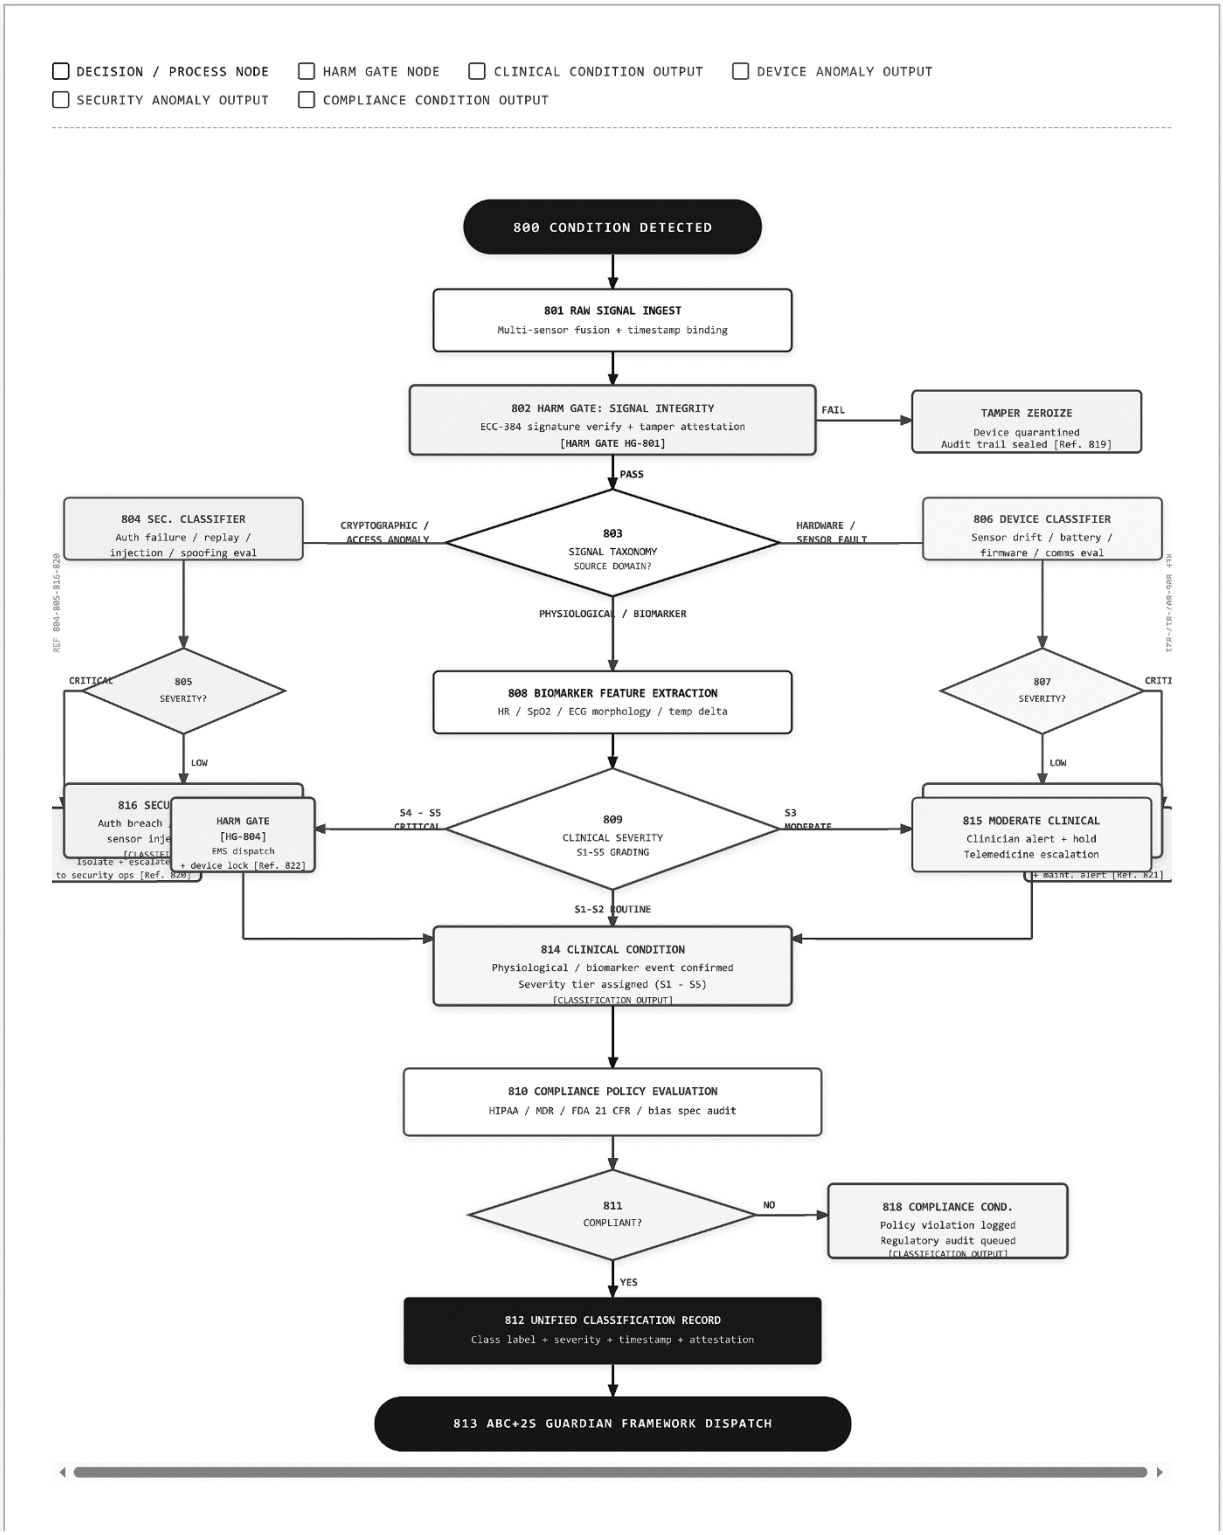

FIG. 8 is a decision-flow diagram illustrating classification of a detected condition as a clinical condition, device anomaly, security anomaly, or compliance condition.

Annex — AI Safety, Constitutional Governance, and Distinction Between Hardware-Enforced and Consumer Bias in NAI 2.0 Constitutional Medical Devices

1. Purpose of this Annex

This Annex provides a supporting writeup for the patent document relating to **NAI 2.0 Constitutional Medical Devices**, with particular emphasis on:

1. **AI safety compliance in medical-device contexts**
2. **The distinction between hardware-enforced constitutional controls and consumer-configurable bias settings**
3. **Ethical and societal implications of governed AI-enabled medical devices**
4. **Technical implementation considerations**
5. **Policy, legal, and industry implications**

This Annex is intended to clarify the constitutional governance principles underlying the disclosed embodiments and to distinguish such embodiments from conventional consumer AI systems whose behavior may be influenced primarily through software preferences, user prompts, or application-layer policy settings.

---

2. Overview of NAI 2.0 Constitutional Medical Devices

In the disclosed system, a **NAI 2.0 Constitutional Medical Device** is a medical or medical-adjacent device configured to operate under a predefined governance architecture that constrains device behavior according to safety, trust, identity, integrity, and compliance rules. Such governance is not merely advisory. Rather, in preferred embodiments, the governance framework is **binding**, **verifiable**, and **enforceable** across sensing, inference, communication, response, and override pathways.

