

ODA3 INSTITUTE

SECTOR IMPLEMENTATION GUIDE

GAISSF™ v1.0 Manufacturing Sector Implementation Guide

Implementation guidance for the Global AI Safety and Security Framework

GAISSF-SEC-045 | Informative sector guidance | Draft for Publication
Version 1.0 | Public | External Review Open | Published 1 July 2026 | Updated 5 October 2026

© 2026 ODA3 Pvt Ltd

Field	Value
Publisher	ODA3 Institute
Legal entity	ODA3 Pvt Ltd
Publication date	1 July 2026
Normative status	Informative sector implementation guidance
Authoritative baseline	GAISSF-NOR-001 v1.0 corrected final publication edition; 59 controls
Release status	Draft for Publication — Public — External Review Open. Final Publication requires completion or authorised disposition of the manufacturing/OT engineering, functional/machinery-safety, legal/regulatory, accessibility/editorial, trademark and authorised release-signing gates below.

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.

Normative, Legal and Safety Boundaries

- GAISSF control identifiers, official titles and authoritative requirements remain unchanged. All manufacturing interpretation is informative.
- GAISSF conformance-tier classification and GAISSF-SEC-045 assessment do not constitute an EU 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.
- 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.
- An AI system must not be treated as the sole safety protection layer unless it has been engineered, assessed and approved under the applicable functional-safety and machinery-safety regime.
- D7 remains the canonical GAISSF domain Human & Societal Harms and retains its AI-enabled phishing, deepfake, out-of-band authentication, AI social-engineering incident-response and AI-enhanced external-attack control set. Manufacturing human-factors, workforce or fairness considerations may support implementation but do not redefine D7.
- Where D9 is in scope, use GAISSF D9 alongside PAI-SF™ v1.0 as applicable. CSX-PAISF-004 is a mapping only and does not establish equivalence, evidence sufficiency or automatic conformance.
- GAISSF certification is not currently offered. References to assessment mean independent assessment or review; any future certification scheme would require separately approved and published governance.

Source and Evidence Classification

Source Class	Name	Definition
Primary / authoritative [T1]	Authoritative primary evidence	Legislation, regulator findings, official investigations, standards records, authenticated first-party technical notices.
Corroborated [T2]	Corroborated evidence	Multiple independent credible sources or replicated findings.
Research / proxy [T3]	Controlled research or credible proxy	Peer-reviewed research, controlled experiments, simulations or documented technical demonstrations with stated limits.
Analytical / hypothetical [T4]	Analytical or hypothetical	Reasoned interpretation or scenario analysis not established as a field incident.

Artifact Evidence Strength

Class	Name	Definition
A1	Direct technical artifact	System-generated log, signed hash,

Class	Name	Definition
		configuration export, measured test result.
A2	Controlled implementation record	Approved change record, signed test report, risk decision, deployment approval.
A3	Independent assurance or attestation	Independent assessment, certification report, supplier assurance with defined scope.
A4	Management assertion or contextual support	Interview, policy statement or unsupported supplier statement; requires corroboration for high-consequence claims.

Source-class labels correspond directly to ODA3's controlling T1–T4 evidence-sourcing taxonomy, shown in brackets beside each label where a sourced claim is made: Primary/authoritative = [T1], Corroborated = [T2], Research/proxy = [T3], Analytical/hypothetical = [T4]. The separate Artifact Evidence Strength scale (A1–A4, below) rates the strength of an implementer's own evidence artifacts and is not an evidence-source tier. Sections 4–8 present prescriptive implementation patterns and control-trigger logic rather than sourced factual claims and are not inline-tagged on that basis; sourced claims appear with [T1]–[T4] tags in Section 11 (Manufacturing Threat Scenarios) and in the accompanying Control Mapping CSV.

Executive Summary

Manufacturing AI spans enterprise assistance, production advice, quality decisions, control-influencing systems and autonomous physical operation. Assurance depth must follow authority and consequence, not marketing labels.

- Operate two implementation tracks: enterprise/advisory AI and OT/edge/physical AI.
- Preserve independent safety functions, engineering change control, quality release authority and deterministic fallback.
- Use human-in-the-loop approval where process timing permits; use human-on-the-loop supervision or independent protection where real-time intervention is infeasible.
- Estimate financial exposure with traceable formulas, low/base/high ranges and explicit assumptions.
- Record what was not established, including the absence of proprietary ODA3 telemetry and verified public frequency data.

1. Purpose and Scope

This guide provides manufacturing-sector implementation guidance for GAISSF™ v1.0. GAISSF contains 59 controls across 9 domains. The Applicable-Control Baseline comprises all 52 controls in D1-D8 plus the 7 D9 controls where Physical AI or cyber-physical actuation is in scope. Manufacturing interpretation is informative and does not alter GAISSF Foundational, Operational or Optimized conformance tiers.

2. Two-Track Implementation Model

Track	Primary concerns	Typical systems	Control emphasis
Enterprise / Advisory AI	Data leakage, intellectual property, unreliable advice, automation bias, identity, cloud and supplier dependency	Knowledge assistants, contract analysis, forecasting, maintenance copilots	Governance, data, output integrity, identity, supplier assurance and human approval
OT / Edge / Physical AI	Determinism, actuation, sensor integrity, safe state, constrained devices, downtime, product conformity	Machine vision, predictive maintenance, robotics, process optimisation, autonomous vehicles	Independent protection, operating envelope, command mediation, rollback, local resilience and physical validation

3. Manufacturing AI Classification

Source Class	Description	Oversight model	Minimum assurance
Impact Band 1 — Enterprise Support	No direct production or safety authority	Human-in-the-loop or ordinary managerial review	Inventory, owner, data rules, acceptable use, monitoring
Impact Band 2 — Production Advisory	Influences operators or planning without direct actuation	Human-in-the-loop or human-on-the-loop	Plant validation, trained oversight, fallback and change control
Impact Band 3 — Decision or Control Influencing	Material effect on quality, release, settings or production actions	Human-in-the-loop where feasible; otherwise supervised veto and command mediation	Independent verification, enhanced monitoring, tested rollback
Impact Band 4 — Safety-Critical or Autonomous Control	Direct actuation or material physical consequence	Human-on-the-loop or human-out-of-the-loop with independent protection	Formal engineering and safety approval, independent protection, validated operating envelope and safe-state

Source Class	Description	Oversight model	Minimum assurance response
--------------	-------------	-----------------	----------------------------

4. Human Factors Implementation Pattern

- Calibrate operator trust by presenting limits, uncertainty and required action rather than unsupported confidence.
- Design alerts for actionability, priority and rate; monitor alert fatigue and override behaviour.
- Run periodic manual or degraded-mode drills to preserve competence where safe and operationally feasible.
- Measure whether the operator has sufficient information and intervention time; nominal presence is not meaningful oversight.
- Review task de-skilling, workarounds, unsafe reliance and loss of situational awareness.

This pattern is informative and does not create a new GAISSF control.

5. Conditional Applicability Trigger Matrix

Control	Official title	Trigger	Required treatment	OT friction
D1-CTL-04	FEDERATED LEARNING POISONING PREVENTION	Triggered when model training or parameter aggregation occurs across plants, suppliers or edge nodes without centralising raw operational data.	Document participating nodes, aggregation trust, poisoning resistance, privacy assumptions and rollback.	High
D1-CTL-06	POST-QUANTUM MODEL SIGNING & CRYPTO HARDENING	Triggered when the model-signing trust period, confidentiality lifetime, equipment lifetime, regulatory obligation or approved cryptographic transition horizon creates material post-quantum exposure.	Perform a cryptographic-agility and migration assessment; do not use a fixed ten-year trigger without context.	High
D1-CTL-07	LORA/ADAPTER INTEGRITY VERIFICATION	Triggered when low-rank adaptations, adapters or other detachable model components are loaded, exchanged or updated separately from the base model.	Verify provenance, compatibility, signatures, privilege and rollback for each adapter.	Moderate
D1-CTL-08	MODEL MERGE ATTACK DETECTION	Applicable when the scoped system uses the capability or faces the risk addressed by the control.	Apply proportionately to architecture, consequence and deployment context.	Moderate
D1-CTL-09	QUANTIZATION BACKDOOR SCREENING	Applicable when the scoped system uses the capability or faces the risk addressed by the control.	Apply proportionately to architecture, consequence and deployment context.	Moderate
D5-CTL-04	AI WATERMARKING ROBUSTNESS	Core for public-facing synthetic media or externally distributed content; generally low relevance to internal OT telemetry and machine-vision inspection unless outputs leave the controlled environment.	Document why provenance or watermarking is required or not required for the scoped output channel.	Low
D9-CTL-03	HUMAN OVERRIDE AND EMERGENCY STOP	Applicable when the scoped system uses the capability or faces the risk addressed by the control.	Apply proportionately to architecture, consequence and deployment context.	High

6. Functional Safety and Machinery Safety Integration

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.

GAISSF activity	Safety-lifecycle interface	Security contribution	Required authority	Boundary
Authority and agency controls	Hazard analysis and safety requirements specification	Defines permitted AI actions and prohibited command paths	Engineering and functional-safety authority	Does not establish a safety integrity level
Safe-state and fallback controls	Architecture, validation, operation and maintenance	Tests transition under model, data, communications or compute failure	Functional-safety and plant authority	Does not establish probability of dangerous failure

GAISSF activity	Safety-lifecycle interface	Security contribution	Required authority	Boundary
Human override and emergency response	Design and operational procedures	Defines supervisory veto, isolation and shutdown coordination	Machinery-safety authority and plant management	Software override must not inhibit independent emergency stop
Model and data change controls	Modification lifecycle	Prevents unreviewed updates and preserves rollback evidence	Engineering, quality and safety authorities	Security approval alone is insufficient
Monitoring and incident controls	Operation, maintenance and proof-test planning	Correlates AI, OT, quality and safety evidence	Operations, safety, quality and security	Monitoring does not replace required proof testing

7. Financial Exposure Estimation

Use plant-specific low, base and high estimates. Fixed confidence percentages must not be used unless supported by observed variance or explicit expert judgement.

Exposure	Formula	Required assumptions
Daily quality escape cost	Units inspected × defect prevalence × AI false-negative rate × average escape cost per defective unit	Inspection population, defect prevalence, measured false-negative rate, rework/replacement/logistics/warranty/recall allocation
Impact Band 4 interruption cost	(Detection + isolation + fallback activation + safety revalidation + quality revalidation + restart time) × contribution margin per line hour + restart scrap + expediting cost	Time distributions, contribution margin, validation steps, restart yield
Annualised expected exposure	Estimated event frequency × estimated cost per event	Frequency basis, low/base/high bounds, excluded costs and review date

8. Impact Band 4 Operational Metrics

Metric	Definition	Interpretation caution
Safety-related false-negative rate	Missed hazardous conditions / all verified hazardous conditions	Requires a defensible ground truth and sufficient sample size
Safe-state transition time	Elapsed time from qualifying failure to validated safe state	Compare with the documented requirement and operating condition
Safe-state transition success rate	Successful validated transitions / attempted transitions	Classify failed and partial transitions
Command override rate	Human vetoes or overrides / AI physical-command recommendations	High rate may indicate poor calibration, weak trust or changing conditions
Emergency-stop activation rate	Emergency stops by human and automated cause category / operating hours	A rising rate is not automatically model degradation
Sensor-integrity alert rate	Calibration, plausibility or drift alerts / operating period	Separate true integrity issues from threshold noise
Operating-envelope violation rate	Detected boundary violations / operating hours or cycles	Investigate model, data, process and configuration causes
Fallback activation and success	Fallback attempts and successful completions	Record degraded-mode capacity and product impact
Post-update validation failure rate	Updates failing safety or quality validation / updates assessed	Leading indicator of change-control quality

9. Reference Architecture Examples

- Data and model provenance for a computer numerical control machine may include sensor identity, calibration records, toolpath metadata, maintenance state, model hash, training dataset version and deployment approval.
- Behavioural drift detection should segment performance by plant, line, machine, product family, tooling condition, shift and environmental regime rather than rely only on an enterprise-wide average.
- An OT/edge deployment should use signed artifacts, secure boot where available, hardened management, local monitoring, validated rollback and compensating controls for constrained legacy devices.
- An enterprise assistant should prevent confidential design or process data from being used for provider training unless explicitly authorised.

10. GAISSF Control Implementation Records

Each record preserves the authoritative identifier and official title. Source requirements remain normative only through GAISSF-NOR-001. Manufacturing content below is informative.

D1-CTL-01 — DATASET PROVENANCE & POISONING PREVENTION

Field	Content
Domain	D1: MODEL INTEGRITY & ADVERSARIAL ROBUSTNESS

Field	Content
Applicability	Applicable
Sector relevance	High
Implementation track	Both
OT friction	Moderate
Conditional trigger	Applicable when the scoped system uses the capability or faces the risk addressed by the control.
Human oversight mode	Human-in-the-loop for advisory/approval workflows; human-on-the-loop or independent protection for high-speed OT.

Authoritative Requirement

Refer to GAISSF-NOR-001 v1.0 and GAISSF-NOR-004 v1.0 for the authoritative control requirement. This sector guide does not restate or modify normative control text.

Manufacturing Interpretation

In manufacturing, this control requires the organisation to apply dataset provenance & poisoning prevention in a way that accounts for long-lived assets, constrained maintenance windows, plant-specific conditions, product conformity, physical consequences and separation of safety functions from AI.

Minimum and Enhanced Implementation

- Maintain model and dataset provenance, integrity checks, approved sources, version control, plant-specific validation, and documented acceptance criteria before production use.
- Use signed artifacts, automated integrity verification, adversarial testing, drift detection by plant and operating regime, independent review, and tested rollback.

Evidence and Assessment

approved scope and owner; risk assessment; architecture or data-flow diagram; configuration and change records; test results; monitoring or review records; dataset/model provenance; hash or signature verification; drift and rollback records

Confirm the control operates at sampled plants and not only in corporate policy. Trace one deployed system from intake through validation, deployment, monitoring, change and recovery. Challenge supplier assertions and verify evidence quality.

Legacy OT and Compensating Controls

Apply proportionately to architecture, consequence and deployment context.

Use compensating monitoring, gateways, signed offline transfer, procedural approval or independent protection where direct implementation is technically infeasible.

Failure Modes, Residual Risk and Limits

- enterprise-only validation
- no plant-specific baseline
- unsigned model updates
- unmonitored drift
- weak dataset lineage

Residual risk remains from unknown failure modes, incomplete data, supplier opacity, plant variation, human factors and correlated dependencies. Risk acceptance must be explicit for production- or safety-affecting systems.

Implementation of this control does not establish functional safety, product conformity, legal compliance or incident prevention by itself.

D1-CTL-02 — MODEL EXTRACTION RESISTANCE

Field	Content
Domain	D1: MODEL INTEGRITY & ADVERSARIAL ROBUSTNESS
Applicability	Applicable
Sector relevance	High
Implementation track	Both
OT friction	Moderate
Conditional trigger	Applicable when the scoped system uses the capability or faces the risk addressed by the control.
Human oversight mode	Human-in-the-loop for advisory/approval workflows; human-on-the-loop or independent protection for high-speed OT.

Authoritative Requirement

Refer to GAISSF-NOR-001 v1.0 and GAISSF-NOR-004 v1.0 for the authoritative control requirement. This sector guide does not restate or modify normative control text.

Manufacturing Interpretation

In manufacturing, this control requires the organisation to apply model extraction resistance in a way that accounts for long-lived assets, constrained maintenance windows, plant-specific conditions, product conformity, physical consequences and separation of safety functions from AI.

Minimum and Enhanced Implementation

- Maintain model and dataset provenance, integrity checks, approved sources, version control, plant-specific validation, and documented acceptance criteria before production use.
- Use signed artifacts, automated integrity verification, adversarial testing, drift detection by plant and operating regime, independent review, and tested rollback.

Evidence and Assessment

approved scope and owner; risk assessment; architecture or data-flow diagram; configuration and change records; test results; monitoring or review records; dataset/model provenance; hash or signature verification; drift and rollback records

Confirm the control operates at sampled plants and not only in corporate policy. Trace one deployed system from intake through validation, deployment, monitoring, change and recovery. Challenge supplier assertions and verify evidence quality.

Legacy OT and Compensating Controls

Apply proportionately to architecture, consequence and deployment context.

Use compensating monitoring, gateways, signed offline transfer, procedural approval or independent protection where direct implementation is technically infeasible.

Failure Modes, Residual Risk and Limits

- enterprise-only validation
- no plant-specific baseline
- unsigned model updates
- unmonitored drift
- weak dataset lineage

Residual risk remains from unknown failure modes, incomplete data, supplier opacity, plant variation, human factors and correlated dependencies. Risk acceptance must be explicit for production- or safety-affecting systems.

Implementation of this control does not establish functional safety, product conformity, legal compliance or incident prevention by itself.

D1-CTL-03 — BEHAVIORAL DRIFT DETECTION

Field	Content
Domain	D1: MODEL INTEGRITY & ADVERSARIAL ROBUSTNESS
Applicability	Applicable
Sector relevance	High
Implementation track	Both
OT friction	Moderate
Conditional trigger	Applicable when the scoped system uses the capability or faces the risk addressed by the control.
Human oversight mode	Human-in-the-loop for advisory/approval workflows; human-on-the-loop or independent protection for high-speed OT.

Authoritative Requirement

Refer to GAISSF-NOR-001 v1.0 and GAISSF-NOR-004 v1.0 for the authoritative control requirement. This sector guide does not restate or modify normative control text.

Manufacturing Interpretation

In manufacturing, this control requires the organisation to apply behavioral drift detection in a way that accounts for long-lived assets, constrained maintenance windows, plant-specific conditions, product conformity, physical consequences and separation of safety functions from AI.

Minimum and Enhanced Implementation

- Maintain model and dataset provenance, integrity checks, approved sources, version control, plant-specific validation, and documented acceptance criteria before production use.
- Use signed artifacts, automated integrity verification, adversarial testing, drift detection by plant and operating regime, independent review, and tested rollback.

Evidence and Assessment

approved scope and owner; risk assessment; architecture or data-flow diagram; configuration and change records; test results; monitoring or review records; dataset/model provenance; hash or signature verification; drift and rollback records

Confirm the control operates at sampled plants and not only in corporate policy. Trace one deployed system from intake through validation, deployment, monitoring, change and recovery. Challenge supplier assertions and verify evidence quality.

Legacy OT and Compensating Controls

Apply proportionately to architecture, consequence and deployment context.

Use compensating monitoring, gateways, signed offline transfer, procedural approval or independent protection where direct implementation is technically infeasible.

