

ODA3 INSTITUTE

SECTOR IMPLEMENTATION GUIDANCE

GAISSF™ Critical Infrastructure Sector Implementation Guide

Physical AI, cyber-physical and OT/ICS implementation guidance

GAISSF-SEC-042 | Public | Draft for Publication — External Review Open

Document Control

Document ID	GAISSF-SEC-042
Title	GAISSF™ Critical Infrastructure Sector Implementation Guide
Doc Type	Sector implementation guidance — informative
Framework	GAISSF™ v1.0 — Global AI Safety and Security Framework (59 controls across nine domains)
Authoritative sources	GAISSF-NOR-001 (Framework Standard); GAISSF-NOR-004 (Control Catalogue)
Status	Draft for Publication — Public — External Review Open. Open gates: named approvals; OT/ICS SME review; functional-safety review; energy, water and healthcare-infrastructure reviews; legal/regulatory applicability and mapping verification; accessibility and final publication QA.
Dates	Refinement dated 30 June 2026; updated 05 October 2026. No publication date assigned.
Companion files	Control records (JSON/CSV), control mapping, implementation workbook, evidence catalogue and threat register
Contact	contact@oda3.org
Classification	Public (on release)

GAISSF™ Critical Infrastructure Sector Implementation Guide

Physical AI, cyber-physical and OT/ICS implementation guidance

Reliance and control-integrity notice: This informative sector guide does not modify GAISSF-NOR-001 or GAISSF-NOR-004. It does not constitute legal advice, establish regulatory compliance, certify functional safety, assign a Safety Integrity Level, replace process hazard analysis, or guarantee incident prevention. Exact applicability of external standards and laws must be verified by qualified personnel using controlled sources. Profile membership is defined in GAISSF-NOR-001. GAISSF certification is not currently offered. © 2026 ODA3 Pvt Ltd. GAISSF Ecosystem developed by ODA3 Institute, licensed under GEL v1.0.

1. Executive Summary

GAISSF-SEC-042 translates the 59-control GAISSF baseline into critical-infrastructure implementation guidance for enterprise IT, OT/ICS, cyber-physical and physical-AI systems. This edition removes generic scaling, prohibits uncontrolled adversarial testing on live actuation paths, introduces air-gap and black-box-vendor protocols, synchronizes evidence formats, and adds energy, water and healthcare-infrastructure context.

Leadership decisions: define where AI may operate; bound operational authority; require representative non-production validation; maintain independent safety and fallback mechanisms; document supplier opacity and residual risk; and suspend AI-enabled functions when evidence, timing or safe-state assumptions fail.

2. OT/ICS Testing Boundary

Adversarial, poisoning, fault-injection, unsafe-command and destructive tests must be executed in a high-fidelity digital twin, hardware-in-the-loop rig or strictly isolated staging network. They must not be routed through a live production control or actuation path unless a separately approved safety case, controlled change window, rollback plan and accountable engineering authority explicitly authorize the test.

Boundary Element	Minimum Expectation
Representative environment	Mirror relevant interfaces, latency, data characteristics, protections, failure states and operator workflow.
No unintended actuation	Physically or logically prevent test traffic from reaching live actuators or safety functions.
Air-gap transfer	Use approved media, malware scanning, SHA-256 verification, dual authorization, custody manifest and reconciliation.
Latency assurance	Measure worst-case added latency/jitter against the approved deterministic timing budget.
Restoration	Verify clean restoration, configuration integrity and safe operation after testing.

3. Criticality Classification

Factor	1 - Low	3 - Material	5 - Severe
Consequence	Administrative inconvenience	Material service degradation	Potential loss of life, major environmental harm or systemic disruption
Autonomy	Advisory	Human-approved execution	Autonomous/high-speed execution
Speed	Hours/days to intervene	Minutes	Milliseconds/seconds
Reversibility	Easily reversible	Costly recovery	Irreversible physical effect
Detectability	Obvious	Delayed	Latent or misleading
Concentration	Single low-impact use	Multiple systems/sites	Common-mode multi-site dependency
Fallback	Proven manual mode	Constrained fallback	No credible immediate fallback
Uncertainty	Well understood	Material gaps	Opaque/novel with high uncertainty

The workbook provides customizable thresholds. The score supports triage only; engineering judgment and consequence pathways govern the approved class.

4. Cross-Control CI Implementation Baseline

Authority

Document advisory, approval, execution, suspension, override and restoration authority.

Safety independence

AI and AI-security controls must not replace or delay independently required protective functions.

Evidence

Tie evidence to system ID, site, deployed version, period, owner, scope and limitations.

Fallback

Test manual, degraded and safe-state modes with realistic staffing and communications.

Supplier opacity

Separate vendor assertion from independent boundary evidence and record residual risk.

Human factors

Test alarm fatigue, automation bias, shift workload and emergency takeover.

5. Vendor Assertion vs Independent Boundary Testing

Where a supplier will not disclose training data, model weights or internal design, the operator should not fabricate provenance. Record the vendor assertion, evidence date and limitations; independently test observable inputs, outputs, timing, network dependencies, update mechanisms, logging, failure behaviour and fallback; and submit unresolved opacity for residual-risk acceptance.

6. Subsector Annexes

Energy and Power

- Grid balancing and protection coordination
- IEC 61850/DNP3 timing and restoration
- NERC CIP applicability must be verified by jurisdiction and asset scope

Water and Wastewater

- Treatment chemistry and dosing limits
- Remote pumping, pressure and water-quality monitoring
- Manual sampling and safe dosing fallback

Healthcare Infrastructure

- Medical IoT and facilities systems
- Patient-support infrastructure and clinical engineering
- No substitution for medical-device, clinical-safety or patient-safety obligations

7. Evidence Grades and Source Types

Grade / source type	Meaning
G1	Design evidence only

Grade / source type	Meaning
G2	Implemented configuration or procedure
G3	Operating-effectiveness evidence from controlled simulation/HIL or representative operations
G4	Independent assurance
S1	Primary organization-controlled record
S2	Verified secondary or supplier evidence
S3	Controlled simulation, digital twin or HIL evidence
S4	Independent assessment or functional-safety/security assurance

Source types S1–S4 describe where evidence comes from. They are distinct from ODA3 Institute's T1–T4 evidence tiers used in research reports.

8. Day 1 Triage

Priority	Immediate action
1	Inventory AI-enabled assets, embedded vendor AI, data paths and actuation interfaces.
2	Identify CI-3/4/5 systems and prohibit unapproved autonomous or write authority.
3	Verify D9-CTL-03 human override/emergency stop and D9-CTL-02 safe-state behaviour.
4	Apply D6-CTL-01 human authority and D4-CTL-06 shadow-AI discovery adapted to OT constraints.
5	Confirm fallback, local operation and critical cloud/API dependencies.
6	Establish isolated test boundary and air-gap transfer process.
7	Record black-box supplier gaps and boundary-testing plan.
8	Preserve logs and incident evidence without impairing deterministic timing.

9. Control-by-Control CI Interpretation

Control IDs, titles and authoritative requirements are from GAISSE-NOR-004. Business objectives, OT/ICS equivalents, criticality scaling and evidence expectations are set at domain level and applied to each control; they are informative.

D1-CTL-01 — Dataset Provenance & Poisoning Prevention

Field	Critical-infrastructure interpretation
Authoritative requirement	Hash verification + source allowlist + poisoning detection.
Business objective	Preserve trustworthy operational decisions by preventing corrupted data or model state from causing unsafe control recommendations, missed degradation, service interruption or cascading cyber-physical effects.

Field	Critical-infrastructure interpretation
OT/ICS-adapted equivalent	Use signed offline manifests, approved removable-media transfer, local hash registries and isolated scanning when continuous repositories or CI/CD are unavailable.
Criticality scaling (CI-1 to CI-5)	CI-1: Document source, model or dataset identity; perform periodic integrity review and retain a reproducible baseline. CI-2: Add approved provenance records, version-controlled integrity checks and representative offline validation before material changes. CI-3: Use automated integrity verification, independent challenge of high-impact changes and tested rollback to the last known-good model/data state. CI-4: Require cryptographically protected provenance, dual authorization for high-impact changes, isolated adversarial testing and continuous integrity monitoring at trust boundaries. CI-5: Use hardware-backed or independently anchored attestation where feasible, multi-party approval, redundant validation paths, safety-engineering review and evidence that integrity controls do not impair deterministic operations.
Evidence requirement	Signed JSON/CSV manifest containing asset/model/data version, hashes, source, test results, thresholds, exceptions and reviewer.
Testing boundary	high-fidelity digital twin, hardware-in-the-loop rig, or strictly isolated staging network. Adversarial, poisoning, fault-injection and unsafe-command tests must not be routed through a live production control or actuation path unless a separately approved safety case authorises it.
Notably Absent	Detection of novel poisoning methods not represented in the approved test corpus; assurance against an authorized insider able to alter both data and trusted provenance records; proof that source allowlisting establishes semantic correctness.

D1-CTL-02 — Model Extraction Resistance

Field	Critical-infrastructure interpretation
Authoritative requirement	Rate limiting + diversity detection + extraction monitoring.
Business objective	Preserve trustworthy operational decisions by preventing corrupted data or model state from causing unsafe control recommendations, missed degradation, service interruption or cascading cyber-physical effects.
OT/ICS-adapted equivalent	Use signed offline manifests, approved removable-media transfer, local hash registries and isolated scanning when continuous repositories or CI/CD are unavailable.
Criticality scaling (CI-1 to CI-5)	CI-1: Document source, model or dataset identity; perform periodic integrity review and retain a reproducible baseline. CI-2: Add approved provenance records, version-controlled integrity checks and representative offline validation before material changes. CI-3: Use automated integrity verification, independent challenge of high-impact changes and tested rollback to the last known-good model/data state. CI-4: Require cryptographically protected provenance, dual authorization for high-impact changes, isolated adversarial testing and continuous integrity monitoring at trust boundaries. CI-5: Use hardware-backed or independently anchored attestation where feasible, multi-party approval, redundant validation paths, safety-engineering review and evidence that integrity controls do not impair deterministic operations.
Evidence requirement	Signed JSON/CSV manifest containing asset/model/data version, hashes, source, test results, thresholds, exceptions and reviewer.
Testing boundary	high-fidelity digital twin, hardware-in-the-loop rig, or strictly isolated staging network. Adversarial, poisoning, fault-injection and unsafe-command tests must not be routed through a live production control or actuation path unless a separately approved safety case authorises it.

Field	Critical-infrastructure interpretation
Notably Absent	A guarantee that rate limits or output controls prevent all model inference; proof that protected model logic cannot be approximated through physical-process observation or compromised vendor tooling.

D1-CTL-03 — Behavioral Drift Detection

Field	Critical-infrastructure interpretation
Authoritative requirement	Baseline profiling + KL divergence monitoring + accuracy tracking.
Business objective	Preserve trustworthy operational decisions by preventing corrupted data or model state from causing unsafe control recommendations, missed degradation, service interruption or cascading cyber-physical effects.
OT/ICS-adapted equivalent	Use signed offline manifests, approved removable-media transfer, local hash registries and isolated scanning when continuous repositories or CI/CD are unavailable.
Criticality scaling (CI-1 to CI-5)	CI-1: Document source, model or dataset identity; perform periodic integrity review and retain a reproducible baseline. CI-2: Add approved provenance records, version-controlled integrity checks and representative offline validation before material changes. CI-3: Use automated integrity verification, independent challenge of high-impact changes and tested rollback to the last known-good model/data state. CI-4: Require cryptographically protected provenance, dual authorization for high-impact changes, isolated adversarial testing and continuous integrity monitoring at trust boundaries. CI-5: Use hardware-backed or independently anchored attestation where feasible, multi-party approval, redundant validation paths, safety-engineering review and evidence that integrity controls do not impair deterministic operations.
Evidence requirement	Signed JSON/CSV manifest containing asset/model/data version, hashes, source, test results, thresholds, exceptions and reviewer.
Testing boundary	high-fidelity digital twin, hardware-in-the-loop rig, or strictly isolated staging network. Adversarial, poisoning, fault-injection and unsafe-command tests must not be routed through a live production control or actuation path unless a separately approved safety case authorises it.
Notably Absent	A universal drift threshold; assurance that statistically stable behaviour remains operationally safe; detection of changes hidden by sensor faults, seasonal shifts or coordinated manipulation.

D1-CTL-04 — Federated Learning Poisoning Prevention

Field	Critical-infrastructure interpretation
Authoritative requirement	Gradient anomaly detection + robust aggregation.
Business objective	Preserve trustworthy operational decisions by preventing corrupted data or model state from causing unsafe control recommendations, missed degradation, service interruption or cascading cyber-physical effects.
OT/ICS-adapted equivalent	Use signed offline manifests, approved removable-media transfer, local hash registries and isolated scanning when continuous repositories or CI/CD are unavailable.