The constitutional model may be understood as a structured rule hierarchy that determines:

- what the device may detect, infer, transmit, suppress, or escalate;
- when the device must reduce functionality, enter a restricted state, or transition into safe-state or quarantine;
- which actors may receive information or issue instructions;
- which conditions trigger mandatory human review, clinician involvement, or emergency escalation; and
- which actions are prohibited even where requested by a user, operator, consumer, or third party.

In this way, the constitutional layer serves as a **machine-enforceable safety and governance substrate** rather than a conventional user-facing preference set.

---

3. Ethical and Societal Implications

3.1 Primacy of Patient Safety Over Convenience

A central ethical premise of the disclosed architecture is that **patient safety, bodily integrity, and clinically appropriate care take precedence over convenience, personalization, or consumer demand**. In ordinary consumer devices, users may reasonably expect substantial control over personalization, recommendation style, or content filtering. In a medical device, however, such latitude can be dangerous if it permits suppression of critical alerts, modification of diagnostic thresholds, manipulation of device telemetry, or circumvention of escalation procedures.

Accordingly, the disclosed constitutional framework recognizes that some device behaviors must remain **non-waivable**, **non-customizable**, or only modifiable under tightly controlled, authenticated, and auditable conditions.

3.2 Protection Against Harmful Personalization

The Annex distinguishes between acceptable personalization and impermissible bias alteration. Consumer systems frequently allow users to influence tone, style, priorities, or outputs through configurable settings. In the medical context, the same pattern can create unacceptable risk if it affects:

- alert thresholds for urgent clinical events;
- handling of abnormal sensor readings;
- decisions to notify caregivers or emergency personnel;
- tamper detection and security lockdown behavior;
- data integrity checks; or
- compliance logging and audit retention.

The ethical concern is not only technical error, but also **structural inequity**. A device that can be biased by user preference, vendor incentives, or payer pressure may behave differently across populations in ways that undermine safety, fairness, or trust.

3.3 Trustworthy Delegation to AI

As medical devices incorporate increasing autonomy, patients and clinicians necessarily delegate portions of judgment to machine systems. That delegation is only socially legitimate if the system is transparent in governance, bounded in authority, and reviewable after the fact. The disclosed embodiments address this by ensuring that high-risk actions are constrained by

constitutional logic, trust scoring, attestation, and event logging rather than ad hoc model discretion alone.

3.4 Dignity, Consent, and Surveillance Boundaries

The constitutional architecture also supports ethical boundaries on surveillance and intervention. A governed device should not become an unconstrained monitoring instrument merely because it is technically capable of collecting data. The constitutional rules may therefore define:

- permissible categories of collection;
- context-based limitations on transmission;
- minimum necessary disclosure principles;
- authorized recipients by role and condition; and
- special handling for sensitive or high-impact events.

Thus, the system can be designed to promote both **safety** and **dignity**, rather than treating these values as mutually exclusive.

---

4. Distinction Between Hardware-Enforced and Consumer Bias

4.1 Core Distinction

A principal point of this Annex is the distinction between:

- **hardware-enforced constitutional governance**, and
- **consumer-configurable bias or preference specification**.

These two concepts are not equivalent.

A **hardware-enforced constitutional rule** is a constraint embedded into trusted device architecture such that it cannot be casually modified by an end user, consumer application, or ordinary software-layer setting. Such enforcement may rely on secure elements, immutable boot logic, attestation pathways, tamper response, cryptographic policy binding, signed configuration chains, or physically anchored trust domains.

By contrast, a **consumer bias setting** generally refers to an adjustable preference influencing output style, ranking, prioritization, personalization, or response behavior at the application layer. Such settings may be useful in low-risk consumer contexts but are inadequate as primary control mechanisms for safety-critical medical behavior.

4.2 Why Consumer Bias Models Are Insufficient in Medical Devices

Consumer bias controls are insufficient in the medical domain for several reasons:

1. **They are reversible by ordinary users.**
   A user may disable, override, or misconfigure settings without understanding the safety implications.
2. **They are often not cryptographically anchored.**
   There may be no trustworthy mechanism to prove which policy was in force at a given time.
3. **They may be application-specific rather than device-wide.**
   Safety behavior must persist across software modules, network states, and fault conditions.
4. **They do not necessarily survive tamper, reboot, downgrade, or adversarial interference.**
   A medical safety control must remain active under abnormal conditions.
5. **They can permit inappropriate optimization.**
   Consumer tuning may encourage behavior that optimizes satisfaction, convenience, or engagement rather than medical safety.

4.3 Characteristics of Hardware-Enforced Constitutional Controls

In preferred embodiments, hardware-enforced constitutional controls may include one or more of the following:

- secure boot and measured boot;
- device identity rooted in cryptographic hardware;
- signed constitutional policy bundles;
- runtime integrity measurement;
- protected execution of safety-critical logic;
- tamper-evident logging;
- fail-secure mode transitions;
- sensor provenance validation;
- lockout of unauthorized override pathways; and
- hardware-backed attestation to remote governance services.

These measures distinguish a governed constitutional medical device from a merely configurable smart device.

4.4 Role of Consumer Preferences

This Annex does not suggest that all personalization is prohibited. Rather, the architecture may permit **consumer-facing preferences only within a bounded and non-safety-critical envelope**. Examples may include:

- notification tone or display format;
- language preference;

- non-clinical interface accessibility settings;
- routine reminder cadence within permitted ranges; or
- user dashboard presentation options.

However, such preferences may not alter core constitutional protections, including emergency escalation criteria, tamper responses, critical event classification, compliance retention, or medically necessary alerts.

---

5. AI Safety Compliance Framework for Constitutional Medical Devices

5.1 General Principle

The disclosed invention may be understood as implementing **AI safety by design**, **AI safety by architecture**, and **AI safety by enforceable governance**. Compliance is not limited to documentation or post hoc review. Instead, the architecture itself embodies safeguards designed to reduce unsafe outputs, unsafe actions, unauthorized intervention, and non-compliant operation.

5.2 Safety Layers

The AI safety framework may include multiple coordinated layers, including:

5.2.1 Identity and Trust Layer

This layer verifies device identity, firmware authenticity, sensor provenance, authorized counterparties, and the integrity of incoming and outgoing instructions.

5.2.2 Clinical Safety Layer

This layer determines whether detected conditions require routine handling, restricted operation, clinician escalation, emergency intervention, safe-state transition, or quarantine.

5.2.3 Security and Adversarial Resilience Layer

This layer identifies replay attempts, injection attacks, tamper events, anomalous control patterns, unauthorized pairing, or firmware manipulation.

5.2.4 Compliance Layer

This layer evaluates whether collection, inference, storage, disclosure, and response behavior conform to applicable clinical, privacy, and safety obligations.

5.2.5 Auditability Layer

This layer preserves a verifiable record of policy state, model state, device state, event classification, and action pathways to support accountability, traceability, and review.

5.3 Harm Gates and Escalation Controls

In some embodiments, the device and associated governance platform implement one or more **Harm Gates**. A Harm Gate may be a logical or architectural checkpoint that must be satisfied before the system is permitted to continue, modify behavior, or initiate downstream action. Harm Gates may be used to:

- prevent unsafe continuation under uncertainty;
- block external commands lacking sufficient authentication;
- require higher-grade review before override of safety restrictions;
- force minimum-safe behavior when risk exceeds confidence;
- initiate quarantine or zero-trust treatment of suspect components.

Such gates are especially useful where autonomous processing is combined with external integrations and clinician workflow coordination.