Failure Modes, Residual Risk and Limits

- enterprise-only validation
- no plant-specific baseline
- unsigned model updates
- unmonitored drift
- weak dataset lineage

Residual risk remains from unknown failure modes, incomplete data, supplier opacity, plant variation, human factors and correlated dependencies. Risk acceptance must be explicit for production- or safety-affecting systems.

Implementation of this control does not establish functional safety, product conformity, legal compliance or incident prevention by itself.

D1-CTL-04 — FEDERATED LEARNING POISONING PREVENTION

Field	Content
Domain	D1: MODEL INTEGRITY & ADVERSARIAL ROBUSTNESS
Applicability	Conditionally applicable
Sector relevance	Conditional
Implementation track	Both
OT friction	High
Conditional trigger	Triggered when model training or parameter aggregation occurs across plants, suppliers or edge nodes without centralising raw operational data.
Human oversight mode	Human-in-the-loop for advisory/approval workflows; human-on-the-loop or independent protection for high-speed OT.

Authoritative Requirement

Refer to GAISSF-NOR-001 v1.0 and GAISSF-NOR-004 v1.0 for the authoritative control requirement. This sector guide does not restate or modify normative control text.

Manufacturing Interpretation

In manufacturing, this control requires the organisation to apply federated learning poisoning prevention in a way that accounts for long-lived assets, constrained maintenance windows, plant-specific conditions, product conformity, physical consequences and separation of safety functions from AI.

Minimum and Enhanced Implementation

- Maintain model and dataset provenance, integrity checks, approved sources, version control, plant-specific validation, and documented acceptance criteria before production use.
- Use signed artifacts, automated integrity verification, adversarial testing, drift detection by plant and operating regime, independent review, and tested rollback.

Evidence and Assessment

approved scope and owner; risk assessment; architecture or data-flow diagram; configuration and change records; test results; monitoring or review records; dataset/model provenance; hash or signature verification; drift and rollback records

Confirm the control operates at sampled plants and not only in corporate policy. Trace one deployed system from intake through validation, deployment, monitoring, change and recovery. Challenge supplier assertions and verify evidence quality.

Legacy OT and Compensating Controls

Document participating nodes, aggregation trust, poisoning resistance, privacy assumptions and rollback.

Use compensating monitoring, gateways, signed offline transfer, procedural approval or independent protection where direct implementation is technically infeasible.

Failure Modes, Residual Risk and Limits

- enterprise-only validation
- no plant-specific baseline
- unsigned model updates
- unmonitored drift
- weak dataset lineage

Residual risk remains from unknown failure modes, incomplete data, supplier opacity, plant variation, human factors and correlated dependencies. Risk acceptance must be explicit for production- or safety-affecting systems.

Implementation of this control does not establish functional safety, product conformity, legal compliance or incident prevention by itself.

D1-CTL-05 — EMBEDDING SPACE ROBUSTNESS

Field	Content
Domain	D1: MODEL INTEGRITY & ADVERSARIAL ROBUSTNESS
Applicability	Applicable
Sector relevance	High
Implementation track	Both
OT friction	Moderate
Conditional trigger	Applicable when the scoped system uses the capability or faces the risk addressed by the control.
Human oversight mode	Human-in-the-loop for advisory/approval workflows; human-on-the-loop or independent protection for high-speed OT.

Authoritative Requirement

Refer to GAISSF-NOR-001 v1.0 and GAISSF-NOR-004 v1.0 for the authoritative control requirement. This sector guide does not restate or modify normative control text.

Manufacturing Interpretation

In manufacturing, this control requires the organisation to apply embedding space robustness in a way that accounts for long-lived assets, constrained maintenance windows, plant-specific conditions, product conformity, physical consequences and separation of safety functions from AI.

Minimum and Enhanced Implementation

- Maintain model and dataset provenance, integrity checks, approved sources, version control, plant-specific validation, and documented acceptance criteria before production use.
- Use signed artifacts, automated integrity verification, adversarial testing, drift detection by plant and operating regime, independent review, and tested rollback.

Evidence and Assessment

approved scope and owner; risk assessment; architecture or data-flow diagram; configuration and change records; test results; monitoring or review records; dataset/model provenance; hash or signature verification; drift and rollback records

Confirm the control operates at sampled plants and not only in corporate policy. Trace one deployed system from intake through validation, deployment, monitoring, change and recovery. Challenge supplier assertions and verify evidence quality.

Legacy OT and Compensating Controls

Apply proportionately to architecture, consequence and deployment context.

Use compensating monitoring, gateways, signed offline transfer, procedural approval or independent protection where direct implementation is technically infeasible.

Failure Modes, Residual Risk and Limits

- enterprise-only validation
- no plant-specific baseline
- unsigned model updates
- unmonitored drift
- weak dataset lineage

Residual risk remains from unknown failure modes, incomplete data, supplier opacity, plant variation, human factors and correlated dependencies. Risk acceptance must be explicit for production- or safety-affecting systems.

Implementation of this control does not establish functional safety, product conformity, legal compliance or incident prevention by itself.

D1-CTL-06 — POST-QUANTUM MODEL SIGNING & CRYPTO HARDENING

Field	Content
Domain	D1: MODEL INTEGRITY & ADVERSARIAL ROBUSTNESS
Applicability	Conditionally applicable
Sector relevance	Conditional
Implementation track	Both
OT friction	High
Conditional trigger	Triggered when the model-signing trust period, confidentiality lifetime, equipment lifetime, regulatory obligation or approved cryptographic transition horizon creates material post-quantum exposure.
Human oversight mode	Human-in-the-loop for advisory/approval workflows; human-on-the-loop or independent protection for high-speed OT.

Authoritative Requirement

Refer to GAISSF-NOR-001 v1.0 and GAISSF-NOR-004 v1.0 for the authoritative control requirement. This sector guide does not restate or modify normative control text.

Manufacturing Interpretation

In manufacturing, this control requires the organisation to apply post-quantum model signing & crypto hardening in a way that accounts for long-lived assets, constrained maintenance windows, plant-specific conditions, product conformity, physical consequences and separation of safety functions from AI.

Minimum and Enhanced Implementation

- Maintain model and dataset provenance, integrity checks, approved sources, version control, plant-specific validation, and documented acceptance criteria before production use.
- Use signed artifacts, automated integrity verification, adversarial testing, drift detection by plant and operating regime, independent review, and tested rollback.

Evidence and Assessment

approved scope and owner; risk assessment; architecture or data-flow diagram; configuration and change records; test results; monitoring or review records; dataset/model provenance; hash or signature verification; drift and rollback records

Confirm the control operates at sampled plants and not only in corporate policy. Trace one deployed system from intake through validation, deployment, monitoring, change and recovery. Challenge supplier assertions and verify evidence quality.

Legacy OT and Compensating Controls

Perform a cryptographic-agility and migration assessment; do not use a fixed ten-year trigger without context.

Use compensating monitoring, gateways, signed offline transfer, procedural approval or independent protection where direct implementation is technically infeasible.

Failure Modes, Residual Risk and Limits

- enterprise-only validation
- no plant-specific baseline
- unsigned model updates
- unmonitored drift
- weak dataset lineage

Residual risk remains from unknown failure modes, incomplete data, supplier opacity, plant variation, human factors and correlated dependencies. Risk acceptance must be explicit for production- or safety-affecting systems.

Implementation of this control does not establish functional safety, product conformity, legal compliance or incident prevention by itself.

D1-CTL-07 — LORA/ADAPTER INTEGRITY VERIFICATION

Field	Content
Domain	D1: MODEL INTEGRITY & ADVERSARIAL ROBUSTNESS
Applicability	Conditionally applicable
Sector relevance	Conditional
Implementation track	Both
OT friction	Moderate
Conditional trigger	Triggered when low-rank adaptations, adapters or other detachable model components are loaded, exchanged or updated separately from the base model.
Human oversight mode	Human-in-the-loop for advisory/approval workflows; human-on-the-loop or independent protection for high-speed OT.

Authoritative Requirement

Refer to GAISSF-NOR-001 v1.0 and GAISSF-NOR-004 v1.0 for the authoritative control requirement. This sector guide does not restate or modify normative control text.

Manufacturing Interpretation

In manufacturing, this control requires the organisation to apply lora/adapter integrity verification in a way that accounts for long-lived assets, constrained maintenance windows, plant-specific conditions, product conformity, physical consequences and separation of safety functions from AI.

Minimum and Enhanced Implementation

- Maintain model and dataset provenance, integrity checks, approved sources, version control, plant-specific validation, and documented acceptance criteria before production use.
- Use signed artifacts, automated integrity verification, adversarial testing, drift detection by plant and operating regime, independent review, and tested rollback.

Evidence and Assessment

approved scope and owner; risk assessment; architecture or data-flow diagram; configuration and change records; test results; monitoring or review records; dataset/model provenance; hash or signature verification; drift and rollback records

Confirm the control operates at sampled plants and not only in corporate policy. Trace one deployed system from intake through validation, deployment, monitoring, change and recovery. Challenge supplier assertions and verify evidence quality.

Legacy OT and Compensating Controls

Verify provenance, compatibility, signatures, privilege and rollback for each adapter.

Use compensating monitoring, gateways, signed offline transfer, procedural approval or independent protection where direct implementation is technically infeasible.

Failure Modes, Residual Risk and Limits

- enterprise-only validation
- no plant-specific baseline
- unsigned model updates
- unmonitored drift
- weak dataset lineage

Residual risk remains from unknown failure modes, incomplete data, supplier opacity, plant variation, human factors and correlated dependencies. Risk acceptance must be explicit for production- or safety-affecting systems.

Implementation of this control does not establish functional safety, product conformity, legal compliance or incident prevention by itself.

D1-CTL-08 — MODEL MERGE ATTACK DETECTION

Field	Content
Domain	D1: MODEL INTEGRITY & ADVERSARIAL ROBUSTNESS
Applicability	Conditionally applicable
Sector relevance	Conditional
Implementation track	Both
OT friction	Moderate

Field	Content
Conditional trigger	Applicable when the scoped system uses the capability or faces the risk addressed by the control.
Human oversight mode	Human-in-the-loop for advisory/approval workflows; human-on-the-loop or independent protection for high-speed OT.

Authoritative Requirement

Refer to GAISSF-NOR-001 v1.0 and GAISSF-NOR-004 v1.0 for the authoritative control requirement. This sector guide does not restate or modify normative control text.

Manufacturing Interpretation

In manufacturing, this control requires the organisation to apply model merge attack detection in a way that accounts for long-lived assets, constrained maintenance windows, plant-specific conditions, product conformity, physical consequences and separation of safety functions from AI.

Minimum and Enhanced Implementation

- Maintain model and dataset provenance, integrity checks, approved sources, version control, plant-specific validation, and documented acceptance criteria before production use.
- Use signed artifacts, automated integrity verification, adversarial testing, drift detection by plant and operating regime, independent review, and tested rollback.

Evidence and Assessment

approved scope and owner; risk assessment; architecture or data-flow diagram; configuration and change records; test results; monitoring or review records; dataset/model provenance; hash or signature verification; drift and rollback records

Confirm the control operates at sampled plants and not only in corporate policy. Trace one deployed system from intake through validation, deployment, monitoring, change and recovery. Challenge supplier assertions and verify evidence quality.

Legacy OT and Compensating Controls

Apply proportionately to architecture, consequence and deployment context.

Use compensating monitoring, gateways, signed offline transfer, procedural approval or independent protection where direct implementation is technically infeasible.

Failure Modes, Residual Risk and Limits

- enterprise-only validation
- no plant-specific baseline
- unsigned model updates
- unmonitored drift
- weak dataset lineage

Residual risk remains from unknown failure modes, incomplete data, supplier opacity, plant variation, human factors and correlated dependencies. Risk acceptance must be explicit for production- or safety-affecting systems.

Implementation of this control does not establish functional safety, product conformity, legal compliance or incident prevention by itself.

D1-CTL-09 — QUANTIZATION BACKDOOR SCREENING

Field	Content
Domain	D1: MODEL INTEGRITY & ADVERSARIAL ROBUSTNESS
Applicability	Conditionally applicable
Sector relevance	Conditional
Implementation track	Both
OT friction	Moderate
Conditional trigger	Applicable when the scoped system uses the capability or faces the risk addressed by the control.
Human oversight mode	Human-in-the-loop for advisory/approval workflows; human-on-the-loop or independent protection for high-speed OT.

Authoritative Requirement

Refer to GAISSF-NOR-001 v1.0 and GAISSF-NOR-004 v1.0 for the authoritative control requirement. This sector guide does not restate or modify normative control text.

Manufacturing Interpretation

In manufacturing, this control requires the organisation to apply quantization backdoor screening in a way that accounts for long-lived assets, constrained maintenance windows, plant-specific conditions, product conformity, physical consequences and separation of safety functions from AI.

Minimum and Enhanced Implementation

- Maintain model and dataset provenance, integrity checks, approved sources, version control, plant-specific validation, and documented acceptance criteria before production use.
- Use signed artifacts, automated integrity verification, adversarial testing, drift detection by plant and operating regime, independent review, and tested rollback.

Evidence and Assessment

approved scope and owner; risk assessment; architecture or data-flow diagram; configuration and change records; test results; monitoring or review records; dataset/model provenance; hash or signature verification; drift and rollback records

Confirm the control operates at sampled plants and not only in corporate policy. Trace one deployed system from intake through validation, deployment, monitoring, change and recovery. Challenge supplier assertions and verify evidence quality.

Legacy OT and Compensating Controls

Apply proportionately to architecture, consequence and deployment context.

Use compensating monitoring, gateways, signed offline transfer, procedural approval or independent protection where direct implementation is technically infeasible.

Failure Modes, Residual Risk and Limits

- enterprise-only validation
- no plant-specific baseline
- unsigned model updates
- unmonitored drift
- weak dataset lineage

Residual risk remains from unknown failure modes, incomplete data, supplier opacity, plant variation, human factors and correlated dependencies. Risk acceptance must be explicit for production- or safety-affecting systems.

Implementation of this control does not establish functional safety, product conformity, legal compliance or incident prevention by itself.

D2-CTL-01 — DIRECT PROMPT INJECTION PREVENTION

Field	Content
Domain	D2: RUNTIME SECURITY & ADVERSARIAL DEFENSE
Applicability	Applicable
Sector relevance	High
Implementation track	Both
OT friction	Moderate
Conditional trigger	Applicable when the scoped system uses the capability or faces the risk addressed by the control.
Human oversight mode	Human-in-the-loop for advisory/approval workflows; human-on-the-loop or independent protection for high-speed OT.

Authoritative Requirement

Refer to GAISSF-NOR-001 v1.0 and GAISSF-NOR-004 v1.0 for the authoritative control requirement. This sector guide does not restate or modify normative control text.

Manufacturing Interpretation

In manufacturing, this control requires the organisation to apply direct prompt injection prevention in a way that accounts for long-lived assets, constrained maintenance windows, plant-specific conditions, product conformity, physical consequences and separation of safety functions from AI.

Minimum and Enhanced Implementation

- Validate and constrain inputs, isolate untrusted content, enforce tool permissions, log security-relevant interactions, and require human approval for production-affecting actions.
- Deploy context separation, deterministic policy enforcement, adversarial testing, runtime anomaly detection, content provenance checks, and automatic containment.

Evidence and Assessment

approved scope and owner; risk assessment; architecture or data-flow diagram; configuration and change records; test results; monitoring or review records; input/output filtering tests; tool permission logs; adversarial test results

Confirm the control operates at sampled plants and not only in corporate policy. Trace one deployed system from intake through validation, deployment, monitoring, change and recovery. Challenge supplier assertions and verify evidence quality.

Legacy OT and Compensating Controls

Apply proportionately to architecture, consequence and deployment context.

Use compensating monitoring, gateways, signed offline transfer, procedural approval or independent protection where direct implementation is technically infeasible.

Failure Modes, Residual Risk and Limits

- treating all plant data as trusted
- unrestricted tool calls
- missing context separation
- no production-action approval gate

Residual risk remains from unknown failure modes, incomplete data, supplier opacity, plant variation, human factors and correlated dependencies. Risk acceptance must be explicit for production- or safety-affecting systems.

Implementation of this control does not establish functional safety, product conformity, legal compliance or incident prevention by itself.

D2-CTL-02 — INDIRECT PROMPT INJECTION PREVENTION

Field	Content
Domain	D2: RUNTIME SECURITY & ADVERSARIAL DEFENSE
Applicability	Applicable
Sector relevance	High
Implementation track	Both
OT friction	Moderate
Conditional trigger	Applicable when the scoped system uses the capability or faces the risk addressed by the control.
Human oversight mode	Human-in-the-loop for advisory/approval workflows; human-on-the-loop or independent protection for high-speed OT.

Authoritative Requirement

Refer to GAISSF-NOR-001 v1.0 and GAISSF-NOR-004 v1.0 for the authoritative control requirement. This sector guide does not restate or modify normative control text.

Manufacturing Interpretation

In manufacturing, this control requires the organisation to apply indirect prompt injection prevention in a way that accounts for long-lived assets, constrained maintenance windows, plant-specific conditions, product conformity, physical consequences and separation of safety functions from AI.

Minimum and Enhanced Implementation

- Validate and constrain inputs, isolate untrusted content, enforce tool permissions, log security-relevant interactions, and require human approval for production-affecting actions.
- Deploy context separation, deterministic policy enforcement, adversarial testing, runtime anomaly detection, content provenance checks, and automatic containment.

Evidence and Assessment

approved scope and owner; risk assessment; architecture or data-flow diagram; configuration and change records; test results; monitoring or review records; input/output filtering tests; tool permission logs; adversarial test results

Confirm the control operates at sampled plants and not only in corporate policy. Trace one deployed system from intake through validation, deployment, monitoring, change and recovery. Challenge supplier assertions and verify evidence quality.

Legacy OT and Compensating Controls

Apply proportionately to architecture, consequence and deployment context.

Use compensating monitoring, gateways, signed offline transfer, procedural approval or independent protection where direct implementation is technically infeasible.

Failure Modes, Residual Risk and Limits

- treating all plant data as trusted
- unrestricted tool calls
- missing context separation
- no production-action approval gate

Residual risk remains from unknown failure modes, incomplete data, supplier opacity, plant variation, human factors and correlated dependencies. Risk acceptance must be explicit for production- or safety-affecting systems.

Implementation of this control does not establish functional safety, product conformity, legal compliance or incident prevention by itself.

D2-CTL-03 — JAILBREAK RESISTANCE TESTING

Field	Content
Domain	D2: RUNTIME SECURITY & ADVERSARIAL DEFENSE
Applicability	Applicable
Sector relevance	High
Implementation track	Both
OT friction	Moderate
Conditional trigger	Applicable when the scoped system uses the capability or faces the risk addressed by the control.
Human oversight mode	Human-in-the-loop for advisory/approval workflows; human-on-the-loop or independent protection for high-speed OT.

