

ODA3 INSTITUTE

EXECUTIVE BRIEF

GAISSF™ v1.0 Manufacturing AI Assurance

Executive brief for the manufacturing sector

GAISSF-SEC-045 | Executive brief | Draft for Publication

Version 1.0 | Public | External Review Open | Published 1 July 2026 | Updated 5 October 2026

© 2026 ODA3 Pvt Ltd

Methodology Note

ODA3 Institute is an independent AI assurance and evidence infrastructure institute founded in March 2026. It does not maintain a historical proprietary client incident dataset covering manufacturing AI events. Sector interpretations and threat scenarios are derived from authoritative standards, public disclosures, regulatory and official materials, peer-reviewed or academic research, controlled technical demonstrations and explicitly labelled analytical scenarios. Absence of proprietary telemetry limits frequency estimation and must not be interpreted as evidence that an event has or has not occurred.

This brief contains no local T1-T4 evidence-tier model. Detailed source/evidence classifications and limitations are maintained in the accompanying technical guide without redefining ODA3's controlling T1-T4 evidence-sourcing taxonomy.

The Executive Decision

Manufacturing AI risk is determined by operational authority and consequence. An office assistant and an autonomous production system should not be governed as though they create the same exposure.

AI category	Primary business exposure	Executive question
Enterprise support	Intellectual-property leakage, unreliable advice, workforce misuse	Are data rules, ownership and supplier terms adequate?
Production advisory	Maintenance error, scheduling disruption, automation bias	Can operators challenge the output and continue safely without it?
Quality or control influencing	Quality escapes, scrap, recall and contractual exposure	Is independent verification operating before product release or process change?
Safety-critical or autonomous	Worker harm, prolonged downtime, equipment damage and regulatory exposure	Are independent protection, safe-state response and validated recovery demonstrably effective?

Board Risk Lens

- Authority: What may the system recommend, approve or execute?
- Consequence: Can failure affect people, equipment, product conformity or production continuity?
- Reversibility: Can operations return to a validated state without extended requalification?
- Dependency: Can the plant continue if a cloud, supplier or model service is unavailable?
- Evidence: Has operating effectiveness been demonstrated at the plant rather than only described in policy?

Financial and Operational Exposure

Financial estimates must use plant-specific assumptions, low/base/high scenarios and an explicit confidence basis. Overall equipment effectiveness is an operational indicator, not a currency measure.

Exposure	Formula	Management use
Quality escape	Units inspected × defect prevalence × AI false-negative rate × average escape cost per defective unit	Estimate rework, warranty, logistics, contractual and recall exposure.
Impact Band 4 interruption	Total time to detect, isolate, fall back, revalidate and restart × contribution margin per line hour + restart scrap and expediting	Estimate production interruption and recovery exposure.
Annualised exposure	Estimated event frequency × estimated cost per event	Compare risk-reduction options on a common basis.

Required Assumptions

- Data source and period
- Low, base and high input values

- Confidence basis and known uncertainty
- Excluded costs and second-order effects
- Owner and review date
- Sensitivity to supplier opacity, product mix and production volume

Impact Band 4 Leading Indicators

Indicator	Executive interpretation
Safe-state transition success and time	Tests whether the system reaches the validated safe condition within the required period.
Human command-override rate	May indicate weak trust, poor calibration or changed operating conditions.
Emergency-stop activations by cause	Requires cause classification; an increase is not automatically model failure.
Operating-envelope violations	Signals that model, process, sensor or configuration assumptions may no longer hold.
Fallback activation and success	Tests whether production can continue or stop safely when AI is unavailable.
Post-update validation failures	Leading indicator of weak change control and supplier release quality.

Governance, Safety and Legal Boundaries

GAISSF™ v1.0 contains 59 controls across 9 domains. The Applicable-Control Baseline is all 52 D1-D8 controls plus D9 where Physical AI or cyber-physical actuation is in scope. D7 remains Human & Societal Harms with its canonical social-engineering control set. Where D9 is in scope, use GAISSF D9 alongside PAI-SF™ v1.0 as applicable; CSX-PAISF-004 is mapping only and does not establish equivalence, evidence sufficiency or automatic conformance.

- Assign an accountable executive and plant owner for every material AI use case.
- Do not allow AI to bypass independent safety functions, engineering change control or quality release authority.
- Use human approval before action where timing permits; use supervisory veto or independent protection where high-speed processes make real-time intervention infeasible.
- Require evidence-backed supplier commitments for provenance, updates, remote access, incident notification, continuity and exit.
- Treat GAISSF tiers as an operational assurance model, not a legal classification.

Functional-Safety Boundary

GAISSF security controls do not calculate, demonstrate or satisfy probability of failure on demand, probability of dangerous failure per hour, safety integrity level, performance level or equivalent functional-safety metrics. Security approval does not replace machinery- or process-safety lifecycle approval.

European Union Artificial Intelligence Act Boundary

GAISSF tier classification and GAISSF-SEC-045 assessment do not constitute an European Union Artificial Intelligence Act legal classification or conformity assessment. They do not replace technical documentation, quality management, risk management, human oversight, registration, post-market monitoring, fundamental-rights impact assessment or other obligations where those requirements apply.

Five Board-Level Indicators

Indicator	Why it matters
Impact Band 3 and Impact Band 4 systems with approved owners	Tests accountability for material systems.
Systems with current fallback and recovery tests	Tests whether production can continue or stop safely.

Indicator	Why it matters
Quality escapes and unsafe-output events involving AI	Connects technical performance to product and safety consequences.
Supplier assurance and remote-access exceptions	Surfaces concentration and third-party exposure.
Overdue model validation and uncontrolled changes	Identifies lifecycle deterioration before an incident.

Management Actions for the Next 90 Days

Period	Required action	Evidence expected
Days 1–30	Complete inventory and classify production-affecting AI; stop or approve ungoverned operational-technology (OT) connections.	Inventory, accountable owners, system boundaries and tier decisions.
Days 31–60	Validate safety, quality, command and supplier boundaries; define financial exposure ranges.	Risk records, supplier evidence, fallback design and low/base/high estimates.
Days 61–90	Run one incident, fallback, product-quarantine and validated restart exercise.	Exercise record, findings, corrective actions and executive risk acceptance.

Notably Absent

- ODA3 Institute does not rely on a proprietary historical manufacturing-incident telemetry corpus for the claims in this guide.
- Public reporting does not support reliable estimates of the frequency of AI-driven operational-technology disruption.
- No claim is made that GAISSF controls replace functional-safety certification.
- No claim is made that GAISSF tiers establish legal classification or conformity.
- Plausible kinetic sabotage, model poisoning and cross-plant propagation scenarios are not treated as confirmed incidents.

Executive Approval Checklist

- Material AI use cases and owners are known.
- Impact Band 3 and Impact Band 4 safety, quality and production boundaries are approved.
- Fallback, rollback and restart have been tested.
- Financial exposure assumptions and ranges are documented.
- Supplier access and continuity exceptions are visible.
- Evidence gaps and notably absent facts are disclosed.

Publication Status

The technical and executive publications are Draft for Publication — Public — External Review Open; not approved as Final Publication. Final Publication requires completion or authorized disposition of applicable manufacturing/OT/safety/legal/accessibility/editorial review gates.