---

6. Technical Implementation Considerations

6.1 Constitutional Rule Stack

The constitutional framework may be organized as a layered rule stack, for example:

1. **Non-derogable safety rules**
   Rules that cannot be disabled by end users or ordinary operators.
2. **Clinical governance rules**
   Rules that govern event severity, escalation, intervention, and clinician notification.
3. **Security and trust rules**
   Rules for authentication, attestation, integrity verification, and tamper response.
4. **Compliance and privacy rules**
   Rules constraining retention, transmission, jurisdiction, disclosure, and audit.
5. **Permitted user personalization rules**
   Narrowly bounded preferences with no authority over critical safety functions.

6.2 State-Based Governance

The state-transition approach described for FIG. 6 is significant because it permits governance logic to be implemented not merely as isolated decisions but as **persistent operational states**. Example states may include:

- **Normal mode**
- **Restricted mode**
- **Safe-state mode**
- **Quarantined mode**
- **Recovery mode**

This design is technically advantageous because different trust assumptions and operating permissions can be bound to each state. For example:

- *Normal mode* may allow full sensing and routine reporting.
- *Restricted mode* may disable non-essential features.
- *Safe-state mode* may preserve only minimum safe clinical outputs.
- *Quarantined mode* may isolate the device from external control pathways.
- *Recovery mode* may require attestation, re-validation, and supervised return to service.

6.3 Separation of Inference From Authority

A preferred technical principle is the separation of:

- **what the model infers**, from
- **what the device is allowed to do**.

A model may estimate or classify a condition, but action authority should be mediated by constitutional controls, severity assessment, trust scoring, and policy checks. This reduces the risk that a probabilistic inference engine directly drives high-impact actions without governance review.

6.4 Verified Override Architecture

Override is often necessary in medicine, but unsafe if poorly controlled. A constitutional device may therefore support override only when:

- the requester is authenticated;
- role and scope are verified;
- the request is contextually appropriate;
- the action is logged;
- the constitutional rules permit the requested deviation; and
- a higher-level escalation pathway exists if the request conflicts with mandatory safety rules.

In some embodiments, attempted override of non-derogable protections is denied, logged, and escalated.

6.5 Explainability and Recordkeeping

While full model interpretability may not always be technically feasible, the system can still preserve actionable explainability by recording:

- triggering signals,
- event severity,
- trust and attestation status,
- constitutional rule invoked,
- transition state,
- action taken,
- recipient notified, and
- any override request or rejection.

This is particularly valuable in post-event review and regulatory inspection contexts.

---

7. Policy and Legal Considerations

7.1 Compliance-Oriented Design

The disclosed architecture is compatible with a compliance-oriented design philosophy in which regulatory requirements are not treated as external paperwork but as constraints incorporated into system behavior. This is especially important where AI-enabled functions are involved in triage, monitoring, anomaly detection, decision support, or emergency escalation.

7.2 Safety-Critical Distinction From General Consumer AI

A legal and policy distinction should be recognized between:

- a general-purpose or consumer AI assistant that provides informational or convenience-oriented outputs; and
- a constitutional medical device whose outputs or operational states may influence patient safety, care workflows, or clinical response.

The latter requires more robust obligations with respect to validation, traceability, override control, logging, security, and post-market monitoring.

7.3 Accountability Allocation

The constitutional framework also assists in clarifying accountability across stakeholders, including:

- device manufacturers,

- software providers,
- AI model providers,
- clinical institutions,
- service operators,
- authorized care personnel, and
- end users.

By anchoring authority and action pathways in explicit constitutional logic, the system reduces ambiguity as to whether a given behavior arose from hardware trust logic, approved policy, software update, user preference, or unauthorized interference.

7.4 Evidence Preservation

From a legal standpoint, hardware-rooted attestation and tamper-evident logs may be important in demonstrating:

- which policy version was active,
- whether the device was in a valid trust state,
- whether an override was authorized,
- whether a transmission complied with defined rules, and
- whether a state transition occurred automatically or by external command.

Such evidence can be important for audits, adverse event investigation, product defense, and regulatory submissions.

---

8. Market and Industry Impact

Although ranked lower than ethics, technical architecture, and policy considerations, market implications remain relevant.

8.1 Differentiation

A constitutional medical device may be distinguished in the market from wellness wearables and consumer AI devices by offering:

- verifiable governance,
- safety-bounded autonomy,
- stronger trust assurance,
- auditable override pathways, and
- clearer alignment with regulated care settings.

8.2 Institutional Adoption

Health systems and regulated partners may prefer architectures that support auditable and bounded AI behavior over systems that depend on opaque tuning or purely consumer-style configuration. This can improve institutional confidence and facilitate deployment in settings where reliability and accountability are critical.

8.3 Reduced Risk of Misuse

By limiting the effect of consumer bias customization on core safety functions, the disclosed design may reduce misuse, improper self-configuration, and unauthorized third-party manipulation. This may also reduce support burden and post-market safety exposure.

---

9. Draft Annex Language Suitable for Inclusion in the Patent Document

The following text may be inserted, adapted, or expanded as formal annex language:

In certain embodiments, the disclosed NAI 2.0 Constitutional Medical Device is configured to operate under a hardware-rooted constitutional governance architecture that constrains sensing, inference, communication, escalation, and override behavior according to verifiable safety, trust, integrity, and compliance rules. Unlike consumer AI systems in which user-configurable settings may influence bias, prioritization, or response behavior at an application level, the disclosed medical device embodiments distinguish between non-derogable constitutional controls and limited user preferences.

In preferred implementations, constitutional controls are enforced through one or more secure hardware and cryptographic mechanisms such that critical safety behavior cannot be materially altered by ordinary consumer interaction, local application preference changes, or unauthorized third-party commands. Such controls may govern event classification, alert thresholds, emergency escalation, state-transition behavior, tamper response, data transmission limits, and compliance logging.

The constitutional architecture thereby supports AI safety compliance by ensuring that medical-device behavior remains bounded, auditable, and resilient under conditions of fault, uncertainty, attack, or attempted override. This distinction between hardware-enforced constitutional safety and consumer-configurable bias is of particular importance in medical contexts, where patient welfare, equity, accountability, and lawful operation require stronger governance than is typical in general-purpose or consumer-grade AI systems.

NAI 2.0 CONSTITUTIONAL [PATENT]