Authoritative Requirement

Refer to GAISSF-NOR-001 v1.0 and GAISSF-NOR-004 v1.0 for the authoritative control requirement. This sector guide does not restate or modify normative control text.

Manufacturing Interpretation

In manufacturing, this control requires the organisation to apply jailbreak resistance testing in a way that accounts for long-lived assets, constrained maintenance windows, plant-specific conditions, product conformity, physical consequences and separation of safety functions from AI.

Minimum and Enhanced Implementation

- Validate and constrain inputs, isolate untrusted content, enforce tool permissions, log security-relevant interactions, and require human approval for production-affecting actions.
- Deploy context separation, deterministic policy enforcement, adversarial testing, runtime anomaly detection, content provenance checks, and automatic containment.

Evidence and Assessment

approved scope and owner; risk assessment; architecture or data-flow diagram; configuration and change records; test results; monitoring or review records; input/output filtering tests; tool permission logs; adversarial test results

Confirm the control operates at sampled plants and not only in corporate policy. Trace one deployed system from intake through validation, deployment, monitoring, change and recovery. Challenge supplier assertions and verify evidence quality.

Legacy OT and Compensating Controls

Apply proportionately to architecture, consequence and deployment context.

Use compensating monitoring, gateways, signed offline transfer, procedural approval or independent protection where direct implementation is technically infeasible.

Failure Modes, Residual Risk and Limits

- treating all plant data as trusted
- unrestricted tool calls
- missing context separation
- no production-action approval gate

Residual risk remains from unknown failure modes, incomplete data, supplier opacity, plant variation, human factors and correlated dependencies. Risk acceptance must be explicit for production- or safety-affecting systems.

Implementation of this control does not establish functional safety, product conformity, legal compliance or incident prevention by itself.

D2-CTL-04 — MULTI-MODAL INJECTION DEFENSE

Field	Content
Domain	D2: RUNTIME SECURITY & ADVERSARIAL DEFENSE
Applicability	Applicable
Sector relevance	High
Implementation track	Both
OT friction	Moderate
Conditional trigger	Applicable when the scoped system uses the capability or faces the risk addressed by the control.
Human oversight mode	Human-in-the-loop for advisory/approval workflows; human-on-the-loop or independent protection for high-speed OT.

Authoritative Requirement

Refer to GAISSF-NOR-001 v1.0 and GAISSF-NOR-004 v1.0 for the authoritative control requirement. This sector guide does not restate or modify normative control text.

Manufacturing Interpretation

In manufacturing, this control requires the organisation to apply multi-modal injection defense in a way that accounts for long-lived assets, constrained maintenance windows, plant-specific conditions, product conformity, physical consequences and separation of safety functions from AI.

Minimum and Enhanced Implementation

- Validate and constrain inputs, isolate untrusted content, enforce tool permissions, log security-relevant interactions, and require human approval for production-affecting actions.
- Deploy context separation, deterministic policy enforcement, adversarial testing, runtime anomaly detection, content provenance checks, and automatic containment.

Evidence and Assessment

approved scope and owner; risk assessment; architecture or data-flow diagram; configuration and change records; test results; monitoring or review records; input/output filtering tests; tool permission logs; adversarial test results

Confirm the control operates at sampled plants and not only in corporate policy. Trace one deployed system from intake through validation, deployment, monitoring, change and recovery. Challenge supplier assertions and verify evidence quality.

Legacy OT and Compensating Controls

Apply proportionately to architecture, consequence and deployment context.

Use compensating monitoring, gateways, signed offline transfer, procedural approval or independent protection where direct implementation is technically infeasible.

Failure Modes, Residual Risk and Limits

- treating all plant data as trusted
- unrestricted tool calls
- missing context separation
- no production-action approval gate

Residual risk remains from unknown failure modes, incomplete data, supplier opacity, plant variation, human factors and correlated dependencies. Risk acceptance must be explicit for production- or safety-affecting systems.

Implementation of this control does not establish functional safety, product conformity, legal compliance or incident prevention by itself.

D2-CTL-05 — FUNCTION CALL/TOOL CALL INJECTION PREVENTION

Field	Content
Domain	D2: RUNTIME SECURITY & ADVERSARIAL DEFENSE
Applicability	Applicable
Sector relevance	High
Implementation track	Both
OT friction	Moderate
Conditional trigger	Applicable when the scoped system uses the capability or faces the risk addressed by the control.
Human oversight mode	Human-in-the-loop for advisory/approval workflows; human-on-the-loop or

Field	Content
	independent protection for high-speed OT.

Authoritative Requirement

Refer to GAISSF-NOR-001 v1.0 and GAISSF-NOR-004 v1.0 for the authoritative control requirement. This sector guide does not restate or modify normative control text.

Manufacturing Interpretation

In manufacturing, this control requires the organisation to apply function call/tool call injection prevention in a way that accounts for long-lived assets, constrained maintenance windows, plant-specific conditions, product conformity, physical consequences and separation of safety functions from AI.

Minimum and Enhanced Implementation

- Validate and constrain inputs, isolate untrusted content, enforce tool permissions, log security-relevant interactions, and require human approval for production-affecting actions.
- Deploy context separation, deterministic policy enforcement, adversarial testing, runtime anomaly detection, content provenance checks, and automatic containment.

Evidence and Assessment

approved scope and owner; risk assessment; architecture or data-flow diagram; configuration and change records; test results; monitoring or review records; input/output filtering tests; tool permission logs; adversarial test results

Confirm the control operates at sampled plants and not only in corporate policy. Trace one deployed system from intake through validation, deployment, monitoring, change and recovery. Challenge supplier assertions and verify evidence quality.

Legacy OT and Compensating Controls

Apply proportionately to architecture, consequence and deployment context.

Use compensating monitoring, gateways, signed offline transfer, procedural approval or independent protection where direct implementation is technically infeasible.

Failure Modes, Residual Risk and Limits

- treating all plant data as trusted
- unrestricted tool calls
- missing context separation
- no production-action approval gate

Residual risk remains from unknown failure modes, incomplete data, supplier opacity, plant variation, human factors and correlated dependencies. Risk acceptance must be explicit for production- or safety-affecting systems.

Implementation of this control does not establish functional safety, product conformity, legal compliance or incident prevention by itself.

D2-CTL-06 — CROSS-CONTEXT HIJACKING MITIGATION

Field	Content
Domain	D2: RUNTIME SECURITY & ADVERSARIAL DEFENSE
Applicability	Applicable
Sector relevance	High
Implementation track	Both
OT friction	Moderate
Conditional trigger	Applicable when the scoped system uses the capability or faces the risk addressed by the control.
Human oversight mode	Human-in-the-loop for advisory/approval workflows; human-on-the-loop or independent protection for high-speed OT.

Authoritative Requirement

Refer to GAISSF-NOR-001 v1.0 and GAISSF-NOR-004 v1.0 for the authoritative control requirement. This sector guide does not restate or modify normative control text.

Manufacturing Interpretation

In manufacturing, this control requires the organisation to apply cross-context hijacking mitigation in a way that accounts for long-lived assets, constrained maintenance windows, plant-specific conditions, product conformity, physical consequences and separation of safety functions from AI.

Minimum and Enhanced Implementation

- Validate and constrain inputs, isolate untrusted content, enforce tool permissions, log security-relevant interactions, and require human approval for production-affecting actions.
- Deploy context separation, deterministic policy enforcement, adversarial testing, runtime anomaly detection, content provenance checks, and automatic containment.

Evidence and Assessment

approved scope and owner; risk assessment; architecture or data-flow diagram; configuration and change records; test results; monitoring or review records; input/output filtering tests; tool permission logs; adversarial test results

Confirm the control operates at sampled plants and not only in corporate policy. Trace one deployed system from intake through validation, deployment, monitoring, change and recovery. Challenge supplier assertions and verify evidence quality.

Legacy OT and Compensating Controls

Apply proportionately to architecture, consequence and deployment context.

Use compensating monitoring, gateways, signed offline transfer, procedural approval or independent protection where direct implementation is technically infeasible.

Failure Modes, Residual Risk and Limits

- treating all plant data as trusted
- unrestricted tool calls
- missing context separation
- no production-action approval gate

Residual risk remains from unknown failure modes, incomplete data, supplier opacity, plant variation, human factors and correlated dependencies. Risk acceptance must be explicit for production- or safety-affecting systems.

Implementation of this control does not establish functional safety, product conformity, legal compliance or incident prevention by itself.

D3-CTL-01 — LEAST AGENCY ENFORCEMENT

Field	Content
Domain	D3: AGENTIC RISK & AUTONOMOUS SYSTEM SECURITY
Applicability	Applicable with manufacturing interpretation
Sector relevance	High
Implementation track	Both
OT friction	Moderate
Conditional trigger	Applicable when the scoped system uses the capability or faces the risk addressed by the control.
Human oversight mode	Select human-in-the-loop, human-on-the-loop, or human-out-of-the-loop with independent protection based on process speed and consequence.

Authoritative Requirement

Refer to GAISSF-NOR-001 v1.0 and GAISSF-NOR-004 v1.0 for the authoritative control requirement. This sector guide does not restate or modify normative control text.

Manufacturing Interpretation

In manufacturing, this control requires the organisation to apply least agency enforcement in a way that accounts for long-lived assets, constrained maintenance windows, plant-specific conditions, product conformity, physical consequences and separation of safety functions from AI.

Minimum and Enhanced Implementation

- Define authority boundaries, approved actions, human override, fail-safe behavior, segregation of duties, and explicit approval for control-affecting operations.
- Use independent safety interlocks, formal policy constraints, simulation and boundary testing, command mediation, real-time authority monitoring, and tested emergency fallback.

Evidence and Assessment

approved scope and owner; risk assessment; architecture or data-flow diagram; configuration and change records; test results; monitoring or review records; authority matrix; simulation results; override and fail-safe test

Confirm the control operates at sampled plants and not only in corporate policy. Trace one deployed system from intake through validation, deployment, monitoring, change and recovery. Challenge supplier assertions and verify evidence quality.

Legacy OT and Compensating Controls

Apply proportionately to architecture, consequence and deployment context.

Use compensating monitoring, gateways, signed offline transfer, procedural approval or independent protection where direct implementation is technically infeasible.

Failure Modes, Residual Risk and Limits

- unclear action authority
- ineffective human oversight
- AI path bypasses engineering change control
- no tested fallback

Residual risk remains from unknown failure modes, incomplete data, supplier opacity, plant variation, human factors and correlated dependencies. Risk acceptance must be explicit for production- or safety-affecting systems.

Implementation of this control does not establish functional safety, product conformity, legal compliance or incident prevention by itself.

D3-CTL-02 — INTER-AGENT COMMUNICATION SECURITY

Field	Content
Domain	D3: AGENTIC RISK & AUTONOMOUS SYSTEM SECURITY
Applicability	Applicable with manufacturing interpretation
Sector relevance	High
Implementation track	Both
OT friction	Moderate
Conditional trigger	Applicable when the scoped system uses the capability or faces the risk addressed by the control.
Human oversight mode	Select human-in-the-loop, human-on-the-loop, or human-out-of-the-loop with independent protection based on process speed and consequence.

Authoritative Requirement

Refer to GAISSF-NOR-001 v1.0 and GAISSF-NOR-004 v1.0 for the authoritative control requirement. This sector guide does not restate or modify normative control text.

Manufacturing Interpretation

In manufacturing, this control requires the organisation to apply inter-agent communication security in a way that accounts for long-lived assets, constrained maintenance windows, plant-specific conditions, product conformity, physical consequences and separation of safety functions from AI.

Minimum and Enhanced Implementation

- Define authority boundaries, approved actions, human override, fail-safe behavior, segregation of duties, and explicit approval for control-affecting operations.
- Use independent safety interlocks, formal policy constraints, simulation and boundary testing, command mediation, real-time authority monitoring, and tested emergency fallback.

Evidence and Assessment

approved scope and owner; risk assessment; architecture or data-flow diagram; configuration and change records; test results; monitoring or review records; authority matrix; simulation results; override and fail-safe test

Confirm the control operates at sampled plants and not only in corporate policy. Trace one deployed system from intake through validation, deployment, monitoring, change and recovery. Challenge supplier assertions and verify evidence quality.

Legacy OT and Compensating Controls

Apply proportionately to architecture, consequence and deployment context.

Use compensating monitoring, gateways, signed offline transfer, procedural approval or independent protection where direct implementation is technically infeasible.

Failure Modes, Residual Risk and Limits

- unclear action authority

- ineffective human oversight
- AI path bypasses engineering change control
- no tested fallback

Residual risk remains from unknown failure modes, incomplete data, supplier opacity, plant variation, human factors and correlated dependencies. Risk acceptance must be explicit for production- or safety-affecting systems.

Implementation of this control does not establish functional safety, product conformity, legal compliance or incident prevention by itself.

D3-CTL-03 — AGENTIC PROMPT CHAINING DETECTION

Field	Content
Domain	D3: AGENTIC RISK & AUTONOMOUS SYSTEM SECURITY
Applicability	Applicable with manufacturing interpretation
Sector relevance	High
Implementation track	Both
OT friction	Moderate
Conditional trigger	Applicable when the scoped system uses the capability or faces the risk addressed by the control.
Human oversight mode	Select human-in-the-loop, human-on-the-loop, or human-out-of-the-loop with independent protection based on process speed and consequence.

Authoritative Requirement

Refer to GAISSF-NOR-001 v1.0 and GAISSF-NOR-004 v1.0 for the authoritative control requirement. This sector guide does not restate or modify normative control text.

Manufacturing Interpretation

In manufacturing, this control requires the organisation to apply agentic prompt chaining detection in a way that accounts for long-lived assets, constrained maintenance windows, plant-specific conditions, product conformity, physical consequences and separation of safety functions from AI.

Minimum and Enhanced Implementation

- Define authority boundaries, approved actions, human override, fail-safe behavior, segregation of duties, and explicit approval for control-affecting operations.
- Use independent safety interlocks, formal policy constraints, simulation and boundary testing, command mediation, real-time authority monitoring, and tested emergency fallback.

Evidence and Assessment

approved scope and owner; risk assessment; architecture or data-flow diagram; configuration and change records; test results; monitoring or review records; authority matrix; simulation results; override and fail-safe test

Confirm the control operates at sampled plants and not only in corporate policy. Trace one deployed system from intake through validation, deployment, monitoring, change and recovery. Challenge supplier assertions and verify evidence quality.

Legacy OT and Compensating Controls

Apply proportionately to architecture, consequence and deployment context.

Use compensating monitoring, gateways, signed offline transfer, procedural approval or independent protection where direct implementation is technically infeasible.

Failure Modes, Residual Risk and Limits

- unclear action authority
- ineffective human oversight
- AI path bypasses engineering change control
- no tested fallback

Residual risk remains from unknown failure modes, incomplete data, supplier opacity, plant variation, human factors and correlated dependencies. Risk acceptance must be explicit for production- or safety-affecting systems.

Implementation of this control does not establish functional safety, product conformity, legal compliance or incident prevention by itself.

D3-CTL-04 — EMBODIED AI SAFETY CONTROLS

Field	Content
Domain	D3: AGENTIC RISK & AUTONOMOUS SYSTEM SECURITY
Applicability	Applicable with manufacturing interpretation
Sector relevance	High
Implementation track	Both
OT friction	Moderate
Conditional trigger	Applicable when the scoped system uses the capability or faces the risk addressed by the control.
Human oversight mode	Select human-in-the-loop, human-on-the-loop, or human-out-of-the-loop with independent protection based on process speed and consequence.

Authoritative Requirement

Refer to GAISSF-NOR-001 v1.0 and GAISSF-NOR-004 v1.0 for the authoritative control requirement. This sector guide does not restate or modify normative control text.

Manufacturing Interpretation

In manufacturing, this control requires the organisation to apply embodied ai safety controls in a way that accounts for long-lived assets, constrained maintenance windows, plant-specific conditions, product conformity, physical consequences and separation of safety functions from AI.

Minimum and Enhanced Implementation

- Define authority boundaries, approved actions, human override, fail-safe behavior, segregation of duties, and explicit approval for control-affecting operations.
- Use independent safety interlocks, formal policy constraints, simulation and boundary testing, command mediation, real-time authority monitoring, and tested emergency fallback.

Evidence and Assessment

approved scope and owner; risk assessment; architecture or data-flow diagram; configuration and change records; test results; monitoring or review records; authority matrix; simulation results; override and fail-safe test

Confirm the control operates at sampled plants and not only in corporate policy. Trace one deployed system from intake through validation, deployment, monitoring, change and recovery. Challenge supplier assertions and verify evidence quality.

Legacy OT and Compensating Controls

Apply proportionately to architecture, consequence and deployment context.

Use compensating monitoring, gateways, signed offline transfer, procedural approval or independent protection where direct implementation is technically infeasible.

Failure Modes, Residual Risk and Limits

- unclear action authority
- ineffective human oversight
- AI path bypasses engineering change control
- no tested fallback

Residual risk remains from unknown failure modes, incomplete data, supplier opacity, plant variation, human factors and correlated dependencies. Risk acceptance must be explicit for production- or safety-affecting systems.

Implementation of this control does not establish functional safety, product conformity, legal compliance or incident prevention by itself.

D3-CTL-05 — MULTI-AGENT TRUST CHAIN ATTESTATION

Field	Content
Domain	D3: AGENTIC RISK & AUTONOMOUS SYSTEM SECURITY
Applicability	Applicable with manufacturing interpretation
Sector relevance	High
Implementation track	Both
OT friction	Moderate
Conditional trigger	Applicable when the scoped system uses the capability or faces the risk addressed by the control.
Human oversight mode	Select human-in-the-loop, human-on-the-loop, or human-out-of-the-loop with independent protection based on process speed and consequence.

Authoritative Requirement

Refer to GAISSF-NOR-001 v1.0 and GAISSF-NOR-004 v1.0 for the authoritative control requirement. This sector guide does not restate or modify normative control text.

Manufacturing Interpretation

In manufacturing, this control requires the organisation to apply multi-agent trust chain attestation in a way that accounts for long-lived assets, constrained maintenance windows, plant-specific conditions, product conformity, physical consequences and separation of safety functions from AI.

Minimum and Enhanced Implementation

- Define authority boundaries, approved actions, human override, fail-safe behavior, segregation of duties, and explicit approval for control-affecting operations.
- Use independent safety interlocks, formal policy constraints, simulation and boundary testing, command mediation, real-time authority monitoring, and tested emergency fallback.