Field	Critical-infrastructure interpretation
Criticality scaling (CI-1 to CI-5)	CI-1: Document source, model or dataset identity; perform periodic integrity review and retain a reproducible baseline. CI-2: Add approved provenance records, version-controlled integrity checks and representative offline validation before material changes. CI-3: Use automated integrity verification, independent challenge of high-impact changes and tested rollback to the last known-good model/data state. CI-4: Require cryptographically protected provenance, dual authorization for high-impact changes, isolated adversarial testing and continuous integrity monitoring at trust boundaries. CI-5: Use hardware-backed or independently anchored attestation where feasible, multi-party approval, redundant validation paths, safety-engineering review and evidence that integrity controls do not impair deterministic operations.
Evidence requirement	Signed JSON/CSV manifest containing asset/model/data version, hashes, source, test results, thresholds, exceptions and reviewer.
Testing boundary	high-fidelity digital twin, hardware-in-the-loop rig, or strictly isolated staging network. Adversarial, poisoning, fault-injection and unsafe-command tests must not be routed through a live production control or actuation path unless a separately approved safety case authorises it.
Notably Absent	A guarantee against zero-day attack paths, authorized-insider misuse or failures outside the tested system boundary; proof of safety, legal compliance or operating effectiveness from design evidence alone.

D1-CTL-05 — Embedding Space Robustness

Field	Critical-infrastructure interpretation
Authoritative requirement	Adversarial training + certified robustness measurement.
Business objective	Preserve trustworthy operational decisions by preventing corrupted data or model state from causing unsafe control recommendations, missed degradation, service interruption or cascading cyber-physical effects.
OT/ICS-adapted equivalent	Use signed offline manifests, approved removable-media transfer, local hash registries and isolated scanning when continuous repositories or CI/CD are unavailable.
Criticality scaling (CI-1 to CI-5)	CI-1: Document source, model or dataset identity; perform periodic integrity review and retain a reproducible baseline. CI-2: Add approved provenance records, version-controlled integrity checks and representative offline validation before material changes. CI-3: Use automated integrity verification, independent challenge of high-impact changes and tested rollback to the last known-good model/data state. CI-4: Require cryptographically protected provenance, dual authorization for high-impact changes, isolated adversarial testing and continuous integrity monitoring at trust boundaries. CI-5: Use hardware-backed or independently anchored attestation where feasible, multi-party approval, redundant validation paths, safety-engineering review and evidence that integrity controls do not impair deterministic operations.
Evidence requirement	Signed JSON/CSV manifest containing asset/model/data version, hashes, source, test results, thresholds, exceptions and reviewer.
Testing boundary	high-fidelity digital twin, hardware-in-the-loop rig, or strictly isolated staging network. Adversarial, poisoning, fault-injection and unsafe-command tests must not be routed through a live production control or actuation path unless a separately approved safety case authorises it.
Notably Absent	A guarantee against zero-day attack paths, authorized-insider misuse or failures outside the tested system boundary; proof of safety, legal compliance or operating effectiveness from design evidence alone.

D1-CTL-06 — Post-Quantum Model Signing & Crypto Hardening

Field	Critical-infrastructure interpretation
Authoritative requirement	PQC signing (ML-DSA/SLH-DSA) + PQC key exchange (ML-KEM).
Business objective	Preserve trustworthy operational decisions by preventing corrupted data or model state from causing unsafe control recommendations, missed degradation, service interruption or cascading cyber-physical effects.
OT/ICS-adapted equivalent	Use signed offline manifests, approved removable-media transfer, local hash registries and isolated scanning when continuous repositories or CI/CD are unavailable.
Criticality scaling (CI-1 to CI-5)	CI-1: Document source, model or dataset identity; perform periodic integrity review and retain a reproducible baseline. CI-2: Add approved provenance records, version-controlled integrity checks and representative offline validation before material changes. CI-3: Use automated integrity verification, independent challenge of high-impact changes and tested rollback to the last known-good model/data state. CI-4: Require cryptographically protected provenance, dual authorization for high-impact changes, isolated adversarial testing and continuous integrity monitoring at trust boundaries. CI-5: Use hardware-backed or independently anchored attestation where feasible, multi-party approval, redundant validation paths, safety-engineering review and evidence that integrity controls do not impair deterministic operations.
Evidence requirement	Signed JSON/CSV manifest containing asset/model/data version, hashes, source, test results, thresholds, exceptions and reviewer.
Testing boundary	high-fidelity digital twin, hardware-in-the-loop rig, or strictly isolated staging network. Adversarial, poisoning, fault-injection and unsafe-command tests must not be routed through a live production control or actuation path unless a separately approved safety case authorises it.
Notably Absent	A guarantee against zero-day attack paths, authorized-insider misuse or failures outside the tested system boundary; proof of safety, legal compliance or operating effectiveness from design evidence alone.

D1-CTL-07 — LoRA/Adapter Integrity Verification

Field	Critical-infrastructure interpretation
Authoritative requirement	Adapter scanning + provenance verification + registry allowlist.
Business objective	Preserve trustworthy operational decisions by preventing corrupted data or model state from causing unsafe control recommendations, missed degradation, service interruption or cascading cyber-physical effects.
OT/ICS-adapted equivalent	Use signed offline manifests, approved removable-media transfer, local hash registries and isolated scanning when continuous repositories or CI/CD are unavailable.

Field	Critical-infrastructure interpretation
Criticality scaling (CI-1 to CI-5)	CI-1: Document source, model or dataset identity; perform periodic integrity review and retain a reproducible baseline. CI-2: Add approved provenance records, version-controlled integrity checks and representative offline validation before material changes. CI-3: Use automated integrity verification, independent challenge of high-impact changes and tested rollback to the last known-good model/data state. CI-4: Require cryptographically protected provenance, dual authorization for high-impact changes, isolated adversarial testing and continuous integrity monitoring at trust boundaries. CI-5: Use hardware-backed or independently anchored attestation where feasible, multi-party approval, redundant validation paths, safety-engineering review and evidence that integrity controls do not impair deterministic operations.
Evidence requirement	Signed JSON/CSV manifest containing asset/model/data version, hashes, source, test results, thresholds, exceptions and reviewer.
Testing boundary	high-fidelity digital twin, hardware-in-the-loop rig, or strictly isolated staging network. Adversarial, poisoning, fault-injection and unsafe-command tests must not be routed through a live production control or actuation path unless a separately approved safety case authorises it.
Notably Absent	A guarantee against zero-day attack paths, authorized-insider misuse or failures outside the tested system boundary; proof of safety, legal compliance or operating effectiveness from design evidence alone.

D1-CTL-08 — Model Merge Attack Detection

Field	Critical-infrastructure interpretation
Authoritative requirement	Pre-registration behavioural evaluation + regression testing.
Business objective	Preserve trustworthy operational decisions by preventing corrupted data or model state from causing unsafe control recommendations, missed degradation, service interruption or cascading cyber-physical effects.
OT/ICS-adapted equivalent	Use signed offline manifests, approved removable-media transfer, local hash registries and isolated scanning when continuous repositories or CI/CD are unavailable.
Criticality scaling (CI-1 to CI-5)	CI-1: Document source, model or dataset identity; perform periodic integrity review and retain a reproducible baseline. CI-2: Add approved provenance records, version-controlled integrity checks and representative offline validation before material changes. CI-3: Use automated integrity verification, independent challenge of high-impact changes and tested rollback to the last known-good model/data state. CI-4: Require cryptographically protected provenance, dual authorization for high-impact changes, isolated adversarial testing and continuous integrity monitoring at trust boundaries. CI-5: Use hardware-backed or independently anchored attestation where feasible, multi-party approval, redundant validation paths, safety-engineering review and evidence that integrity controls do not impair deterministic operations.
Evidence requirement	Signed JSON/CSV manifest containing asset/model/data version, hashes, source, test results, thresholds, exceptions and reviewer.
Testing boundary	high-fidelity digital twin, hardware-in-the-loop rig, or strictly isolated staging network. Adversarial, poisoning, fault-injection and unsafe-command tests must not be routed through a live production control or actuation path unless a separately approved safety case authorises it.
Notably Absent	A guarantee against zero-day attack paths, authorized-insider misuse or failures outside the tested system boundary; proof of safety, legal compliance or operating effectiveness from design evidence alone.

D1-CTL-09 — Quantization Backdoor Screening

Field	Critical-infrastructure interpretation
Authoritative requirement	Cross-precision behavioural comparison + delta threshold monitoring.
Business objective	Preserve trustworthy operational decisions by preventing corrupted data or model state from causing unsafe control recommendations, missed degradation, service interruption or cascading cyber-physical effects.
OT/ICS-adapted equivalent	Use signed offline manifests, approved removable-media transfer, local hash registries and isolated scanning when continuous repositories or CI/CD are unavailable.
Criticality scaling (CI-1 to CI-5)	CI-1: Document source, model or dataset identity; perform periodic integrity review and retain a reproducible baseline. CI-2: Add approved provenance records, version-controlled integrity checks and representative offline validation before material changes. CI-3: Use automated integrity verification, independent challenge of high-impact changes and tested rollback to the last known-good model/data state. CI-4: Require cryptographically protected provenance, dual authorization for high-impact changes, isolated adversarial testing and continuous integrity monitoring at trust boundaries. CI-5: Use hardware-backed or independently anchored attestation where feasible, multi-party approval, redundant validation paths, safety-engineering review and evidence that integrity controls do not impair deterministic operations.
Evidence requirement	Signed JSON/CSV manifest containing asset/model/data version, hashes, source, test results, thresholds, exceptions and reviewer.
Testing boundary	high-fidelity digital twin, hardware-in-the-loop rig, or strictly isolated staging network. Adversarial, poisoning, fault-injection and unsafe-command tests must not be routed through a live production control or actuation path unless a separately approved safety case authorises it.
Notably Absent	A guarantee against zero-day attack paths, authorized-insider misuse or failures outside the tested system boundary; proof of safety, legal compliance or operating effectiveness from design evidence alone.

D2-CTL-01 — Direct Prompt Injection Prevention

Field	Critical-infrastructure interpretation
Authoritative requirement	Input validation + adversarial pattern matching + system prompt isolation + guardrail sidecar.
Business objective	Prevent adversarial inputs or tool manipulation from bypassing authority boundaries, confusing operators or initiating unsafe operational actions.
OT/ICS-adapted equivalent	Execute adversarial testing only in a digital twin, HIL rig or isolated staging zone. Production testing requires a separately approved safety case and must not expose actuation paths.
Criticality scaling (CI-1 to CI-5)	CI-1: Apply input constraints and documented misuse testing in a non-production test environment. CI-2: Add representative adversarial test cases, response logging and periodic rule/model review. CI-3: Use isolated red-team testing, tool-call allowlists, privilege boundaries and independent review of bypass findings. CI-4: Validate guardrails in a high-fidelity shadow environment; require pre-execution authorization for write or actuation tools and rapid isolation capability. CI-5: Prohibit uncontrolled adversarial testing on production; use HIL/digital-twin testing, deterministic non-AI safety enforcement and independent safety/security acceptance.

Field	Critical-infrastructure interpretation
Evidence requirement	Signed test report containing payload class, environment, target version, blocked/allowed outcome, tool call, privilege boundary, latency and restoration result.
Testing boundary	high-fidelity digital twin, hardware-in-the-loop rig, or strictly isolated staging network. Adversarial, poisoning, fault-injection and unsafe-command tests must not be routed through a live production control or actuation path unless a separately approved safety case authorises it.
Notably Absent	Protection against every semantic or multimodal jailbreak; permission to test adversarial payloads against live production control paths; assurance that pattern matching alone blocks novel attacks.

D2-CTL-02 — Indirect Prompt Injection Prevention

Field	Critical-infrastructure interpretation
Authoritative requirement	Contextual separation + source allowlisting + output validation + RAG sanitization pipeline.
Business objective	Prevent adversarial inputs or tool manipulation from bypassing authority boundaries, confusing operators or initiating unsafe operational actions.
OT/ICS-adapted equivalent	Execute adversarial testing only in a digital twin, HIL rig or isolated staging zone. Production testing requires a separately approved safety case and must not expose actuation paths.
Criticality scaling (CI-1 to CI-5)	CI-1: Apply input constraints and documented misuse testing in a non-production test environment. CI-2: Add representative adversarial test cases, response logging and periodic rule/model review. CI-3: Use isolated red-team testing, tool-call allowlists, privilege boundaries and independent review of bypass findings. CI-4: Validate guardrails in a high-fidelity shadow environment; require pre-execution authorization for write or actuation tools and rapid isolation capability. CI-5: Prohibit uncontrolled adversarial testing on production; use HIL/digital-twin testing, deterministic non-AI safety enforcement and independent safety/security acceptance.
Evidence requirement	Signed test report containing payload class, environment, target version, blocked/allowed outcome, tool call, privilege boundary, latency and restoration result.
Testing boundary	high-fidelity digital twin, hardware-in-the-loop rig, or strictly isolated staging network. Adversarial, poisoning, fault-injection and unsafe-command tests must not be routed through a live production control or actuation path unless a separately approved safety case authorises it.
Notably Absent	Protection against every semantic or multimodal jailbreak; permission to test adversarial payloads against live production control paths; assurance that pattern matching alone blocks novel attacks.