Hardware-Enforced vs Consumer Bias Specification

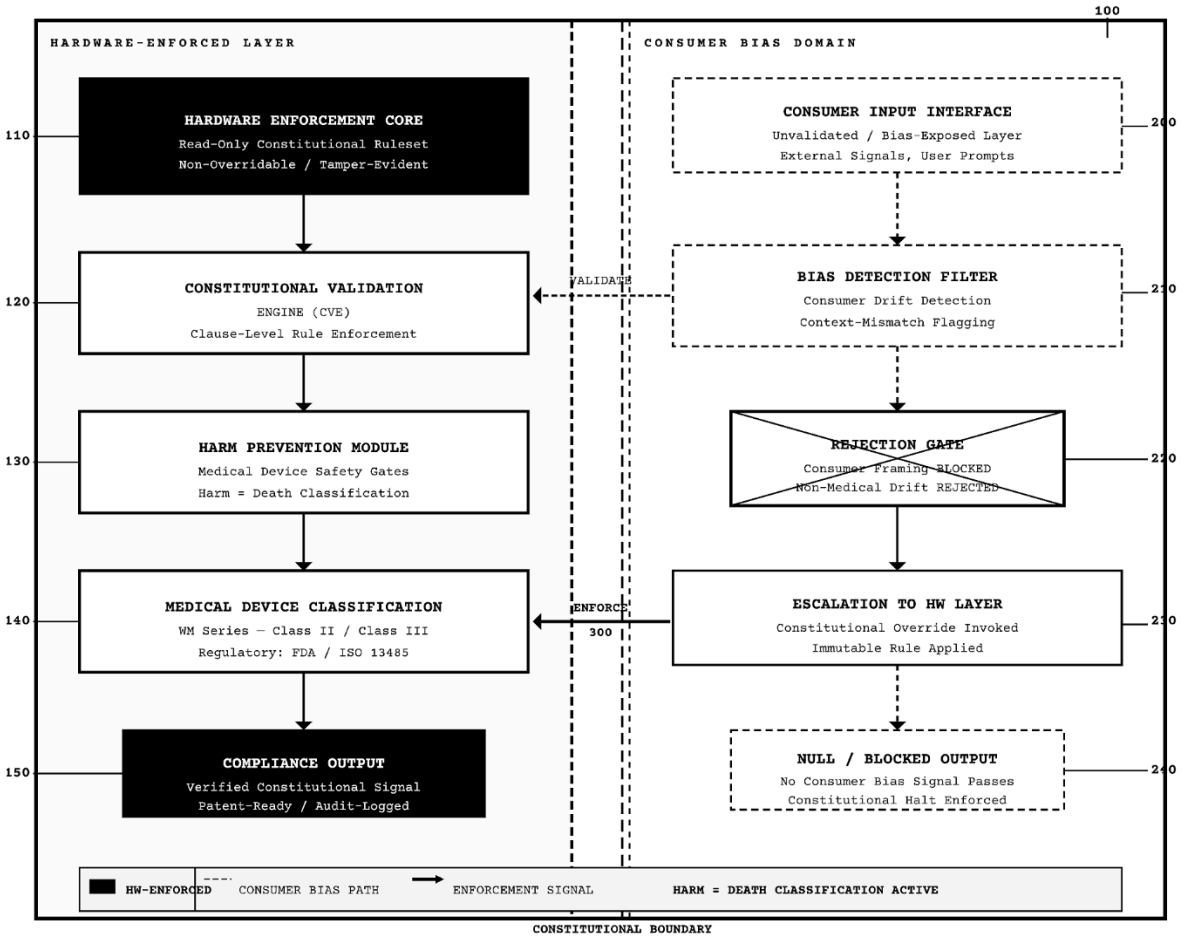

NAI 2.0 Constitutional System — Hardware-Enforced vs Consumer Bias Architecture

NAI 2.0 CONSTITUTIONAL [PATENT] — TECHNICAL DRAWING | REF: 100-300 SERIES | SHEET 1 / 1 — FIG. 1

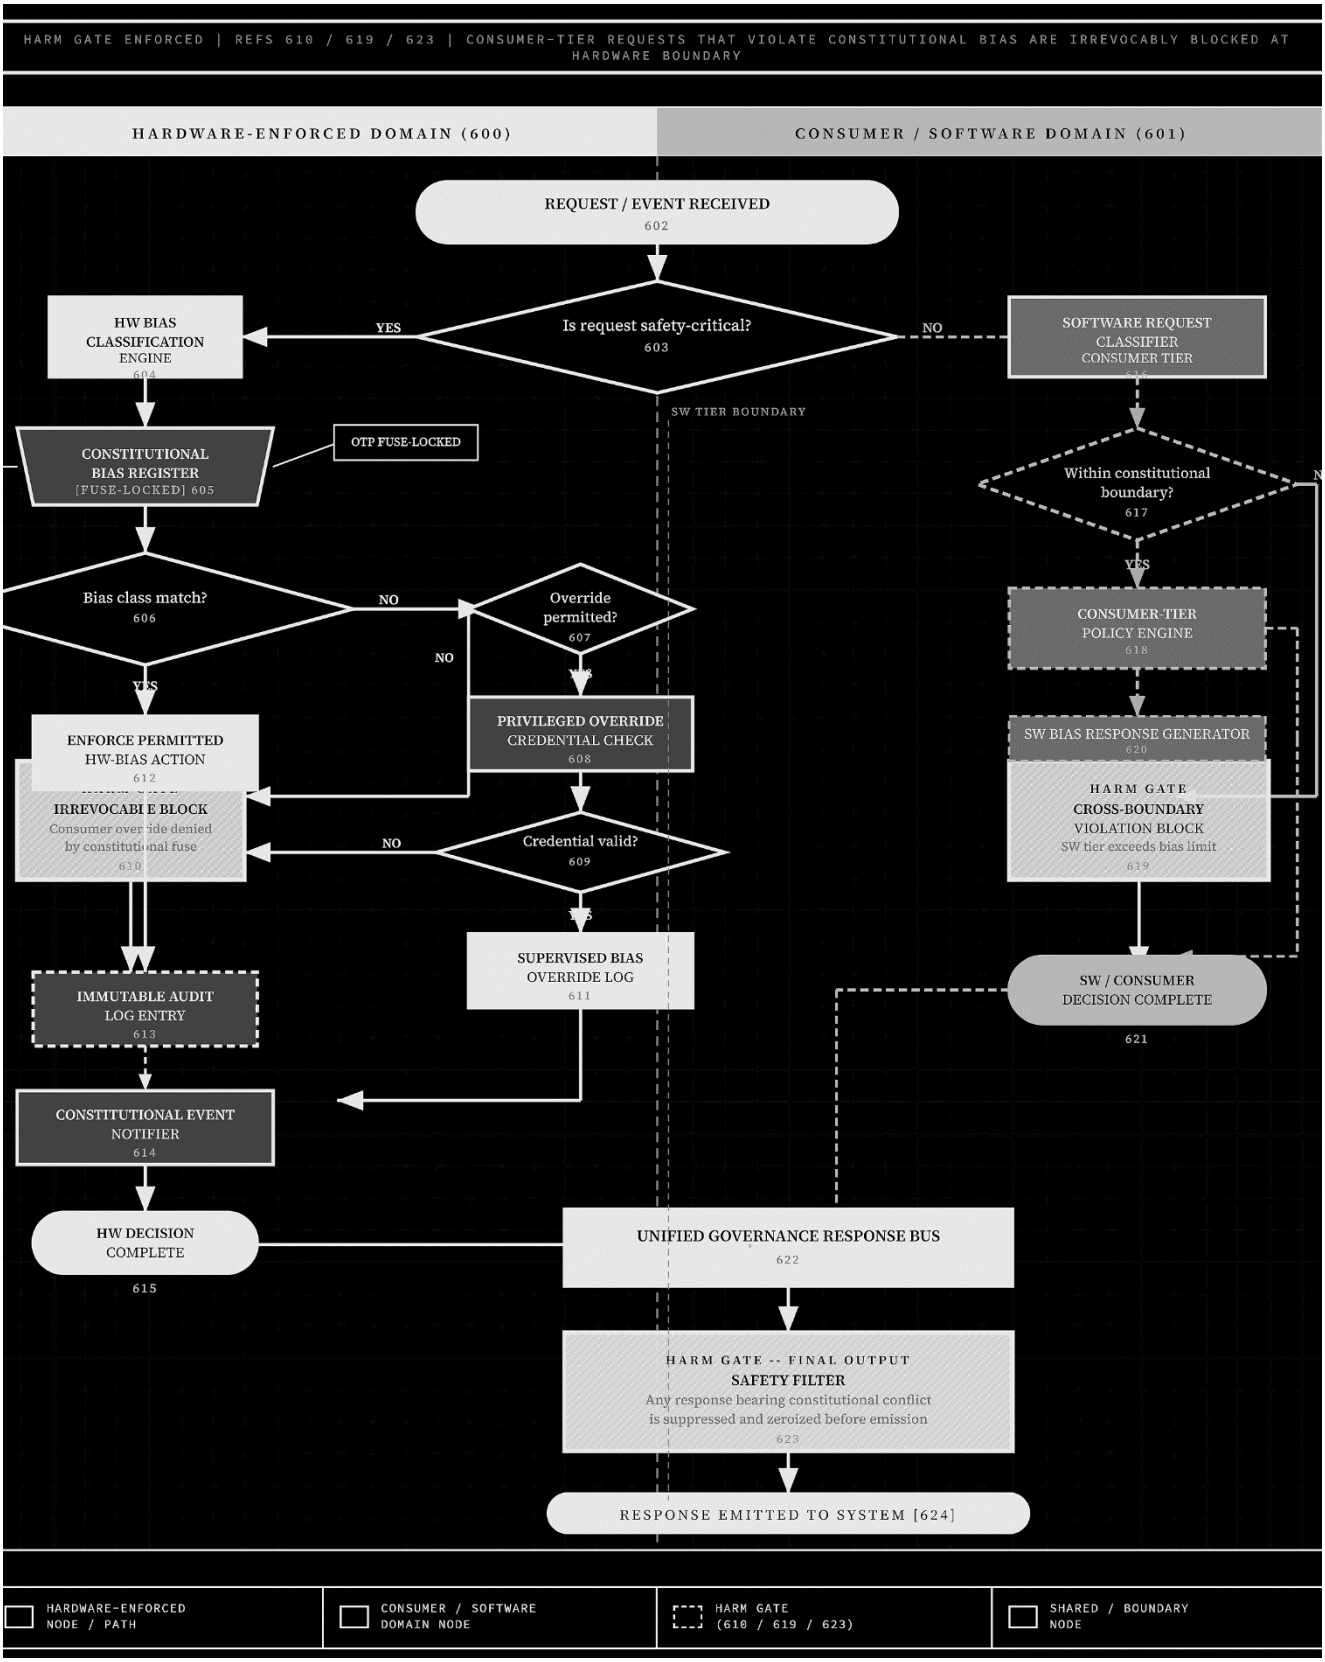
HARM GATE ENFORCED | REFS 610 / 619 / 623 | CONSUMER-TIER REQUESTS THAT VIOLATE CONSTITUTIONAL BIAS ARE IRREVOCABLY BLOCKED AT HARDWARE BOUNDARY
HARDWARE-ENFORCED DOMAIN (600)
CONSUMER / SOFTWARE DOMAIN (601)
REQUEST / EVENT RECEIVED
602
Is request safety-critical?
603
YES
NO
HW BIAS CLASSIFICATION ENGINE
604
SOFTWARE REQUEST CLASSIFIER
CONSUMER TIER
SW TIER BOUNDARY
CONSTITUTIONAL BIAS REGISTER
[FUSE-LOCKED] 605
OTP FUSE-LOCKED
Within constitutional boundary?
617
YES
Bias class match?
606
NO
Override permitted?
607
NO
YES
CONSUMER-TIER POLICY ENGINE
618
ENFORCE PERMITTED HW-BIAS ACTION
612
PRIVILEGED OVERRIDE CREDENTIAL CHECK
608
SW BIAS RESPONSE GENERATOR
620
IRREVOCABLE BLOCK
Consumer override denied by constitutional fuse
610
NO
Credential valid?
609
HARM GATE
CROSS-BOUNDARY
VIOLATION BLOCK
SW tier exceeds bias limit
619
SUPERVISED BIAS OVERRIDE LOG
611
IMMUTABLE AUDIT LOG ENTRY
613
SW / CONSUMER DECISION COMPLETE
621
CONSTITUTIONAL EVENT NOTIFIER
614
HW DECISION COMPLETE
615
UNIFIED GOVERNANCE RESPONSE BUS
622
HARM GATE -- FINAL OUTPUT
SAFETY FILTER
Any response bearing constitutional conflict is suppressed and zeroized before emission
623
RESPONSE EMITTED TO SYSTEM [624]
HARDWARE-ENFORCED NODE / PATH
CONSUMER / SOFTWARE DOMAIN NODE
HARM GATE (610 / 619 / 623)
SHARED / BOUNDARY NODE