Evidence and Assessment

approved scope and owner; risk assessment; architecture or data-flow diagram; configuration and change records; test results; monitoring or review records; authority matrix; simulation results; override and fail-safe test

Confirm the control operates at sampled plants and not only in corporate policy. Trace one deployed system from intake through validation, deployment, monitoring, change and recovery. Challenge supplier assertions and verify evidence quality.

Legacy OT and Compensating Controls

Apply proportionately to architecture, consequence and deployment context.

Use compensating monitoring, gateways, signed offline transfer, procedural approval or independent protection where direct implementation is technically infeasible.

Failure Modes, Residual Risk and Limits

- unclear action authority
- ineffective human oversight
- AI path bypasses engineering change control
- no tested fallback

Residual risk remains from unknown failure modes, incomplete data, supplier opacity, plant variation, human factors and correlated dependencies. Risk acceptance must be explicit for production- or safety-affecting systems.

Implementation of this control does not establish functional safety, product conformity, legal compliance or incident prevention by itself.

D3-CTL-06 — PERSISTENT MEMORY EXFILTRATION PREVENTION

Field	Content
Domain	D3: AGENTIC RISK & AUTONOMOUS SYSTEM SECURITY
Applicability	Applicable with manufacturing interpretation
Sector relevance	High
Implementation track	Both
OT friction	Moderate
Conditional trigger	Applicable when the scoped system uses the capability or faces the risk addressed by the control.
Human oversight mode	Select human-in-the-loop, human-on-the-loop, or human-out-of-the-loop with independent protection based on process speed and consequence.

Authoritative Requirement

Refer to GAISSF-NOR-001 v1.0 and GAISSF-NOR-004 v1.0 for the authoritative control requirement. This sector guide does not restate or modify normative control text.

Manufacturing Interpretation

In manufacturing, this control requires the organisation to apply persistent memory exfiltration prevention in a way that accounts for long-lived assets, constrained maintenance windows, plant-specific conditions, product conformity, physical consequences and separation of safety functions from AI.

Minimum and Enhanced Implementation

- Define authority boundaries, approved actions, human override, fail-safe behavior, segregation of duties, and explicit approval for control-affecting operations.
- Use independent safety interlocks, formal policy constraints, simulation and boundary testing, command mediation, real-time authority monitoring, and tested emergency fallback.

Evidence and Assessment

approved scope and owner; risk assessment; architecture or data-flow diagram; configuration and change records; test results; monitoring or review records; authority matrix; simulation results; override and fail-safe test

Confirm the control operates at sampled plants and not only in corporate policy. Trace one deployed system from intake through validation, deployment, monitoring, change and recovery. Challenge supplier assertions and verify evidence quality.

Legacy OT and Compensating Controls

Apply proportionately to architecture, consequence and deployment context.

Use compensating monitoring, gateways, signed offline transfer, procedural approval or independent protection where direct implementation is technically infeasible.

Failure Modes, Residual Risk and Limits

- unclear action authority
- ineffective human oversight
- AI path bypasses engineering change control
- no tested fallback

Residual risk remains from unknown failure modes, incomplete data, supplier opacity, plant variation, human factors and correlated dependencies. Risk acceptance must be explicit for production- or safety-affecting systems.

Implementation of this control does not establish functional safety, product conformity, legal compliance or incident prevention by itself.

D3-CTL-07 — SECURE MEMORY LIFECYCLE MANAGEMENT

Field	Content
Domain	D3: AGENTIC RISK & AUTONOMOUS SYSTEM SECURITY
Applicability	Applicable with manufacturing interpretation
Sector relevance	High
Implementation track	Both
OT friction	Moderate
Conditional trigger	Applicable when the scoped system uses the capability or faces the risk addressed by the control.
Human oversight mode	Select human-in-the-loop, human-on-the-loop, or human-out-of-the-loop with independent protection based on process speed and consequence.

Authoritative Requirement

Refer to GAISSF-NOR-001 v1.0 and GAISSF-NOR-004 v1.0 for the authoritative control requirement. This sector guide does not restate or modify normative control text.

Manufacturing Interpretation

In manufacturing, this control requires the organisation to apply secure memory lifecycle management in a way that accounts for long-lived assets, constrained maintenance windows, plant-specific conditions, product conformity, physical consequences and separation of safety functions from AI.

Minimum and Enhanced Implementation

- Define authority boundaries, approved actions, human override, fail-safe behavior, segregation of duties, and explicit approval for control-affecting operations.
- Use independent safety interlocks, formal policy constraints, simulation and boundary testing, command mediation, real-time authority monitoring, and tested emergency fallback.

Evidence and Assessment

approved scope and owner; risk assessment; architecture or data-flow diagram; configuration and change records; test results; monitoring or review records; authority matrix; simulation results; override and fail-safe test

Confirm the control operates at sampled plants and not only in corporate policy. Trace one deployed system from intake through validation, deployment, monitoring, change and recovery. Challenge supplier assertions and verify evidence quality.

Legacy OT and Compensating Controls

Apply proportionately to architecture, consequence and deployment context.

Use compensating monitoring, gateways, signed offline transfer, procedural approval or independent protection where direct implementation is technically infeasible.

Failure Modes, Residual Risk and Limits

- unclear action authority
- ineffective human oversight
- AI path bypasses engineering change control
- no tested fallback

Residual risk remains from unknown failure modes, incomplete data, supplier opacity, plant variation, human factors and correlated dependencies. Risk acceptance must be explicit for production- or safety-affecting systems.

Implementation of this control does not establish functional safety, product conformity, legal compliance or incident prevention by itself.

D4-CTL-01 — AI BILL OF MATERIALS (AI BOM) MAINTENANCE

Field	Content
Domain	D4: SUPPLY CHAIN & THIRD-PARTY AI SECURITY
Applicability	Applicable
Sector relevance	High
Implementation track	Both
OT friction	Moderate
Conditional trigger	Applicable when the scoped system uses the capability or faces the risk addressed by the control.
Human oversight mode	Human-in-the-loop for advisory/approval workflows; human-on-the-loop or independent protection for high-speed OT.

Authoritative Requirement

Refer to GAISSF-NOR-001 v1.0 and GAISSF-NOR-004 v1.0 for the authoritative control requirement. This sector guide does not restate or modify normative control text.

Manufacturing Interpretation

In manufacturing, this control requires the organisation to apply ai bill of materials (ai bom) maintenance in a way that accounts for long-lived assets, constrained maintenance windows, plant-specific conditions, product conformity, physical consequences and separation of safety functions from AI.

Minimum and Enhanced Implementation

- Inventory suppliers and dependencies, define security requirements, obtain provenance and update commitments, restrict access, and maintain exit and continuity plans.
- Require signed releases, independent assurance, SBOM/AI-BOM evidence, continuous supplier monitoring, escrow or fallback arrangements, and tested replacement procedures.

Evidence and Assessment

approved scope and owner; risk assessment; architecture or data-flow diagram; configuration and change records; test results; monitoring or review records; contract clauses; supplier evidence; AI-BOM/SBOM; remote-access review

Confirm the control operates at sampled plants and not only in corporate policy. Trace one deployed system from intake through validation, deployment, monitoring, change and recovery. Challenge supplier assertions and verify evidence quality.

Legacy OT and Compensating Controls

Apply proportionately to architecture, consequence and deployment context.

Use compensating monitoring, gateways, signed offline transfer, procedural approval or independent protection where direct implementation is technically infeasible.

Failure Modes, Residual Risk and Limits

- accepting marketing claims as evidence

- persistent vendor access
- no model provenance
- no cloud outage contingency

Residual risk remains from unknown failure modes, incomplete data, supplier opacity, plant variation, human factors and correlated dependencies. Risk acceptance must be explicit for production- or safety-affecting systems.

Implementation of this control does not establish functional safety, product conformity, legal compliance or incident prevention by itself.

D4-CTL-02 — MODEL FILE & ARTIFACT SCANNING

Field	Content
Domain	D4: SUPPLY CHAIN & THIRD-PARTY AI SECURITY
Applicability	Applicable
Sector relevance	High
Implementation track	Both
OT friction	Moderate
Conditional trigger	Applicable when the scoped system uses the capability or faces the risk addressed by the control.
Human oversight mode	Human-in-the-loop for advisory/approval workflows; human-on-the-loop or independent protection for high-speed OT.

Authoritative Requirement

Refer to GAISSF-NOR-001 v1.0 and GAISSF-NOR-004 v1.0 for the authoritative control requirement. This sector guide does not restate or modify normative control text.

Manufacturing Interpretation

In manufacturing, this control requires the organisation to apply model file & artifact scanning in a way that accounts for long-lived assets, constrained maintenance windows, plant-specific conditions, product conformity, physical consequences and separation of safety functions from AI.

Minimum and Enhanced Implementation

- Inventory suppliers and dependencies, define security requirements, obtain provenance and update commitments, restrict access, and maintain exit and continuity plans.
- Require signed releases, independent assurance, SBOM/AI-BOM evidence, continuous supplier monitoring, escrow or fallback arrangements, and tested replacement procedures.

Evidence and Assessment

approved scope and owner; risk assessment; architecture or data-flow diagram; configuration and change records; test results; monitoring or review records; contract clauses; supplier evidence; AI-BOM/SBOM; remote-access review

Confirm the control operates at sampled plants and not only in corporate policy. Trace one deployed system from intake through validation, deployment, monitoring, change and recovery. Challenge supplier assertions and verify evidence quality.

Legacy OT and Compensating Controls

Apply proportionately to architecture, consequence and deployment context.

Use compensating monitoring, gateways, signed offline transfer, procedural approval or independent protection where direct implementation is technically infeasible.

Failure Modes, Residual Risk and Limits

- accepting marketing claims as evidence
- persistent vendor access
- no model provenance
- no cloud outage contingency

Residual risk remains from unknown failure modes, incomplete data, supplier opacity, plant variation, human factors and correlated dependencies. Risk acceptance must be explicit for production- or safety-affecting systems.

Implementation of this control does not establish functional safety, product conformity, legal compliance or incident prevention by itself.

D4-CTL-03 — MODEL HUB & REGISTRY VETTING

Field	Content
Domain	D4: SUPPLY CHAIN & THIRD-PARTY AI SECURITY
Applicability	Applicable
Sector relevance	High
Implementation track	Both
OT friction	Moderate
Conditional trigger	Applicable when the scoped system uses the capability or faces the risk addressed by the control.
Human oversight mode	Human-in-the-loop for advisory/approval workflows; human-on-the-loop or independent protection for high-speed OT.

Authoritative Requirement

Refer to GAISSF-NOR-001 v1.0 and GAISSF-NOR-004 v1.0 for the authoritative control requirement. This sector guide does not restate or modify normative control text.

Manufacturing Interpretation

In manufacturing, this control requires the organisation to apply model hub & registry vetting in a way that accounts for long-lived assets, constrained maintenance windows, plant-specific conditions, product conformity, physical consequences and separation of safety functions from AI.

Minimum and Enhanced Implementation

- Inventory suppliers and dependencies, define security requirements, obtain provenance and update commitments, restrict access, and maintain exit and continuity plans.
- Require signed releases, independent assurance, SBOM/AI-BOM evidence, continuous supplier monitoring, escrow or fallback arrangements, and tested replacement procedures.

Evidence and Assessment

approved scope and owner; risk assessment; architecture or data-flow diagram; configuration and change records; test results; monitoring or review records; contract clauses; supplier evidence; AI-BOM/SBOM; remote-access review

Confirm the control operates at sampled plants and not only in corporate policy. Trace one deployed system from intake through validation, deployment, monitoring, change and recovery. Challenge supplier assertions and verify evidence quality.

Legacy OT and Compensating Controls

Apply proportionately to architecture, consequence and deployment context.

Use compensating monitoring, gateways, signed offline transfer, procedural approval or independent protection where direct implementation is technically infeasible.

Failure Modes, Residual Risk and Limits

- accepting marketing claims as evidence
- persistent vendor access
- no model provenance
- no cloud outage contingency

Residual risk remains from unknown failure modes, incomplete data, supplier opacity, plant variation, human factors and correlated dependencies. Risk acceptance must be explicit for production- or safety-affecting systems.

Implementation of this control does not establish functional safety, product conformity, legal compliance or incident prevention by itself.

D4-CTL-04 — MCP SERVER BEHAVIORAL MONITORING

Field	Content
Domain	D4: SUPPLY CHAIN & THIRD-PARTY AI SECURITY
Applicability	Applicable
Sector relevance	High
Implementation track	Both
OT friction	Moderate
Conditional trigger	Applicable when the scoped system uses the capability or faces the risk addressed by the control.
Human oversight mode	Human-in-the-loop for advisory/approval workflows; human-on-the-loop or independent protection for high-speed OT.

Authoritative Requirement

Refer to GAISSF-NOR-001 v1.0 and GAISSF-NOR-004 v1.0 for the authoritative control requirement. This sector guide does not restate or modify normative control text.

Manufacturing Interpretation

In manufacturing, this control requires the organisation to apply mcp server behavioral monitoring in a way that accounts for long-lived assets, constrained maintenance windows, plant-specific conditions, product conformity, physical consequences and separation of safety functions from AI.

Minimum and Enhanced Implementation

- Inventory suppliers and dependencies, define security requirements, obtain provenance and update commitments, restrict access, and maintain exit and continuity plans.
- Require signed releases, independent assurance, SBOM/AI-BOM evidence, continuous supplier monitoring, escrow or fallback arrangements, and tested replacement procedures.

Evidence and Assessment

approved scope and owner; risk assessment; architecture or data-flow diagram; configuration and change records; test results; monitoring or review records; contract clauses; supplier evidence; AI-BOM/SBOM; remote-access review

Confirm the control operates at sampled plants and not only in corporate policy. Trace one deployed system from intake through validation, deployment, monitoring, change and recovery. Challenge supplier assertions and verify evidence quality.

Legacy OT and Compensating Controls

Apply proportionately to architecture, consequence and deployment context.

Use compensating monitoring, gateways, signed offline transfer, procedural approval or independent protection where direct implementation is technically infeasible.

Failure Modes, Residual Risk and Limits

- accepting marketing claims as evidence
- persistent vendor access
- no model provenance
- no cloud outage contingency

Residual risk remains from unknown failure modes, incomplete data, supplier opacity, plant variation, human factors and correlated dependencies. Risk acceptance must be explicit for production- or safety-affecting systems.

Implementation of this control does not establish functional safety, product conformity, legal compliance or incident prevention by itself.

D4-CTL-05 — THIRD-PARTY AI API SECURITY ASSESSMENT

Field	Content
Domain	D4: SUPPLY CHAIN & THIRD-PARTY AI SECURITY
Applicability	Applicable
Sector relevance	High
Implementation track	Both
OT friction	Moderate
Conditional trigger	Applicable when the scoped system uses the capability or faces the risk addressed by the control.
Human oversight mode	Human-in-the-loop for advisory/approval workflows; human-on-the-loop or independent protection for high-speed OT.

Authoritative Requirement

Refer to GAISSF-NOR-001 v1.0 and GAISSF-NOR-004 v1.0 for the authoritative control requirement. This sector guide does not restate or modify normative control text.

Manufacturing Interpretation

In manufacturing, this control requires the organisation to apply third-party ai api security assessment in a way that accounts for long-lived assets, constrained maintenance windows, plant-specific conditions, product conformity, physical consequences and separation of safety functions from AI.

Minimum and Enhanced Implementation

- Inventory suppliers and dependencies, define security requirements, obtain provenance and update commitments, restrict access, and maintain exit and continuity plans.
- Require signed releases, independent assurance, SBOM/AI-BOM evidence, continuous supplier monitoring, escrow or fallback arrangements, and tested replacement procedures.

Evidence and Assessment

approved scope and owner; risk assessment; architecture or data-flow diagram; configuration and change records; test results; monitoring or review records; contract clauses; supplier evidence; AI-BOM/SBOM; remote-access review

Confirm the control operates at sampled plants and not only in corporate policy. Trace one deployed system from intake through validation, deployment, monitoring, change and recovery. Challenge supplier assertions and verify evidence quality.

Legacy OT and Compensating Controls

Apply proportionately to architecture, consequence and deployment context.

Use compensating monitoring, gateways, signed offline transfer, procedural approval or independent protection where direct implementation is technically infeasible.

Failure Modes, Residual Risk and Limits

- accepting marketing claims as evidence
- persistent vendor access
- no model provenance
- no cloud outage contingency

Residual risk remains from unknown failure modes, incomplete data, supplier opacity, plant variation, human factors and correlated dependencies. Risk acceptance must be explicit for production- or safety-affecting systems.

Implementation of this control does not establish functional safety, product conformity, legal compliance or incident prevention by itself.

D4-CTL-06 — SHADOW AI DISCOVERY & GOVERNANCE

Field	Content
Domain	D4: SUPPLY CHAIN & THIRD-PARTY AI SECURITY
Applicability	Applicable
Sector relevance	High
Implementation track	Both
OT friction	Moderate
Conditional trigger	Applicable when the scoped system uses the capability or faces the risk addressed by the control.
Human oversight mode	Human-in-the-loop for advisory/approval workflows; human-on-the-loop or independent protection for high-speed OT.

Authoritative Requirement

Refer to GAISSF-NOR-001 v1.0 and GAISSF-NOR-004 v1.0 for the authoritative control requirement. This sector guide does not restate or modify normative control text.

Manufacturing Interpretation

In manufacturing, this control requires the organisation to apply shadow ai discovery & governance in a way that accounts for long-lived assets, constrained maintenance windows, plant-specific conditions, product conformity, physical consequences and separation of safety functions from AI.

Minimum and Enhanced Implementation

- Inventory suppliers and dependencies, define security requirements, obtain provenance and update commitments, restrict access, and maintain exit and continuity plans.
- Require signed releases, independent assurance, SBOM/AI-BOM evidence, continuous supplier monitoring, escrow or fallback arrangements, and tested replacement procedures.

Evidence and Assessment

approved scope and owner; risk assessment; architecture or data-flow diagram; configuration and change records; test results; monitoring or review records; contract clauses; supplier evidence; AI-BOM/SBOM; remote-access review

Confirm the control operates at sampled plants and not only in corporate policy. Trace one deployed system from intake through validation, deployment, monitoring, change and recovery. Challenge supplier assertions and verify evidence quality.

Legacy OT and Compensating Controls

Apply proportionately to architecture, consequence and deployment context.

Use compensating monitoring, gateways, signed offline transfer, procedural approval or independent protection where direct implementation is technically infeasible.

Failure Modes, Residual Risk and Limits

- accepting marketing claims as evidence
- persistent vendor access
- no model provenance
- no cloud outage contingency

Residual risk remains from unknown failure modes, incomplete data, supplier opacity, plant variation, human factors and correlated dependencies. Risk acceptance must be explicit for production- or safety-affecting systems.