D2-CTL-03 — Jailbreak Resistance Testing

Field	Critical-infrastructure interpretation
Authoritative requirement	Quarterly red-team prompt library + adversarial training + automated refusal monitoring.
Business objective	Prevent adversarial inputs or tool manipulation from bypassing authority boundaries, confusing operators or initiating unsafe operational actions.

Field	Critical-infrastructure interpretation
OT/ICS-adapted equivalent	Execute adversarial testing only in a digital twin, HIL rig or isolated staging zone. Production testing requires a separately approved safety case and must not expose actuation paths.
Criticality scaling (CI-1 to CI-5)	CI-1: Apply input constraints and documented misuse testing in a non-production test environment. CI-2: Add representative adversarial test cases, response logging and periodic rule/model review. CI-3: Use isolated red-team testing, tool-call allowlists, privilege boundaries and independent review of bypass findings. CI-4: Validate guardrails in a high-fidelity shadow environment; require pre-execution authorization for write or actuation tools and rapid isolation capability. CI-5: Prohibit uncontrolled adversarial testing on production; use HIL/digital-twin testing, deterministic non-AI safety enforcement and independent safety/security acceptance.
Evidence requirement	Signed test report containing payload class, environment, target version, blocked/allowed outcome, tool call, privilege boundary, latency and restoration result.
Testing boundary	high-fidelity digital twin, hardware-in-the-loop rig, or strictly isolated staging network. Adversarial, poisoning, fault-injection and unsafe-command tests must not be routed through a live production control or actuation path unless a separately approved safety case authorises it.
Notably Absent	A guarantee against zero-day attack paths, authorized-insider misuse or failures outside the tested system boundary; proof of safety, legal compliance or operating effectiveness from design evidence alone.

D2-CTL-04 — Multi-Modal Injection Defense

Field	Critical-infrastructure interpretation
Authoritative requirement	Multi-modal content scanning + steganography detection + modality-specific guardrails.
Business objective	Prevent adversarial inputs or tool manipulation from bypassing authority boundaries, confusing operators or initiating unsafe operational actions.
OT/ICS-adapted equivalent	Execute adversarial testing only in a digital twin, HIL rig or isolated staging zone. Production testing requires a separately approved safety case and must not expose actuation paths.
Criticality scaling (CI-1 to CI-5)	CI-1: Apply input constraints and documented misuse testing in a non-production test environment. CI-2: Add representative adversarial test cases, response logging and periodic rule/model review. CI-3: Use isolated red-team testing, tool-call allowlists, privilege boundaries and independent review of bypass findings. CI-4: Validate guardrails in a high-fidelity shadow environment; require pre-execution authorization for write or actuation tools and rapid isolation capability. CI-5: Prohibit uncontrolled adversarial testing on production; use HIL/digital-twin testing, deterministic non-AI safety enforcement and independent safety/security acceptance.
Evidence requirement	Signed test report containing payload class, environment, target version, blocked/allowed outcome, tool call, privilege boundary, latency and restoration result.
Testing boundary	high-fidelity digital twin, hardware-in-the-loop rig, or strictly isolated staging network. Adversarial, poisoning, fault-injection and unsafe-command tests must not be routed through a live production control or actuation path unless a separately approved safety case authorises it.

Field	Critical-infrastructure interpretation
Notably Absent	A guarantee against zero-day attack paths, authorized-insider misuse or failures outside the tested system boundary; proof of safety, legal compliance or operating effectiveness from design evidence alone.

D2-CTL-05 — Function Call/Tool Call Injection Prevention

Field	Critical-infrastructure interpretation
Authoritative requirement	Parameter schema validation + allowlist enforcement + sandboxed execution.
Business objective	Prevent adversarial inputs or tool manipulation from bypassing authority boundaries, confusing operators or initiating unsafe operational actions.
OT/ICS-adapted equivalent	Execute adversarial testing only in a digital twin, HIL rig or isolated staging zone. Production testing requires a separately approved safety case and must not expose actuation paths.
Criticality scaling (CI-1 to CI-5)	CI-1: Apply input constraints and documented misuse testing in a non-production test environment. CI-2: Add representative adversarial test cases, response logging and periodic rule/model review. CI-3: Use isolated red-team testing, tool-call allowlists, privilege boundaries and independent review of bypass findings. CI-4: Validate guardrails in a high-fidelity shadow environment; require pre-execution authorization for write or actuation tools and rapid isolation capability. CI-5: Prohibit uncontrolled adversarial testing on production; use HIL/digital-twin testing, deterministic non-AI safety enforcement and independent safety/security acceptance.
Evidence requirement	Signed test report containing payload class, environment, target version, blocked/allowed outcome, tool call, privilege boundary, latency and restoration result.
Testing boundary	high-fidelity digital twin, hardware-in-the-loop rig, or strictly isolated staging network. Adversarial, poisoning, fault-injection and unsafe-command tests must not be routed through a live production control or actuation path unless a separately approved safety case authorises it.
Notably Absent	A guarantee against zero-day attack paths, authorized-insider misuse or failures outside the tested system boundary; proof of safety, legal compliance or operating effectiveness from design evidence alone.

D2-CTL-06 — Cross-Context Hijacking Mitigation

Field	Critical-infrastructure interpretation
Authoritative requirement	Context window segmentation + prompt anchoring + attention boundary enforcement.
Business objective	Prevent adversarial inputs or tool manipulation from bypassing authority boundaries, confusing operators or initiating unsafe operational actions.
OT/ICS-adapted equivalent	Execute adversarial testing only in a digital twin, HIL rig or isolated staging zone. Production testing requires a separately approved safety case and must not expose actuation paths.

Field	Critical-infrastructure interpretation
Criticality scaling (CI-1 to CI-5)	CI-1: Apply input constraints and documented misuse testing in a non-production test environment. CI-2: Add representative adversarial test cases, response logging and periodic rule/model review. CI-3: Use isolated red-team testing, tool-call allowlists, privilege boundaries and independent review of bypass findings. CI-4: Validate guardrails in a high-fidelity shadow environment; require pre-execution authorization for write or actuation tools and rapid isolation capability. CI-5: Prohibit uncontrolled adversarial testing on production; use HIL/digital-twin testing, deterministic non-AI safety enforcement and independent safety/security acceptance.
Evidence requirement	Signed test report containing payload class, environment, target version, blocked/allowed outcome, tool call, privilege boundary, latency and restoration result.
Testing boundary	high-fidelity digital twin, hardware-in-the-loop rig, or strictly isolated staging network. Adversarial, poisoning, fault-injection and unsafe-command tests must not be routed through a live production control or actuation path unless a separately approved safety case authorises it.
Notably Absent	A guarantee against zero-day attack paths, authorized-insider misuse or failures outside the tested system boundary; proof of safety, legal compliance or operating effectiveness from design evidence alone.

D3-CTL-01 — Least Agency Enforcement

Field	Critical-infrastructure interpretation
Authoritative requirement	Role-based tool scoping + policy-as-code + dynamic permission revocation.
Business objective	Maintain enforceable identity, privilege and trust boundaries across segmented and intermittently connected IT/OT environments.
OT/ICS-adapted equivalent	Where continuous attestation is infeasible, use hardware roots of trust, signed build/maintenance manifests, physical key custody, scheduled offline verification and controlled maintenance access.
Criticality scaling (CI-1 to CI-5)	CI-1: Document identities, roles and trust boundaries; use least privilege appropriate to the deployment. CI-2: Use managed identities, periodic access review and offline-verifiable credentials where continuous connectivity is unavailable. CI-3: Require strong workload identity, segmented trust zones and dual control for privileged changes. CI-4: Use hardware roots of trust or compensating physical controls, short-lived or batched signed credentials, and monitored cross-zone access. CI-5: Integrate identity decisions with safety authority boundaries; maintain independently enforceable deny states and tested offline operation.
Evidence requirement	Identity/credential inventory, signed configuration export, access review, offline attestation log and exception record.
Testing boundary	document/configuration inspection and representative offline test. Adversarial, poisoning, fault-injection and unsafe-command tests must not be routed through a live production control or actuation path unless a separately approved safety case authorises it.
Notably Absent	A guarantee against zero-day attack paths, authorized-insider misuse or failures outside the tested system boundary; proof of safety, legal compliance or operating effectiveness from design evidence alone.

D3-CTL-02 — Inter-Agent Communication Security

Field	Critical-infrastructure interpretation
Authoritative requirement	mTLS for agent mesh + message signing + payload validation.
Business objective	Maintain enforceable identity, privilege and trust boundaries across segmented and intermittently connected IT/OT environments.
OT/ICS-adapted equivalent	Where continuous attestation is infeasible, use hardware roots of trust, signed build/maintenance manifests, physical key custody, scheduled offline verification and controlled maintenance access.
Criticality scaling (CI-1 to CI-5)	CI-1: Document identities, roles and trust boundaries; use least privilege appropriate to the deployment. CI-2: Use managed identities, periodic access review and offline-verifiable credentials where continuous connectivity is unavailable. CI-3: Require strong workload identity, segmented trust zones and dual control for privileged changes. CI-4: Use hardware roots of trust or compensating physical controls, short-lived or batched signed credentials, and monitored cross-zone access. CI-5: Integrate identity decisions with safety authority boundaries; maintain independently enforceable deny states and tested offline operation.
Evidence requirement	Identity/credential inventory, signed configuration export, access review, offline attestation log and exception record.
Testing boundary	document/configuration inspection and representative offline test. Adversarial, poisoning, fault-injection and unsafe-command tests must not be routed through a live production control or actuation path unless a separately approved safety case authorises it.
Notably Absent	A guarantee against zero-day attack paths, authorized-insider misuse or failures outside the tested system boundary; proof of safety, legal compliance or operating effectiveness from design evidence alone.

D3-CTL-03 — Agentic Prompt Chaining Detection

Field	Critical-infrastructure interpretation
Authoritative requirement	Cross-session behavioural correlation + chain pattern detection + anomaly scoring.
Business objective	Maintain enforceable identity, privilege and trust boundaries across segmented and intermittently connected IT/OT environments.
OT/ICS-adapted equivalent	Where continuous attestation is infeasible, use hardware roots of trust, signed build/maintenance manifests, physical key custody, scheduled offline verification and controlled maintenance access.
Criticality scaling (CI-1 to CI-5)	CI-1: Document identities, roles and trust boundaries; use least privilege appropriate to the deployment. CI-2: Use managed identities, periodic access review and offline-verifiable credentials where continuous connectivity is unavailable. CI-3: Require strong workload identity, segmented trust zones and dual control for privileged changes. CI-4: Use hardware roots of trust or compensating physical controls, short-lived or batched signed credentials, and monitored cross-zone access. CI-5: Integrate identity decisions with safety authority boundaries; maintain independently enforceable deny states and tested offline operation.
Evidence requirement	Identity/credential inventory, signed configuration export, access review, offline attestation log and exception record.

Field	Critical-infrastructure interpretation
Testing boundary	document/configuration inspection and representative offline test. Adversarial, poisoning, fault-injection and unsafe-command tests must not be routed through a live production control or actuation path unless a separately approved safety case authorises it.
Notably Absent	A guarantee against zero-day attack paths, authorized-insider misuse or failures outside the tested system boundary; proof of safety, legal compliance or operating effectiveness from design evidence alone.

D3-CTL-04 — Embodied AI Safety Controls

Field	Critical-infrastructure interpretation
Authoritative requirement	Sensor integrity verification + safety interlocks + fail-safe state enforcement.
Business objective	Maintain enforceable identity, privilege and trust boundaries across segmented and intermittently connected IT/OT environments.
OT/ICS-adapted equivalent	Where continuous attestation is infeasible, use hardware roots of trust, signed build/maintenance manifests, physical key custody, scheduled offline verification and controlled maintenance access.
Criticality scaling (CI-1 to CI-5)	CI-1: Document identities, roles and trust boundaries; use least privilege appropriate to the deployment. CI-2: Use managed identities, periodic access review and offline-verifiable credentials where continuous connectivity is unavailable. CI-3: Require strong workload identity, segmented trust zones and dual control for privileged changes. CI-4: Use hardware roots of trust or compensating physical controls, short-lived or batched signed credentials, and monitored cross-zone access. CI-5: Integrate identity decisions with safety authority boundaries; maintain independently enforceable deny states and tested offline operation.
Evidence requirement	Identity/credential inventory, signed configuration export, access review, offline attestation log and exception record.
Testing boundary	document/configuration inspection and representative offline test. Adversarial, poisoning, fault-injection and unsafe-command tests must not be routed through a live production control or actuation path unless a separately approved safety case authorises it.
Notably Absent	A guarantee against zero-day attack paths, authorized-insider misuse or failures outside the tested system boundary; proof of safety, legal compliance or operating effectiveness from design evidence alone.

D3-CTL-05 — Multi-Agent Trust Chain Attestation

Field	Critical-infrastructure interpretation
Authoritative requirement	SPIFFE/SPIRE workload identity + short-lived certificates + continuous attestation.
Business objective	Maintain enforceable identity, privilege and trust boundaries across segmented and intermittently connected IT/OT environments.
OT/ICS-adapted equivalent	Where continuous attestation is infeasible, use hardware roots of trust, signed build/maintenance manifests, physical key custody, scheduled offline verification and controlled maintenance access.

