ODA3 INSTITUTE

SECTOR CALIBRATION PACKAGE

PAI-SF™ Autonomous Vehicles Sector Guide

Physical AI Security Framework v1.0 — Sector Calibration Package

ODA3-2026-07-SCP-PAISF-013 | Public-Facing — Sector Calibration Package

Document Control

Doc ID	ODA3-2026-07-SCP-PAISF-013
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™ Autonomous Vehicles Sector Guide

Physical AI Security Framework v1.0 — Sector Calibration Package

Sector Position

This guide calibrates PAI-SF™ v1.0 for autonomous vehicles and mobility systems where AI-enabled sensing, inference, autonomy, and actuation may affect people, vehicles, infrastructure, or public and controlled operating environments. It is sector guidance, not a safety certification, deployment permission, 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 autonomous vehicles, ADAS-adjacent mobility systems, automated shuttles, delivery robots, yard tractors, autonomous trucks, mining or port vehicles, and closed-site mobility systems whose behavior depends on models, perception components, localization, sensor fusion, or adaptive autonomy. It is intended for security, safety, fleet, engineering, compliance, and assurance teams preparing an initial PAI-SF scoping exercise.

Included	Use With Care	Out of Scope
AI-enabled mobility systems with direct or indirect steering, braking, throttle, path, or route-control consequence.	Driver-assistance systems where the AI recommends action but the human remains continuously responsible for physical control.	Replacement of vehicle safety engineering, type approval, roadworthiness, transport regulation, or deployment authorization.
Security conditions that may create or worsen physical consequences in vehicles, fleets, and operating environments.	General vehicle cybersecurity with no AI-driven perception, autonomy, or physical command path. Functional-safety and SOTIF 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 autonomous vehicles because they sit closest to perception, autonomy, command, fleet update, and physical-consequence paths.

Cluster	Primary PAI-SF Controls	Sector Interpretation
Perception and operating design domain	PAI-SF-SEN-001 to SEN-003; ENV-001 to ENV-003	Sensor calibration, adversarial/physical-world input detection, environment assumption validation, and change response across road, site, fleet, weather, lighting, and traffic contexts.
Autonomy boundary and control handoff	PAI-SF-AUT-001 to AUT-004; HOV-001 to HOV-003	Authorized autonomy levels, ODD enforcement, degraded authority states, fallback handoff, driver/operator override, and reset integrity after intervention.
Safe-state and degraded-mode behavior	PAI-SF-SAF-001 to SAF-004; DEG-001	Credible fallback behavior for perception loss, model drift, localization degradation, communications loss, and conditions where immediate stop may itself introduce risk, plus defined degraded-mode entry/exit criteria.
Actuation and command safety	PAI-SF-ACT-001 to ACT-003	Independent constraints on steering, braking, throttle, path, speed, timing, command replay, and actuator divergence monitoring.
Telemetry, incident recovery, and fleet evidence	PAI-SF-TEL-001 to TEL-004; IRP-001 to IRP-003	Tamper-evident logs, event reconstruction, fleet incident containment, custody of evidence, and human-in-the-loop return-to-service.
Updates, validation, and safety interface	PAI-SF-SUP-001 to SUP-003; TST-001 to TST-003; FSI-001 to FSI-004	Staged model/firmware rollout, regression testing, simulation-to-reality validation, ISO 26262 / ISO 21448 interface awareness, and unresolved AI-security gap escalation.

3. Domain Calibration

This table translates the 12-domain PAI-SF architecture into autonomous-vehicle and mobility-system review language. It does not change the canonical control catalogue.

No.	PAI-SF Domain	Autonomous Vehicle / Mobility Calibration
1	Sensor and Perception Integrity	Validate sensor calibration, occlusion behavior, adversarial physical-world inputs, localization integrity, and perception