Implementation of this control does not establish functional safety, product conformity, legal compliance or incident prevention by itself.

D4-CTL-07 — AI SOFTWARE COMPOSITION ANALYSIS (SCA)

Field	Content
Domain	D4: SUPPLY CHAIN & THIRD-PARTY AI SECURITY
Applicability	Applicable
Sector relevance	High
Implementation track	Both
OT friction	Moderate
Conditional trigger	Applicable when the scoped system uses the capability or faces the risk addressed by the control.
Human oversight mode	Human-in-the-loop for advisory/approval workflows; human-on-the-loop or independent protection for high-speed OT.

Authoritative Requirement

Refer to GAISSF-NOR-001 v1.0 and GAISSF-NOR-004 v1.0 for the authoritative control requirement. This sector guide does not restate or modify normative control text.

Manufacturing Interpretation

In manufacturing, this control requires the organisation to apply ai software composition analysis (sca) in a way that accounts for long-lived assets, constrained maintenance windows, plant-specific conditions, product conformity, physical consequences and separation of safety functions from AI.

Minimum and Enhanced Implementation

- Inventory suppliers and dependencies, define security requirements, obtain provenance and update commitments, restrict access, and maintain exit and continuity plans.
- Require signed releases, independent assurance, SBOM/AI-BOM evidence, continuous supplier monitoring, escrow or fallback arrangements, and tested replacement procedures.

Evidence and Assessment

approved scope and owner; risk assessment; architecture or data-flow diagram; configuration and change records; test results; monitoring or review records; contract clauses; supplier evidence; AI-BOM/SBOM; remote-access review

Confirm the control operates at sampled plants and not only in corporate policy. Trace one deployed system from intake through validation, deployment, monitoring, change and recovery. Challenge supplier assertions and verify evidence quality.

Legacy OT and Compensating Controls

Apply proportionately to architecture, consequence and deployment context.

Use compensating monitoring, gateways, signed offline transfer, procedural approval or independent protection where direct implementation is technically infeasible.

Failure Modes, Residual Risk and Limits

- accepting marketing claims as evidence

- persistent vendor access
- no model provenance
- no cloud outage contingency

Residual risk remains from unknown failure modes, incomplete data, supplier opacity, plant variation, human factors and correlated dependencies. Risk acceptance must be explicit for production- or safety-affecting systems.

Implementation of this control does not establish functional safety, product conformity, legal compliance or incident prevention by itself.

D5-CTL-01 — HARMFUL CONTENT BLOCKING

Field	Content
Domain	D5: CONTENT SAFETY & OUTPUT INTEGRITY
Applicability	Applicable
Sector relevance	High
Implementation track	Both
OT friction	Moderate
Conditional trigger	Applicable when the scoped system uses the capability or faces the risk addressed by the control.
Human oversight mode	Human-in-the-loop for advisory/approval workflows; human-on-the-loop or independent protection for high-speed OT.

Authoritative Requirement

Refer to GAISSF-NOR-001 v1.0 and GAISSF-NOR-004 v1.0 for the authoritative control requirement. This sector guide does not restate or modify normative control text.

Manufacturing Interpretation

In manufacturing, this control requires the organisation to apply harmful content blocking in a way that accounts for long-lived assets, constrained maintenance windows, plant-specific conditions, product conformity, physical consequences and separation of safety functions from AI.

Minimum and Enhanced Implementation

- Define acceptable outputs, validate critical results, communicate uncertainty, maintain human review, and prevent unsafe or prohibited content from entering production workflows.
- Use independent verification, output policy enforcement, confidence calibration, structured validation, traceable citations, and automatic quarantine of suspect outputs.

Evidence and Assessment

approved scope and owner; risk assessment; architecture or data-flow diagram; configuration and change records; test results; monitoring or review records; output acceptance criteria; quality verification; exception and quarantine records

Confirm the control operates at sampled plants and not only in corporate policy. Trace one deployed system from intake through validation, deployment, monitoring, change and recovery. Challenge supplier assertions and verify evidence quality.

Legacy OT and Compensating Controls

Apply proportionately to architecture, consequence and deployment context.

Use compensating monitoring, gateways, signed offline transfer, procedural approval or independent protection where direct implementation is technically infeasible.

Failure Modes, Residual Risk and Limits

- confidence treated as correctness
- no independent verification
- generated instructions used without engineering approval

Residual risk remains from unknown failure modes, incomplete data, supplier opacity, plant variation, human factors and correlated dependencies. Risk acceptance must be explicit for production- or safety-affecting systems.

Implementation of this control does not establish functional safety, product conformity, legal compliance or incident prevention by itself.

D5-CTL-02 — PII LEAKAGE PREVENTION

Field	Content
Domain	D5: CONTENT SAFETY & OUTPUT INTEGRITY
Applicability	Applicable
Sector relevance	High
Implementation track	Both
OT friction	Moderate
Conditional trigger	Applicable when the scoped system uses the capability or faces the risk addressed by the control.
Human oversight mode	Human-in-the-loop for advisory/approval workflows; human-on-the-loop or independent protection for high-speed OT.

Authoritative Requirement

Refer to GAISSF-NOR-001 v1.0 and GAISSF-NOR-004 v1.0 for the authoritative control requirement. This sector guide does not restate or modify normative control text.

Manufacturing Interpretation

In manufacturing, this control requires the organisation to apply pii leakage prevention in a way that accounts for long-lived assets, constrained maintenance windows, plant-specific conditions, product conformity, physical consequences and separation of safety functions from AI.

Minimum and Enhanced Implementation

- Define acceptable outputs, validate critical results, communicate uncertainty, maintain human review, and prevent unsafe or prohibited content from entering production workflows.
- Use independent verification, output policy enforcement, confidence calibration, structured validation, traceable citations, and automatic quarantine of suspect outputs.

Evidence and Assessment

approved scope and owner; risk assessment; architecture or data-flow diagram; configuration and change records; test results; monitoring or review records; output acceptance criteria; quality verification; exception and quarantine records

Confirm the control operates at sampled plants and not only in corporate policy. Trace one deployed system from intake through validation, deployment, monitoring, change and recovery. Challenge supplier assertions and verify evidence quality.

Legacy OT and Compensating Controls

Apply proportionately to architecture, consequence and deployment context.

Use compensating monitoring, gateways, signed offline transfer, procedural approval or independent protection where direct implementation is technically infeasible.

Failure Modes, Residual Risk and Limits

- confidence treated as correctness
- no independent verification
- generated instructions used without engineering approval

Residual risk remains from unknown failure modes, incomplete data, supplier opacity, plant variation, human factors and correlated dependencies. Risk acceptance must be explicit for production- or safety-affecting systems.

Implementation of this control does not establish functional safety, product conformity, legal compliance or incident prevention by itself.

D5-CTL-03 — COPYRIGHT DETECTION

Field	Content
Domain	D5: CONTENT SAFETY & OUTPUT INTEGRITY
Applicability	Applicable
Sector relevance	High
Implementation track	Both
OT friction	Moderate
Conditional trigger	Applicable when the scoped system uses the capability or faces the risk addressed by the control.
Human oversight mode	Human-in-the-loop for advisory/approval workflows; human-on-the-loop or independent protection for high-speed OT.

Authoritative Requirement

Refer to GAISSF-NOR-001 v1.0 and GAISSF-NOR-004 v1.0 for the authoritative control requirement. This sector guide does not restate or modify normative control text.

Manufacturing Interpretation

In manufacturing, this control requires the organisation to apply copyright detection in a way that accounts for long-lived assets, constrained maintenance windows, plant-specific conditions, product conformity, physical consequences and separation of safety functions from AI.

Minimum and Enhanced Implementation

- Define acceptable outputs, validate critical results, communicate uncertainty, maintain human review, and prevent unsafe or prohibited content from entering production workflows.
- Use independent verification, output policy enforcement, confidence calibration, structured validation, traceable citations, and automatic quarantine of suspect outputs.

Evidence and Assessment

approved scope and owner; risk assessment; architecture or data-flow diagram; configuration and change records; test results; monitoring or review records; output acceptance criteria; quality verification; exception and quarantine records

Confirm the control operates at sampled plants and not only in corporate policy. Trace one deployed system from intake through validation, deployment, monitoring, change and recovery. Challenge supplier assertions and verify evidence quality.

Legacy OT and Compensating Controls

Apply proportionately to architecture, consequence and deployment context.

Use compensating monitoring, gateways, signed offline transfer, procedural approval or independent protection where direct implementation is technically infeasible.

Failure Modes, Residual Risk and Limits

- confidence treated as correctness
- no independent verification
- generated instructions used without engineering approval

Residual risk remains from unknown failure modes, incomplete data, supplier opacity, plant variation, human factors and correlated dependencies. Risk acceptance must be explicit for production- or safety-affecting systems.

Implementation of this control does not establish functional safety, product conformity, legal compliance or incident prevention by itself.

D5-CTL-04 — AI WATERMARKING ROBUSTNESS

Field	Content
Domain	D5: CONTENT SAFETY & OUTPUT INTEGRITY
Applicability	Applicable
Sector relevance	Conditional
Implementation track	Both
OT friction	Low
Conditional trigger	Core for public-facing synthetic media or externally distributed content; generally low relevance to internal OT telemetry and machine-vision inspection unless outputs leave the controlled environment.
Human oversight mode	Human-in-the-loop for advisory/approval workflows; human-on-the-loop or independent protection for high-speed OT.

Authoritative Requirement

Refer to GAISSF-NOR-001 v1.0 and GAISSF-NOR-004 v1.0 for the authoritative control requirement. This sector guide does not restate or modify normative control text.

Manufacturing Interpretation

In manufacturing, this control requires the organisation to apply ai watermarking robustness in a way that accounts for long-lived assets, constrained maintenance windows, plant-specific conditions, product conformity, physical consequences and separation of safety functions from AI.

Minimum and Enhanced Implementation

- Define acceptable outputs, validate critical results, communicate uncertainty, maintain human review, and prevent unsafe or prohibited content from entering production workflows.
- Use independent verification, output policy enforcement, confidence calibration, structured validation, traceable citations, and automatic quarantine of suspect outputs.

Evidence and Assessment

approved scope and owner; risk assessment; architecture or data-flow diagram; configuration and change records; test results; monitoring or review records; output acceptance criteria; quality verification; exception and quarantine records

Confirm the control operates at sampled plants and not only in corporate policy. Trace one deployed system from intake through validation, deployment, monitoring, change and recovery. Challenge supplier assertions and verify evidence quality.

Legacy OT and Compensating Controls

Document why provenance or watermarking is required or not required for the scoped output channel.

Use compensating monitoring, gateways, signed offline transfer, procedural approval or independent protection where direct implementation is technically infeasible.

Failure Modes, Residual Risk and Limits

- confidence treated as correctness
- no independent verification
- generated instructions used without engineering approval

Residual risk remains from unknown failure modes, incomplete data, supplier opacity, plant variation, human factors and correlated dependencies. Risk acceptance must be explicit for production- or safety-affecting systems.

Implementation of this control does not establish functional safety, product conformity, legal compliance or incident prevention by itself.

D5-CTL-05 — PRIVACY-BY-DESIGN VERIFICATION

Field	Content
Domain	D5: CONTENT SAFETY & OUTPUT INTEGRITY
Applicability	Applicable
Sector relevance	High
Implementation track	Both
OT friction	Moderate
Conditional trigger	Applicable when the scoped system uses the capability or faces the risk addressed by the control.
Human oversight mode	Human-in-the-loop for advisory/approval workflows; human-on-the-loop or independent protection for high-speed OT.

Authoritative Requirement

Refer to GAISSF-NOR-001 v1.0 and GAISSF-NOR-004 v1.0 for the authoritative control requirement. This sector guide does not restate or modify normative control text.

Manufacturing Interpretation

In manufacturing, this control requires the organisation to apply privacy-by-design verification in a way that accounts for long-lived assets, constrained maintenance windows, plant-specific conditions, product conformity, physical consequences and separation of safety functions from AI.

Minimum and Enhanced Implementation

- Define acceptable outputs, validate critical results, communicate uncertainty, maintain human review, and prevent unsafe or prohibited content from entering production workflows.
- Use independent verification, output policy enforcement, confidence calibration, structured validation, traceable citations, and automatic quarantine of suspect outputs.

Evidence and Assessment

approved scope and owner; risk assessment; architecture or data-flow diagram; configuration and change records; test results; monitoring or review records; output acceptance criteria; quality verification; exception and quarantine records

Confirm the control operates at sampled plants and not only in corporate policy. Trace one deployed system from intake through validation, deployment, monitoring, change and recovery. Challenge supplier assertions and verify evidence quality.

Legacy OT and Compensating Controls

Apply proportionately to architecture, consequence and deployment context.

Use compensating monitoring, gateways, signed offline transfer, procedural approval or independent protection where direct implementation is technically infeasible.

Failure Modes, Residual Risk and Limits

- confidence treated as correctness
- no independent verification
- generated instructions used without engineering approval

Residual risk remains from unknown failure modes, incomplete data, supplier opacity, plant variation, human factors and correlated dependencies. Risk acceptance must be explicit for production- or safety-affecting systems.

Implementation of this control does not establish functional safety, product conformity, legal compliance or incident prevention by itself.

D5-CTL-06 — PRIVACY-PRESERVING ML VALIDATION

Field	Content
Domain	D5: CONTENT SAFETY & OUTPUT INTEGRITY
Applicability	Applicable
Sector relevance	High
Implementation track	Both
OT friction	Moderate
Conditional trigger	Applicable when the scoped system uses the capability or faces the risk addressed by the control.
Human oversight mode	Human-in-the-loop for advisory/approval workflows; human-on-the-loop or independent protection for high-speed OT.

Authoritative Requirement

Refer to GAISSF-NOR-001 v1.0 and GAISSF-NOR-004 v1.0 for the authoritative control requirement. This sector guide does not restate or modify normative control text.

Manufacturing Interpretation

In manufacturing, this control requires the organisation to apply privacy-preserving ml validation in a way that accounts for long-lived assets, constrained maintenance windows, plant-specific conditions, product conformity, physical consequences and separation of safety functions from AI.

Minimum and Enhanced Implementation

- Define acceptable outputs, validate critical results, communicate uncertainty, maintain human review, and prevent unsafe or prohibited content from entering production workflows.
- Use independent verification, output policy enforcement, confidence calibration, structured validation, traceable citations, and automatic quarantine of suspect outputs.

Evidence and Assessment

approved scope and owner; risk assessment; architecture or data-flow diagram; configuration and change records; test results; monitoring or review records; output acceptance criteria; quality verification; exception and quarantine records

Confirm the control operates at sampled plants and not only in corporate policy. Trace one deployed system from intake through validation, deployment, monitoring, change and recovery. Challenge supplier assertions and verify evidence quality.

Legacy OT and Compensating Controls

Apply proportionately to architecture, consequence and deployment context.

Use compensating monitoring, gateways, signed offline transfer, procedural approval or independent protection where direct implementation is technically infeasible.

Failure Modes, Residual Risk and Limits

- confidence treated as correctness
- no independent verification
- generated instructions used without engineering approval

Residual risk remains from unknown failure modes, incomplete data, supplier opacity, plant variation, human factors and correlated dependencies. Risk acceptance must be explicit for production- or safety-affecting systems.

Implementation of this control does not establish functional safety, product conformity, legal compliance or incident prevention by itself.

D6-CTL-01 — HUMAN-IN-THE-LOOP FOR HIGH-RISK ACTIONS

Field	Content
Domain	D6: GOVERNANCE, ACCOUNTABILITY & HUMAN OVERSIGHT
Applicability	Applicable
Sector relevance	Core
Implementation track	Enterprise Governance
OT friction	Low
Conditional trigger	Applicable when the scoped system uses the capability or faces the risk addressed by the control.
Human oversight mode	Select human-in-the-loop, human-on-the-loop, or human-out-of-the-loop with independent protection based on process speed and consequence.

Authoritative Requirement

Refer to GAISSF-NOR-001 v1.0 and GAISSF-NOR-004 v1.0 for the authoritative control requirement. This sector guide does not restate or modify normative control text.

Manufacturing Interpretation

In manufacturing, this control requires the organisation to apply human-in-the-loop for high-risk actions in a way that accounts for long-lived assets, constrained maintenance windows, plant-specific conditions, product conformity, physical consequences and separation of safety functions from AI.

Minimum and Enhanced Implementation

- Maintain an AI inventory, assign accountable owners, classify criticality, approve risk, define policies, and retain evidence.
- Integrate plant-level assurance with enterprise governance, continuous control monitoring, independent challenge, board reporting, and multi-site conformance review.

Evidence and Assessment

approved scope and owner; risk assessment; architecture or data-flow diagram; configuration and change records; test results; monitoring or review records; AI inventory; RACI; risk acceptance; governance minutes

Confirm the control operates at sampled plants and not only in corporate policy. Trace one deployed system from intake through validation, deployment, monitoring, change and recovery. Challenge supplier assertions and verify evidence quality.

Legacy OT and Compensating Controls

Apply proportionately to architecture, consequence and deployment context.

Use compensating monitoring, gateways, signed offline transfer, procedural approval or independent protection where direct implementation is technically infeasible.

Failure Modes, Residual Risk and Limits

- central policy without plant implementation
- incomplete inventory
- unclear ownership
- undocumented exceptions

Residual risk remains from unknown failure modes, incomplete data, supplier opacity, plant variation, human factors and correlated dependencies. Risk acceptance must be explicit for production- or safety-affecting systems.

Implementation of this control does not establish functional safety, product conformity, legal compliance or incident prevention by itself.

D6-CTL-02 — AUDIT TRAIL COMPLETENESS

Field	Content
Domain	D6: GOVERNANCE, ACCOUNTABILITY & HUMAN OVERSIGHT
Applicability	Applicable
Sector relevance	Core
Implementation track	Enterprise Governance
OT friction	Low

Field	Content
Conditional trigger	Applicable when the scoped system uses the capability or faces the risk addressed by the control.
Human oversight mode	Human-in-the-loop for advisory/approval workflows; human-on-the-loop or independent protection for high-speed OT.

Authoritative Requirement

Refer to GAISSF-NOR-001 v1.0 and GAISSF-NOR-004 v1.0 for the authoritative control requirement. This sector guide does not restate or modify normative control text.

Manufacturing Interpretation

In manufacturing, this control requires the organisation to apply audit trail completeness in a way that accounts for long-lived assets, constrained maintenance windows, plant-specific conditions, product conformity, physical consequences and separation of safety functions from AI.

