ODA3 INSTITUTE

SECTOR CALIBRATION PACKAGE

PAI-SF™ Medical Robotics Sector Guide

Physical AI Security Framework v1.0 — Sector Calibration Package

ODA3-2026-07-SCP-PAISF-015 | Public-Facing — Sector Calibration Package

Document Control

Doc ID	ODA3-2026-07-SCP-PAISF-015
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™ Medical Robotics Sector Guide

Physical AI Security Framework v1.0 — Sector Calibration Package

Sector Position

This guide calibrates PAI-SF™ v1.0 for medical robotics where AI-enabled perception, clinical-context interpretation, supervised autonomy, actuation, instrument control, or physical movement may affect patients, clinicians, staff, facilities, or care environments. It is sector guidance, not medical-device approval, clinical validation, or an 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 surgical robots, rehabilitation robots, hospital logistics robots, AI-enabled assistive systems, lab automation with physical handling, imaging-guided robotic systems, bedside or care-environment robots, and related medical or clinical-support systems whose behavior depends on models, perception components, sensor fusion, adaptive autonomy, or command pathways. It is intended for security, clinical engineering, safety, quality, compliance, product, and assurance teams preparing an initial PAI-SF scoping exercise.

Included	Use With Care	Out of Scope
AI-enabled medical, clinical-support, assistive, surgical, rehabilitation, laboratory, or hospital robots with sensing, inference, autonomy, motion, instrument, or physical-environment consequence.	Conventional medical-device, quality-system, clinical-risk, usability, and product-safety artifacts used as evidence for the AI-security assurance layer. Systems where AI supports planning, image interpretation, workflow recommendation, or clinician decision support but does not directly command physical motion.	Medical-device regulatory approval, clinical effectiveness, treatment protocols, patient-care decisions, or general healthcare cybersecurity issues with no embodied AI or physical command path. Replacement of clinical validation, medical-device risk management, quality management, safety acceptance, or regulator submissions.
Security conditions that may create		

Included	Use With Care	Out of Scope
or worsen physical consequence for patients, clinicians, staff, equipment, specimens, facilities, or clinical environments.		

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 medical robotics because they sit closest to the patient, clinician, instrument, command, and clinical-workflow consequence path.

Cluster	Primary PAI-SF Controls	Sector Interpretation
Patient-proximity perception	PAI-SF-SEN-001 to SEN-003; ENV-001 to ENV-003	Sensor calibration, occlusion, patient positioning, tissue/anatomy recognition where applicable, staff proximity, lighting, sterile-field constraints, and environment-change assumptions.
Clinical autonomy boundary	PAI-SF-AUT-001 to AUT-004	Authorized autonomy level, clinical workflow boundary, supervised mode, task envelope, procedure phase, handoff criteria, and restrictions on adaptive behavior near patients.
Kinetic command and safe state	PAI-SF-ACT-001 to ACT-003; SAF-001 to SAF-004; DEG-001	Independent limits on force, torque, speed, trajectory, timing, replay, instrument/payload state, pause/hold behavior, and credible safe-state transitions.
Human override and clinical intervention	PAI-SF-HOV-001 to HOV-003	Clinician override, emergency stop, manual takeover, escalation path, operator readiness, alarm handling, and reset criteria after intervention.
Telemetry, incident response, and recovery	PAI-SF-TEL-001 to TEL-004; IRP-001 to IRP-003	Tamper-evident logs for perception, autonomy mode, command state, override, alarms, clinical-context events, physical recovery, and HITL return-to-service.
Update, validation, and safety interface	PAI-SF-ATT-001 to ATT-003; SUP-001 to SUP-003; TST-001 to TST-003; FSI-001 to FSI-004	Device/software/model provenance, controlled update pathways, simulation and bench testing, clinical-workflow regression, safety-case linkage, and unresolved AI-security gap escalation.