No.	PAI-SF Domain	Autonomous Vehicle / Mobility Calibration
		disagreement handling across the vehicle's declared operating context.
2	Actuator and Kinetic Command Safety	Treat steering, braking, throttle, path, speed, timing, and replay controls as security-assurance evidence because command-path compromise can create physical consequence.
3	Edge Hardware and Model Attestation	Verify vehicle/edge controller, ECU, firmware, model, sensor-fusion, and perception-stack provenance where runtime behavior depends on local compute.
4	Autonomy Boundaries	Document and enforce authorized autonomy levels, ODD boundaries, route/site restrictions, transition conditions, and reduced-autonomy states.
5	Safe-State and Degraded-Mode Behavior	Define and test fallback behavior for perception degradation, model drift, localization uncertainty, actuator disagreement, and degraded connectivity.
6	Human Override and Intervention	Require clear driver, safety operator, remote operator, or fleet-operations intervention paths where the system design includes human fallback.
7	Runtime Monitoring and Telemetry	Preserve tamper-evident evidence for perception, localization, autonomy mode, command, override, safety-interface, and incident-reconstruction events.
8	Incident Response and Physical Recovery	Prepare fleet and vehicle-level containment, preservation, investigation, software rollback, and human-in-the-loop revalidation procedures.
9	Supply Chain and Update Assurance	Control model, map, firmware, sensor, ECU, dependency, and over-the-air update paths with staged rollout and physical regression evidence.
10	Simulation, Testing, and Validation	Use scenario, closed-course, hardware-in-the-loop, fleet replay, adversarial, and regression testing where behavior depends on context.
11	Functional Safety Interface	Map AI-security findings into automotive safety-case processes, including ISO 26262 and ISO 21448/SOTIF discussions where applicable, without claiming clause-level compliance.

No.	PAI-SF Domain	Autonomous Vehicle / Mobility Calibration
12	Physical Operating Environment	Maintain a validated account of roads, routes, sites, weather, lighting, traffic, geofencing, charging/maintenance environments, and other operating assumptions.

4. Starter Evidence Pack

A first-pass review should collect enough evidence to determine whether the vehicle, fleet, or mobility 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 ODD scope	Vehicle/platform inventory, autonomy level, ODD statement, deployment routes or sites, geofence rules, driver/operator role, and physical operating assumptions.
Perception and localization assurance	Sensor calibration records, perception test cases, localization confidence handling, occlusion/adversarial-input tests, and drift/anomaly monitoring.
Actuation and fallback controls	Steering/braking/throttle bounds, command-path architecture, replay/timing controls, actuator-health monitoring, fallback-state definitions, and degraded-mode test evidence.
Handoff and human intervention	Driver/operator handoff logic, override architecture, remote/fleet-operations escalation rules, training records, and reset criteria after intervention.
Telemetry and incident readiness	Event data map, tamper-evidence controls, retention and custody rules, incident triage process, fleet containment playbook, and return-to-service evidence.
Updates, validation, and safety interface	Model/map/firmware provenance, OTA rollout controls, regression suite evidence, scenario library, safety-case mapping, and unresolved AI-security gap register.

5. Integration Notes

PAI-SF should be applied beside existing automotive cybersecurity, validation, functional-safety, and SOTIF practices. For this sector, ISO 26262 can support functional-safety interface review; ISO 21448/SOTIF can support intended-functionality and scenario reasoning; IEC 61508 concepts may inform broader safety lifecycle discussions in adjacent mobility systems. 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, adaptive, or fleet-updated behavior.

Boundary Discipline: A safety case, SOTIF analysis, validation campaign, cybersecurity case, or road-testing package 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/map updates, ODD

boundary changes, command-path compromise, telemetry gaps, incident reconstruction, and return-to-service risks are addressed.

6. Notably Absent

Absent Claim	Clarification
No autonomous-vehicle safety certification claim	This guide does not certify compliance with ISO 26262, ISO 21448/SOTIF, type approval, roadworthiness, transport law, or local vehicle-safety requirements.
No deployment permission claim	The guide does not establish permission to test or operate on public roads, private sites, ports, warehouses, campuses, mines, or controlled mobility corridors.
No replacement for automotive engineering	Functional safety, product safety, validation engineering, cybersecurity engineering, and regulatory approval processes remain necessary and separate.
No scoring or certification logic	The guide does not publish assessment scoring, maturity thresholds, auditor procedures, or certification-decision criteria.
No universal ODD pattern	Passenger AVs, ADAS-equipped vehicles, shuttles, trucks, yard tractors, delivery robots, and closed-site mobility systems 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, road-test authorization, deployment approval, and certification language remain outside this guide and should be handled through the controlled assessment and conformance-documentation set when authorized.