

ODA3 INSTITUTE

ASSESSMENT INTAKE QUESTIONNAIRE

PAI-SF™ Assessment Intake Questionnaire

Physical AI Security Framework v1.0 — Assessment Support Artifact

ODA3-2026-07-AIQ-PAISF-017 | Public-Facing — Assessment Support Artifact

© ODA3 Pvt Ltd

Document Control

Doc ID	ODA3-2026-07-AIQ-PAISF-017
Doc Type	Assessment Intake Questionnaire
Framework Tag	PAI-SF™ v1.0
Audience	Assessment requester, system owner, security lead, AI governance lead, assurance team
Companion Document	ODA3-2026-07-NOR-PAISF-005; ODA3-2026-07-CAT-PAISF-006; ODA3-2026-07-QSG-PAISF-008
Governing Licence	GAISSF Ecosystem Licence (GEL v1.0)
Access Classification	Open / Public intake aid; scoring, certification, and auditor procedures 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 — Assessment Support Artifact

PAI-SF™ Assessment Intake Questionnaire

Physical AI Security Framework v1.0 — Assessment Support Artifact

Purpose

This questionnaire captures the minimum information needed to decide whether a physically embodied AI system can be scoped for PAI-SF™ assessment planning. It supports intake and triage only; it does not disclose scoring logic, maturity thresholds, auditor procedures, certification decisions, or test protocols.

This questionnaire is a triage tool for the **Kinetic Zero Trust** boundary PAI-SF assesses: whether AI-generated physical actions in the system under review are validated against independent safety boundaries, or whether they are trusted on model confidence alone. Section 5's red-flag screen is built around detecting the latter.

1. Requester and System Profile

Complete this section before collecting technical evidence. The goal is to identify the accountable system, owner, deployment stage, and reason for assessment intake.

Item	Intake Prompt	Response / Notes
Organization / business unit	Name the legal entity, operating unit, site, or programme requesting review.	
System or product name	Provide the system name, version, deployment stage, and whether the system is in pilot, production, laboratory, or pre-deployment review.	

Item	Intake Prompt	Response / Notes
System category	Identify the closest category: robotics, industrial automation, autonomous vehicle, drone/UAS, medical robotics, smart infrastructure, facility automation, other physical AI.	
Primary owner	Name the accountable business, security, engineering, safety, compliance, or operations owner.	
Assessment purpose	Describe the reason for intake: internal readiness, partner review, procurement, incident follow-up, assurance planning, regulator-facing preparation, or pre-commercial review.	

2. Physical-AI Scope Screen

Use this section to determine whether PAI-SF is materially relevant. If there is no credible path from AI behavior to physical consequence, another ODA3 framework or ordinary AI governance review may be more appropriate.

Item	Intake Prompt	Response / Notes
Physical consequence path	Can AI-enabled sensing, inference, autonomy, decision support, or control logic influence movement, force, process control, infrastructure state, environment, access, payload, vehicle behavior, or human safety?	
Operating environment	Describe the physical environment: factory, road, site, airspace, clinical setting, public space, building, utility environment, campus, warehouse, laboratory, or mixed deployment.	
Human proximity	Identify whether people may be nearby, co-working, riding, operating, supervising, receiving care, entering the controlled space, or affected indirectly.	
Autonomy mode	State the highest expected autonomy mode and whether a human can approve, pause, modify, override, or reset operation.	
Deployment boundary	List the geography, site, operating design domain, zone, route, airspace, room, cell, facility, or asset boundary in which the system is allowed to operate.	

3. 12-Domain Triage

Mark each domain as In Scope, Possibly In Scope, or Not Evident. The short prompt is intentionally diagnostic, not exhaustive.

