ODA3 INSTITUTE

GLOSSARY

Physical AI Security Framework Glossary

Standalone Terminology Reference for PAI-SF™ v1.0

ODA3-2026-07-GLO-PAISF-009 | Public-Facing — Publication-Supporting Reference

Document Control

Doc ID	ODA3-2026-07-GLO-PAISF-009
Doc Type	Glossary
Framework Tag	PAI-SF™ v1.0
Governing Licence	GAISSF Ecosystem Licence (GEL v1.0)
Access Classification	Open / Public — terminology reference; commercial use remains subject to GEL v1.0 licence terms
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-Supporting Reference

Physical AI Security Framework Glossary

Standalone Terminology Reference for PAI-SF™ v1.0

Purpose

This glossary standardizes PAI-SF terminology for public framework references, implementation planning, and downstream documents. It defines terms as used inside PAI-SF; it does not create certification authority, functional-safety certification, regulatory recognition, or commercial delivery rights.

How to Use This Glossary

Use this document as the terminology anchor for the Framework Standard, Control Catalogue, Quick Start Guide, sector guides, assessment templates, and website references. Where a term has a narrower PAI-SF meaning than its general industry meaning, the PAI-SF usage note should control.

Terms are intentionally scoped to physical-AI security assurance. Functional safety, product safety, hazard analysis, and sector-specific safety certification retain their own meanings under their respective standards and regulatory regimes.

Domain Vocabulary

PAI-SF Domain	Term Definition	Primary Reference
Sensor and Perception Integrity	Domain 1. Assurance over sensor trust, cross-modal validation, calibration drift, and adversarial physical-world input detection.	PAI-SF-SEN controls
Actuator and Kinetic Command Safety	Domain 2. Assurance over physical command execution, kinematic limits, command freshness, and actuator divergence.	PAI-SF-ACT controls
Edge Hardware and Model Attestation	Domain 3. Assurance that edge hardware, firmware, models, and	PAI-SF-ATT controls

PAI-SF Domain	Term Definition	Primary Reference
	attestation keys remain authorized and protected.	
Autonomy Boundaries	Domain 4. Governance of autonomy level, operating-domain authorization, graduated boundary response, and handoff integrity.	PAI-SF-AUT controls
Safe-State and Degraded-Mode Behavior	Domain 5. Assurance that credible failure modes trigger safe or reduced-risk physical behavior through defined and validated mechanisms.	PAI-SF-SAF / PAI-SF-DEG controls
Human Override and Intervention	Domain 6. Assurance that human intervention paths exist, are usable under pressure, and do not permit silent unsafe resumption.	PAI-SF-HOV controls
Runtime Monitoring and Telemetry	Domain 7. Assurance over telemetry integrity, monitoring independence, threat-model coverage, retention, and evidentiary custody.	PAI-SF-TEL controls
Incident Response and Physical Recovery	Domain 8. Physical-AI extension of incident response, including containment, evidence preservation, and return-to-service revalidation.	PAI-SF-IRP controls
Supply Chain and Update Assurance	Domain 9. Assurance over physical components, firmware/model updates, staged rollout, regression testing, and third-party component trust.	PAI-SF-SUP controls
Simulation, Testing, and Validation	Domain 10. Assurance from hardware-in-the-loop testing, adversarial physical-world testing, and sim-to-real gap quantification.	PAI-SF-TST controls
Functional Safety Interface	Domain 11. Security-assurance interface with functional safety practice; it identifies AI-specific safety-case gaps without claiming safety certification.	PAI-SF-FSI controls
Physical Operating Environment	Domain 12. Specification and monitoring of the validated environment in which the physical-AI system is allowed to operate.	PAI-SF-ENV controls

Core Glossary

A

Adversarial Physical-World Input — A deliberately crafted physical stimulus, object, signal, placement, or environmental condition intended to manipulate perception or system behavior. *Usage note: Used for threat modeling and SEN/TST evidence; not asserted as widely exploited at scale.*

AI-IRF™ — ODA3 Institute's AI Incident Response Framework, used for preparation, detection/analysis, containment, recovery, and learning after AI incidents. *Usage note: PAI-SF extends it with physical containment, safe-state management, and HITL revalidation.*

AI-Specific Hazard — A hazard whose cause, trigger, or escalation path arises from learned-model behavior, autonomy, perception, inference, or AI-enabled control logic. *Usage note: Relevant to Functional Safety Interface controls.*

Assessment — An evidence-informed review of implementation against PAI-SF controls, methodology, and stated scope. *Usage note: PAI-SF v1.0 does not claim current certification authority.*