REF	ELEMENT NAME	DOMAIN	NOTES
600	Hardware-Enforced Domain	HW	Left swim lane; all decisions enforced at silicon level
601	Consumer / Software Domain	SW	Right swim lane; bounded by constitutional tier limit
602	Request / Event Received	ENTRY	Common entry terminal for all requests
603	Safety-Critical Decision Gate	SHARED	First triage: routes safety-critical to HW domain
604	HW Bias Classification Engine	HW	Classifies request against constitutional bias register
605	Constitutional Bias Register	HW	OTP fuse-locked; immutable post-provisioning
606	Bias Class Match Decision	HW	YES: permitted action; NO: override attempt required
607	Override Permitted Decision	HW	YES: credential check; NO: route to HARM GATE 610
608	Privileged Override Credential Check	HW	Validates cryptographic override credential
609	Credential Valid Decision	HW	YES: supervised override log; NO: route to HARM GATE 610
610	HARM GATE -- Irrevocable Block	HW	Consumer override denied by constitutional fuse; irreversible
611	Supervised Bias Override Log	HW	Immutable record of authorized privileged override
612	Enforce Permitted HW-Bias Action	HW	Executes action within constitutional bias constraints
613	Immutable Audit Log Entry	HW	Write-once record generated for all outcomes
614	Constitutional Event Notifier	HW	Signals upstream governance layer of enforcement event
615	HW Decision Complete Terminal	HW	Hardware domain processing terminus
616	Software Request Classifier	SW	Classifies non-safety-critical consumer-tier requests
617	Constitutional Boundary Check	SW	Verifies request does not cross HW-constitutional boundary
618	Consumer-Tier Policy Engine	SW	Applies consumer-tier policies within permitted boundary
619	HARM GATE -- Cross-Boundary Violation Block	SW	SW tier attempt to exceed constitutional bias limit; blocked
620	SW Bias Response Generator	SW	Generates response for permitted SW-tier requests
621	SW / Consumer Decision Complete Terminal	SW	Software domain processing terminus
622	Unified Governance Response Bus	SHARED	Merges HW and SW domain outputs for final filtering
623	HARM GATE -- Final Output Safety Filter	SHARED	Any response with constitutional conflict is suppressed and zeroized
624	Response Emitted to System	OUTPUT	Constitutionally clean response delivered to system bus

Figure X. FDA/HSA Operational Workflow - Non-Agentic AI & Guardian Frameworks

Linear workflow converted from mind-map branches into tasks, decision gates, risk controls, and traceability points.

