ODA3 INSTITUTE

ASSESSMENT READINESS CHECKLIST

PAI-SF™ Assessment Readiness Checklist

Physical AI Security Framework v1.0 — Assessment Support Artifact

ODA3-2026-07-ARC-PAISF-019 | Public-Facing — Assessment Support Artifact

Document Control

Doc ID	ODA3-2026-07-ARC-PAISF-019
Doc Type	Assessment Readiness Checklist
Framework Tag	PAI-SF™ v1.0
Audience	System owner, assessor, security lead, AI governance lead, safety lead, assurance coordinator
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
Governing Licence	GAISSF Ecosystem Licence (GEL v1.0)
Access Classification	Open / Public assessment support; 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 Readiness Checklist

Physical AI Security Framework v1.0 — Assessment Support Artifact

Purpose

This checklist helps determine whether a physically embodied AI system is ready to enter PAI-SF™ assessment planning. It supports readiness triage only; it does not disclose scoring logic, maturity thresholds, auditor procedures, certification decisions, or controlled test protocols.

Readiness, throughout this checklist, means readiness to demonstrate **Kinetic Zero Trust** — PAI-SF's core assurance principle that AI-generated physical actions must be validated against independent safety boundaries, not trusted on model confidence alone. A system that cannot yet produce evidence of that validation is not ready for assessment planning, regardless of how complete its other documentation is.

1. Readiness State Definitions

State	Use This State When
Ready	Scope is defined, accountable owners are identified, starter evidence is available, and no known blocker prevents assessment planning.
Partially Ready	Scope or evidence exists but has material gaps, unresolved owners, incomplete operating-boundary definition, or pending supplier information.
Not Ready	Physical consequence path, system boundary, command path, safety/degraded behavior, telemetry, or update exposure cannot yet be described.
Escalate	A red flag suggests the assessment plan may require

State	Use This State When
	controlled evidence handling, specialist review, incident review, or executive decision before intake proceeds.

2. Assessment Entry Gate

Mark each gate as Ready, Partial, Missing, or Escalate. Use the notes column for evidence owner, date, blocker, or agreed next step.

Gate	Readiness Check	Status / Notes
System identity	System name, version, deployment stage, owner, operator role, and assessment purpose are known.	
Physical consequence path	The organization can explain how AI-enabled sensing, inference, autonomy, or command logic can affect physical state or human proximity.	
Operating boundary	The deployment zone, operating design domain, route, airspace, cell, facility, public-space exposure, or environmental assumptions are documented.	
Command and override path	Command flow, authorization, interlocks, emergency stop, local/remote override, handoff, reset, and return-to-service paths are identified.	
Safe-state behavior	Safe-state and degraded-mode triggers, fallback logic, manual-mode transition, and recovery validation are documented at least at starter level.	
Telemetry coverage	Logs can reconstruct sensor state, autonomy mode, command path, override, degraded state, incident event, and physical recovery activity.	
Supplier/update exposure	Suppliers, integrators, firmware/model/ruleset/map updates, rollout, rollback, and regression-test evidence are known.	
Evidence handling	Sensitive artifacts, regulated data, credentials, source code, and operational records have an agreed secure-handling path or redaction plan.	

3. Domain Readiness Sweep

A domain can be marked ready for assessment planning when the requester can identify the relevant owner and produce starter evidence. This is not a domain score or maturity rating.