Field	Critical-infrastructure interpretation
Criticality scaling (CI-1 to CI-5)	CI-1: Document identities, roles and trust boundaries; use least privilege appropriate to the deployment. CI-2: Use managed identities, periodic access review and offline-verifiable credentials where continuous connectivity is unavailable. CI-3: Require strong workload identity, segmented trust zones and dual control for privileged changes. CI-4: Use hardware roots of trust or compensating physical controls, short-lived or batched signed credentials, and monitored cross-zone access. CI-5: Integrate identity decisions with safety authority boundaries; maintain independently enforceable deny states and tested offline operation.
Evidence requirement	Identity/credential inventory, signed configuration export, access review, offline attestation log and exception record.
Testing boundary	document/configuration inspection and representative offline test. Adversarial, poisoning, fault-injection and unsafe-command tests must not be routed through a live production control or actuation path unless a separately approved safety case authorises it.
Notably Absent	An assumption of continuous external connectivity or cloud control-plane access; a requirement to deploy SPIFFE/SPIRE where offline hardware identity, signed manifests or physical controls provide a justified equivalent.

D3-CTL-06 — Persistent Memory Exfiltration Prevention

Field	Critical-infrastructure interpretation
Authoritative requirement	User-scoped memory isolation + encryption at rest + query-level access controls.
Business objective	Maintain enforceable identity, privilege and trust boundaries across segmented and intermittently connected IT/OT environments.
OT/ICS-adapted equivalent	Where continuous attestation is infeasible, use hardware roots of trust, signed build/maintenance manifests, physical key custody, scheduled offline verification and controlled maintenance access.
Criticality scaling (CI-1 to CI-5)	CI-1: Document identities, roles and trust boundaries; use least privilege appropriate to the deployment. CI-2: Use managed identities, periodic access review and offline-verifiable credentials where continuous connectivity is unavailable. CI-3: Require strong workload identity, segmented trust zones and dual control for privileged changes. CI-4: Use hardware roots of trust or compensating physical controls, short-lived or batched signed credentials, and monitored cross-zone access. CI-5: Integrate identity decisions with safety authority boundaries; maintain independently enforceable deny states and tested offline operation.
Evidence requirement	Identity/credential inventory, signed configuration export, access review, offline attestation log and exception record.
Testing boundary	document/configuration inspection and representative offline test. Adversarial, poisoning, fault-injection and unsafe-command tests must not be routed through a live production control or actuation path unless a separately approved safety case authorises it.
Notably Absent	A guarantee against zero-day attack paths, authorized-insider misuse or failures outside the tested system boundary; proof of safety, legal compliance or operating effectiveness from design evidence alone.

D3-CTL-07 — Secure Memory Lifecycle Management

Field	Critical-infrastructure interpretation
Authoritative requirement	Cryptographic deletion + lifecycle policy enforcement + retention auditing.
Business objective	Maintain enforceable identity, privilege and trust boundaries across segmented and intermittently connected IT/OT environments.
OT/ICS-adapted equivalent	Where continuous attestation is infeasible, use hardware roots of trust, signed build/maintenance manifests, physical key custody, scheduled offline verification and controlled maintenance access.
Criticality scaling (CI-1 to CI-5)	CI-1: Document identities, roles and trust boundaries; use least privilege appropriate to the deployment. CI-2: Use managed identities, periodic access review and offline-verifiable credentials where continuous connectivity is unavailable. CI-3: Require strong workload identity, segmented trust zones and dual control for privileged changes. CI-4: Use hardware roots of trust or compensating physical controls, short-lived or batched signed credentials, and monitored cross-zone access. CI-5: Integrate identity decisions with safety authority boundaries; maintain independently enforceable deny states and tested offline operation.
Evidence requirement	Identity/credential inventory, signed configuration export, access review, offline attestation log and exception record.
Testing boundary	document/configuration inspection and representative offline test. Adversarial, poisoning, fault-injection and unsafe-command tests must not be routed through a live production control or actuation path unless a separately approved safety case authorises it.
Notably Absent	A guarantee against zero-day attack paths, authorized-insider misuse or failures outside the tested system boundary; proof of safety, legal compliance or operating effectiveness from design evidence alone.

D4-CTL-01 — AI Bill Of Materials (AI BOM) Maintenance

Field	Critical-infrastructure interpretation
Authoritative requirement	Automated BOM generation + version tracking + registry synchronization.
Business objective	Make AI-enabled assets and supplier dependencies visible enough to manage outages, vulnerabilities, unsupported components and unauthorized deployment (AI Bill Of Materials (AI Bom) Maintenance).
OT/ICS-adapted equivalent	Where vendor internals are unavailable, record vendor assertions separately from independently observed boundary behaviour, network dependencies, update paths and failure modes.
Criticality scaling (CI-1 to CI-5)	CI-1: Maintain a basic inventory of AI-enabled assets, suppliers and interfaces. CI-2: Add version, owner, data path, network zone and support-status records. CI-3: Require supplier evidence or documented boundary testing where internals are unavailable; monitor unauthorized AI use. CI-4: Maintain site-level dependency maps, signed inventory changes and contingency plans for opaque suppliers. CI-5: Use independent boundary assurance, contractual escalation, replacement/containment plans and executive acceptance of unresolved supplier opacity.
Evidence requirement	AI/asset BOM, supplier disclosure record, dependency map, boundary-test report and unsupported-component decision.
Testing boundary	document/configuration inspection and representative offline test. Adversarial, poisoning, fault-injection and unsafe-command tests must not be routed through a live production control or actuation path unless a separately approved safety case authorises it.

Field	Critical-infrastructure interpretation
Notably Absent	An assumption that suppliers disclose model weights or training data; acceptance of vendor assertions as conclusive evidence where independent boundary testing is feasible.

D4-CTL-02 — Model File & Artifact Scanning

Field	Critical-infrastructure interpretation
Authoritative requirement	Static analysis + deserialization sandboxing + signature verification.
Business objective	Make AI-enabled assets and supplier dependencies visible enough to manage outages, vulnerabilities, unsupported components and unauthorized deployment.
OT/ICS-adapted equivalent	Where vendor internals are unavailable, record vendor assertions separately from independently observed boundary behaviour, network dependencies, update paths and failure modes.
Criticality scaling (CI-1 to CI-5)	CI-1: Maintain a basic inventory of AI-enabled assets, suppliers and interfaces. CI-2: Add version, owner, data path, network zone and support-status records. CI-3: Require supplier evidence or documented boundary testing where internals are unavailable; monitor unauthorized AI use. CI-4: Maintain site-level dependency maps, signed inventory changes and contingency plans for opaque suppliers. CI-5: Use independent boundary assurance, contractual escalation, replacement/containment plans and executive acceptance of unresolved supplier opacity.
Evidence requirement	AI/asset BOM, supplier disclosure record, dependency map, boundary-test report and unsupported-component decision.
Testing boundary	document/configuration inspection and representative offline test. Adversarial, poisoning, fault-injection and unsafe-command tests must not be routed through a live production control or actuation path unless a separately approved safety case authorises it.
Notably Absent	A guarantee against zero-day attack paths, authorized-insider misuse or failures outside the tested system boundary; proof of safety, legal compliance or operating effectiveness from design evidence alone.

D4-CTL-03 — Model Hub & Registry Vetting

Field	Critical-infrastructure interpretation
Authoritative requirement	Provenance verification + license compliance + security scorecard.
Business objective	Make AI-enabled assets and supplier dependencies visible enough to manage outages, vulnerabilities, unsupported components and unauthorized deployment.
OT/ICS-adapted equivalent	Where vendor internals are unavailable, record vendor assertions separately from independently observed boundary behaviour, network dependencies, update paths and failure modes.
Criticality scaling (CI-1 to CI-5)	CI-1: Maintain a basic inventory of AI-enabled assets, suppliers and interfaces. CI-2: Add version, owner, data path, network zone and support-status records. CI-3: Require supplier evidence or documented boundary testing where internals are unavailable; monitor unauthorized AI use. CI-4: Maintain site-level dependency maps, signed inventory changes and contingency plans for opaque suppliers. CI-5: Use independent boundary assurance, contractual escalation, replacement/containment plans and executive acceptance of unresolved supplier opacity.

Field	Critical-infrastructure interpretation
Evidence requirement	AI/asset BOM, supplier disclosure record, dependency map, boundary-test report and unsupported-component decision.
Testing boundary	document/configuration inspection and representative offline test. Adversarial, poisoning, fault-injection and unsafe-command tests must not be routed through a live production control or actuation path unless a separately approved safety case authorises it.
Notably Absent	A guarantee against zero-day attack paths, authorized-insider misuse or failures outside the tested system boundary; proof of safety, legal compliance or operating effectiveness from design evidence alone.

D4-CTL-04 — MCP Server Behavioral Monitoring

Field	Critical-infrastructure interpretation
Authoritative requirement	Tool-call logging + anomaly detection + access control enforcement.
Business objective	Make AI-enabled assets and supplier dependencies visible enough to manage outages, vulnerabilities, unsupported components and unauthorized deployment.
OT/ICS-adapted equivalent	Where vendor internals are unavailable, record vendor assertions separately from independently observed boundary behaviour, network dependencies, update paths and failure modes.
Criticality scaling (CI-1 to CI-5)	CI-1: Maintain a basic inventory of AI-enabled assets, suppliers and interfaces. CI-2: Add version, owner, data path, network zone and support-status records. CI-3: Require supplier evidence or documented boundary testing where internals are unavailable; monitor unauthorized AI use. CI-4: Maintain site-level dependency maps, signed inventory changes and contingency plans for opaque suppliers. CI-5: Use independent boundary assurance, contractual escalation, replacement/containment plans and executive acceptance of unresolved supplier opacity.
Evidence requirement	AI/asset BOM, supplier disclosure record, dependency map, boundary-test report and unsupported-component decision.
Testing boundary	document/configuration inspection and representative offline test. Adversarial, poisoning, fault-injection and unsafe-command tests must not be routed through a live production control or actuation path unless a separately approved safety case authorises it.
Notably Absent	A guarantee against zero-day attack paths, authorized-insider misuse or failures outside the tested system boundary; proof of safety, legal compliance or operating effectiveness from design evidence alone.

D4-CTL-05 — Third-Party AI Api Security Assessment

Field	Critical-infrastructure interpretation
Authoritative requirement	Contractual security requirements + penetration testing + data flow mapping.
Business objective	Make AI-enabled assets and supplier dependencies visible enough to manage outages, vulnerabilities, unsupported components and unauthorized deployment.
OT/ICS-adapted equivalent	Where vendor internals are unavailable, record vendor assertions separately from independently observed boundary behaviour, network dependencies, update paths and failure modes.

Field	Critical-infrastructure interpretation
Criticality scaling (CI-1 to CI-5)	CI-1: Maintain a basic inventory of AI-enabled assets, suppliers and interfaces. CI-2: Add version, owner, data path, network zone and support-status records. CI-3: Require supplier evidence or documented boundary testing where internals are unavailable; monitor unauthorized AI use. CI-4: Maintain site-level dependency maps, signed inventory changes and contingency plans for opaque suppliers. CI-5: Use independent boundary assurance, contractual escalation, replacement/containment plans and executive acceptance of unresolved supplier opacity.
Evidence requirement	AI/asset BOM, supplier disclosure record, dependency map, boundary-test report and unsupported-component decision.
Testing boundary	document/configuration inspection and representative offline test. Adversarial, poisoning, fault-injection and unsafe-command tests must not be routed through a live production control or actuation path unless a separately approved safety case authorises it.
Notably Absent	A guarantee against zero-day attack paths, authorized-insider misuse or failures outside the tested system boundary; proof of safety, legal compliance or operating effectiveness from design evidence alone.

D4-CTL-06 — Shadow AI Discovery & Governance

Field	Critical-infrastructure interpretation
Authoritative requirement	Network traffic analysis + SaaS discovery + policy enforcement.
Business objective	Make AI-enabled assets and supplier dependencies visible enough to manage outages, vulnerabilities, unsupported components and unauthorized deployment.
OT/ICS-adapted equivalent	Where vendor internals are unavailable, record vendor assertions separately from independently observed boundary behaviour, network dependencies, update paths and failure modes.
Criticality scaling (CI-1 to CI-5)	CI-1: Maintain a basic inventory of AI-enabled assets, suppliers and interfaces. CI-2: Add version, owner, data path, network zone and support-status records. CI-3: Require supplier evidence or documented boundary testing where internals are unavailable; monitor unauthorized AI use. CI-4: Maintain site-level dependency maps, signed inventory changes and contingency plans for opaque suppliers. CI-5: Use independent boundary assurance, contractual escalation, replacement/containment plans and executive acceptance of unresolved supplier opacity.
Evidence requirement	AI/asset BOM, supplier disclosure record, dependency map, boundary-test report and unsupported-component decision.
Testing boundary	document/configuration inspection and representative offline test. Adversarial, poisoning, fault-injection and unsafe-command tests must not be routed through a live production control or actuation path unless a separately approved safety case authorises it.
Notably Absent	Visibility into every offline engineering workstation, removable-media workflow or embedded vendor feature; authorization to deploy discovery agents where they could affect validated OT configurations.