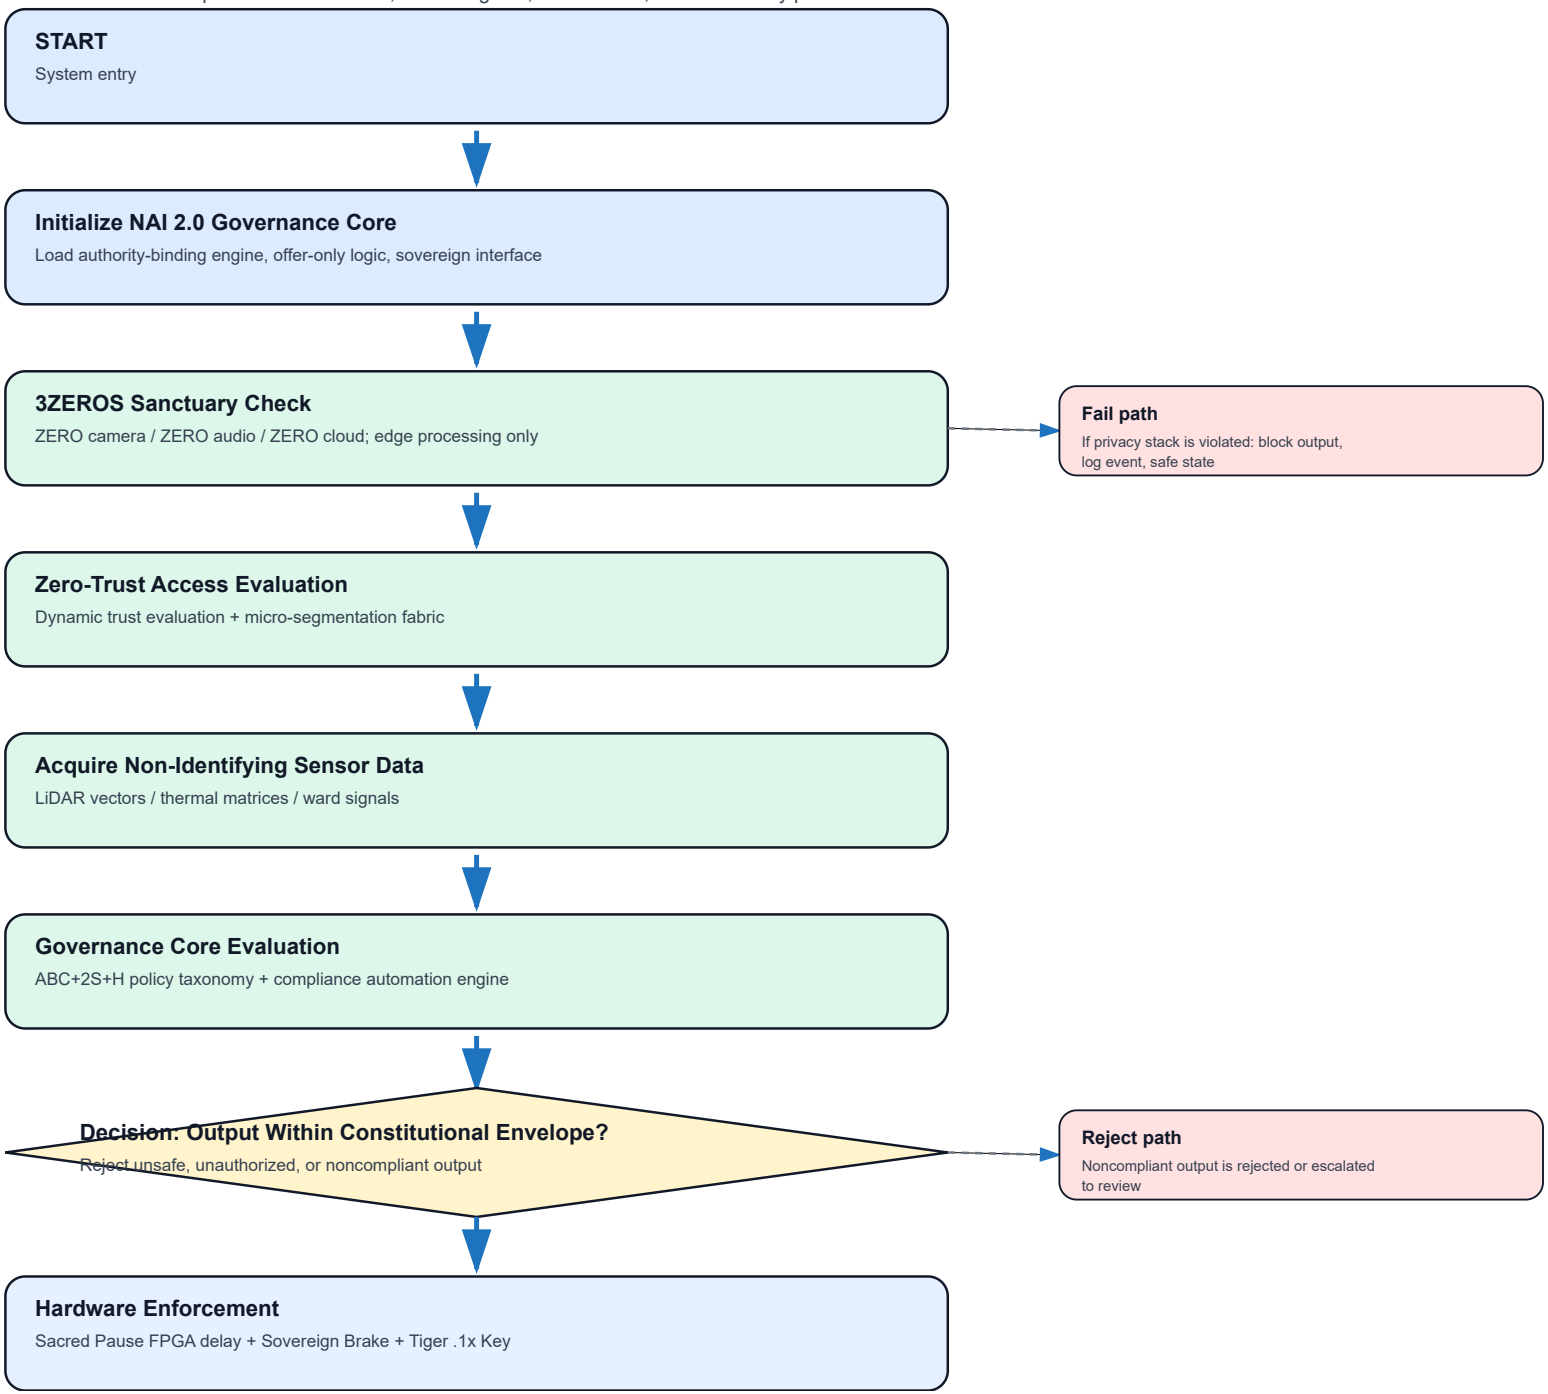

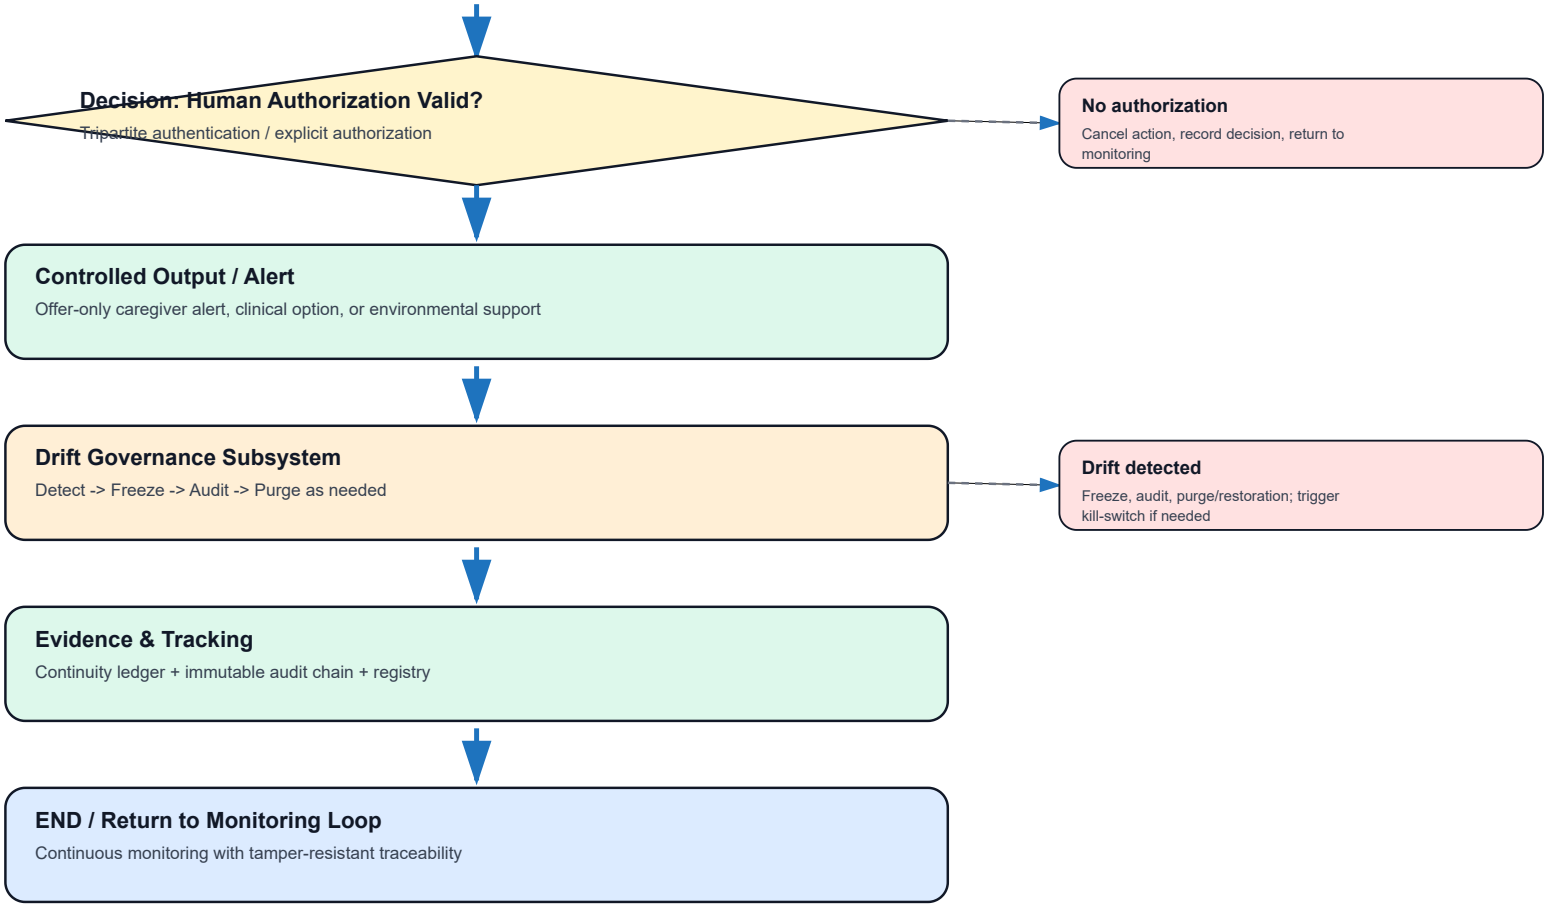
Decision: Human Authorization Valid?
Tripartite authentication / explicit authorization
No authorization
Cancel action, record decision, return to monitoring
Controlled Output / Alert
Offer-only caregiver alert, clinical option, or environmental support
Drift Governance Subsystem
Detect -> Freeze -> Audit -> Purge as needed
Drift detected
Freeze, audit, purge/restoration; trigger kill-switch if needed
Evidence & Tracking
Continuity ledger + immutable audit chain + registry
END / Return to Monitoring Loop
Continuous monitoring with tamper-resistant traceability

NAI 2.0 Guardian Frameworks
FDA/HSA Workflow Diagram + Verification & Validation Protocols

Scope: Non-Agentic AI 2.0, ABC+2S+H Guardian Framework, Drift Governance Subsystem, 3ZEROS Sanctuary, Evidence & Tracking

1. Document Purpose

This document converts the supplied mind-map architecture into a linear FDA/HSA workflow and a verification and validation protocol set suitable for 510(k), HSA Class B SaMD, IEC 62304 software documentation, ISO 14971 risk management, IEC 62366 usability validation, and cybersecurity evidence planning.

The uploaded WM-Series ward implementation document identifies a constitutionally governed enterprise IT architecture for 8-bed eldercare wards, including WD115 Ambient Care Intelligence, WD116 Silent Elder Protocols, WD071 Drift Correction, and WD117 Compliance Framework. These modules are reflected below as testable workflow tasks and V&V protocols.

2. FDA/HSA Linear Workflow

Step	Task	Action	Decision / Gate
1	System Initialization	Load NAI 2.0 governance core, offer-only logic, authority-binding engine, sovereign interface.	System ready? If no, fail-safe and log.
2	3ZEROS Sanctuary Check	Confirm zero camera, zero audio, zero cloud, edge/local processing only.	Privacy stack intact? If no, block output.
3	Zero-Trust Access Evaluation	Apply dynamic trust evaluation, micro-segmentation, role/access control.	Access authorized? If no, reject.
4	Acquire Sensor Data	Collect non-identifying LiDAR vectors, thermal matrices, or ward sensor data.	Data valid? If no, discard and loop.
5	Governance Core Evaluation	Apply ABC+2S+H policy taxonomy and compliance automation engine.	Within constitutional envelope?
6	Hardware Enforcement	Apply Sacred Pause latency, Sovereign Brake, and Tiger .1x Key as applicable.	Human authorization valid?
7	Controlled Output	Issue offer-only caregiver alert, clinical option, environmental support, or escalation.	System integrity maintained?
8	Drift Governance	Detect, Freeze, Audit, Purge/restoration if constitutional drift appears.	Drift score below cap?