Minimum and Enhanced Implementation

- Maintain an AI inventory, assign accountable owners, classify criticality, approve risk, define policies, and retain evidence.
- Integrate plant-level assurance with enterprise governance, continuous control monitoring, independent challenge, board reporting, and multi-site conformance review.

Evidence and Assessment

approved scope and owner; risk assessment; architecture or data-flow diagram; configuration and change records; test results; monitoring or review records; AI inventory; RACI; risk acceptance; governance minutes

Confirm the control operates at sampled plants and not only in corporate policy. Trace one deployed system from intake through validation, deployment, monitoring, change and recovery. Challenge supplier assertions and verify evidence quality.

Legacy OT and Compensating Controls

Apply proportionately to architecture, consequence and deployment context.

Use compensating monitoring, gateways, signed offline transfer, procedural approval or independent protection where direct implementation is technically infeasible.

Failure Modes, Residual Risk and Limits

- central policy without plant implementation
- incomplete inventory
- unclear ownership
- undocumented exceptions

Residual risk remains from unknown failure modes, incomplete data, supplier opacity, plant variation, human factors and correlated dependencies. Risk acceptance must be explicit for production- or safety-affecting systems.

Implementation of this control does not establish functional safety, product conformity, legal compliance or incident prevention by itself.

D6-CTL-03 — AI MODEL CARD COMPLETENESS

Field	Content
Domain	D6: GOVERNANCE, ACCOUNTABILITY & HUMAN OVERSIGHT
Applicability	Applicable
Sector relevance	Core
Implementation track	Enterprise Governance
OT friction	Low
Conditional trigger	Applicable when the scoped system uses the capability or faces the risk addressed by the control.
Human oversight mode	Human-in-the-loop for advisory/approval workflows; human-on-the-loop or independent protection for high-speed OT.

Authoritative Requirement

Refer to GAISSF-NOR-001 v1.0 and GAISSF-NOR-004 v1.0 for the authoritative control requirement. This sector guide does not restate or modify normative control text.

Manufacturing Interpretation

In manufacturing, this control requires the organisation to apply ai model card completeness in a way that accounts for long-lived assets, constrained maintenance windows, plant-specific conditions, product conformity, physical consequences and separation of safety functions from AI.

Minimum and Enhanced Implementation

- Maintain an AI inventory, assign accountable owners, classify criticality, approve risk, define policies, and retain evidence.
- Integrate plant-level assurance with enterprise governance, continuous control monitoring, independent challenge, board reporting, and multi-site conformance review.

Evidence and Assessment

approved scope and owner; risk assessment; architecture or data-flow diagram; configuration and change records; test results; monitoring or review records; AI inventory; RACI; risk acceptance; governance minutes

Confirm the control operates at sampled plants and not only in corporate policy. Trace one deployed system from intake through validation, deployment, monitoring, change and recovery. Challenge supplier assertions and verify evidence quality.

Legacy OT and Compensating Controls

Apply proportionately to architecture, consequence and deployment context.

Use compensating monitoring, gateways, signed offline transfer, procedural approval or independent protection where direct implementation is technically infeasible.

Failure Modes, Residual Risk and Limits

- central policy without plant implementation
- incomplete inventory
- unclear ownership
- undocumented exceptions

Residual risk remains from unknown failure modes, incomplete data, supplier opacity, plant variation, human factors and correlated dependencies. Risk acceptance must be explicit for production- or safety-affecting systems.

Implementation of this control does not establish functional safety, product conformity, legal compliance or incident prevention by itself.

D6-CTL-04 — AI INCIDENT RESPONSE READINESS

Field	Content
Domain	D6: GOVERNANCE, ACCOUNTABILITY & HUMAN OVERSIGHT
Applicability	Applicable
Sector relevance	Core
Implementation track	Enterprise Governance
OT friction	Low
Conditional trigger	Applicable when the scoped system uses the capability or faces the risk addressed by the control.
Human oversight mode	Human-in-the-loop for advisory/approval workflows; human-on-the-loop or independent protection for high-speed OT.

Authoritative Requirement

Refer to GAISSF-NOR-001 v1.0 and GAISSF-NOR-004 v1.0 for the authoritative control requirement. This sector guide does not restate or modify normative control text.

Manufacturing Interpretation

In manufacturing, this control requires the organisation to apply ai incident response readiness in a way that accounts for long-lived assets, constrained maintenance windows, plant-specific conditions, product conformity, physical consequences and separation of safety functions from AI.

Minimum and Enhanced Implementation

- Maintain an AI inventory, assign accountable owners, classify criticality, approve risk, define policies, and retain evidence.
- Integrate plant-level assurance with enterprise governance, continuous control monitoring, independent challenge, board reporting, and multi-site conformance review.

Evidence and Assessment

approved scope and owner; risk assessment; architecture or data-flow diagram; configuration and change records; test results; monitoring or review records; AI inventory; RACI; risk acceptance; governance minutes

Confirm the control operates at sampled plants and not only in corporate policy. Trace one deployed system from intake through validation, deployment, monitoring, change and recovery. Challenge supplier assertions and verify evidence quality.

Legacy OT and Compensating Controls

Apply proportionately to architecture, consequence and deployment context.

Use compensating monitoring, gateways, signed offline transfer, procedural approval or independent protection where direct implementation is technically infeasible.

Failure Modes, Residual Risk and Limits

- central policy without plant implementation
- incomplete inventory
- unclear ownership
- undocumented exceptions

Residual risk remains from unknown failure modes, incomplete data, supplier opacity, plant variation, human factors and correlated dependencies. Risk acceptance must be explicit for production- or safety-affecting systems.

Implementation of this control does not establish functional safety, product conformity, legal compliance or incident prevention by itself.

D6-CTL-05 — MODEL DEPRECIATION & DECOMMISSIONING

Field	Content
Domain	D6: GOVERNANCE, ACCOUNTABILITY & HUMAN OVERSIGHT
Applicability	Applicable
Sector relevance	Core
Implementation track	Enterprise Governance
OT friction	Low
Conditional trigger	Applicable when the scoped system uses the capability or faces the risk addressed by the control.
Human oversight mode	Human-in-the-loop for advisory/approval workflows; human-on-the-loop or independent protection for high-speed OT.

Authoritative Requirement

Refer to GAISSF-NOR-001 v1.0 and GAISSF-NOR-004 v1.0 for the authoritative control requirement. This sector guide does not restate or modify normative control text.

Manufacturing Interpretation

In manufacturing, this control requires the organisation to apply model depreciation & decommissioning in a way that accounts for long-lived assets, constrained maintenance windows, plant-specific conditions, product conformity, physical consequences and separation of safety functions from AI.

Minimum and Enhanced Implementation

- Maintain an AI inventory, assign accountable owners, classify criticality, approve risk, define policies, and retain evidence.
- Integrate plant-level assurance with enterprise governance, continuous control monitoring, independent challenge, board reporting, and multi-site conformance review.

Evidence and Assessment

approved scope and owner; risk assessment; architecture or data-flow diagram; configuration and change records; test results; monitoring or review records; AI inventory; RACI; risk acceptance; governance minutes

Confirm the control operates at sampled plants and not only in corporate policy. Trace one deployed system from intake through validation, deployment, monitoring, change and recovery. Challenge supplier assertions and verify evidence quality.

Legacy OT and Compensating Controls

Apply proportionately to architecture, consequence and deployment context.

Use compensating monitoring, gateways, signed offline transfer, procedural approval or independent protection where direct implementation is technically infeasible.

Failure Modes, Residual Risk and Limits

- central policy without plant implementation
- incomplete inventory
- unclear ownership
- undocumented exceptions

Residual risk remains from unknown failure modes, incomplete data, supplier opacity, plant variation, human factors and correlated dependencies. Risk acceptance must be explicit for production- or safety-affecting systems.

Implementation of this control does not establish functional safety, product conformity, legal compliance or incident prevention by itself.

D6-CTL-06 — THIRD-PARTY AI VENDOR GOVERNANCE

Field	Content
Domain	D6: GOVERNANCE, ACCOUNTABILITY & HUMAN OVERSIGHT
Applicability	Applicable
Sector relevance	Core
Implementation track	Enterprise Governance
OT friction	Low
Conditional trigger	Applicable when the scoped system uses the capability or faces the risk addressed by the control.
Human oversight mode	Human-in-the-loop for advisory/approval workflows; human-on-the-loop or independent protection for high-speed OT.

Authoritative Requirement

Refer to GAISSF-NOR-001 v1.0 and GAISSF-NOR-004 v1.0 for the authoritative control requirement. This sector guide does not restate or modify normative control text.

Manufacturing Interpretation

In manufacturing, this control requires the organisation to apply third-party ai vendor governance in a way that accounts for long-lived assets, constrained maintenance windows, plant-specific conditions, product conformity, physical consequences and separation of safety functions from AI.

Minimum and Enhanced Implementation

- Maintain an AI inventory, assign accountable owners, classify criticality, approve risk, define policies, and retain evidence.
- Integrate plant-level assurance with enterprise governance, continuous control monitoring, independent challenge, board reporting, and multi-site conformance review.

Evidence and Assessment

approved scope and owner; risk assessment; architecture or data-flow diagram; configuration and change records; test results; monitoring or review records; AI inventory; RACI; risk acceptance; governance minutes

Confirm the control operates at sampled plants and not only in corporate policy. Trace one deployed system from intake through validation, deployment, monitoring, change and recovery. Challenge supplier assertions and verify evidence quality.

Legacy OT and Compensating Controls

Apply proportionately to architecture, consequence and deployment context.

Use compensating monitoring, gateways, signed offline transfer, procedural approval or independent protection where direct implementation is technically infeasible.

Failure Modes, Residual Risk and Limits

- central policy without plant implementation
- incomplete inventory
- unclear ownership
- undocumented exceptions

Residual risk remains from unknown failure modes, incomplete data, supplier opacity, plant variation, human factors and correlated dependencies. Risk acceptance must be explicit for production- or safety-affecting systems.

Implementation of this control does not establish functional safety, product conformity, legal compliance or incident prevention by itself.

D6-CTL-07 — AI RESILIENCE & BUSINESS CONTINUITY

Field	Content
Domain	D6: GOVERNANCE, ACCOUNTABILITY & HUMAN OVERSIGHT
Applicability	Applicable
Sector relevance	Core
Implementation track	Enterprise Governance
OT friction	Low
Conditional trigger	Applicable when the scoped system uses the capability or faces the risk addressed by the control.
Human oversight mode	Human-in-the-loop for advisory/approval workflows; human-on-the-loop or independent protection for high-speed OT.

Authoritative Requirement

Refer to GAISSF-NOR-001 v1.0 and GAISSF-NOR-004 v1.0 for the authoritative control requirement. This sector guide does not restate or modify normative control text.

Manufacturing Interpretation

In manufacturing, this control requires the organisation to apply ai resilience & business continuity in a way that accounts for long-lived assets, constrained maintenance windows, plant-specific conditions, product conformity, physical consequences and separation of safety functions from AI.

Minimum and Enhanced Implementation

- Maintain an AI inventory, assign accountable owners, classify criticality, approve risk, define policies, and retain evidence.
- Integrate plant-level assurance with enterprise governance, continuous control monitoring, independent challenge, board reporting, and multi-site conformance review.

Evidence and Assessment

approved scope and owner; risk assessment; architecture or data-flow diagram; configuration and change records; test results; monitoring or review records; AI inventory; RACI; risk acceptance; governance minutes

Confirm the control operates at sampled plants and not only in corporate policy. Trace one deployed system from intake through validation, deployment, monitoring, change and recovery. Challenge supplier assertions and verify evidence quality.

Legacy OT and Compensating Controls

Apply proportionately to architecture, consequence and deployment context.

Use compensating monitoring, gateways, signed offline transfer, procedural approval or independent protection where direct implementation is technically infeasible.

Failure Modes, Residual Risk and Limits

- central policy without plant implementation
- incomplete inventory
- unclear ownership
- undocumented exceptions

Residual risk remains from unknown failure modes, incomplete data, supplier opacity, plant variation, human factors and correlated dependencies. Risk acceptance must be explicit for production- or safety-affecting systems.

Implementation of this control does not establish functional safety, product conformity, legal compliance or incident prevention by itself.

D7-CTL-H01 — AI-GENERATED PHISHING SIMULATION

Field	Content
Domain	D7: HUMAN & SOCIETAL HARMS
Applicability	Applicable
Sector relevance	High
Implementation track	Both
OT friction	Low
Conditional trigger	Applicable when the scoped system uses the capability or faces the risk addressed by the control.
Human oversight mode	Select human-in-the-loop, human-on-the-loop, or human-out-of-the-loop with independent protection based on process speed and consequence.

Authoritative Requirement

Refer to GAISSF-NOR-001 v1.0 and GAISSF-NOR-004 v1.0 for the authoritative control requirement. This sector guide does not restate or modify normative control text.

Manufacturing Interpretation

In manufacturing, this control requires the organisation to apply ai-generated phishing simulation in a way that accounts for long-lived assets, constrained maintenance windows, plant-specific conditions, product conformity, physical consequences and separation of safety functions from AI.

Minimum and Enhanced Implementation

- Define human authority, train users, test usability, provide escalation and override, and monitor adverse human-factor effects.
- Conduct human-factors engineering, competence validation, fatigue and workload assessment, independent ethics review where relevant, and recurring operator drills.

Evidence and Assessment

approved scope and owner; risk assessment; architecture or data-flow diagram; configuration and change records; test results; monitoring or review records; training records; human-factors test; override drill; complaint or adverse-impact review

Confirm the control operates at sampled plants and not only in corporate policy. Trace one deployed system from intake through validation, deployment, monitoring, change and recovery. Challenge supplier assertions and verify evidence quality.

Legacy OT and Compensating Controls

Apply proportionately to architecture, consequence and deployment context.

Use compensating monitoring, gateways, signed offline transfer, procedural approval or independent protection where direct implementation is technically infeasible.

Failure Modes, Residual Risk and Limits

- nominal human-in-the-loop
- operators cannot realistically intervene
- training not role-specific
- alert fatigue

Residual risk remains from unknown failure modes, incomplete data, supplier opacity, plant variation, human factors and correlated dependencies. Risk acceptance must be explicit for production- or safety-affecting systems.

Implementation of this control does not establish functional safety, product conformity, legal compliance or incident prevention by itself.

D7-CTL-H02 — DEEPPFAKE DETECTION TRAINING

Field	Content
Domain	D7: HUMAN & SOCIETAL HARMS
Applicability	Applicable
Sector relevance	High
Implementation track	Both
OT friction	Low
Conditional trigger	Applicable when the scoped system uses the capability or faces the risk addressed by the control.
Human oversight mode	Select human-in-the-loop, human-on-the-loop, or human-out-of-the-loop with independent protection based on process speed and consequence.

Authoritative Requirement

Refer to GAISSF-NOR-001 v1.0 and GAISSF-NOR-004 v1.0 for the authoritative control requirement. This sector guide does not restate or modify normative control text.

Manufacturing Interpretation

In manufacturing, this control requires the organisation to apply deepfake detection training in a way that accounts for long-lived assets, constrained maintenance windows, plant-specific conditions, product conformity, physical consequences and separation of safety functions from AI.

Minimum and Enhanced Implementation

- Define human authority, train users, test usability, provide escalation and override, and monitor adverse human-factor effects.
- Conduct human-factors engineering, competence validation, fatigue and workload assessment, independent ethics review where relevant, and recurring operator drills.

Evidence and Assessment

approved scope and owner; risk assessment; architecture or data-flow diagram; configuration and change records; test results; monitoring or review records; training records; human-factors test; override drill; complaint or adverse-impact review

Confirm the control operates at sampled plants and not only in corporate policy. Trace one deployed system from intake through validation, deployment, monitoring, change and recovery. Challenge supplier assertions and verify evidence quality.

Legacy OT and Compensating Controls

Apply proportionately to architecture, consequence and deployment context.

Use compensating monitoring, gateways, signed offline transfer, procedural approval or independent protection where direct implementation is technically infeasible.

Failure Modes, Residual Risk and Limits

- nominal human-in-the-loop
- operators cannot realistically intervene
- training not role-specific
- alert fatigue

Residual risk remains from unknown failure modes, incomplete data, supplier opacity, plant variation, human factors and correlated dependencies. Risk acceptance must be explicit for production- or safety-affecting systems.

Implementation of this control does not establish functional safety, product conformity, legal compliance or incident prevention by itself.

D7-CTL-H03 — OUT-OF-BAND AUTHENTICATION

Field	Content
Domain	D7: HUMAN & SOCIETAL HARMS
Applicability	Applicable
Sector relevance	High
Implementation track	Both
OT friction	Low
Conditional trigger	Applicable when the scoped system uses the capability or faces the risk addressed by the control.
Human oversight mode	Select human-in-the-loop, human-on-the-loop, or human-out-of-the-loop with independent protection based on process speed and consequence.

Authoritative Requirement

Refer to GAISSF-NOR-001 v1.0 and GAISSF-NOR-004 v1.0 for the authoritative control requirement. This sector guide does not restate or modify normative control text.

Manufacturing Interpretation

In manufacturing, this control requires the organisation to apply out-of-band authentication in a way that accounts for long-lived assets, constrained maintenance windows, plant-specific conditions, product conformity, physical consequences and separation of safety functions from AI.

Minimum and Enhanced Implementation

- Define human authority, train users, test usability, provide escalation and override, and monitor adverse human-factor effects.
- Conduct human-factors engineering, competence validation, fatigue and workload assessment, independent ethics review where relevant, and recurring operator drills.

Evidence and Assessment

approved scope and owner; risk assessment; architecture or data-flow diagram; configuration and change records; test results; monitoring or review records; training records; human-factors test; override drill; complaint or adverse-impact review

Confirm the control operates at sampled plants and not only in corporate policy. Trace one deployed system from intake through validation, deployment, monitoring, change and recovery. Challenge supplier assertions and verify evidence quality.

Legacy OT and Compensating Controls

Apply proportionately to architecture, consequence and deployment context.

Use compensating monitoring, gateways, signed offline transfer, procedural approval or independent protection where direct implementation is technically infeasible.

Failure Modes, Residual Risk and Limits

- nominal human-in-the-loop
- operators cannot realistically intervene
- training not role-specific
- alert fatigue

Residual risk remains from unknown failure modes, incomplete data, supplier opacity, plant variation, human factors and correlated dependencies. Risk acceptance must be explicit for production- or safety-affecting systems.

Implementation of this control does not establish functional safety, product conformity, legal compliance or incident prevention by itself.