Attestation — Evidence that a physical-AI system, edge device, firmware, model, or signing key is the authorized version and has not silently changed. *Usage note: Domain 3 term; also supports telemetry integrity.*

Autonomy Boundary — A documented constraint on when, where, how, and at what level a physical-AI system may act without additional authorization. *Usage note: Different from the physical environment itself; Domain 4 owns authorization decisions.*

C

Chain of Custody — Documented custody, access, retention, and integrity history for telemetry or physical incident evidence. *Usage note: Supports post-incident investigation and possible legal/regulatory use.*

Commercial Use — Use that monetizes, delivers, embeds, or offers PAI-SF material as part of paid services, products, assessments, certification, software, or advisory work. *Usage note: Subject to GEL v1.0 and ODA3 commercial licensing.*

Conformance — A scoped statement that evidence has been reviewed against defined PAI-SF controls and criteria. *Usage note: Detailed conformance record template is a separate backlog item.*

Controlled Asset — A PAI-SF asset intentionally not published as open reference material, such as scoring thresholds, test procedures, auditor guidance, scenario libraries, and certification decision logic. *Usage note: Protects assessment integrity and commercial delivery channels.*

D

Degraded Mode — A reduced-capability operating mode designed to lower physical risk when full operation is not adequately assured. *Usage note: Domain 5 term; not merely slower normal operation.*

E

Evidence Requirement — The documentation, logs, tests, architecture records, or operational artifacts expected to support a control. *Usage note: Field used in the Control Catalogue.*

Evidence Tier — ODA3's corpus-wide marker for source strength: T1 Primary Verified, T2 Secondary Verified, T3 Controlled Simulation/Academic Proxy, T4 Anecdotal/Unverified. *Usage note: Tiering qualifies source strength; it does not by itself prove control adequacy.*

F

Foundational Physical-AI Attestation — Illustrative future assurance pathway representing the entry-level conformance tier: documented system architecture, basic hazard/threat identification, and design-evidence-primary control implementation. *Usage note: Proposed pathway only; the lightest

of the three illustrative levels, alongside Operational Physical-AI Assurance and High-Assurance Physical-AI Evaluation.*

Functional Safety — Engineering discipline concerned with safety-related systems and hazard risk reduction under applicable standards and sector regimes. *Usage note: PAI-SF interfaces with it but does not replace it.*

Functional Safety Interface — The PAI-SF domain that maps AI-specific security assurance concerns into safety-case awareness, standards gap tracking, and exception escalation. *Usage note: Domain 11; not a safety certification claim.*

G

GAISSF Ecosystem — The ODA3 framework ecosystem containing GAISSF, UAIF, AI-IRF, PAI-SF, licensing, publications, and related implementation assets. *Usage note: Governed by GEL v1.0.*

GAISSF™ — ODA3 Institute's Global AI Safety and Security Framework, the top-level framework in the GAISSF Ecosystem. *Usage note: PAI-SF operates alongside GAISSF, UAIF, and AI-IRF.*

GEL v1.0 — GAISSF Ecosystem Licence v1.0, the ODA3-controlled source-available licence governing public-use and commercial-use boundaries. *Usage note: Not an OSI-approved open-source licence; commercial use requires the applicable ODA3 authorization.*

H

Hardware-in-the-Loop (HITL) — Testing or revalidation where real hardware components participate in the test loop rather than relying only on simulation. *Usage note: Used before deployment and before return to service after incidents.*

High-Assurance Physical-AI Evaluation — Illustrative future assurance pathway for higher-consequence physical-AI systems with deeper evidence, independence, and testing expectations. *Usage note: Proposed pathway only; not current certification authority. The highest of the three illustrative levels.*

Human Override — A mechanism by which an authorized human can intervene in or halt physical-AI operation through a path independent enough for the consequence profile. *Usage note: Domain 6 term.*

I

Incident Response and Physical Recovery — PAI-SF's physical-AI-specific extension of incident response covering physical containment, evidence preservation, HITL revalidation, and return to service. *Usage note: Domain 8 term.*

Internal Use — Non-commercial organizational use permitted under GEL v1.0 for internal evaluation, governance, and implementation planning within licence limits. *Usage note: Does not authorize paid third-party delivery unless separately licensed.*

K

Kinematic Bounding — Independent constraints on force, torque, speed, position, motion envelope, or other physical movement characteristics. *Usage note: Domain 2 term.*

Kinetic Zero Trust — PAI-SF's principle that AI-generated physical action should not be trusted solely because the model is authorized, confident, or operating in a familiar context. *Usage note: Requires independent validation against safety boundaries, operational constraints, sensor confidence, environment, and override expectations.*

