ODA3 INSTITUTE

EXECUTIVE SUMMARY

PAI-SF™ v1.0 Executive Summary

Physical AI Security Framework — Normative-Adjacent Summary

ODA3-2026-07-EXS-PAISF-007 | Public-Facing — Publication Support Document

Document Control

Doc ID	ODA3-2026-07-EXS-PAISF-007
Doc Type	Executive Summary
Framework Tag	PAI-SF™ v1.0
Governing Licence	GAISSF Ecosystem Licence (GEL v1.0)
Access Classification	Open / Public — normative-adjacent summary content; no Controlled-asset material
Publication State	Publication Candidate — pending formal publish action
Publish Date / Review Cycle	01 August 2026 (content finalized); publication date to be confirmed
Classification	Public-Facing — Publication Support Document

PAI-SF™ v1.0 Executive Summary

Physical AI Security Framework — Normative-Adjacent Summary

Executive Position

PAI-SF™ defines the AI-security assurance layer for systems where model outputs, sensor interpretation, autonomy decisions, and actuation can create physical consequence. It does not replace functional safety, product safety, sector regulation, or certification regimes for aviation, automotive, medical devices, industrial systems, or robotics.

Kinetic Zero Trust is the core assurance principle underlying this framework: AI-generated physical actions should not be trusted solely because the model is authorized, confident, or operating within a known deployment context. Physical commands must be validated against independent safety boundaries, operational constraints, sensor confidence, environmental conditions, and human-override expectations before execution.

Purpose and Audience

This executive summary provides a concise orientation to PAI-SF™ v1.0 for CISOs, Security Architects, AI Governance Leads, Compliance Officers, standards participants, safety-facing technical leaders, and risk committees. It is not a substitute for the Framework Standard or Control Catalogue; it is the short citable entry point for understanding why the framework exists and how it should be used.

PAI-SF™ sits inside the GAISSF Ecosystem as ODA3 Institute's fourth core framework: GAISSF™ governs and assures AI systems broadly, UAIF™ classifies AI incidents, AI-IRF™ structures response, and PAI-SF™ addresses physical-AI security assurance where cyber, model, autonomy, sensor, and actuator behavior meet physical consequence.

Why PAI-SF Exists

- Existing AI governance frameworks generally stop before physical actuation, where model behavior becomes movement, force, process control, or environmental manipulation.

- Functional safety and product-safety regimes remain essential, but they were not designed to cover adversarial model behavior, sensor manipulation, runtime drift, or AI-specific command-path assurance as security questions.
- PAI-SF™ fills this gap by defining controls, evidence expectations, and assurance boundaries for embodied and cyber-physical AI systems.

Scope

PAI-SF™ applies to AI-enabled systems whose outputs directly or indirectly cause, guide, constrain, or fail to prevent physical action. Covered contexts include robotics, industrial automation, drones, autonomous vehicles, medical robotics, smart infrastructure, embodied AI agents, and AI-connected IoT or OT systems with physical effectors.

Pure software AI risks remain in GAISSF™/UAIF™/AI-IRF™ unless they create a foreseeable and direct physical, kinetic, operational, infrastructure, or human-safety consequence.

Canonical Architecture

PAI-SF™ v1.0 contains 12 domains and 41 controls. The domains are intended to be read together as an assurance chain: perception, attestation, autonomy boundaries, command safety, safe state, override, telemetry, recovery, supply chain, validation, safety interface, and operating environment.

No.	Domain	Executive Meaning
1	Sensor and Perception Integrity	Validates sensor trustworthiness before model outputs drive physical action.
2	Actuator and Kinetic Command Safety	Constrains commands, timing, replay risk, and actuator divergence.
3	Edge Hardware and Model Attestation	Establishes boot, model, hardware, and attestation-key trust.
4	Autonomy Boundaries	Defines and enforces permitted autonomy levels and operating-design-domain changes.
5	Safe-State and Degraded-Mode Behavior	Specifies safe-state triggers, fallback control, and degraded operation.
6	Human Override and Intervention	Maintains independent physical override, operator readiness, and reset integrity.
7	Runtime Monitoring and Telemetry	Makes telemetry trustworthy, independent, coverage-tested, and retained.
8	Incident Response and Physical Recovery	Adds physical containment, evidence preservation, and return-to-service revalidation.
9	Supply Chain and Update Assurance	Covers component provenance, staged updates, and third-party assurance.
10	Simulation, Testing, and Validation	Addresses hardware-in-the-loop, adversarial testing, and sim-to-real gap evidence.

No.	Domain	Executive Meaning
11	Functional Safety Interface	Interfaces with safety cases without claiming to replace safety certification.
12	Physical Operating Environment	Defines and monitors the physical context in which the system may operate.

Evidence and Assurance Posture

PAI-SF™ is evidence-driven and candid about the limits of the current public record. The framework does not claim proprietary incident telemetry, field-validated market adoption, regulatory recognition, or elimination of physical harm. Many threat patterns are derived from public research, controlled testing, adjacent OT/robotics evidence, and systematic architecture analysis rather than mature operational incident datasets.

Conformance should therefore be treated as evidence-informed confidence, not as a universal safety guarantee. Evidence quality depends on relevance, independence, currency, coverage, and integrity; it is not satisfied by policy statements alone.

What PAI-SF Does Not Claim

- It does not replace functional safety, product safety, safety engineering, hazard analysis, or sector-specific safety certification.
- It does not claim recognition or endorsement by regulators, standards bodies, notified bodies, laboratories, or certification authorities.
- It does not claim ODA3 Institute currently operates physical testing laboratories or proprietary physical-AI incident datasets.
- GAISSF™ Domain 9 control content is cited via the verified crosswalk (ODA3-2026-07-CSX-PAISF-004), not reproduced independently here.

Document Set Relationship

This summary should be cited alongside the PAI-SF™ Framework Standard and the standalone PAI-SF™ Control Catalogue. The Technical + Compliance Report remains the fuller narrative and methodology source; the Framework Standard separates the normative core; the Control Catalogue separates the 41-control register for implementation, assessment, and reference.

Recommended Use

- Use this summary as the introduction for executive readers, standards reviewers, and partner discussions.
- Use the Framework Standard as the normative anchor for scope, principles, architecture, lifecycle, and conformance boundaries.
- Use the Control Catalogue for control-by-control implementation and evidence planning.
- Use future sector guides, assessment methodology, conformance templates, and scoring rubrics only after their controlled/public status is deliberately assigned.

Notably Absent

Notably absent from PAI-SF™ v1.0 are public certification decision rules, finalized scoring thresholds, sector-specific implementation profiles, a public telemetry/incident JSON schema, and field-validated claims of adoption. These are documentation and validation gaps to be closed through the maturity roadmap, not facts to imply through marketing language.