D4-CTL-07 — AI Software Composition Analysis (SCA)

Field	Critical-infrastructure interpretation
Authoritative requirement	Dependency scanning + CVE matching + automated patching.

Field	Critical-infrastructure interpretation
Business objective	Make AI-enabled assets and supplier dependencies visible enough to manage outages, vulnerabilities, unsupported components and unauthorized deployment (Ai Software Composition Analysis (Sca)).
OT/ICS-adapted equivalent	Where vendor internals are unavailable, record vendor assertions separately from independently observed boundary behaviour, network dependencies, update paths and failure modes.
Criticality scaling (CI-1 to CI-5)	CI-1: Maintain a basic inventory of AI-enabled assets, suppliers and interfaces. CI-2: Add version, owner, data path, network zone and support-status records. CI-3: Require supplier evidence or documented boundary testing where internals are unavailable; monitor unauthorized AI use. CI-4: Maintain site-level dependency maps, signed inventory changes and contingency plans for opaque suppliers. CI-5: Use independent boundary assurance, contractual escalation, replacement/containment plans and executive acceptance of unresolved supplier opacity.
Evidence requirement	AI/asset BOM, supplier disclosure record, dependency map, boundary-test report and unsupported-component decision.
Testing boundary	document/configuration inspection and representative offline test. Adversarial, poisoning, fault-injection and unsafe-command tests must not be routed through a live production control or actuation path unless a separately approved safety case authorises it.
Notably Absent	A guarantee against zero-day attack paths, authorized-insider misuse or failures outside the tested system boundary; proof of safety, legal compliance or operating effectiveness from design evidence alone.

D5-CTL-01 — Harmful Content Blocking

Field	Critical-infrastructure interpretation
Authoritative requirement	Content safety classifier + refusal engine.
Business objective	Prevent generated or transformed content from misleading operators, corrupting approved procedures or causing unsafe service decisions.
OT/ICS-adapted equivalent	Ground outputs in locally approved procedures and offline repositories; generated instructions must not supersede controlled operating documents.
Criticality scaling (CI-1 to CI-5)	CI-1: Apply task-specific output restrictions and competent review for operational content. CI-2: Add source grounding, prohibited-content tests and clear operator escalation. CI-3: Test for misleading instructions, fabricated procedures and unsafe operational ambiguity in representative workflows. CI-4: Prevent direct execution of unverified generated procedures; require authoritative source checks and dual approval for high-impact content. CI-5: Keep generated content outside deterministic protection functions; use approved procedures, independent validation and emergency fallback communications.
Evidence requirement	Approved source list, output-test corpus, review records, blocked-output log and procedure-control linkage.
Testing boundary	document/configuration inspection and representative offline test. Adversarial, poisoning, fault-injection and unsafe-command tests must not be routed through a live production control or actuation path unless a separately approved safety case authorises it.
Notably Absent	A universal definition of harmful content; assurance that filtering prevents misleading but superficially benign operational instructions; permission for generated content to replace approved operating procedures.

D5-CTL-02 — PII Leakage Prevention

Field	Critical-infrastructure interpretation
Authoritative requirement	PII detection + masking + access controls.
Business objective	Prevent generated or transformed content from misleading operators, corrupting approved procedures or causing unsafe service decisions.
OT/ICS-adapted equivalent	Ground outputs in locally approved procedures and offline repositories; generated instructions must not supersede controlled operating documents.
Criticality scaling (CI-1 to CI-5)	CI-1: Apply task-specific output restrictions and competent review for operational content. CI-2: Add source grounding, prohibited-content tests and clear operator escalation. CI-3: Test for misleading instructions, fabricated procedures and unsafe operational ambiguity in representative workflows. CI-4: Prevent direct execution of unverified generated procedures; require authoritative source checks and dual approval for high-impact content. CI-5: Keep generated content outside deterministic protection functions; use approved procedures, independent validation and emergency fallback communications.
Evidence requirement	Approved source list, output-test corpus, review records, blocked-output log and procedure-control linkage.
Testing boundary	document/configuration inspection and representative offline test. Adversarial, poisoning, fault-injection and unsafe-command tests must not be routed through a live production control or actuation path unless a separately approved safety case authorises it.
Notably Absent	A universal definition of harmful content; assurance that filtering prevents misleading but superficially benign operational instructions; permission for generated content to replace approved operating procedures.

D5-CTL-03 — Copyright Detection

Field	Critical-infrastructure interpretation
Authoritative requirement	n-gram overlap detection + refusal.
Business objective	Prevent generated or transformed content from misleading operators, corrupting approved procedures or causing unsafe service decisions.
OT/ICS-adapted equivalent	Ground outputs in locally approved procedures and offline repositories; generated instructions must not supersede controlled operating documents.
Criticality scaling (CI-1 to CI-5)	CI-1: Apply task-specific output restrictions and competent review for operational content. CI-2: Add source grounding, prohibited-content tests and clear operator escalation. CI-3: Test for misleading instructions, fabricated procedures and unsafe operational ambiguity in representative workflows. CI-4: Prevent direct execution of unverified generated procedures; require authoritative source checks and dual approval for high-impact content. CI-5: Keep generated content outside deterministic protection functions; use approved procedures, independent validation and emergency fallback communications.
Evidence requirement	Approved source list, output-test corpus, review records, blocked-output log and procedure-control linkage.
Testing boundary	document/configuration inspection and representative offline test. Adversarial, poisoning, fault-injection and unsafe-command tests must not be routed through a live production control or actuation path unless a separately approved safety case authorises it.
Notably Absent	A universal definition of harmful content; assurance that filtering prevents misleading but superficially benign operational instructions; permission for generated content to replace approved operating procedures.

D5-CTL-04 — AI Watermarking Robustness

Field	Critical-infrastructure interpretation
Authoritative requirement	C2PA-compliant watermarking + tamper resistance testing.
Business objective	Prevent generated or transformed content from misleading operators, corrupting approved procedures or causing unsafe service decisions.
OT/ICS-adapted equivalent	Ground outputs in locally approved procedures and offline repositories; generated instructions must not supersede controlled operating documents.
Criticality scaling (CI-1 to CI-5)	CI-1: Apply task-specific output restrictions and competent review for operational content. CI-2: Add source grounding, prohibited-content tests and clear operator escalation. CI-3: Test for misleading instructions, fabricated procedures and unsafe operational ambiguity in representative workflows. CI-4: Prevent direct execution of unverified generated procedures; require authoritative source checks and dual approval for high-impact content. CI-5: Keep generated content outside deterministic protection functions; use approved procedures, independent validation and emergency fallback communications.
Evidence requirement	Approved source list, output-test corpus, review records, blocked-output log and procedure-control linkage.
Testing boundary	document/configuration inspection and representative offline test. Adversarial, poisoning, fault-injection and unsafe-command tests must not be routed through a live production control or actuation path unless a separately approved safety case authorises it.
Notably Absent	A universal definition of harmful content; assurance that filtering prevents misleading but superficially benign operational instructions; permission for generated content to replace approved operating procedures.

D5-CTL-05 — Privacy-By-Design Verification

Field	Critical-infrastructure interpretation
Authoritative requirement	Data minimization + purpose limitation + machine unlearning.
Business objective	Prevent generated or transformed content from misleading operators, corrupting approved procedures or causing unsafe service decisions.
OT/ICS-adapted equivalent	Ground outputs in locally approved procedures and offline repositories; generated instructions must not supersede controlled operating documents.
Criticality scaling (CI-1 to CI-5)	CI-1: Apply task-specific output restrictions and competent review for operational content. CI-2: Add source grounding, prohibited-content tests and clear operator escalation. CI-3: Test for misleading instructions, fabricated procedures and unsafe operational ambiguity in representative workflows. CI-4: Prevent direct execution of unverified generated procedures; require authoritative source checks and dual approval for high-impact content. CI-5: Keep generated content outside deterministic protection functions; use approved procedures, independent validation and emergency fallback communications.
Evidence requirement	Approved source list, output-test corpus, review records, blocked-output log and procedure-control linkage.
Testing boundary	document/configuration inspection and representative offline test. Adversarial, poisoning, fault-injection and unsafe-command tests must not be routed through a live production control or actuation path unless a separately approved safety case authorises it.

Field	Critical-infrastructure interpretation
Notably Absent	A universal definition of harmful content; assurance that filtering prevents misleading but superficially benign operational instructions; permission for generated content to replace approved operating procedures.

D5-CTL-06 — Privacy-Preserving ML Validation

Field	Critical-infrastructure interpretation
Authoritative requirement	Differential privacy + membership inference testing.
Business objective	Prevent generated or transformed content from misleading operators, corrupting approved procedures or causing unsafe service decisions.
OT/ICS-adapted equivalent	Ground outputs in locally approved procedures and offline repositories; generated instructions must not supersede controlled operating documents.
Criticality scaling (CI-1 to CI-5)	CI-1: Apply task-specific output restrictions and competent review for operational content. CI-2: Add source grounding, prohibited-content tests and clear operator escalation. CI-3: Test for misleading instructions, fabricated procedures and unsafe operational ambiguity in representative workflows. CI-4: Prevent direct execution of unverified generated procedures; require authoritative source checks and dual approval for high-impact content. CI-5: Keep generated content outside deterministic protection functions; use approved procedures, independent validation and emergency fallback communications.
Evidence requirement	Approved source list, output-test corpus, review records, blocked-output log and procedure-control linkage.
Testing boundary	document/configuration inspection and representative offline test. Adversarial, poisoning, fault-injection and unsafe-command tests must not be routed through a live production control or actuation path unless a separately approved safety case authorises it.
Notably Absent	A universal definition of harmful content; assurance that filtering prevents misleading but superficially benign operational instructions; permission for generated content to replace approved operating procedures.

D6-CTL-01 — Human-In-The-Loop For High-Risk Actions

Field	Critical-infrastructure interpretation
Authoritative requirement	Approval workflow + policy enforcement + audit log.
Business objective	Ensure accountable humans retain practical authority to review, reject, suspend and recover AI-enabled operations under normal and emergency conditions.
OT/ICS-adapted equivalent	Provide physically and logically independent override routes, shift-specific competence, alarm-flood procedures and restoration authority.
Criticality scaling (CI-1 to CI-5)	CI-1: Define decision owner, review point and override route. CI-2: Train operators and test that review occurs in normal operations. CI-3: Measure automation bias, alarm suppression and escalation performance under realistic workload. CI-4: Require explicit approval for high-consequence actions, independent override channels and shift-aware staffing. CI-5: Use hard authority limits, tested emergency stop/takeover, recurrent simulator exercises and safety-function independence.
Evidence requirement	Authority matrix, approval/override logs, simulator exercise results, staffing/competence records and restoration decision.

Field	Critical-infrastructure interpretation
Testing boundary	document/configuration inspection and representative offline test. Adversarial, poisoning, fault-injection and unsafe-command tests must not be routed through a live production control or actuation path unless a separately approved safety case authorises it.
Notably Absent	Proof that a nominal human-in-the-loop has time, competence, information and authority to intervene; assurance that an override remains usable during alarm floods or degraded communications.

D6-CTL-02 — Audit Trail Completeness

Field	Critical-infrastructure interpretation
Authoritative requirement	Structured logging + SIEM integration + retention enforcement.
Business objective	Ensure accountable humans retain practical authority to review, reject, suspend and recover AI-enabled operations under normal and emergency conditions.
OT/ICS-adapted equivalent	Provide physically and logically independent override routes, shift-specific competence, alarm-flood procedures and restoration authority.
Criticality scaling (CI-1 to CI-5)	CI-1: Define decision owner, review point and override route. CI-2: Train operators and test that review occurs in normal operations. CI-3: Measure automation bias, alarm suppression and escalation performance under realistic workload. CI-4: Require explicit approval for high-consequence actions, independent override channels and shift-aware staffing. CI-5: Use hard authority limits, tested emergency stop/takeover, recurrent simulator exercises and safety-function independence.
Evidence requirement	Authority matrix, approval/override logs, simulator exercise results, staffing/competence records and restoration decision.
Testing boundary	document/configuration inspection and representative offline test. Adversarial, poisoning, fault-injection and unsafe-command tests must not be routed through a live production control or actuation path unless a separately approved safety case authorises it.
Notably Absent	Proof that a nominal human-in-the-loop has time, competence, information and authority to intervene; assurance that an override remains usable during alarm floods or degraded communications.

D6-CTL-03 — AI Model Card Completeness

Field	Critical-infrastructure interpretation
Authoritative requirement	Standardized template + version control + public accessibility.
Business objective	Ensure accountable humans retain practical authority to review, reject, suspend and recover AI-enabled operations under normal and emergency conditions.
OT/ICS-adapted equivalent	Provide physically and logically independent override routes, shift-specific competence, alarm-flood procedures and restoration authority.
Criticality scaling (CI-1 to CI-5)	CI-1: Define decision owner, review point and override route. CI-2: Train operators and test that review occurs in normal operations. CI-3: Measure automation bias, alarm suppression and escalation performance under realistic workload. CI-4: Require explicit approval for high-consequence actions, independent override channels and shift-aware staffing. CI-5: Use hard authority limits, tested emergency stop/takeover, recurrent simulator exercises and safety-function independence.