M

Model Drift — Gradual divergence between expected and observed learned-model behavior over time, context, data distribution, or operating conditions. *Usage note: PAI-SF treats drift as distinct from discrete hardware or software faults.*

N

Notably Absent — ODA3 discipline requiring explicit documentation of claims, evidence, partnerships, recognition, or capabilities that are not present. *Usage note: Used to avoid inflated threat, certification, or market-recognition claims.*

O

Operational Design Domain (ODD) — The defined operating conditions, geography, environment, tasks, and constraints under which a system is intended and validated to operate. *Usage note: Domain 4 governs authorization to change it; Domain 12 describes the environment.*

Operational Physical-AI Assurance — Illustrative assurance pathway for systems with evidence-backed operational controls across relevant domains. *Usage note: Proposed pathway only pending detailed scoring methodology. The middle of the three illustrative levels.*

P

PAI-SF™ — Physical AI Security Framework, ODA3 Institute's framework for AI-security assurance where sensing, inference, autonomy, and actuation can produce physical consequences. *Usage note: Fourth core framework in the GAISSF Ecosystem.*

Physical AI — An AI-enabled system whose outputs directly or indirectly cause, guide, constrain, or fail to prevent physical action where erroneous, adversarial, or degraded operation could cause harm. *Usage note: Core scope term.*

Physical Consequence — A kinetic, environmental, infrastructure, operational, or human-safety effect caused or influenced by AI-enabled action or failure to act. *Usage note: Scope trigger for PAI-SF relevance.*

Physical Effector Path — The causal path from AI output or control logic to physical movement, force, process control, environmental manipulation, or other real-world effect. *Usage note: Distinguishes PAI-SF systems from pure software AI systems.*

Physical Operating Environment — The actual and specified physical context in which a system operates, including site, environmental conditions, human proximity, and spatial constraints. *Usage note: Domain 12 term.*

Public-Use Reference — PAI-SF terminology, domain names, control IDs, high-level objectives, and selected vocabulary made available for citation and internal use under GEL v1.0. *Usage note: Not a grant of unrestricted commercial use.*

R

Return to Service — The controlled process of allowing a physical-AI system to resume operation after incident, override, update, or recovery activity. *Usage note: Requires relevant revalidation evidence.*

S

Safe State — A defined condition or behavior that reduces physical risk for a credible failure mode, such as halt, minimal-risk maneuver, reduced capability, or human handover. *Usage note: Domain 5 term; must be physically achievable.*

Safety Case — Structured argument and evidence used in safety engineering to justify that a system is acceptably safe for its intended use. *Usage note: PAI-SF may feed AI-specific hazards into it but does not replace it.*

Scoring / Maturity Rubric — A detailed method for assigning maturity or assurance level based on evidence and control performance. *Usage note: Controlled backlog item; not fully published in this glossary.*

Simulation-to-Reality Gap — The difference between modeled or simulated system behavior and observed behavior in real physical operation. *Usage note: Domain 10 term.*

T

Telemetry Integrity — Assurance that runtime telemetry has not been silently altered, suppressed, or made unreliable for investigation or assurance. *Usage note: Domain 7 term.*

Threat Model — A structured account of plausible threats, adversaries, failure modes, assumptions, and evidence maturity for the system and deployment context. *Usage note: PAI-SF expects calibration to consequence and evidence maturity.*

U

UAIF™ — ODA3 Institute's Unified AI Incident Framework, used for classifying AI incidents. *Usage note: PAI-SF extends it with physical-consequence categories.*

Terms Not Used as Claims

Term / Phrase	Boundary
Certification	This glossary does not claim that ODA3 currently certifies PAI-SF implementations.
Regulatory recognition	Terms in this document do not imply recognition by a regulator or standards body.
Safety guarantee	PAI-SF terminology does not mean elimination of physical harm or replacement of functional safety practice.
Commercial authorization	Public availability of terms does not authorize commercial service delivery without the applicable ODA3 licence/partner status.
GAISSF D9 control titles	This glossary does not itself restate individual GAISSF Domain 9 control names; the verified control-by-control crosswalk is published separately as ODA3-2026-07-CSX-PAISF-004.

Notably Absent

> This glossary does not assert confirmed market adoption, regulator endorsement, third-party certification authority, or proprietary physical-AI incident datasets. It also does not publish controlled assessment scoring, auditor guidance, scenario libraries, or certification decision logic.

Next Document Dependency

The next backlog item should be the Change History / Release Notes document. It will preserve version-control discipline across the PAI-SF publication-supporting set now that the Framework Standard, Control Catalogue, Executive Summary, Quick Start Guide, and Glossary exist as separate artifacts.