ODA3 INSTITUTE

SECTOR CALIBRATION PACKAGE

PAI-SF™ Robotics and Industrial Automation Sector Guide

Physical AI Security Framework v1.0 — Sector Calibration Package

ODA3-2026-07-SCP-PAISF-012 | Public-Facing — Sector Calibration Package

Document Control

Doc ID	ODA3-2026-07-SCP-PAISF-012
Doc Type	Sector Calibration Package
Framework Tag	PAI-SF™ v1.0
Governing Licence	GAISSF Ecosystem Licence (GEL v1.0)
Access Classification	Open / Public sector guidance; commercial delivery 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 — Sector Calibration Package

PAI-SF™ Robotics and Industrial Automation Sector Guide

Physical AI Security Framework v1.0 — Sector Calibration Package

Sector Position

This guide calibrates PAI-SF™ v1.0 for robotics and industrial automation environments where AI-enabled sensing, inference, autonomy, and actuation may affect people, equipment, production assets, or physical operating environments. It is sector guidance, not a safety certification or assessment-scoring method, and it operates under **Kinetic Zero Trust** — PAI-SF's core assurance principle that AI-generated physical action must be validated against independent safety boundaries, not trusted on model confidence alone.

1. Scope and Use

The guide applies to industrial robots, collaborative robots, autonomous mobile robots, warehouse and logistics automation, AI-assisted production lines, PLC- or controller-integrated automation, and robotic systems whose behavior depends on models, perception components, or adaptive autonomy. It is intended for security, OT, safety, robotics, compliance, and assurance teams preparing an initial PAI-SF scoping exercise.

Included	Use With Care	Out of Scope
AI-enabled robots, automation cells, mobile robots, human-robot collaboration zones, and AI-assisted production lines.	Legacy automation where AI only advises humans and has no direct or indirect actuation path. General factory cybersecurity with no AI-driven sensing, inference, autonomy, or actuation path.	Replacement of machinery safety, robotics safety, functional safety, or local regulatory obligations.
Security conditions that may create or worsen physical consequences.	Functional-safety artifacts used as evidence for the AI-security assurance layer.	

2. Priority Control Clusters

All 12 PAI-SF domains remain relevant to sector scoping. The following clusters should usually be treated as first-pass priorities for robotics and industrial automation because they sit closest to the physical consequence path.

Cluster	Primary PAI-SF Controls	Sector Interpretation
Sensor and perception assurance	PAI-SF-SEN-001, SEN-002, SEN-003	Sensor calibration, occlusion handling, and adversarial-input detection — the most common real-world cause of robotics safety incidents (e.g., failure to detect a person in a collaborative cell).
Actuation and motion control	PAI-SF-ACT-001, ACT-002, ACT-003	Independent motion bounding, timing/replay protection, actuator divergence monitoring.
Autonomy boundary and handoff	PAI-SF-AUT-001, AUT-002, AUT-003, AUT-004	Clear autonomy levels, living operating domain, graduated boundary enforcement, reliable handoff.
Safe-state and degraded mode	PAI-SF-SAF-001, SAF-002, SAF-003, SAF-004, DEG-001	Credible failure modes, drift-triggered safe state, independence from compromised compute, fallback control law, and degraded-mode entry/exit criteria.
Human intervention	PAI-SF-HOV-001, HOV-002, HOV-003	Physical override, operator familiarity, reset integrity after intervention.
Telemetry and recovery	PAI-SF-TEL-001 to TEL-004; IRP-001 to IRP-003	Tamper-evident monitoring, coverage validation, custody, containment, evidence capture, HITL return-to-service.
Updates, tests, and safety interface	PAI-SF-SUP-001 to SUP-003; TST-001 to TST-003; FSI-001 to FSI-004	Component provenance, regression tests, adversarial simulation, functional-safety mapping, gap escalation.

3. Domain Calibration

This table translates the 12-domain PAI-SF architecture into robotics and industrial automation review language. It does not change the canonical control catalogue.