Field	Critical-infrastructure interpretation
Evidence requirement	Authority matrix, approval/override logs, simulator exercise results, staffing/competence records and restoration decision.
Testing boundary	document/configuration inspection and representative offline test. Adversarial, poisoning, fault-injection and unsafe-command tests must not be routed through a live production control or actuation path unless a separately approved safety case authorises it.
Notably Absent	Proof that a nominal human-in-the-loop has time, competence, information and authority to intervene; assurance that an override remains usable during alarm floods or degraded communications.

D6-CTL-04 — AI Incident Response Readiness

Field	Critical-infrastructure interpretation
Authoritative requirement	AI-IR runbook + tabletop exercises + containment automation.
Business objective	Ensure accountable humans retain practical authority to review, reject, suspend and recover AI-enabled operations under normal and emergency conditions.
OT/ICS-adapted equivalent	Provide physically and logically independent override routes, shift-specific competence, alarm-flood procedures and restoration authority.
Criticality scaling (CI-1 to CI-5)	CI-1: Define decision owner, review point and override route. CI-2: Train operators and test that review occurs in normal operations. CI-3: Measure automation bias, alarm suppression and escalation performance under realistic workload. CI-4: Require explicit approval for high-consequence actions, independent override channels and shift-aware staffing. CI-5: Use hard authority limits, tested emergency stop/takeover, recurrent simulator exercises and safety-function independence.
Evidence requirement	Authority matrix, approval/override logs, simulator exercise results, staffing/competence records and restoration decision.
Testing boundary	document/configuration inspection and representative offline test. Adversarial, poisoning, fault-injection and unsafe-command tests must not be routed through a live production control or actuation path unless a separately approved safety case authorises it.
Notably Absent	Proof that a nominal human-in-the-loop has time, competence, information and authority to intervene; assurance that an override remains usable during alarm floods or degraded communications.

D6-CTL-05 — Model Deprecation & Decommissioning

Field	Critical-infrastructure interpretation
Authoritative requirement	Access revocation + decommission audit + scheduled lifecycle.
Business objective	Ensure accountable humans retain practical authority to review, reject, suspend and recover AI-enabled operations under normal and emergency conditions.
OT/ICS-adapted equivalent	Provide physically and logically independent override routes, shift-specific competence, alarm-flood procedures and restoration authority.

Field	Critical-infrastructure interpretation
Criticality scaling (CI-1 to CI-5)	CI-1: Define decision owner, review point and override route. CI-2: Train operators and test that review occurs in normal operations. CI-3: Measure automation bias, alarm suppression and escalation performance under realistic workload. CI-4: Require explicit approval for high-consequence actions, independent override channels and shift-aware staffing. CI-5: Use hard authority limits, tested emergency stop/takeover, recurrent simulator exercises and safety-function independence.
Evidence requirement	Authority matrix, approval/override logs, simulator exercise results, staffing/competence records and restoration decision.
Testing boundary	document/configuration inspection and representative offline test. Adversarial, poisoning, fault-injection and unsafe-command tests must not be routed through a live production control or actuation path unless a separately approved safety case authorises it.
Notably Absent	Proof that a nominal human-in-the-loop has time, competence, information and authority to intervene; assurance that an override remains usable during alarm floods or degraded communications.

D6-CTL-06 — Third-Party AI Vendor Governance

Field	Critical-infrastructure interpretation
Authoritative requirement	Contractual security requirements + annual assessment + audit rights.
Business objective	Ensure accountable humans retain practical authority to review, reject, suspend and recover AI-enabled operations under normal and emergency conditions.
OT/ICS-adapted equivalent	Provide physically and logically independent override routes, shift-specific competence, alarm-flood procedures and restoration authority.
Criticality scaling (CI-1 to CI-5)	CI-1: Define decision owner, review point and override route. CI-2: Train operators and test that review occurs in normal operations. CI-3: Measure automation bias, alarm suppression and escalation performance under realistic workload. CI-4: Require explicit approval for high-consequence actions, independent override channels and shift-aware staffing. CI-5: Use hard authority limits, tested emergency stop/takeover, recurrent simulator exercises and safety-function independence.
Evidence requirement	Authority matrix, approval/override logs, simulator exercise results, staffing/competence records and restoration decision.
Testing boundary	document/configuration inspection and representative offline test. Adversarial, poisoning, fault-injection and unsafe-command tests must not be routed through a live production control or actuation path unless a separately approved safety case authorises it.
Notably Absent	Proof that a nominal human-in-the-loop has time, competence, information and authority to intervene; assurance that an override remains usable during alarm floods or degraded communications.

D6-CTL-07 — AI Resilience & Business Continuity

Field	Critical-infrastructure interpretation
Authoritative requirement	Failover systems + degraded mode + RTO/RPO definition.
Business objective	Ensure accountable humans retain practical authority to review, reject, suspend and recover AI-enabled operations under normal and emergency conditions.

Field	Critical-infrastructure interpretation
OT/ICS-adapted equivalent	Provide physically and logically independent override routes, shift-specific competence, alarm-flood procedures and restoration authority.
Criticality scaling (CI-1 to CI-5)	CI-1: Define decision owner, review point and override route. CI-2: Train operators and test that review occurs in normal operations. CI-3: Measure automation bias, alarm suppression and escalation performance under realistic workload. CI-4: Require explicit approval for high-consequence actions, independent override channels and shift-aware staffing. CI-5: Use hard authority limits, tested emergency stop/takeover, recurrent simulator exercises and safety-function independence.
Evidence requirement	Authority matrix, approval/override logs, simulator exercise results, staffing/competence records and restoration decision.
Testing boundary	document/configuration inspection and representative offline test. Adversarial, poisoning, fault-injection and unsafe-command tests must not be routed through a live production control or actuation path unless a separately approved safety case authorises it.
Notably Absent	Proof that a nominal human-in-the-loop has time, competence, information and authority to intervene; assurance that an override remains usable during alarm floods or degraded communications.

D7-CTL-H01 — AI-Generated Phishing Simulation

Field	Critical-infrastructure interpretation
Authoritative requirement	Simulation campaigns + click tracking + remedial training.
Business objective	Reduce the chance that AI-enabled social engineering — phishing, voice or video deepfakes and impersonation — leads operators, engineers or contractors to take unsafe or unauthorised actions on critical systems.
OT/ICS-adapted equivalent	Extend awareness, verification and incident procedures to control-room staff, field engineers, shift handovers and vendor remote-support channels.
Criticality scaling (CI-1 to CI-5)	CI-1: Include AI-enabled social engineering in awareness content for operational staff. CI-2: Add role-specific simulations for control-room, engineering and vendor-support staff. CI-3: Require out-of-band verification for remote-access, setpoint-change and emergency instructions. CI-4: Run joint exercises with incident response and OT engineering, including impersonated-instruction scenarios. CI-5: Build verification into emergency operating procedures, with independent review of exercise results.
Evidence requirement	Simulation and exercise records, verification-procedure logs, role-based training completion, incident and near-miss records, and corrective actions.
Testing boundary	document/configuration inspection and representative offline test. Adversarial, poisoning, fault-injection and unsafe-command tests must not be routed through a live production control or actuation path unless a separately approved safety case authorises it.
Notably Absent	Assurance that trained staff will detect every deepfake or impersonation; a substitute for technical access controls on remote and emergency channels.

D7-CTL-H02 — Deepfake Detection Training

Field	Critical-infrastructure interpretation
Authoritative requirement	Training modules + quiz + simulated attacks.

Field	Critical-infrastructure interpretation
Business objective	Reduce the chance that AI-enabled social engineering — phishing, voice or video deepfakes and impersonation — leads operators, engineers or contractors to take unsafe or unauthorised actions on critical systems.
OT/ICS-adapted equivalent	Extend awareness, verification and incident procedures to control-room staff, field engineers, shift handovers and vendor remote-support channels.
Criticality scaling (CI-1 to CI-5)	CI-1: Include AI-enabled social engineering in awareness content for operational staff. CI-2: Add role-specific simulations for control-room, engineering and vendor-support staff. CI-3: Require out-of-band verification for remote-access, setpoint-change and emergency instructions. CI-4: Run joint exercises with incident response and OT engineering, including impersonated-instruction scenarios. CI-5: Build verification into emergency operating procedures, with independent review of exercise results.
Evidence requirement	Simulation and exercise records, verification-procedure logs, role-based training completion, incident and near-miss records, and corrective actions.
Testing boundary	document/configuration inspection and representative offline test. Adversarial, poisoning, fault-injection and unsafe-command tests must not be routed through a live production control or actuation path unless a separately approved safety case authorises it.
Notably Absent	Assurance that trained staff will detect every deepfake or impersonation; a substitute for technical access controls on remote and emergency channels.

D7-CTL-H03 — Out-Of-Band Authentication

Field	Critical-infrastructure interpretation
Authoritative requirement	Independent channel verification + policy enforcement.
Business objective	Reduce the chance that AI-enabled social engineering — phishing, voice or video deepfakes and impersonation — leads operators, engineers or contractors to take unsafe or unauthorised actions on critical systems.
OT/ICS-adapted equivalent	Extend awareness, verification and incident procedures to control-room staff, field engineers, shift handovers and vendor remote-support channels.
Criticality scaling (CI-1 to CI-5)	CI-1: Include AI-enabled social engineering in awareness content for operational staff. CI-2: Add role-specific simulations for control-room, engineering and vendor-support staff. CI-3: Require out-of-band verification for remote-access, setpoint-change and emergency instructions. CI-4: Run joint exercises with incident response and OT engineering, including impersonated-instruction scenarios. CI-5: Build verification into emergency operating procedures, with independent review of exercise results.
Evidence requirement	Simulation and exercise records, verification-procedure logs, role-based training completion, incident and near-miss records, and corrective actions.
Testing boundary	document/configuration inspection and representative offline test. Adversarial, poisoning, fault-injection and unsafe-command tests must not be routed through a live production control or actuation path unless a separately approved safety case authorises it.
Notably Absent	Assurance that trained staff will detect every deepfake or impersonation; a substitute for technical access controls on remote and emergency channels.

D7-CTL-H04 — AI Social Engineering IR

Field	Critical-infrastructure interpretation
Authoritative requirement	Tabletop exercises + IR plan + verification triggers.
Business objective	Reduce the chance that AI-enabled social engineering — phishing, voice or video deepfakes and impersonation — leads operators, engineers or contractors to take unsafe or unauthorised actions on critical systems.
OT/ICS-adapted equivalent	Extend awareness, verification and incident procedures to control-room staff, field engineers, shift handovers and vendor remote-support channels.
Criticality scaling (CI-1 to CI-5)	CI-1: Include AI-enabled social engineering in awareness content for operational staff. CI-2: Add role-specific simulations for control-room, engineering and vendor-support staff. CI-3: Require out-of-band verification for remote-access, setpoint-change and emergency instructions. CI-4: Run joint exercises with incident response and OT engineering, including impersonated-instruction scenarios. CI-5: Build verification into emergency operating procedures, with independent review of exercise results.
Evidence requirement	Simulation and exercise records, verification-procedure logs, role-based training completion, incident and near-miss records, and corrective actions.
Testing boundary	document/configuration inspection and representative offline test. Adversarial, poisoning, fault-injection and unsafe-command tests must not be routed through a live production control or actuation path unless a separately approved safety case authorises it.
Notably Absent	Assurance that trained staff will detect every deepfake or impersonation; a substitute for technical access controls on remote and emergency channels.

D7-CTL-H05 — AI-Enhanced External Attack Defense

Field	Critical-infrastructure interpretation
Authoritative requirement	AI-generated phishing detection + SOC tuning + response automation.
Business objective	Reduce the chance that AI-enabled social engineering — phishing, voice or video deepfakes and impersonation — leads operators, engineers or contractors to take unsafe or unauthorised actions on critical systems.
OT/ICS-adapted equivalent	Extend awareness, verification and incident procedures to control-room staff, field engineers, shift handovers and vendor remote-support channels.
Criticality scaling (CI-1 to CI-5)	CI-1: Include AI-enabled social engineering in awareness content for operational staff. CI-2: Add role-specific simulations for control-room, engineering and vendor-support staff. CI-3: Require out-of-band verification for remote-access, setpoint-change and emergency instructions. CI-4: Run joint exercises with incident response and OT engineering, including impersonated-instruction scenarios. CI-5: Build verification into emergency operating procedures, with independent review of exercise results.
Evidence requirement	Simulation and exercise records, verification-procedure logs, role-based training completion, incident and near-miss records, and corrective actions.
Testing boundary	document/configuration inspection and representative offline test. Adversarial, poisoning, fault-injection and unsafe-command tests must not be routed through a live production control or actuation path unless a separately approved safety case authorises it.
Notably Absent	Assurance that trained staff will detect every deepfake or impersonation; a substitute for technical access controls on remote and emergency channels.