No.	PAI-SF Domain	Readiness Check	Mark
1	Sensor and Perception Integrity	Sensor inventory, calibration, drift/disagreement handling, perception test evidence, and known blind spots can be produced.	
2	Actuator and Kinetic Command Safety	Physical command bounds, interlocks, timing/replay protections, divergence monitoring, and force/speed/process constraints can be produced.	
3	Edge Hardware and Model Attestation	Edge-device inventory, firmware/image provenance, model-package identity, attestation evidence, and field update records can be produced.	
4	Autonomy Boundaries	Authorized tasks, autonomy modes, operating-domain definitions, escalation criteria, and boundary-change approvals can be produced.	
5	Safe-State and Degraded-Mode Behavior	Safe-state definitions, degraded-mode triggers, manual transition, fallback procedure, and recovery validation can be produced.	
6	Human Override and Intervention	Operator authority, emergency stop, intervention logs, handoff procedure, reset criteria, and training evidence can be produced.	
7	Runtime Monitoring and Telemetry	Telemetry coverage, tamper-evidence, retention, sensor-confidence logs, autonomy-state logs, command logs, and safety-interface events can be produced.	
8	Incident Response and Physical Recovery	Incident playbooks, physical containment procedure, evidence	

No.	PAI-SF Domain	Readiness Check	Mark
		preservation, recovery checklist, and return-to-service validation can be produced.	
9	Supply Chain and Update Assurance	Vendor/integrator inventory, component provenance, update approvals, staged rollout, rollback, supplier attestations, and regression tests can be produced.	
10	Simulation, Testing, and Validation	Scenario library summary, simulation or hardware-in-the-loop evidence, physical-world tests, acceptance criteria, and sim-to-real limitations can be produced.	
11	Functional Safety Interface	Safety-case links, unresolved AI-security gap register, exception route, safety-engineering owner, and escalation evidence can be produced.	
12	Physical Operating Environment	Site maps or operating-zone definitions, environmental assumptions, obstacles, occupancy, maintenance zones, and environment-change handling can be produced.	

4. Evidence Preparation Checklist

Preparation Item	Expected Readiness Status / Owner
Evidence index prepared	List artifacts by title, owner, date, version, sensitivity, source system, and whether the evidence is implemented, tested, planned, inherited, or vendor-asserted.
Version basis identified	Record system, model, firmware, controller, map, rule, operating procedure, and deployment version used for readiness review.
Secure submission route agreed	Define whether artifacts will be shared as full documents, excerpts, screenshots, summaries, redacted copies, controlled-room review, or evidence index only.
Owner availability confirmed	Confirm accountable people for engineering, AI/model, safety, operations, incident response, supplier management, and compliance are available for follow-up.

Preparation Item	Expected Readiness Status / Owner
Assumptions logged	Separate confirmed evidence from planned remediation, operating assumptions, supplier assertions, inherited platform claims, and unknowns.

5. Escalation Blockers

If any blocker is present, do not treat the intake as routine. Escalate the assessment plan before requesting broad evidence or making external readiness statements.

Blocker	Why It Blocks Routine Readiness
Unknown physical command path	The team cannot describe where AI behavior can affect physical movement, access, process state, infrastructure state, or human proximity.
No tested degraded-mode behavior	There is no evidence for perception loss, model drift, communications loss, command compromise, operator loss, or degraded environment response.
Override cannot be demonstrated	Human override, emergency stop, handoff, reset, or return-to-service procedures are undocumented or cannot be exercised.
Telemetry cannot support incident reconstruction	Logs do not preserve enough evidence about sensor state, autonomy mode, command path, override, degraded mode, and physical effect.
Updates can silently alter physical behavior	Firmware, model, ruleset, map, or controller updates can change behavior without approval, staged rollout, rollback, or regression evidence.
Evidence handling is unresolved	Requested artifacts contain secrets, regulated data, personal data, source code, or sensitive operational detail without an agreed handling plan.

6. Notably Absent

Absent Claim	Clarification
No certification decision	This checklist does not produce PAI-SF certification, compliance status, conformance statement, maturity rating, or public assurance claim.
No scoring publication	It does not publish scoring thresholds, weighting, pass/fail criteria, maturity levels, auditor procedures, controlled scenarios, or certification-decision logic.
No regulatory approval claim	Readiness does not establish safety certification, medical-device approval, aviation approval, road authorization, building-code compliance, infrastructure approval, or legal compliance.
No unnecessary secret disclosure	The checklist does not require source code, credentials, trade secrets, regulated data, personal data, or sensitive operational records unless secure handling is agreed.

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.