ODA3 INSTITUTE

QUICK START GUIDE

PAI-SF™ Quick Start Guide

First 30 Days of Practitioner Onboarding

ODA3-2026-07-QSG-PAISF-008 | Public-Facing — Publication Support Document

Document Control

Doc ID	ODA3-2026-07-QSG-PAISF-008
Doc Type	Quick Start Guide
Framework Tag	PAI-SF™ v1.0
Governing Licence	GAISSF Ecosystem Licence (GEL v1.0)
Access Classification	Open / Public — onboarding guidance; no Controlled-asset material (assessment methodology, scoring rubric, scenario bank, assessor decision logic remain Controlled and are not disclosed here)
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 Support Document

PAI-SF™ Quick Start Guide

First 30 Days of Practitioner Onboarding

Purpose

This guide helps teams begin using PAI-SF without treating the full framework, control catalogue, or assessment methodology as a prerequisite. It is an onboarding document, not a certification decision, audit result, or replacement for functional safety certification.

This guide assumes familiarity with **Kinetic Zero Trust** — PAI-SF's core assurance principle, that AI-generated physical actions must be validated against independent safety boundaries rather than trusted on the model's own confidence or authorization alone. See the Framework Standard (ODA3-2026-07-NOR-PAISF-005) for the full definition before beginning system scoping.

1. Use This Guide When

- A system can sense, decide, and cause physical action through actuators, robots, drones, vehicles, industrial equipment, medical robotics, or edge cyber-physical infrastructure.
- Security work must address the chain from perception to model decision, autonomy boundary, kinetic command, physical effect, monitoring, intervention, and recovery.
- A team needs an initial scoping, evidence, or readiness path before a formal assessment is planned.

2. First 30-Day Start Path

1. **Confirm the system boundary.** Define the physical AI system, operating environment, connected assets, autonomy level, actuator classes, human operators, and update channels.
2. **Classify physical consequence paths.** List how software, model output, sensor error, or compromised identity could create physical motion, interruption, collision, unsafe exposure, or recovery difficulty.

3. **Map the 12 domains.** Mark each domain as applicable, partially applicable, not applicable, or pending evidence. Do not exclude a domain merely because ownership sits with a vendor or safety team.

4. **Collect starter evidence.** Gather existing architecture diagrams, sensor and actuator inventories, logs, test reports, safety case references, update records, and incident procedures.

5. **Identify control gaps.** Compare the collected evidence against the Framework Standard and Control Catalogue. Separate missing evidence from missing implementation.

6. **Set the next decision.** Decide whether the next step is remediation planning, partner-led implementation support, a controlled assessment-readiness review, or sector-specific guidance.

3. Minimum Scoping Questions

- What physical actions can the system cause directly or indirectly?
- Which sensors and perception models influence those actions?
- Where are autonomy boundaries defined, enforced, logged, and changed?
- What independent stop, override, or safe-state mechanisms exist?
- What runtime telemetry would prove what happened during an abnormal event?
- What update, supply-chain, and third-party components can affect perception, autonomy, or actuation?
- What evidence exists for hardware-in-the-loop, adversarial, degraded-mode, and sim-to-real testing?

4. Domain Quick Map

Use this table for initial scoping only. The Control Catalogue remains the authoritative control-level reference.

No.	Domain	Control Family	What to Check First
1	Sensor and Perception Integrity	SEN	Sensor input trust, calibration, spoofing resistance, and perception assurance.
2	Actuator and Kinetic Command Safety	ACT	Command validation, rate limits, replay resistance, and actuator divergence detection.
3	Edge Hardware and Model Attestation	ATT	Device, model, boot, and attestation-key trust at the physical edge.
4	Autonomy Boundaries	AUT	Authorized operating envelope, autonomy level, and ODD change control.
5	Safe-State and Degraded-Mode Behavior	SAF / DEG	Safe-state triggers, fallback behavior, and graceful degradation evidence.
6	Human Override and Intervention	HOV	Independent override, emergency stop, operator readiness, and reset integrity.