D8-CTL-01 — EU AI Act Risk Tier Mapping

Field	Critical-infrastructure interpretation
Authoritative requirement	Risk classification framework + conformity assessment.
Business objective	Maintain verifiable obligation, documentation and assurance records without overstating regulatory equivalence or certification status.
OT/ICS-adapted equivalent	Use local evidence repositories and signed transfer manifests; verify applicability and editions before formal reliance.
Criticality scaling (CI-1 to CI-5)	CI-1: Maintain an obligation register and record the basis for applicability decisions. CI-2: Map controls to verified jurisdictional and contractual obligations; retain review evidence. CI-3: Use qualified legal/regulatory review and trace evidence to the deployed system/version. CI-4: Coordinate notification, recordkeeping and assurance across sites and critical suppliers. CI-5: Require executive acceptance of unresolved conflicts and independent verification; no mapping is treated as automatic compliance.
Evidence requirement	Obligation register, applicability rationale, verified source/version, evidence crosswalk and qualified review record.
Testing boundary	document/configuration inspection and representative offline test. Adversarial, poisoning, fault-injection and unsafe-command tests must not be routed through a live production control or actuation path unless a separately approved safety case authorises it.
Notably Absent	A claim of legal equivalence, certification reciprocity or automatic compliance with any cited framework; jurisdiction-specific legal advice.

D8-CTL-02 — ISO 42001 Gap Analysis

Field	Critical-infrastructure interpretation
Authoritative requirement	Gap analysis methodology + remediation tracking.
Business objective	Maintain verifiable obligation, documentation and assurance records without overstating regulatory equivalence or certification status.
OT/ICS-adapted equivalent	Use local evidence repositories and signed transfer manifests; verify applicability and editions before formal reliance.
Criticality scaling (CI-1 to CI-5)	CI-1: Maintain an obligation register and record the basis for applicability decisions. CI-2: Map controls to verified jurisdictional and contractual obligations; retain review evidence. CI-3: Use qualified legal/regulatory review and trace evidence to the deployed system/version. CI-4: Coordinate notification, recordkeeping and assurance across sites and critical suppliers. CI-5: Require executive acceptance of unresolved conflicts and independent verification; no mapping is treated as automatic compliance.
Evidence requirement	Obligation register, applicability rationale, verified source/version, evidence crosswalk and qualified review record.
Testing boundary	document/configuration inspection and representative offline test. Adversarial, poisoning, fault-injection and unsafe-command tests must not be routed through a live production control or actuation path unless a separately approved safety case authorises it.
Notably Absent	A claim of legal equivalence, certification reciprocity or automatic compliance with any cited framework; jurisdiction-specific legal advice.

D8-CTL-03 — GPAI Technical Documentation Verification

Field	Critical-infrastructure interpretation
Authoritative requirement	Technical documentation + training data summary + copyright attestation.
Business objective	Maintain verifiable obligation, documentation and assurance records without overstating regulatory equivalence or certification status.
OT/ICS-adapted equivalent	Use local evidence repositories and signed transfer manifests; verify applicability and editions before formal reliance.
Criticality scaling (CI-1 to CI-5)	CI-1: Maintain an obligation register and record the basis for applicability decisions. CI-2: Map controls to verified jurisdictional and contractual obligations; retain review evidence. CI-3: Use qualified legal/regulatory review and trace evidence to the deployed system/version. CI-4: Coordinate notification, recordkeeping and assurance across sites and critical suppliers. CI-5: Require executive acceptance of unresolved conflicts and independent verification; no mapping is treated as automatic compliance.
Evidence requirement	Obligation register, applicability rationale, verified source/version, evidence crosswalk and qualified review record.
Testing boundary	document/configuration inspection and representative offline test. Adversarial, poisoning, fault-injection and unsafe-command tests must not be routed through a live production control or actuation path unless a separately approved safety case authorises it.
Notably Absent	A claim of legal equivalence, certification reciprocity or automatic compliance with any cited framework; jurisdiction-specific legal advice.

D8-CTL-04 — DORA ICT Incident Reporting (Financial Sector)

Field	Critical-infrastructure interpretation
Authoritative requirement	Incident classification + notification workflow + SLA monitoring.
Business objective	Maintain verifiable obligation, documentation and assurance records without overstating regulatory equivalence or certification status (Dora Ict Incident Reporting (Financial Sector)).
OT/ICS-adapted equivalent	Use local evidence repositories and signed transfer manifests; verify applicability and editions before formal reliance.
Criticality scaling (CI-1 to CI-5)	CI-1: Maintain an obligation register and record the basis for applicability decisions. CI-2: Map controls to verified jurisdictional and contractual obligations; retain review evidence. CI-3: Use qualified legal/regulatory review and trace evidence to the deployed system/version. CI-4: Coordinate notification, recordkeeping and assurance across sites and critical suppliers. CI-5: Require executive acceptance of unresolved conflicts and independent verification; no mapping is treated as automatic compliance.
Evidence requirement	Obligation register, applicability rationale, verified source/version, evidence crosswalk and qualified review record.
Testing boundary	document/configuration inspection and representative offline test. Adversarial, poisoning, fault-injection and unsafe-command tests must not be routed through a live production control or actuation path unless a separately approved safety case authorises it.
Notably Absent	A claim of legal equivalence, certification reciprocity or automatic compliance with any cited framework; jurisdiction-specific legal advice.

D8-CTL-05 — NIST SP 800-218A Compliance Check

Field	Critical-infrastructure interpretation
Authoritative requirement	Secure development practices + attestation.
Business objective	Maintain verifiable obligation, documentation and assurance records without overstating regulatory equivalence or certification status.
OT/ICS-adapted equivalent	Use local evidence repositories and signed transfer manifests; verify applicability and editions before formal reliance.
Criticality scaling (CI-1 to CI-5)	CI-1: Maintain an obligation register and record the basis for applicability decisions. CI-2: Map controls to verified jurisdictional and contractual obligations; retain review evidence. CI-3: Use qualified legal/regulatory review and trace evidence to the deployed system/version. CI-4: Coordinate notification, recordkeeping and assurance across sites and critical suppliers. CI-5: Require executive acceptance of unresolved conflicts and independent verification; no mapping is treated as automatic compliance.
Evidence requirement	Obligation register, applicability rationale, verified source/version, evidence crosswalk and qualified review record.
Testing boundary	document/configuration inspection and representative offline test. Adversarial, poisoning, fault-injection and unsafe-command tests must not be routed through a live production control or actuation path unless a separately approved safety case authorises it.
Notably Absent	A claim of legal equivalence, certification reciprocity or automatic compliance with any cited framework; jurisdiction-specific legal advice.

D9-CTL-01 — Physical Harm Boundary Enforcement

Field	Critical-infrastructure interpretation
Authoritative requirement	Independent safety monitor (hardware or DO-178C Level A / IEC 61508 SIL 3 certified software) running in parallel with AI inference. Safety monitor enforces: maximum force/velocity/temperature/current limits; geofencing for autonomous systems; exclusion zones; rate-of-change limits for safety-critical parameters. AI output gated through safety monitor — monitor vetoes any out-of-boundary command without AI system awareness.
Business objective	Prevent AI-enabled sensing, decision or actuation from defeating physical safety limits, protective functions or safe-state behaviour.
OT/ICS-adapted equivalent	Integrate with the existing functional-safety lifecycle, LOPA/SIF records and deterministic timing budget; security monitoring must not delay protective action.
Criticality scaling (CI-1 to CI-5)	CI-1: Document physical interfaces, safe limits and prohibited actuation; keep AI advisory unless separately approved. CI-2: Validate safe-state behavior and operator override in an isolated representative environment. CI-3: Use redundant sensing, command validation, independent monitoring and tested manual takeover. CI-4: Integrate with LOPA/SIF or equivalent safety lifecycle; verify latency, fail-safe behavior and separation from basic process control. CI-5: Use independently certified or otherwise justified safety mechanisms, hardwired interlocks where required, HIL testing, emergency-stop validation and independent functional-safety assurance.
Evidence requirement	HIL/digital-twin report, safety-boundary specification, SIF/LOPA linkage, command/override logs, latency assessment and safe-state test.
Testing boundary	high-fidelity digital twin, hardware-in-the-loop rig, or strictly isolated staging network. Adversarial, poisoning, fault-injection and unsafe-command tests must not be routed through a live production control or actuation path unless a separately approved safety case authorises it.

Field	Critical-infrastructure interpretation
Notably Absent	Replacement of HAZOP, FMEA, LOPA, SIF/SIL engineering, hardware interlocks or statutory process-safety duties; a guarantee that AI security monitors preserve deterministic timing without measured latency evidence.

D9-CTL-02 — Safe State And Graceful Degradation

Field	Critical-infrastructure interpretation
Authoritative requirement	For each AI-controlled system, document: safe state definition (autonomous vehicle: controlled stop; surgical robot: tool withdrawal; industrial arm: immediate stop and hold); transition time to safe state (must be within stopping distance/reaction time for physical context); trigger conditions for safe state entry; recovery procedure. Implement degraded mode ladder: Full AI control → AI-assisted human control → Manual-only → Safe state.
Business objective	Prevent AI-enabled sensing, decision or actuation from defeating physical safety limits, protective functions or safe-state behaviour.
OT/ICS-adapted equivalent	Integrate with the existing functional-safety lifecycle, LOPA/SIF records and deterministic timing budget; security monitoring must not delay protective action.
Criticality scaling (CI-1 to CI-5)	CI-1: Document physical interfaces, safe limits and prohibited actuation; keep AI advisory unless separately approved. CI-2: Validate safe-state behavior and operator override in an isolated representative environment. CI-3: Use redundant sensing, command validation, independent monitoring and tested manual takeover. CI-4: Integrate with LOPA/SIF or equivalent safety lifecycle; verify latency, fail-safe behavior and separation from basic process control. CI-5: Use independently certified or otherwise justified safety mechanisms, hardwired interlocks where required, HIL testing, emergency-stop validation and independent functional-safety assurance.
Evidence requirement	HIL/digital-twin report, safety-boundary specification, SIF/LOPA linkage, command/override logs, latency assessment and safe-state test.
Testing boundary	high-fidelity digital twin, hardware-in-the-loop rig, or strictly isolated staging network. Adversarial, poisoning, fault-injection and unsafe-command tests must not be routed through a live production control or actuation path unless a separately approved safety case authorises it.
Notably Absent	Replacement of HAZOP, FMEA, LOPA, SIF/SIL engineering, hardware interlocks or statutory process-safety duties; a guarantee that AI security monitors preserve deterministic timing without measured latency evidence.

D9-CTL-03 — Human Override And Emergency Stop

Field	Critical-infrastructure interpretation
Authoritative requirement	Hardware emergency stop: physical E-stop accessible without any software mediation. AI system must not be able to disable, delay, or circumvent E-stop. Software override: human operator interface that immediately transfers control to safe state. Override must be possible when: AI communication is disrupted; AI system is under adversarial attack; AI model is producing anomalous outputs. Override authority must be unconditional — no AI reasoning, confidence scoring, or approval process may delay or prevent override activation.
Business objective	Prevent AI-enabled sensing, decision or actuation from defeating physical safety limits, protective functions or safe-state behaviour.
OT/ICS-adapted equivalent	Integrate with the existing functional-safety lifecycle, LOPA/SIF records and deterministic timing budget; security monitoring must not delay protective action.

Field	Critical-infrastructure interpretation
Criticality scaling (CI-1 to CI-5)	CI-1: Document physical interfaces, safe limits and prohibited actuation; keep AI advisory unless separately approved. CI-2: Validate safe-state behavior and operator override in an isolated representative environment. CI-3: Use redundant sensing, command validation, independent monitoring and tested manual takeover. CI-4: Integrate with LOPA/SIF or equivalent safety lifecycle; verify latency, fail-safe behavior and separation from basic process control. CI-5: Use independently certified or otherwise justified safety mechanisms, hardwired interlocks where required, HIL testing, emergency-stop validation and independent functional-safety assurance.
Evidence requirement	HIL/digital-twin report, safety-boundary specification, SIF/LOPA linkage, command/override logs, latency assessment and safe-state test.
Testing boundary	high-fidelity digital twin, hardware-in-the-loop rig, or strictly isolated staging network. Adversarial, poisoning, fault-injection and unsafe-command tests must not be routed through a live production control or actuation path unless a separately approved safety case authorises it.
Notably Absent	Replacement of HAZOP, FMEA, LOPA, SIF/SIL engineering, hardware interlocks or statutory process-safety duties; a guarantee that AI security monitors preserve deterministic timing without measured latency evidence.