No.	PAI-SF Domain	Triage Prompt	Scope Mark
1	Sensor and Perception Integrity	Sensors, perception models, calibration, drift, spoofing, disagreement, confidence, and environmental assumptions.	
2	Actuator and Kinetic Command Safety	Movement, force, speed, torque, payload, steering, braking, valves, gates, barriers, process controls, timing, replay, and command bounds.	
3	Edge Hardware and Model Attestation	Edge devices, controllers, embedded compute, firmware, model packages, attestation keys, and local runtime trust.	
4	Autonomy Boundaries	Authorized task envelope, operating domain, autonomy level, expansion controls, boundary detection, and escalation rules.	
5	Safe-State and Degraded-Mode Behavior	Safe stop, reduced-risk operation, manual mode, degraded mode, fallback behavior, independence from compromised compute, and recovery criteria.	
6	Human Override and Intervention	Operator authority, emergency stop, remote/local override, handoff, intervention readiness, reset integrity, and role accountability.	
7	Runtime Monitoring and Telemetry	Logs and monitoring for sensor confidence, autonomy state, command path, override events, safety interfaces, evidence custody, and retention.	
8	Incident Response and Physical Recovery	Containment, evidence preservation, physical recovery, return-to-service validation, and post-incident learning.	
9	Supply Chain and Update	Vendor components,	

No.	PAI-SF Domain	Triage Prompt	Scope Mark
	Assurance	integrators, firmware/model updates, staged rollout, rollback, regression testing, and third-party dependencies.	
10	Simulation, Testing, and Validation	Scenario tests, hardware-in-the-loop, adversarial physical-world testing, digital twin, sim-to-real limits, and acceptance criteria.	
11	Functional Safety Interface	Safety case links, unresolved AI-security gaps, exception escalation, safety engineering interface, and standards boundary awareness.	
12	Physical Operating Environment	Validated environment, weather, lighting, floor/road/site conditions, occupancy, obstacles, maintenance zones, and environment-change response.	

4. Starter Evidence Availability

For each evidence area, identify whether the requested material is Available, Partially Available, Missing, or Unknown. Do not submit controlled secrets unless a secure review channel has been agreed.

Item	Intake Prompt	Response / Notes
Architecture and asset inventory	System diagram; AI components; controllers; sensors; actuators; cloud/edge dependencies; third-party components; operator roles.	
Physical command path	Command flow; command authorization; constraints; interlocks; timing/replay controls; emergency stop or safe-state mechanisms.	
Sensor and model assurance	Sensor inventory; calibration records; model/version records; perception tests; drift monitoring; known blind spots or degraded conditions.	
Autonomy and boundary controls	Operating domain definition; autonomy modes; boundary-change approvals; handoff criteria; escalation rules; manual-control process.	

Item	Intake Prompt	Response / Notes
Telemetry and incident readiness	Log map; monitoring coverage; retention; tamper-evidence controls; incident playbooks; physical recovery and return-to-service criteria.	
Supplier and update records	Supplier list; firmware/model update records; rollout/rollback process; maintenance windows; regression-test evidence; integrator responsibilities.	

5. Initial Red-Flag Screen

Flag any condition below for deeper review before a normal assessment plan is accepted. A flag does not mean failure; it means the intake should not be treated as routine.

Flag	Why It Matters	Present? / Notes
Unbounded physical command path	AI output can influence movement, access, environment, process control, or payload behavior without an independently bounded command layer.	
No tested safe-state or degraded mode	The system has no validated response for loss of perception, model drift, communications failure, operator loss, or command-path compromise.	
Unclear human override	Operators cannot clearly pause, override, recover, or reset the system with accountable authority and evidence capture.	
Opaque update path	Firmware, model, rule, map, or controller updates can alter physical behavior without staged rollout, rollback, or regression evidence.	
Telemetry gap	Incident reconstruction would not preserve enough evidence about sensor state, autonomy mode, command path, override, and physical effect.	

6. Notably Absent

Absent Claim	Clarification
No certification decision	Completing this intake questionnaire does not produce certification, compliance status, maturity rating, or conformance statement.
No scoring or unpublished control logic	The questionnaire does not publish scoring thresholds, weighting, maturity criteria, auditor procedures, scenario

Absent Claim	Clarification
	libraries, or test protocols.
No regulatory approval claim	The intake does not establish safety certification, medical-device approval, aviation approval, road authorization, building-code compliance, or infrastructure approval.
No substitute for system-specific review	A completed questionnaire is only a scoping artifact; evidence still requires system-specific validation.