Domain	PAI-SF Domain	Robotics / Industrial Automation Calibration
1	Sensor and Perception Integrity	Verify sensor calibration, occlusion assumptions, perception drift, and disagreement handling in cells, lines, warehouses, and mobile work zones.
2	Actuator and Kinetic Command Safety	Treat force, speed, torque, path, timing, and replay controls as first-order security-assurance evidence,

Domain	PAI-SF Domain	Robotics / Industrial Automation Calibration
		not only safety-engineering artifacts.
3	Edge Hardware and Model Attestation	Validate controller, edge device, firmware, and model-provenance claims where robot behavior depends on local compute.
4	Autonomy Boundaries	Document authorized autonomy levels and enforce zone, task, payload, and operating-mode boundaries with change control.
5	Safe-State and Degraded-Mode Behavior	Define and test safe-state behavior for loss of perception, model drift, command-path compromise, and degraded automation modes.
6	Human Override and Intervention	Require independent override paths, role-based readiness, and reset controls after operator intervention.
7	Runtime Monitoring and Telemetry	Preserve tamper-evident evidence for sensor, command, autonomy, override, and safety-interface events.
8	Incident Response and Physical Recovery	Prepare containment and return-to-service playbooks for robot cells, mobile robots, and automation lines.
9	Supply Chain and Update Assurance	Control firmware, model, controller, component, and vendor update paths, including regression testing after changes.
10	Simulation, Testing, and Validation	Use hardware-in-the-loop, digital twin, adversarial, and scenario testing where behavior can change with context.
11	Functional Safety Interface	Map AI-security assurance findings into ISO 10218, ISO/TS 15066, IEC 61508, and site safety-case processes where applicable.
12	Physical Operating Environment	Maintain a validated description of work envelope, fixtures, floor conditions, lighting, human co-presence, payloads, and environmental constraints.

4. Starter Evidence Pack

A first-pass review should collect enough evidence to determine whether the system can be scoped for deeper PAI-SF assessment. The list below is intentionally a starter pack, not a complete audit request list.

Evidence Area	Examples to Collect
System and operating scope	Robot or automation asset inventory, operating modes,

Evidence Area	Examples to Collect
	work envelopes, payload classes, autonomy level, and validated physical operating environment.
Motion and command controls	Kinematic bounds, command-path architecture, replay/timing controls, actuator-health monitoring, emergency-stop and override architecture.
Perception and model assurance	Sensor calibration records, perception test cases, drift monitoring, model/version provenance, edge attestation records.
Safe-state and handoff records	Safe-state definitions, degraded-mode logic, operator handoff procedures, override training evidence, post-override reset criteria.
Telemetry and incident readiness	Telemetry coverage map, tamper-evidence controls, retention/custody rules, containment playbook, HITL revalidation procedure.
Update and safety interface	Vendor and component provenance, update regression test evidence, safety-case mapping, unresolved AI-specific hazard register and escalation decisions.

5. Integration Notes

PAI-SF should be applied beside existing industrial security, robotics safety, and functional-safety practices. For this sector, IEC 62443-style zone and conduit thinking can support network and OT segmentation evidence; ISO 10218 and ISO/TS 15066 can support robot and collaborative-robot safety context; IEC 61508 and related safety cases can support functional-safety interface review. PAI-SF's role is to ask whether AI-specific sensing, inference, autonomy, command, telemetry, update, and recovery conditions are adequately assured where those practices may not fully address learned or adaptive behavior.

Boundary Discipline: A safety case, robot risk assessment, or OT security audit may provide useful evidence, but it should not be treated as automatic PAI-SF coverage. The review must still test whether AI-specific perception drift, model updates, autonomy boundary changes, command-path attacks, telemetry gaps, and return-to-service risks are addressed.

6. Notably Absent

Absent Claim	Clarification
No robotics safety certification claim	This guide does not certify compliance with ISO 10218, ISO/TS 15066, IEC 61508, local machinery-safety law, or site safety acceptance requirements.
No substitution for OT security	IEC 62443-style OT controls remain necessary. PAI-SF calibrates AI-security assurance where sensing, inference, autonomy, and actuation interact.
No scoring or certification logic	The guide does not publish assessment scoring, maturity thresholds, auditor procedures, or certification-decision criteria.
No universal implementation pattern	Industrial robots, collaborative robots, AMRs, PLC-integrated lines, and AI-assisted process automation require system-specific scoping.

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 guide.

7. Next Use

This sector guide should be used with the PAI-SF Framework Standard, Control Catalogue, Quick Start Guide, and Standards Crosswalk. Assessment planning, scoring, maturity treatment, auditor qualification, and certification language remain outside this guide and should be handled through the controlled assessment and conformance-documentation set when authorized.