3. Domain Calibration

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

No.	PAI-SF Domain	Medical Robotics Calibration
1	Sensor and Perception Integrity	Validate cameras, force sensors, encoders, imaging feeds, navigation inputs, instrument tracking, patient-positioning signals, and sensor disagreement under clinical, sterile, lighting, occlusion, and proximity conditions.
2	Actuator and Kinetic Command Safety	Treat force, torque, speed, trajectory, instrument motion, timing, replay, hold, and release behavior as security-assurance evidence because compromise may create direct patient or staff consequence.
3	Edge Hardware and Model Attestation	Verify controller, edge compute, firmware, model, instrument module, imaging integration, and attestation-key provenance where clinical behavior depends on local compute.
4	Autonomy Boundaries	Document and enforce clinical task, procedure phase, supervision, anatomy/site, treatment-delivery, dosage/payload, and operator-handoff boundaries.
5	Safe-State and Degraded-Mode Behavior	Define and test credible degraded states such as pause, hold, lockout, instrument retract where appropriate, manual mode, alarm escalation, and controlled shutdown.
6	Human Override and Intervention	Require independent clinician/operator override paths, emergency controls, training evidence, intervention escalation, and reset integrity after manual takeover.
7	Runtime Monitoring and Telemetry	Preserve tamper-evident telemetry for perception confidence, autonomy mode, command state, actuator state, override events, alarm events, update state, and incident reconstruction.
8	Incident Response and Physical Recovery	Prepare containment, patient-safety escalation, evidence preservation, device quarantine, rollback, service checks, and human-in-the-loop return-to-service procedures.
9	Supply Chain and Update Assurance	Control device, instrument, sensor, firmware, model, third-party component, clinical integration, and update paths with staged regression evidence.
10	Simulation, Testing, and Validation	Use simulation, bench testing, hardware-in-the-loop, phantom/test-fixture, workflow replay, adversarial physical-world, and degraded-

No.	PAI-SF Domain	Medical Robotics Calibration
		sensor testing.
11	Functional Safety Interface	Map AI-security findings into applicable medical-device risk management, clinical safety, product safety, quality-management, and functional-safety processes without claiming regulatory clearance.
12	Physical Operating Environment	Maintain a validated account of operating room, treatment room, bedside, rehabilitation, lab, home-care, sterile field, patient position, staff movement, lighting, network, and emergency assumptions.

4. Starter Evidence Pack

A first-pass review should collect enough evidence to determine whether the device, robot, care-environment platform, laboratory system, or clinical-support 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 clinical-use scope	Device/platform inventory, intended clinical workflow, autonomy level, procedure or care setting, patient-proximity assumptions, operator role, and physical operating assumptions.
Perception and environment assurance	Sensor calibration records, imaging/navigation confidence handling, patient-positioning assumptions, occlusion tests, sterile-field constraints, lighting limits, and sensor-disagreement responses.
Command, safe-state, and override	Motion/force architecture, trajectory limits, timing/replay controls, pause/hold/retract logic where applicable, emergency stop, clinician override path, and post-override reset criteria.
Telemetry and incident readiness	Log map, autonomy-mode telemetry, command and actuator telemetry, alarm records, tamper-evidence controls, retention/custody rules, and incident reconstruction procedure.
Update and component assurance	Firmware/model/device provenance, instrument or module inventory, third-party integration list, staged update process, clinical-workflow regression suite, and rollback records.
Validation and safety interface	Simulation, bench, fixture, workflow replay, degraded-sensor tests, risk-management mapping, safety-case references, and unresolved AI-security gap register.

5. Integration Notes

PAI-SF should be applied beside medical-device risk management, quality management, clinical safety, cybersecurity, usability engineering, product safety, maintenance, clinical governance, and facility safety practices. For this sector, existing product, clinical, and quality artifacts may support evidence about intended use, operating limits, alarms, maintenance controls, operator responsibilities, and post-market processes. PAI-SF's role is narrower: it asks whether AI-specific sensing, inference, autonomy, command, telemetry, update, and recovery conditions are adequately assured where conventional medical-device processes may not fully address learned, adaptive, or field-updated behavior.

Boundary Discipline: A regulatory submission, quality-system record, clinical validation file, risk-management report, usability file, or cybersecurity file 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, autonomy boundary changes, command compromise, telemetry gaps, incident reconstruction, update regression, and return-to-service risks are addressed.

6. Notably Absent

Absent Claim	Clarification
No medical-device approval claim	This guide does not establish medical-device clearance, approval, conformity assessment, clinical validation, quality-system compliance, or acceptance by any health regulator.
No clinical safety certification claim	The guide does not certify that a medical robot is safe, effective, clinically appropriate, or fit for a procedure, patient group, facility, or care pathway.
No replacement for medical-device engineering	Clinical risk management, usability engineering, cybersecurity, quality management, product safety, clinical evaluation, post-market surveillance, and regulator submissions remain necessary and separate.
No treatment or clinical recommendation	This document is technical sector guidance for AI-security assurance. It does not provide medical advice, clinical protocol, treatment guidance, or patient-care decision criteria.
No scoring or certification logic	The guide does not publish assessment scoring, maturity thresholds, auditor procedures, test protocols, or certification-decision criteria.
No universal medical robotics profile	Surgical robots, rehabilitation robots, hospital logistics robots, diagnostic-assistance platforms, therapy systems, lab automation, and home-care robots 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, clinical validation, regulator submission, and certification language remain outside this

guide and should be handled through the controlled assessment and conformance-documentation set when authorized.