ODA3 INSTITUTE

PUBLIC FAQ

PAI-SF™ Public FAQ

Physical AI Security Framework v1.0 — Public Explanatory Reference

ODA3-2026-07-FAQ-PAISF-021 | Public-Facing — Public Explanatory Reference

© ODA3 Pvt Ltd

Document Control

Doc ID	ODA3-2026-07-FAQ-PAISF-021
Doc Type	Public FAQ
Framework Tag	PAI-SF™ v1.0
Audience	Enterprise leaders, security teams, AI governance leads, operators, partners, and assessment requesters
Companion Document	ODA3-2026-07-NOR-PAISF-005; ODA3-2026-07-CAT-PAISF-006; ODA3-2026-07-AIQ-PAISF-017; ODA3-2026-07-ERL-PAISF-018; ODA3-2026-07-ARC-PAISF-019; ODA3-2026-07-CHT-PAISF-020
Governing Licence	GAISSF Ecosystem Licence (GEL v1.0)
Access Classification	Open / Public explanatory reference; scoring, certification, auditor procedures, and controlled scenarios remain Controlled
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 — Public Explanatory Reference

PAI-SF™ Public FAQ

Physical AI Security Framework v1.0 — Public Explanatory Reference

Purpose

This FAQ explains how PAI-SF™ should be understood by public readers, internal adopters, and assessment requesters. It is explanatory only; it does not publish certification logic, scoring methods, auditor procedures, or controlled test content.

1. Frequently Asked Questions

1. What is PAI-SF™? PAI-SF™ is ODA3 Institute's Physical AI Security Framework for assessing AI-enabled systems that can affect physical state, movement, access, infrastructure, operating environments, or human proximity. It focuses on the chain from sensing and model behavior through autonomy boundaries, physical command, monitoring, intervention, and recovery.

2. What is Kinetic Zero Trust? Kinetic Zero Trust is PAI-SF's core assurance principle: an AI system's own confidence, authorization, or familiarity with its deployment context is never, by itself, sufficient to authorize a physical action. Every physical command must be validated against independent safety boundaries, sensor confidence, operational constraints, environmental conditions, and human-override expectations before it executes. It is the single idea that most distinguishes PAI-SF from ordinary AI governance frameworks, which generally stop short of physical actuation.

3. Why is PAI-SF™ separate from GAISSF™? GAISSF™ remains the broader AI safety and security assurance framework. PAI-SF™ is separate because physically embodied AI needs additional assurance language for sensors, actuators, degraded modes, human override, physical recovery, safety-interface handoff, and real-world operating environments.

4. Which systems are in scope? Typical examples include robots, autonomous vehicles, drones and uncrewed systems, medical or assistive robotics, industrial automation, smart infrastructure, physical access systems, and AI-enabled equipment where model-driven behavior can produce physical effects. Scope depends on actual behavior, not product category alone.

5. Does PAI-SF™ certify a system as safe? No. Public PAI-SF™ documents do not certify safety, compliance, conformance, maturity, regulatory approval, or operational readiness. Certification and assessment decisions, where offered, use controlled procedures and authorized delivery channels.

6. Does PAI-SF™ replace functional-safety or sector regulation? No. PAI-SF™ is not a substitute for functional-safety engineering, aviation approval, medical-device approval, road authorization, machinery-safety acceptance, building-code compliance, workplace-safety obligations, or legal review. It helps identify AI-security assurance issues that must interface with those processes.

7. Can my organization use PAI-SF™ internally? Public PAI-SF™ materials may be used for internal learning, scoping, readiness review, and non-commercial reference within the published licence boundaries. Commercial implementation, audit, certification, training delivery, or client-facing services require the applicable ODA3 authorization.

8. Can consultants use PAI-SF™ for clients? Client-facing implementation, audit, certification, or commercial advisory delivery must be conducted through authorized ODA3 commercial or partner channels. This preserves evidence quality, assessor consistency, revenue integrity, and trust in the certification ecosystem.

9. What evidence does PAI-SF™ typically request? Starter evidence usually covers system scope, sensor/perception integrity, command path, autonomy boundaries, degraded modes, override routes, telemetry, incident recovery, update assurance, validation evidence, safety-interface ownership, and operating-environment assumptions.

10. Does PAI-SF™ publish scoring or maturity levels? No. Public artifacts deliberately omit scoring thresholds, weighting, pass/fail criteria, maturity levels, auditor procedures, certification-decision criteria, and controlled test scenarios.

11. How should teams start? Start with the Assessment Intake Questionnaire, Evidence Request List, Assessment Readiness Checklist, and Cheat Sheet. Use these to define system scope, physical consequence paths, evidence availability, unresolved assumptions, and escalation blockers before requesting deeper assessment.

2. Domain Logic in One View

Domains	Assurance Lens	Core Question
1–3	Sensing, command, and edge trust	Can the system trust what it senses, bound what it commands, and trust what it runs on?
4–6	Autonomy, safe state, and intervention	What may it decide alone, and how can people stop or redirect it?
7–9	Telemetry, recovery, and updates	Can events be reconstructed, recovered, and protected from behavior-changing updates?
10–12	Validation, safety interface, and environment	Have realistic scenarios, safety handoffs, and operating assumptions been tested and owned?

3. Notably Absent

Absent Claim	Clarification
No certification claim	This FAQ does not create PAI-SF certification, compliance status, maturity rating, conformance statement, or public assurance language.
No regulatory approval	It does not establish aviation, medical-device, vehicle, machinery, infrastructure, building-code, workplace-safety, or legal approval.
No scoring disclosure	It does not publish thresholds, weighting, pass/fail rules, maturity levels, auditor procedures, controlled scenarios, or certification logic.
No full assessment request	It is explanatory. Use the intake questionnaire, evidence request list, readiness checklist, and controlled assessment materials for assessment planning.

On GAISSF Domain 9: GAISSF Domain 9 control titles are cited via the verified crosswalk (ODA3-2026-07-CSX-PAISF-004), not reproduced independently in this artifact.