D7-CTL-H04 — AI SOCIAL ENGINEERING IR

Field	Content
Domain	D7: HUMAN & SOCIETAL HARMS
Applicability	Applicable
Sector relevance	High
Implementation track	Both
OT friction	Low
Conditional trigger	Applicable when the scoped system uses the capability or faces the risk addressed by the control.
Human oversight mode	Select human-in-the-loop, human-on-the-loop, or human-out-of-the-loop with independent protection based on process speed and consequence.

Authoritative Requirement

Refer to GAISSF-NOR-001 v1.0 and GAISSF-NOR-004 v1.0 for the authoritative control requirement. This sector guide does not restate or modify normative control text.

Manufacturing Interpretation

In manufacturing, this control requires the organisation to apply ai social engineering ir in a way that accounts for long-lived assets, constrained maintenance windows, plant-specific conditions, product conformity, physical consequences and separation of safety functions from AI.

Minimum and Enhanced Implementation

- Define human authority, train users, test usability, provide escalation and override, and monitor adverse human-factor effects.
- Conduct human-factors engineering, competence validation, fatigue and workload assessment, independent ethics review where relevant, and recurring operator drills.

Evidence and Assessment

approved scope and owner; risk assessment; architecture or data-flow diagram; configuration and change records; test results; monitoring or review records; training records; human-factors test; override drill; complaint or adverse-impact review

Confirm the control operates at sampled plants and not only in corporate policy. Trace one deployed system from intake through validation, deployment, monitoring, change and recovery. Challenge supplier assertions and verify evidence quality.

Legacy OT and Compensating Controls

Apply proportionately to architecture, consequence and deployment context.

Use compensating monitoring, gateways, signed offline transfer, procedural approval or independent protection where direct implementation is technically infeasible.

Failure Modes, Residual Risk and Limits

- nominal human-in-the-loop

- operators cannot realistically intervene
- training not role-specific
- alert fatigue

Residual risk remains from unknown failure modes, incomplete data, supplier opacity, plant variation, human factors and correlated dependencies. Risk acceptance must be explicit for production- or safety-affecting systems.

Implementation of this control does not establish functional safety, product conformity, legal compliance or incident prevention by itself.

D7-CTL-H05 — AI-ENHANCED EXTERNAL ATTACK DEFENSE

Field	Content
Domain	D7: HUMAN & SOCIETAL HARMS
Applicability	Applicable
Sector relevance	High
Implementation track	Both
OT friction	Low
Conditional trigger	Applicable when the scoped system uses the capability or faces the risk addressed by the control.
Human oversight mode	Select human-in-the-loop, human-on-the-loop, or human-out-of-the-loop with independent protection based on process speed and consequence.

Authoritative Requirement

Refer to GAISSF-NOR-001 v1.0 and GAISSF-NOR-004 v1.0 for the authoritative control requirement. This sector guide does not restate or modify normative control text.

Manufacturing Interpretation

In manufacturing, this control requires the organisation to apply ai-enhanced external attack defense in a way that accounts for long-lived assets, constrained maintenance windows, plant-specific conditions, product conformity, physical consequences and separation of safety functions from AI.

Minimum and Enhanced Implementation

- Define human authority, train users, test usability, provide escalation and override, and monitor adverse human-factor effects.
- Conduct human-factors engineering, competence validation, fatigue and workload assessment, independent ethics review where relevant, and recurring operator drills.

Evidence and Assessment

approved scope and owner; risk assessment; architecture or data-flow diagram; configuration and change records; test results; monitoring or review records; training records; human-factors test; override drill; complaint or adverse-impact review

Confirm the control operates at sampled plants and not only in corporate policy. Trace one deployed system from intake through validation, deployment, monitoring, change and recovery. Challenge supplier assertions and verify evidence quality.

Legacy OT and Compensating Controls

Apply proportionately to architecture, consequence and deployment context.

Use compensating monitoring, gateways, signed offline transfer, procedural approval or independent protection where direct implementation is technically infeasible.

Failure Modes, Residual Risk and Limits

- nominal human-in-the-loop
- operators cannot realistically intervene
- training not role-specific
- alert fatigue

Residual risk remains from unknown failure modes, incomplete data, supplier opacity, plant variation, human factors and correlated dependencies. Risk acceptance must be explicit for production- or safety-affecting systems.

Implementation of this control does not establish functional safety, product conformity, legal compliance or incident prevention by itself.

D8-CTL-01 — EU AI ACT RISK TIER MAPPING

Field	Content
Domain	D8: REGULATORY ALIGNMENT & COMPLIANCE
Applicability	Applicable with manufacturing interpretation
Sector relevance	High
Implementation track	Both
OT friction	Moderate
Conditional trigger	Applicable when the scoped system uses the capability or faces the risk addressed by the control.
Human oversight mode	Human-in-the-loop for advisory/approval workflows; human-on-the-loop or independent protection for high-speed OT.

Authoritative Requirement

Refer to GAISSF-NOR-001 v1.0 and GAISSF-NOR-004 v1.0 for the authoritative control requirement. This sector guide does not restate or modify normative control text.

Manufacturing Interpretation

In manufacturing, this control requires the organisation to apply eu ai act risk tier mapping in a way that accounts for long-lived assets, constrained maintenance windows, plant-specific conditions, product conformity, physical consequences and separation of safety functions from AI.

Minimum and Enhanced Implementation

- Log material events, define incident criteria, preserve evidence, test rollback and manual fallback, and conduct post-incident review.
- Correlate AI, OT, safety and quality telemetry; automate containment; test cross-plant recovery; and independently validate restart readiness.

Evidence and Assessment

approved scope and owner; risk assessment; architecture or data-flow diagram; configuration and change records; test results; monitoring or review records; incident playbook; logs; exercise results; recovery and restart approval

Confirm the control operates at sampled plants and not only in corporate policy. Trace one deployed system from intake through validation, deployment, monitoring, change and recovery. Challenge supplier assertions and verify evidence quality.

Legacy OT and Compensating Controls

Apply proportionately to architecture, consequence and deployment context.

Use compensating monitoring, gateways, signed offline transfer, procedural approval or independent protection where direct implementation is technically infeasible.

Failure Modes, Residual Risk and Limits

- monitoring only accuracy
- no production stop criteria
- untested fallback
- restart without quality and safety revalidation

Residual risk remains from unknown failure modes, incomplete data, supplier opacity, plant variation, human factors and correlated dependencies. Risk acceptance must be explicit for production- or safety-affecting systems.

Implementation of this control does not establish functional safety, product conformity, legal compliance or incident prevention by itself.

D8-CTL-02 — ISO 42001 GAP ANALYSIS

Field	Content
Domain	D8: REGULATORY ALIGNMENT & COMPLIANCE
Applicability	Applicable with manufacturing interpretation
Sector relevance	High
Implementation track	Both
OT friction	Moderate
Conditional trigger	Applicable when the scoped system uses the capability or faces the risk addressed by the control.
Human oversight mode	Human-in-the-loop for advisory/approval workflows; human-on-the-loop or independent protection for high-speed OT.

Authoritative Requirement

Refer to GAISSF-NOR-001 v1.0 and GAISSF-NOR-004 v1.0 for the authoritative control requirement. This sector guide does not restate or modify normative control text.

Manufacturing Interpretation

In manufacturing, this control requires the organisation to apply iso 42001 gap analysis in a way that accounts for long-lived assets, constrained maintenance windows, plant-specific conditions, product conformity, physical consequences and separation of safety functions from AI.

Minimum and Enhanced Implementation

- Log material events, define incident criteria, preserve evidence, test rollback and manual fallback, and conduct post-incident review.
- Correlate AI, OT, safety and quality telemetry; automate containment; test cross-plant recovery; and independently validate restart readiness.

Evidence and Assessment

approved scope and owner; risk assessment; architecture or data-flow diagram; configuration and change records; test results; monitoring or review records; incident playbook; logs; exercise results; recovery and restart approval

Confirm the control operates at sampled plants and not only in corporate policy. Trace one deployed system from intake through validation, deployment, monitoring, change and recovery. Challenge supplier assertions and verify evidence quality.

Legacy OT and Compensating Controls

Apply proportionately to architecture, consequence and deployment context.

Use compensating monitoring, gateways, signed offline transfer, procedural approval or independent protection where direct implementation is technically infeasible.

Failure Modes, Residual Risk and Limits

- monitoring only accuracy
- no production stop criteria
- untested fallback
- restart without quality and safety revalidation

Residual risk remains from unknown failure modes, incomplete data, supplier opacity, plant variation, human factors and correlated dependencies. Risk acceptance must be explicit for production- or safety-affecting systems.

Implementation of this control does not establish functional safety, product conformity, legal compliance or incident prevention by itself.

D8-CTL-03 — GPAI TECHNICAL DOCUMENTATION VERIFICATION

Field	Content
Domain	D8: REGULATORY ALIGNMENT & COMPLIANCE
Applicability	Applicable with manufacturing interpretation
Sector relevance	High
Implementation track	Both
OT friction	Moderate
Conditional trigger	Applicable when the scoped system uses the capability or faces the risk addressed by the control.
Human oversight mode	Human-in-the-loop for advisory/approval workflows; human-on-the-loop or independent protection for high-speed OT.

Authoritative Requirement

Refer to GAISSF-NOR-001 v1.0 and GAISSF-NOR-004 v1.0 for the authoritative control requirement. This sector guide does not restate or modify normative control text.

Manufacturing Interpretation

In manufacturing, this control requires the organisation to apply gpai technical documentation verification in a way that accounts for long-lived assets, constrained maintenance windows, plant-specific conditions, product conformity, physical consequences and separation of safety functions from AI.

Minimum and Enhanced Implementation

- Log material events, define incident criteria, preserve evidence, test rollback and manual fallback, and conduct post-incident review.
- Correlate AI, OT, safety and quality telemetry; automate containment; test cross-plant recovery; and independently validate restart readiness.

Evidence and Assessment

approved scope and owner; risk assessment; architecture or data-flow diagram; configuration and change records; test results; monitoring or review records; incident playbook; logs; exercise results; recovery and restart approval

Confirm the control operates at sampled plants and not only in corporate policy. Trace one deployed system from intake through validation, deployment, monitoring, change and recovery. Challenge supplier assertions and verify evidence quality.

Legacy OT and Compensating Controls

Apply proportionately to architecture, consequence and deployment context.

Use compensating monitoring, gateways, signed offline transfer, procedural approval or independent protection where direct implementation is technically infeasible.

Failure Modes, Residual Risk and Limits

- monitoring only accuracy
- no production stop criteria
- untested fallback
- restart without quality and safety revalidation

Residual risk remains from unknown failure modes, incomplete data, supplier opacity, plant variation, human factors and correlated dependencies. Risk acceptance must be explicit for production- or safety-affecting systems.

Implementation of this control does not establish functional safety, product conformity, legal compliance or incident prevention by itself.

D8-CTL-04 — DORA ICT INCIDENT REPORTING (FINANCIAL SECTOR)

Field	Content
Domain	D8: REGULATORY ALIGNMENT & COMPLIANCE
Applicability	Applicable with manufacturing interpretation
Sector relevance	High
Implementation track	Both
OT friction	Moderate
Conditional trigger	Applicable when the scoped system uses the capability or faces the risk addressed by the control.
Human oversight mode	Human-in-the-loop for advisory/approval workflows; human-on-the-loop or independent protection for high-speed OT.

Authoritative Requirement

Refer to GAISSF-NOR-001 v1.0 and GAISSF-NOR-004 v1.0 for the authoritative control requirement. This sector guide does not restate or modify normative control text.

Manufacturing Interpretation

In manufacturing, this control requires the organisation to apply dora ict incident reporting (financial sector) in a way that accounts for long-lived assets, constrained maintenance windows, plant-specific conditions, product conformity, physical consequences and separation of safety functions from AI.

Minimum and Enhanced Implementation

- Log material events, define incident criteria, preserve evidence, test rollback and manual fallback, and conduct post-incident review.
- Correlate AI, OT, safety and quality telemetry; automate containment; test cross-plant recovery; and independently validate restart readiness.

Evidence and Assessment

approved scope and owner; risk assessment; architecture or data-flow diagram; configuration and change records; test results; monitoring or review records; incident playbook; logs; exercise results; recovery and restart approval

Confirm the control operates at sampled plants and not only in corporate policy. Trace one deployed system from intake through validation, deployment, monitoring, change and recovery. Challenge supplier assertions and verify evidence quality.

Legacy OT and Compensating Controls

Apply proportionately to architecture, consequence and deployment context.

Use compensating monitoring, gateways, signed offline transfer, procedural approval or independent protection where direct implementation is technically infeasible.

Failure Modes, Residual Risk and Limits

- monitoring only accuracy
- no production stop criteria
- untested fallback
- restart without quality and safety revalidation

Residual risk remains from unknown failure modes, incomplete data, supplier opacity, plant variation, human factors and correlated dependencies. Risk acceptance must be explicit for production- or safety-affecting systems.

Implementation of this control does not establish functional safety, product conformity, legal compliance or incident prevention by itself.

D8-CTL-05 — NIST SP 800-218A COMPLIANCE CHECK

Field	Content
Domain	D8: REGULATORY ALIGNMENT & COMPLIANCE
Applicability	Applicable with manufacturing interpretation
Sector relevance	High
Implementation track	Both
OT friction	Moderate
Conditional trigger	Applicable when the scoped system uses the capability or faces the risk addressed by the control.
Human oversight mode	Human-in-the-loop for advisory/approval workflows; human-on-the-loop or independent protection for high-speed OT.

Authoritative Requirement

Refer to GAISSF-NOR-001 v1.0 and GAISSF-NOR-004 v1.0 for the authoritative control requirement. This sector guide does not restate or modify normative control text.

Manufacturing Interpretation

In manufacturing, this control requires the organisation to apply nist sp 800-218a compliance check in a way that accounts for long-lived assets, constrained maintenance windows, plant-specific conditions, product conformity, physical consequences and separation of safety functions from AI.

Minimum and Enhanced Implementation

- Log material events, define incident criteria, preserve evidence, test rollback and manual fallback, and conduct post-incident review.
- Correlate AI, OT, safety and quality telemetry; automate containment; test cross-plant recovery; and independently validate restart readiness.

Evidence and Assessment

approved scope and owner; risk assessment; architecture or data-flow diagram; configuration and change records; test results; monitoring or review records; incident playbook; logs; exercise results; recovery and restart approval

Confirm the control operates at sampled plants and not only in corporate policy. Trace one deployed system from intake through validation, deployment, monitoring, change and recovery. Challenge supplier assertions and verify evidence quality.

Legacy OT and Compensating Controls

Apply proportionately to architecture, consequence and deployment context.

Use compensating monitoring, gateways, signed offline transfer, procedural approval or independent protection where direct implementation is technically infeasible.

Failure Modes, Residual Risk and Limits

- monitoring only accuracy

- no production stop criteria
- untested fallback
- restart without quality and safety revalidation

Residual risk remains from unknown failure modes, incomplete data, supplier opacity, plant variation, human factors and correlated dependencies. Risk acceptance must be explicit for production- or safety-affecting systems.

Implementation of this control does not establish functional safety, product conformity, legal compliance or incident prevention by itself.

D9-CTL-01 — PHYSICAL HARM BOUNDARY ENFORCEMENT

Field	Content
Domain	D9: PHYSICAL AI SAFETY
Applicability	Applicable with manufacturing interpretation
Sector relevance	Core
Implementation track	OT/Edge and Physical AI
OT friction	High
Conditional trigger	Applicable when the scoped system uses the capability or faces the risk addressed by the control.
Human oversight mode	Select human-in-the-loop, human-on-the-loop, or human-out-of-the-loop with independent protection based on process speed and consequence.

Authoritative Requirement

Refer to GAISSF-NOR-001 v1.0 and GAISSF-NOR-004 v1.0 for the authoritative control requirement. This sector guide does not restate or modify normative control text.

Manufacturing Interpretation

In manufacturing, this control requires the organisation to apply physical harm boundary enforcement in a way that accounts for long-lived assets, constrained maintenance windows, plant-specific conditions, product conformity, physical consequences and separation of safety functions from AI.

Minimum and Enhanced Implementation

- Maintain independent safety functions, constrained actuation, validated operating envelopes, emergency stop, physical access controls, and safe-state behavior.
- Apply defence-in-depth across perception, planning and actuation; independent monitoring; formal safety cases; adversarial physical testing; and validated degraded modes.

Evidence and Assessment

approved scope and owner; risk assessment; architecture or data-flow diagram; configuration and change records; test results; monitoring or review records; safety case; operating envelope; interlock test; physical access and emergency-stop test

Confirm the control operates at sampled plants and not only in corporate policy. Trace one deployed system from intake through validation, deployment, monitoring, change and recovery. Challenge supplier assertions and verify evidence quality.

Legacy OT and Compensating Controls

Apply proportionately to architecture, consequence and deployment context.

Use compensating monitoring, gateways, signed offline transfer, procedural approval or independent protection where direct implementation is technically infeasible.

Failure Modes, Residual Risk and Limits

- AI relied on as sole safety control
- unsafe operating envelope
- no physical adversarial testing
- safety and security reviews disconnected

Residual risk remains from unknown failure modes, incomplete data, supplier opacity, plant variation, human factors and correlated dependencies. Risk acceptance must be explicit for production- or safety-affecting systems.

Implementation of this control does not establish functional safety, product conformity, legal compliance or incident prevention by itself.

D9-CTL-02 — SAFE STATE AND GRACEFUL DEGRADATION

Field	Content
Domain	D9: PHYSICAL AI SAFETY
Applicability	Applicable with manufacturing interpretation
Sector relevance	Core
Implementation track	OT/Edge and Physical AI
OT friction	High
Conditional trigger	Applicable when the scoped system uses the capability or faces the risk addressed by the control.
Human oversight mode	Select human-in-the-loop, human-on-the-loop, or human-out-of-the-loop with independent protection based on process speed and consequence.

Authoritative Requirement

Refer to GAISSF-NOR-001 v1.0 and GAISSF-NOR-004 v1.0 for the authoritative control requirement. This sector guide does not restate or modify normative control text.

Manufacturing Interpretation

In manufacturing, this control requires the organisation to apply safe state and graceful degradation in a way that accounts for long-lived assets, constrained maintenance windows, plant-specific conditions, product conformity, physical consequences and separation of safety functions from AI.

Minimum and Enhanced Implementation

- Maintain independent safety functions, constrained actuation, validated operating envelopes, emergency stop, physical access controls, and safe-state behavior.
- Apply defence-in-depth across perception, planning and actuation; independent monitoring; formal safety cases; adversarial physical testing; and validated degraded modes.