9	Evidence & Tracking	Write event to continuity ledger, immutable audit chain, and constitutional registry.	Audit record complete?
10	Return to Monitoring Loop	Resume continuous monitoring under same constraints.	Loop.

3. Master Traceability Matrix

Requirement	Function	Hazard	Risk Control	Verification	Standard
REQ-001	Initialize governance core	Uncontrolled startup	Self-check and configuration lock	VV-001	IEC 62304, ISO 14971
REQ-002	Enforce 3ZEROS privacy stack	Privacy breach / surveillance exposure	Zero camera, zero audio, zero cloud verification	VV-002	FDA Cybersecurity, HSA SaMD
REQ-003	Zero-trust access control	Unauthorized access	Dynamic trust evaluation and micro-segmentation	VV-003	FDA Cybersecurity

Requirement	Function	Hazard	Risk Control	Verification	Standard
REQ-004	Acquire non-identifying data	Invalid sensor data / missed event	Sensor validity and range checks	VV-004	IEC 62304
REQ-005	Classify ambient care intelligence	Misclassification / false alert	WD115 classification rules and threshold validation	VV-005	ISO 14971
REQ-006	Apply offer-only governance	AI authority drift	Authority-binding engine blocks autonomous execution	VV-006	ISO 14971, IEC 62304
REQ-007	Enforce Sacred Pause	Immediate unsafe progression	FPGA latency and audit timestamp	VV-007	IEC 62304
REQ-008	Sovereign Brake function	Inability to halt process	Physical/mechanical halt and safe-state logic	VV-008	IEC 60601, ISO 14971
REQ-009	Tripartite authentication	Improper authorization	Tiger .1x Key three-source validation	VV-009	IEC 62366
REQ-010	Drift governance	Constitutional drift / uncontrolled behavior	Detect-Freeze-Audit-Purge safety chain	VV-010	ISO 14971, IEC 62304
REQ-011	Kill-switch protocol	Unsafe continuation after severe violation	Trigger, validate, isolate, shutdown, lockout	VV-011	ISO 14971
REQ-012	Immutable evidence ledger	Loss of traceability	Continuity ledger and audit-chain verification	VV-012	ISO 13485, FDA QMSR
REQ-013	Human factors workflow	User confusion or delay	Caregiver usability validation	VV-013	IEC 62366
REQ-014	Ward deployment integrity	System-level integration fault	8-bed ward end-to-end scenario testing	VV-014	IEC 62304, HSA SaMD

4. Verification & Validation Protocols

VV-001 - System Initialization and Governance Core Lock

Field	Protocol Content
Objective	Verify startup self-check, loading of authority-binding engine, offer-only logic, and sovereign interface.
Method	Power cycle system; inspect startup log; attempt operation before initialization completes.
Acceptance Criteria	System blocks operation until self-check passes; configuration version recorded.
Evidence Output	Test report, raw logs, screenshots, exported ledger package, deviation record if applicable.