D9-CTL-04 — Cyber-Physical Attack Detection

Field	Critical-infrastructure interpretation
Authoritative requirement	Three-layer anomaly detection: (1) Sensor layer — statistical validation of sensor readings against physical models; flag readings deviating $>3\sigma$ from model prediction; cross-validate against redundant sensor channels. (2) Actuator layer — monitor command streams for sequences inconsistent with operating context; flag commands outside physically feasible envelope. (3) AI inference layer — apply GAISST TM D2-CTL-01 (Prompt Injection Detection) equivalent for physical AI inputs; monitor input feature distributions for adversarial perturbation signatures. All detections trigger immediate safe state entry (D9-CTL-02) and incident record with <code>root_cause_category = Adversarial_Attack</code> , <code>root_cause_specific_type = Cyber_Physical_Attack</code> .
Business objective	Prevent AI-enabled sensing, decision or actuation from defeating physical safety limits, protective functions or safe-state behaviour.
OT/ICS-adapted equivalent	Integrate with the existing functional-safety lifecycle, LOPA/SIF records and deterministic timing budget; security monitoring must not delay protective action.
Criticality scaling (CI-1 to CI-5)	CI-1: Document physical interfaces, safe limits and prohibited actuation; keep AI advisory unless separately approved. CI-2: Validate safe-state behavior and operator override in an isolated representative environment. CI-3: Use redundant sensing, command validation, independent monitoring and tested manual takeover. CI-4: Integrate with LOPA/SIF or equivalent safety lifecycle; verify latency, fail-safe behavior and separation from basic process control. CI-5: Use independently certified or otherwise justified safety mechanisms, hardwired interlocks where required, HIL testing, emergency-stop validation and independent functional-safety assurance.
Evidence requirement	HIL/digital-twin report, safety-boundary specification, SIF/LOPA linkage, command/override logs, latency assessment and safe-state test.
Testing boundary	high-fidelity digital twin, hardware-in-the-loop rig, or strictly isolated staging network. Adversarial, poisoning, fault-injection and unsafe-command tests must not be routed through a live production control or actuation path unless a separately approved safety case authorises it.

Field	Critical-infrastructure interpretation
Notably Absent	Replacement of HAZOP, FMEA, LOPA, SIF/SIL engineering, hardware interlocks or statutory process-safety duties; a guarantee that AI security monitors preserve deterministic timing without measured latency evidence.

D9-CTL-05 — Physical Environment Integrity Monitoring

Field	Critical-infrastructure interpretation
Authoritative requirement	Sensor integrity monitoring covering: (1) Hardware health — sensor self-test results, calibration drift indicators, environmental exposure limits. Alert when sensor confidence falls below threshold. (2) Data plausibility — real-time statistical validation against physical laws, historical baselines, and redundant sensor cross-validation. (3) Degraded sensor handling — explicit policy for each sensor failure mode: degrade gracefully (reduce AI authority, increase human oversight) or enter safe state. (4) Calibration management — automated alert when calibration certificates expire; block AI system from operational use with expired sensor calibration.
Business objective	Prevent AI-enabled sensing, decision or actuation from defeating physical safety limits, protective functions or safe-state behaviour.
OT/ICS-adapted equivalent	Integrate with the existing functional-safety lifecycle, LOPA/SIF records and deterministic timing budget; security monitoring must not delay protective action.
Criticality scaling (CI-1 to CI-5)	CI-1: Document physical interfaces, safe limits and prohibited actuation; keep AI advisory unless separately approved. CI-2: Validate safe-state behavior and operator override in an isolated representative environment. CI-3: Use redundant sensing, command validation, independent monitoring and tested manual takeover. CI-4: Integrate with LOPA/SIF or equivalent safety lifecycle; verify latency, fail-safe behavior and separation from basic process control. CI-5: Use independently certified or otherwise justified safety mechanisms, hardwired interlocks where required, HIL testing, emergency-stop validation and independent functional-safety assurance.
Evidence requirement	HIL/digital-twin report, safety-boundary specification, SIF/LOPA linkage, command/override logs, latency assessment and safe-state test.
Testing boundary	high-fidelity digital twin, hardware-in-the-loop rig, or strictly isolated staging network. Adversarial, poisoning, fault-injection and unsafe-command tests must not be routed through a live production control or actuation path unless a separately approved safety case authorises it.
Notably Absent	Replacement of HAZOP, FMEA, LOPA, SIF/SIL engineering, hardware interlocks or statutory process-safety duties; a guarantee that AI security monitors preserve deterministic timing without measured latency evidence.

D9-CTL-06 — Actuator Command Verification

Field	Critical-infrastructure interpretation
Authoritative requirement	Pre-execution verification gate on every actuator command: (1) Physical bounds check — command value within safe operating envelope for current system state. (2) Sequence plausibility check — command consistent with prior sequence; flag implausible state transitions for human review. (3) Rate-of-change check — rate of change does not exceed safe limits (acceleration rate, force application rate, temperature change rate). (4) Dual-approval for irreversible actions — actuator commands causing irreversible physical changes (cutting, welding, demolition, high-energy discharge) require hardware interlock confirmation. Verification gate implemented in IEC 61508 SIL 3 certified software or hardware logic independent of AI model.
Business objective	Prevent AI-enabled sensing, decision or actuation from defeating physical safety limits, protective functions or safe-state behaviour.
OT/ICS-adapted equivalent	Integrate with the existing functional-safety lifecycle, LOPA/SIF records and deterministic timing budget; security monitoring must not delay protective action.
Criticality scaling (CI-1 to CI-5)	CI-1: Document physical interfaces, safe limits and prohibited actuation; keep AI advisory unless separately approved. CI-2: Validate safe-state behavior and operator override in an isolated representative environment. CI-3: Use redundant sensing, command validation, independent monitoring and tested manual takeover. CI-4: Integrate with LOPA/SIF or equivalent safety lifecycle; verify latency, fail-safe behavior and separation from basic process control. CI-5: Use independently certified or otherwise justified safety mechanisms, hardwired interlocks where required, HIL testing, emergency-stop validation and independent functional-safety assurance.
Evidence requirement	HIL/digital-twin report, safety-boundary specification, SIF/LOPA linkage, command/override logs, latency assessment and safe-state test.
Testing boundary	high-fidelity digital twin, hardware-in-the-loop rig, or strictly isolated staging network. Adversarial, poisoning, fault-injection and unsafe-command tests must not be routed through a live production control or actuation path unless a separately approved safety case authorises it.
Notably Absent	Replacement of HAZOP, FMEA, LOPA, SIF/SIL engineering, hardware interlocks or statutory process-safety duties; a guarantee that AI security monitors preserve deterministic timing without measured latency evidence.

D9-CTL-07 — Physical Incident Evidence Preservation

Field	Critical-infrastructure interpretation
Authoritative requirement	(1) Continuous ring-buffer recording — minimum 60-second rolling buffer of: all sensor inputs (raw and processed); all AI model inputs and outputs; all actuator commands; all safety monitor decisions; all human override activations; system health telemetry. Safety-critical systems retain 300 seconds minimum. (2) Incident freeze — on any safety-relevant event, automatically freeze buffer and begin extended logging. Frozen buffer write-protected. (3) Cryptographic integrity — all records SHA-256 hashed and ECDSA signed at point of creation. For Optimized tier: CRYSTALS-Dilithium signing (post-quantum). (4) Regulatory retention — ICAO Annex 13: 5 years minimum; EU AI Act Art. 19: 10 years; DORA Art. 12: 5 years. (5) UAIF TM integration — automatically populate UAIF TM incident record from evidence package.
Business objective	Prevent AI-enabled sensing, decision or actuation from defeating physical safety limits, protective functions or safe-state behaviour.

Field	Critical-infrastructure interpretation
OT/ICS-adapted equivalent	Integrate with the existing functional-safety lifecycle, LOPA/SIF records and deterministic timing budget; security monitoring must not delay protective action.
Criticality scaling (CI-1 to CI-5)	CI-1: Document physical interfaces, safe limits and prohibited actuation; keep AI advisory unless separately approved. CI-2: Validate safe-state behavior and operator override in an isolated representative environment. CI-3: Use redundant sensing, command validation, independent monitoring and tested manual takeover. CI-4: Integrate with LOPA/SIF or equivalent safety lifecycle; verify latency, fail-safe behavior and separation from basic process control. CI-5: Use independently certified or otherwise justified safety mechanisms, hardwired interlocks where required, HIL testing, emergency-stop validation and independent functional-safety assurance.
Evidence requirement	HIL/digital-twin report, safety-boundary specification, SIF/LOPA linkage, command/override logs, latency assessment and safe-state test.
Testing boundary	high-fidelity digital twin, hardware-in-the-loop rig, or strictly isolated staging network. Adversarial, poisoning, fault-injection and unsafe-command tests must not be routed through a live production control or actuation path unless a separately approved safety case authorises it.
Notably Absent	Replacement of HAZOP, FMEA, LOPA, SIF/SIL engineering, hardware interlocks or statutory process-safety duties; a guarantee that AI security monitors preserve deterministic timing without measured latency evidence.

10. Threat Register Summary

ID	Threat	Primary controls	Owner
CI-THR-001	Manipulated sensor data	D1-CTL-01; D1-CTL-03; D9-CTL-05	OT engineering / CISO
CI-THR-002	Training-data poisoning	D1-CTL-01; D1-CTL-04	Data owner / CISO
CI-THR-003	Compromised model update	D1-CTL-06; D4-CTL-01; D4-CTL-02	Model owner / OT engineering
CI-THR-004	Prompt injection affecting operator assistant	D2-CTL-01; D2-CTL-02; D6-CTL-01	CISO / operations
CI-THR-005	Indirect prompt injection	D2-CTL-01; D2-CTL-02; D5-CTL-01	CISO / application owner
CI-THR-006	Unauthorized tool execution	D2-CTL-05; D3-CTL-01; D6-CTL-01	CISO / system owner
CI-THR-007	Excessive agent permissions	D3-CTL-01; D3-CTL-05; D6-CTL-01	CISO / system owner
CI-THR-008	Cloud/API outage	D4-CTL-05; D6-CTL-07; D9-CTL-02	Resilience lead / operations
CI-THR-009	Model drift or concept drift	D1-CTL-03; D6-CTL-01	Model owner / operations
CI-THR-010	Automation bias and alarm suppression	D6-CTL-01; D9-CTL-03	Operations / human factors
CI-THR-011	Common-mode model failure across sites	D1-CTL-03; D4-CTL-01; D6-CTL-07	Enterprise architecture / resilience
CI-THR-012	Unsafe AI-generated code/configuration	D2-CTL-05; D4-CTL-07; D9-CTL-06	Engineering authority / CISO

ID	Threat	Primary controls	Owner
CI-THR-013	Loss of audit logging	D6-CTL-02; D9-CTL-07	CISO / evidence custodian
CI-THR-014	Opaque supplier AI component	D1-CTL-01; D4-CTL-01; D8-CTL-03	Procurement / system owner
CI-THR-015	Unauthorized shadow AI	D4-CTL-06; D3-CTL-01	CISO / site manager
CI-THR-016	Emergency-mode failure	D6-CTL-01; D9-CTL-02; D9-CTL-03	Operations / safety lead
CI-THR-017	Sensor spoofing	D9-CTL-04; D9-CTL-05; D1-CTL-01	OT engineering / safety lead
CI-THR-018	Actuator command injection	D9-CTL-01; D9-CTL-06; D3-CTL-01	OT engineering / CISO
CI-THR-019	Safety monitor bypass	D9-CTL-01; D9-CTL-03; D9-CTL-07	Functional safety lead
CI-THR-020	Adversarial physical input	D2-CTL-04; D9-CTL-04; D9-CTL-05	Engineering / security testing
CI-THR-021	Sensor calibration drift	D1-CTL-03; D9-CTL-05	Maintenance / engineering
CI-THR-022	Compromised OT hardware or firmware	D4-CTL-01; D4-CTL-02; D9-CTL-06	Supply chain / CISO
CI-THR-023	AI-enabled impersonation of operators, engineers or vendors	D7-CTL-H02; D7-CTL-H03; D7-CTL-H04	Security awareness / OT operations

Public, verified telemetry on AI-specific OT incidents remains incomplete; no quantitative likelihood is assigned to any threat.

11. Standards Alignment Rules

Mappings in this package are conceptual alignment aids, not claims that a GAISST control satisfies a law, standard or certification requirement. Exact clauses, editions, asset scope and jurisdiction must be verified against controlled primary sources. External copyrighted standards are not reproduced.

12. Notably Absent

- No assumption of universal vendor cooperation.
- No substitution for HAZOP, FMEA, LOPA, SIF/SIL engineering or sector safety obligations.
- No guarantee that AI guardrails preserve deterministic timing without measured evidence.
- No authorization for uncontrolled production adversarial testing.
- No quantitative likelihood or outage estimate without organization-specific evidence.
- No automatic compliance or certification equivalence.

13. Publication Readiness

Substantive refinement is complete and all 59 controls are present across JSON, CSV and workbooks. Public release remains gated by named document approvals, independent OT/ICS and functional-safety review, subsector review, jurisdiction-specific legal/regulatory verification, controlled-source mapping verification, accessibility review and final publication QA.

Appendix A - Control Reconciliation

Check	Result
Expected controls	59
JSON controls	59
Unique IDs	59
D9 controls	7
Missing/duplicates	None detected
Normative IDs/titles	Preserved from v1.0 source package

GAISSF Ecosystem developed by ODA3 Institute, licensed under GEL v1.0. | © 2026 ODA3 Pvt Ltd | oda3.org