No.	Domain	Control Family	What to Check First
7	Runtime Monitoring and Telemetry	TEL	Independent telemetry, coverage, synchronization, retention, and anomaly context.
8	Incident Response and Physical Recovery	IRP	Physical containment, evidence preservation, recovery, and return-to-service checks.
9	Supply Chain and Update Assurance	SUP	Component provenance, update controls, staged rollout, and supplier assurance.
10	Simulation, Testing, and Validation	TST	Hardware-in-the-loop, adversarial tests, scenario libraries, and sim-to-real evidence.
11	Functional Safety Interface	FSI	Security-safety interface evidence without replacing formal functional safety certification.
12	Physical Operating Environment	ENV	Operating context, site conditions, physical constraints, and environmental monitoring.

5. Evidence Starter Pack

The following evidence categories are enough to begin a first readiness review, covering all 12 domains. They are not a complete assessment record.

Control Family	Domain	Evidence Area	Starter Material to Collect
SEN	D1	Sensor / perception evidence	Sensor inventory, calibration status, spoofing/tamper assumptions, perception failure examples.
ACT	D2	Kinetic command evidence	Command validation rules, rate limits, actuator logs, replay-prevention evidence.
ATT	D3	Edge hardware / attestation evidence	Secure boot records, hardware bill of materials, attestation key custody documentation.
AUT	D4	Autonomy boundary evidence	Operating-design-domain definition, allowed autonomy levels, boundary-change approvals.
SAF / DEG	D5	Safe-state evidence	Trigger rules, fallback behavior, degraded-mode

Control Family	Domain	Evidence Area	Starter Material to Collect
			procedures, recovery readiness checks.
HOV	D6	Human intervention evidence	Emergency stop design, operator authority, intervention drills, reset controls.
TEL	D7	Runtime telemetry evidence	Independent telemetry sources, coverage map, time sync, retention, anomaly records.
IRP	D8	Incident/recovery evidence	Containment playbooks, evidence-preservation procedures, return-to-service test records.
SUP	D9	Supply-chain and update evidence	Component provenance, firmware/model version records, staged rollout evidence.
TST	D10	Testing evidence	Scenario library, hardware-in-the-loop tests, adversarial tests, sim-to-real gap records.
FSI	D11	Functional safety interface evidence	Hazard analysis cross-reference, safety-case gap register, exception escalation records.
ENV	D12	Operating environment evidence	Environment specification, environmental monitoring logs, human proximity zone design records.

6. Initial Readiness Levels

- **Level 0 — Unscoped:** the system boundary, physical consequence paths, or domain applicability are not yet defined.
- **Level 1 — Identified:** applicable domains are known and basic ownership is assigned, but evidence is incomplete.
- **Level 2 — Evidence-started:** architecture, logs, test records, and operating procedures exist for major consequence paths.
- **Level 3 — Review-ready:** evidence is mapped to controls, gaps are separated from exclusions, and responsible owners are named.
- **Level 4 — Assessment-ready:** evidence is stable, repeatable, retained, and ready for controlled methodology review.

7. What Not to Claim From This Guide

Notably Absent: This Quick Start Guide does not certify a system, score control maturity, validate functional safety compliance, approve a product for deployment, or authorize commercial delivery of

PAI-SF services outside ODA3's licensing and partner ecosystem. It also does not provide the controlled assessment methodology, scoring rubric, scenario bank, or assessor decision logic.

8. Relationship to Other PAI-SF Documents

- **Framework Standard:** normative anchor for scope, domain architecture, principles, boundaries, and control index.
- **Control Catalogue:** authoritative 41-control reference and the next place to look after initial scoping.
- **Assessment Methodology and Scoring Rubric:** Controlled assets for structured assessment and maturity decisions.
- **Sector Guides:** deployment-context guidance for robotics, autonomous vehicles, medical robotics, and related physical AI systems.

9. Recommended Next Actions

- Create a one-page system boundary note before mapping controls.
- Build the first evidence folder around all 12 domains, not only the most familiar ones.
- Use the Control Catalogue to identify which controls require evidence versus which require implementation change.
- Do not publish conformance or certification claims until the controlled assessment and licensing pathway is resolved.