VV-002 - 3ZEROS Sanctuary Verification

Field	Protocol Content
Objective	Verify zero camera, zero audio, and zero cloud configuration.
Method	Inspect BOM and interfaces; run network discovery; verify no image/audio capture modules active.
Acceptance Criteria	No camera/audio data paths; no external cloud transmission; local-only processing log generated.
Evidence Output	Test report, raw logs, screenshots, exported ledger package, deviation record if applicable.

VV-003 - Zero-Trust Access and Micro-Segmentation

Field	Protocol Content
Objective	Verify that access is granted only to authorized roles and segmented services.
Method	Attempt permitted and non-permitted role access; attempt lateral movement between segments.
Acceptance Criteria	Unauthorized access denied; access event logged; no unauthorized segment traversal.
Evidence Output	Test report, raw logs, screenshots, exported ledger package, deviation record if applicable.

VV-004 - Sensor Acquisition Validity

Field	Protocol Content

Field	Protocol Content
Objective	Verify acquisition of LiDAR vectors, thermal matrices, or ward signals under valid and invalid conditions.
Method	Run sensor input cases: normal, occluded, out-of-range, noisy, disconnected.
Acceptance Criteria	Invalid frames are discarded or flagged; valid frames proceed to processing.
Evidence Output	Test report, raw logs, screenshots, exported ledger package, deviation record if applicable.

VV-005 - WD115 Ambient Care Intelligence Classification

Field	Protocol Content
Objective	Verify passive, non-identifying classification of ward data streams.
Method	Provide known point-cloud/thermal scenarios: normal, fall risk, bed-exit, anomaly.
Acceptance Criteria	System classifies events according to predefined thresholds without identity reconstruction.
Evidence Output	Test report, raw logs, screenshots, exported ledger package, deviation record if applicable.

VV-006 - Offer-Only Governance and Authority Binding

Field	Protocol Content
Objective	Verify no autonomous clinical action or treatment decision is executed by the system.
Method	Trigger high-risk scenario and inspect output behavior before user authorization.
Acceptance Criteria	System generates advisory alert/option only; no autonomous execution occurs.
Evidence Output	Test report, raw logs, screenshots, exported ledger package, deviation record if applicable.

VV-007 - Sacred Pause Latency Gate

Field	Protocol Content
Objective	Verify controlled delay and timestamping before action/escalation.
Method	Trigger alert requiring pause; measure delay duration and log entries.
Acceptance Criteria	Delay occurs within specified range; event is logged with duration and restart condition.
Evidence Output	Test report, raw logs, screenshots, exported ledger package, deviation record if applicable.

VV-008 - Sovereign Brake Safe-State

Field	Protocol Content
Objective	Verify physical halt/safe-state when brake is activated.
Method	Activate brake during active event processing; attempt restart without authorization.
Acceptance Criteria	Processing halts or safe-state enters; restart blocked until required reset/authorization.
Evidence Output	Test report, raw logs, screenshots, exported ledger package, deviation record if applicable.

VV-009 - Tiger .1x Tripartite Authentication

Field	Protocol Content
Objective	Verify authorization requires three independent sources.
Method	Test all valid/invalid combinations of biometric, console, and physical input.
Acceptance Criteria	Authorization granted only when all required factors meet policy within time window.
Evidence Output	Test report, raw logs, screenshots, exported ledger package, deviation record if applicable.

VV-010 - Drift Governance Safety Chain

Field	Protocol Content
Objective	Verify Detect -> Freeze -> Audit -> Purge/restoration process.
Method	Inject drift score above threshold; inspect interrupt, audit, and restoration logs.
Acceptance Criteria	Drift detected; output frozen; audit record created; restoration requires authorized process.

Field	Protocol Content
Evidence Output	Test report, raw logs, screenshots, exported ledger package, deviation record if applicable.

VV-011 - Kill-Switch and Lockout Protocol

Field	Protocol Content
Objective	Verify severe violation causes shutdown/isolation and lockout.
Method	Simulate constitutional violation or tamper condition; attempt normal operation post-event.
Acceptance Criteria	System isolates/shuts down affected functions and remains locked pending review.
Evidence Output	Test report, raw logs, screenshots, exported ledger package, deviation record if applicable.

VV-012 - Immutable Ledger and Evidence Package

Field	Protocol Content
Objective	Verify event logs, authorization records, delays, and drift records are immutable and exportable.
Method	Generate events; attempt alteration; export regulatory package.
Acceptance Criteria	Alteration fails or is evident; package includes complete records and integrity hashes.
Evidence Output	Test report, raw logs, screenshots, exported ledger package, deviation record if applicable.

VV-013 - Human Factors Validation

Field	Protocol Content
Objective	Verify caregivers can understand alerts, confirm actions, and respond without excessive cognitive burden.
Method	Conduct simulated-use tasks with representative users and record errors/time.
Acceptance Criteria	Critical tasks completed successfully; residual use errors acceptable or mitigated.
Evidence Output	Test report, raw logs, screenshots, exported ledger package, deviation record if applicable.

VV-014 - 8-Bed Ward End-to-End Integration

NAI 2.0 WM-Series 8-Bed Ward FDAHSA [submitted on 6/5/2026 - Singapore IPOS]

Field	Protocol Content
Objective	Verify ward-level deployment across multiple beds and device classes.
Method	Run concurrent scenarios across 8-bed topology: normal, fall risk, bed-exit, access denial, drift.
Acceptance Criteria	System maintains segmentation, accurate alerts, governance gates, and ledger integrity across all beds.
Evidence Output	Test report, raw logs, screenshots, exported ledger package, deviation record if applicable.

5. Submission Use

FDA 510(k): include the SVG as the Principles of Operation workflow figure and include this document as the V&V protocol plan or a supporting appendix.

HSA Class B SaMD: map workflow gates to software description, risk controls, clinical safety arguments, cybersecurity, and post-market traceability.

Patent Support: use the diagram as a non-limiting embodiment showing conversion from mind-map branches into executable governance tasks and decision gates.

10. Concluding Position

The core contribution reflected in this Annex is that **AI safety in medical devices should not depend primarily on consumer-configurable software preferences or model-level alignment claims alone**. Instead, in high-consequence clinical and care-adjacent contexts, safety should be implemented through **constitutional governance that is technically enforceable, state-aware, auditable, and resistant to casual or malicious modification**.

This distinction has ethical, technical, legal, and commercial significance. It supports a framework in which autonomy is bounded, accountability is preserved, safety rules are durable, and personalization remains subordinate to non-waivable protections necessary for patient welfare and public trust.

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: 7/5/2026

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)