Evidence and Assessment

approved scope and owner; risk assessment; architecture or data-flow diagram; configuration and change records; test results; monitoring or review records; safety case; operating envelope; interlock test; physical access and emergency-stop test

Confirm the control operates at sampled plants and not only in corporate policy. Trace one deployed system from intake through validation, deployment, monitoring, change and recovery. Challenge supplier assertions and verify evidence quality.

Legacy OT and Compensating Controls

Apply proportionately to architecture, consequence and deployment context.

Use compensating monitoring, gateways, signed offline transfer, procedural approval or independent protection where direct implementation is technically infeasible.

Failure Modes, Residual Risk and Limits

- AI relied on as sole safety control
- unsafe operating envelope
- no physical adversarial testing
- safety and security reviews disconnected

Residual risk remains from unknown failure modes, incomplete data, supplier opacity, plant variation, human factors and correlated dependencies. Risk acceptance must be explicit for production- or safety-affecting systems.

Implementation of this control does not establish functional safety, product conformity, legal compliance or incident prevention by itself.

D9-CTL-03 — HUMAN OVERRIDE AND EMERGENCY STOP

Field	Content
Domain	D9: PHYSICAL AI SAFETY
Applicability	Applicable with manufacturing interpretation
Sector relevance	Conditional
Implementation track	OT/Edge and Physical AI
OT friction	High
Conditional trigger	Applicable when the scoped system uses the capability or faces the risk addressed by the control.
Human oversight mode	Select human-in-the-loop, human-on-the-loop, or human-out-of-the-loop with independent protection based on process speed and consequence.

Authoritative Requirement

Refer to GAISSF-NOR-001 v1.0 and GAISSF-NOR-004 v1.0 for the authoritative control requirement. This sector guide does not restate or modify normative control text.

Manufacturing Interpretation

In manufacturing, this control requires the organisation to apply human override and emergency stop in a way that accounts for long-lived assets, constrained maintenance windows, plant-specific conditions, product conformity, physical consequences and separation of safety functions from AI.

Minimum and Enhanced Implementation

- Maintain independent safety functions, constrained actuation, validated operating envelopes, emergency stop, physical access controls, and safe-state behavior.
- Apply defence-in-depth across perception, planning and actuation; independent monitoring; formal safety cases; adversarial physical testing; and validated degraded modes.

Evidence and Assessment

approved scope and owner; risk assessment; architecture or data-flow diagram; configuration and change records; test results; monitoring or review records; safety case; operating envelope; interlock test; physical access and emergency-stop test

Confirm the control operates at sampled plants and not only in corporate policy. Trace one deployed system from intake through validation, deployment, monitoring, change and recovery. Challenge supplier assertions and verify evidence quality.

Legacy OT and Compensating Controls

Apply proportionately to architecture, consequence and deployment context.

Use compensating monitoring, gateways, signed offline transfer, procedural approval or independent protection where direct implementation is technically infeasible.

Failure Modes, Residual Risk and Limits

- AI relied on as sole safety control
- unsafe operating envelope
- no physical adversarial testing
- safety and security reviews disconnected

Residual risk remains from unknown failure modes, incomplete data, supplier opacity, plant variation, human factors and correlated dependencies. Risk acceptance must be explicit for production- or safety-affecting systems.

Implementation of this control does not establish functional safety, product conformity, legal compliance or incident prevention by itself.

D9-CTL-04 — CYBER-PHYSICAL ATTACK DETECTION

Field	Content
Domain	D9: PHYSICAL AI SAFETY
Applicability	Applicable with manufacturing interpretation
Sector relevance	Core
Implementation track	OT/Edge and Physical AI
OT friction	High
Conditional trigger	Applicable when the scoped system uses the capability or faces the risk addressed by the control.
Human oversight mode	Select human-in-the-loop, human-on-the-loop, or human-out-of-the-loop with independent protection based on process speed and consequence.

Authoritative Requirement

Refer to GAISSF-NOR-001 v1.0 and GAISSF-NOR-004 v1.0 for the authoritative control requirement. This sector guide does not restate or modify normative control text.

Manufacturing Interpretation

In manufacturing, this control requires the organisation to apply cyber-physical attack detection in a way that accounts for long-lived assets, constrained maintenance windows, plant-specific conditions, product conformity, physical consequences and separation of safety functions from AI.

Minimum and Enhanced Implementation

- Maintain independent safety functions, constrained actuation, validated operating envelopes, emergency stop, physical access controls, and safe-state behavior.
- Apply defence-in-depth across perception, planning and actuation; independent monitoring; formal safety cases; adversarial physical testing; and validated degraded modes.

Evidence and Assessment

approved scope and owner; risk assessment; architecture or data-flow diagram; configuration and change records; test results; monitoring or review records; safety case; operating envelope; interlock test; physical access and emergency-stop test

Confirm the control operates at sampled plants and not only in corporate policy. Trace one deployed system from intake through validation, deployment, monitoring, change and recovery. Challenge supplier assertions and verify evidence quality.

Legacy OT and Compensating Controls

Apply proportionately to architecture, consequence and deployment context.

Use compensating monitoring, gateways, signed offline transfer, procedural approval or independent protection where direct implementation is technically infeasible.

Failure Modes, Residual Risk and Limits

- AI relied on as sole safety control
- unsafe operating envelope
- no physical adversarial testing
- safety and security reviews disconnected

Residual risk remains from unknown failure modes, incomplete data, supplier opacity, plant variation, human factors and correlated dependencies. Risk acceptance must be explicit for production- or safety-affecting systems.

Implementation of this control does not establish functional safety, product conformity, legal compliance or incident prevention by itself.

D9-CTL-05 — PHYSICAL ENVIRONMENT INTEGRITY MONITORING

Field	Content
Domain	D9: PHYSICAL AI SAFETY
Applicability	Applicable with manufacturing interpretation
Sector relevance	Core
Implementation track	OT/Edge and Physical AI
OT friction	High
Conditional trigger	Applicable when the scoped system uses the capability or faces the risk addressed by the control.
Human oversight mode	Select human-in-the-loop, human-on-the-loop, or human-out-of-the-loop with independent protection based on process speed and consequence.

Authoritative Requirement

Refer to GAISSF-NOR-001 v1.0 and GAISSF-NOR-004 v1.0 for the authoritative control requirement. This sector guide does not restate or modify normative control text.

Manufacturing Interpretation

In manufacturing, this control requires the organisation to apply physical environment integrity monitoring in a way that accounts for long-lived assets, constrained maintenance windows, plant-specific conditions, product conformity, physical consequences and separation of safety functions from AI.

Minimum and Enhanced Implementation

- Maintain independent safety functions, constrained actuation, validated operating envelopes, emergency stop, physical access controls, and safe-state behavior.
- Apply defence-in-depth across perception, planning and actuation; independent monitoring; formal safety cases; adversarial physical testing; and validated degraded modes.

Evidence and Assessment

approved scope and owner; risk assessment; architecture or data-flow diagram; configuration and change records; test results; monitoring or review records; safety case; operating envelope; interlock test; physical access and emergency-stop test

Confirm the control operates at sampled plants and not only in corporate policy. Trace one deployed system from intake through validation, deployment, monitoring, change and recovery. Challenge supplier assertions and verify evidence quality.

Legacy OT and Compensating Controls

Apply proportionately to architecture, consequence and deployment context.

Use compensating monitoring, gateways, signed offline transfer, procedural approval or independent protection where direct implementation is technically infeasible.

Failure Modes, Residual Risk and Limits

- AI relied on as sole safety control
- unsafe operating envelope
- no physical adversarial testing
- safety and security reviews disconnected

Residual risk remains from unknown failure modes, incomplete data, supplier opacity, plant variation, human factors and correlated dependencies. Risk acceptance must be explicit for production- or safety-affecting systems.

Implementation of this control does not establish functional safety, product conformity, legal compliance or incident prevention by itself.

D9-CTL-06 — ACTUATOR COMMAND VERIFICATION

Field	Content
Domain	D9: PHYSICAL AI SAFETY
Applicability	Applicable with manufacturing interpretation
Sector relevance	Core
Implementation track	OT/Edge and Physical AI
OT friction	High
Conditional trigger	Applicable when the scoped system uses the capability or faces the risk addressed by the control.
Human oversight mode	Select human-in-the-loop, human-on-the-loop, or human-out-of-the-loop with independent protection based on process speed and consequence.

Authoritative Requirement

Refer to GAISSF-NOR-001 v1.0 and GAISSF-NOR-004 v1.0 for the authoritative control requirement. This sector guide does not restate or modify normative control text.

Manufacturing Interpretation

In manufacturing, this control requires the organisation to apply actuator command verification in a way that accounts for long-lived assets, constrained maintenance windows, plant-specific conditions, product conformity, physical consequences and separation of safety functions from AI.

Minimum and Enhanced Implementation

- Maintain independent safety functions, constrained actuation, validated operating envelopes, emergency stop, physical access controls, and safe-state behavior.
- Apply defence-in-depth across perception, planning and actuation; independent monitoring; formal safety cases; adversarial physical testing; and validated degraded modes.

Evidence and Assessment

approved scope and owner; risk assessment; architecture or data-flow diagram; configuration and change records; test results; monitoring or review records; safety case; operating envelope; interlock test; physical access and emergency-stop test

Confirm the control operates at sampled plants and not only in corporate policy. Trace one deployed system from intake through validation, deployment, monitoring, change and recovery. Challenge supplier assertions and verify evidence quality.

Legacy OT and Compensating Controls

Apply proportionately to architecture, consequence and deployment context.

Use compensating monitoring, gateways, signed offline transfer, procedural approval or independent protection where direct implementation is technically infeasible.

Failure Modes, Residual Risk and Limits

- AI relied on as sole safety control

- unsafe operating envelope
- no physical adversarial testing
- safety and security reviews disconnected

Residual risk remains from unknown failure modes, incomplete data, supplier opacity, plant variation, human factors and correlated dependencies. Risk acceptance must be explicit for production- or safety-affecting systems.

Implementation of this control does not establish functional safety, product conformity, legal compliance or incident prevention by itself.

D9-CTL-07 — PHYSICAL INCIDENT EVIDENCE PRESERVATION

Field	Content
Domain	D9: PHYSICAL AI SAFETY
Applicability	Applicable with manufacturing interpretation
Sector relevance	Core
Implementation track	OT/Edge and Physical AI
OT friction	High
Conditional trigger	Applicable when the scoped system uses the capability or faces the risk addressed by the control.
Human oversight mode	Select human-in-the-loop, human-on-the-loop, or human-out-of-the-loop with independent protection based on process speed and consequence.

Authoritative Requirement

Refer to GAISSF-NOR-001 v1.0 and GAISSF-NOR-004 v1.0 for the authoritative control requirement. This sector guide does not restate or modify normative control text.

Manufacturing Interpretation

In manufacturing, this control requires the organisation to apply physical incident evidence preservation in a way that accounts for long-lived assets, constrained maintenance windows, plant-specific conditions, product conformity, physical consequences and separation of safety functions from AI.

Minimum and Enhanced Implementation

- Maintain independent safety functions, constrained actuation, validated operating envelopes, emergency stop, physical access controls, and safe-state behavior.
- Apply defence-in-depth across perception, planning and actuation; independent monitoring; formal safety cases; adversarial physical testing; and validated degraded modes.

Evidence and Assessment

approved scope and owner; risk assessment; architecture or data-flow diagram; configuration and change records; test results; monitoring or review records; safety case; operating envelope; interlock test; physical access and emergency-stop test

Confirm the control operates at sampled plants and not only in corporate policy. Trace one deployed system from intake through validation, deployment, monitoring, change and recovery. Challenge supplier assertions and verify evidence quality.

Legacy OT and Compensating Controls

Apply proportionately to architecture, consequence and deployment context.

Use compensating monitoring, gateways, signed offline transfer, procedural approval or independent protection where direct implementation is technically infeasible.

Failure Modes, Residual Risk and Limits

- AI relied on as sole safety control
- unsafe operating envelope
- no physical adversarial testing
- safety and security reviews disconnected

Residual risk remains from unknown failure modes, incomplete data, supplier opacity, plant variation, human factors and correlated dependencies. Risk acceptance must be explicit for production- or safety-affecting systems.

Implementation of this control does not establish functional safety, product conformity, legal compliance or incident prevention by itself.

11. Manufacturing Threat Scenarios

ID	Scenario and source class	Actor/cause	Impact Band	Confidence	Evidence status
TS-01	Poisoned predictive-maintenance data [Analytical] [T4]	Malicious insider or compromised data source	Impact Band 3	Moderate	Plausible; not confirmed for a specific site
TS-02	Manipulated machine-vision inspection [Analytical] [T4]	External attacker or insider	Impact Band 3	Moderate	Observed technique; manufacturing occurrence context-dependent
TS-03	Compromised edge inference device [Research / proxy] [T3]	Cybercriminal or supplier compromise	Impact Band 3	Moderate	Credible observed technique
TS-04	Malicious model update [Research / proxy] [T3]	Supplier, attacker or insider	Impact Band 3	Moderate	Credible scenario
TS-05	Cloud AI outage affecting scheduling [Corroborated] [T2]	Cloud/service failure	Impact Band 2	High	Common dependency failure pattern
TS-06	Generative AI leakage of proprietary design data [Analytical] [T4]	Negligent user or provider exposure	Impact Band 2	High	Observed class of risk
TS-07	Adversarial input to robotic vision [Analytical] [T4]	External or insider	Impact Band 4	Low	Plausible; site-specific feasibility required
TS-08	False sensor data drives incorrect optimisation [Research / proxy] [T3]	Compromised sensor or network path	Impact Band 4	Moderate	Credible industrial attack path
TS-09	Compromised supplier model [Research / proxy] [T3]	Supply-chain actor	Impact Band 3	Moderate	Plausible; evidence depends on supplier
TS-10	Unauthorised AI connection to OT [Analytical] [T4]	Employee, contractor or integrator	Impact Band 3	High	Common governance failure pattern
TS-11	Operator overreliance on incorrect maintenance advice [Corroborated] [T2]	Human factors / model error	Impact Band 2	High	Well-established automation-bias mechanism
TS-12	Manipulated digital twin [Analytical] [T4]	Attacker or data integrity failure	Impact Band 3	Moderate	Plausible
TS-13	Quality drift remains undetected [Corroborated] [T2]	Model/data drift	Impact Band 3	High	Common model lifecycle risk
TS-14	AI-enabled insider theft [Analytical] [T4]	Malicious insider	Impact Band 2	Moderate	Credible
TS-15	Ransomware disrupts AI-dependent production [Corroborated] [T2]	Cybercriminal	Impact Band 3	High	Observed industrial threat class
TS-16	Unauthorised threshold modification [Research / proxy] [T3]	Insider or compromised admin	Impact Band 3	Moderate	Credible
TS-17	Compromised autonomous mobile robot [Research / proxy] [T3]	External attacker or supplier	Impact Band 4	Moderate	Plausible/observed component techniques
TS-18	Model rollback failure [Analytical] [T4]	Operational failure	Impact Band 3	Moderate	Credible
TS-19	Defective model propagates across plants [Analytical] [T4]	Central deployment pipeline failure	Impact Band 3	Moderate	Credible
TS-20	Supplier remote-access abuse [Research / proxy] [T3]	Supplier or compromised account	Impact Band 3	High	Observed industrial access pattern
TS-21	Synthetic data creates unsafe blind spots [Analytical] [T4]	Development error	Impact Band 3	Moderate	Plausible
TS-22	Shadow generative AI handles confidential data [Corroborated] [T2]	Negligent employee	Impact Band 2	High	Common enterprise risk
TS-23	AI recommendation bypasses engineering change control [Analytical] [T4]	Process design failure	Impact Band 4	Moderate	Credible
TS-24	Worker-safety detection false negative [Analytical] [T4]	Model/data failure	Impact Band 4	High	Intrinsic performance risk
TS-25	Environmental change collapses performance [Corroborated] [T2]	Distribution shift	Impact Band 3	High	Established ML failure mode

The tags classify the evidentiary basis for the scenario description, not the severity of the scenario. No site-specific access path, occurrence or consequence is established without supporting evidence.

12. EU Artificial Intelligence Act Boundary

GAISSF conformance-tier classification and GAISSF-SEC-045 assessment do not constitute an EU 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.

Manufacturing organisations must determine legal classification from intended purpose, role in the value chain, product integration, applicable annexes, sector legislation and jurisdiction. GAISSF tiers are an operational assurance model, not a legal classification.

13. Notably Absent

- Verified public disclosures establishing the frequency of AI-driven operational-technology disruption in manufacturing were not identified.
- Proprietary internal ODA3 telemetry confirming the frequency of AI-driven operational-technology disruption is absent because ODA3 Institute was founded in March 2026 and does not maintain a historical proprietary client incident dataset.
- Malicious AI-model manipulation, kinetic sabotage, cross-plant propagation or direct safety-system compromise is not treated as confirmed merely because it is technically plausible.
- Any assertion that GAISSF security controls replace or satisfy probability-of-failure metrics required for functional-safety certification is notably absent and expressly disclaimed.
- The review did not establish that reported manufacturing incidents are predominantly attributable to any single AI failure category.

14. Limitations

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.

Public disclosure is incomplete, architectures are often confidential, attribution is difficult, and controlled research does not establish field frequency. External standards, legal obligations and supplier products must be reverified at release and use.

15. Worked Example — Impact Band 2 Predictive Maintenance for a CNC Machine

Stage	Worked example
Intake	Single high-value CNC machine; vibration, temperature and maintenance records; advisory maintenance recommendation; no direct actuation.
Classification	Impact Band 2 because recommendations influence maintenance timing but do not command the machine. Escalate to Impact Band 3 if work orders or production release decisions are automated.
Primary controls	D1-CTL-01 provenance; D1-CTL-03 drift; relevant input/runtime controls; supplier controls; output integrity; governance and monitoring controls.
Threat scenarios	TS-01 poisoned maintenance data; TS-11 operator overreliance; TS-13 undetected drift; TS-25 changed environmental conditions.
Evidence	Sensor calibration, data lineage, model hash, validation by operating regime, maintenance-owner approval, drift dashboard, rollback package and operator training.
Fallback	Use the approved preventive-maintenance schedule and engineering inspection process if the model is unavailable or unreliable.
Financial estimate	Expected avoidable downtime and maintenance exposure using low/base/high event frequency and line-hour contribution margin.

Stage	Worked example
Notably absent	No claim that the model directly controls the machine or that a malicious poisoning event has occurred.

A separately populated worked-example workbook is included in the package.

Publication status: Draft for Publication — Public — External Review Open; not approved as Final Publication. Final Publication requires completion or authorised disposition of the manufacturing/OT engineering, functional or machinery safety, legal/regulatory, accessibility/editorial and authorized final-publication-release gates.