

ODA3 INSTITUTE

SECTOR IMPLEMENTATION GUIDE

GAISSF™ v1.0 Aerospace Sector Implementation Guide

Implementation guidance for the Global AI Safety and Security Framework

GAISSF-SEC-050 | Informative sector guidance | Draft for Publication
Version 1.0 | Public | External Review Open | Updated 6 October 2026

© 2026 ODA3 Pvt Ltd

Publication Status and Precedence

This public-draft edition is reconciled to the current GAISSF-NOR-001/NOR-004 v1.0 baseline. GAISSF-NOR-001 remains authoritative for normative control wording; this sector guide is informative and does not alter control membership or establish legal, regulatory, airworthiness, spaceworthiness or certification outcomes.

Where GAISSF-SEC-050 conflicts with a normative GAISSF publication, the normative publication prevails. External aerospace materials are references only; this edition makes no clause-level equivalence, means-of-compliance or regulatory-acceptance claim. GAISSF certification is not currently offered. External Review Open applies to the specialist reviews required before Final Publication.

Executive Summary

GAISSF-SEC-050 translates the 59 GAISSF™ v1.0 controls across nine domains into aerospace implementation guidance for aviation, space, launch, manufacturing, MRO, ground segments and enterprise AI. The Applicable-Control Baseline comprises all 52 controls in D1-D8 plus the seven D9 controls where Physical AI or cyber-physical actuation is in scope.

Master Notably Absent

- No legal advice or determination of jurisdictional applicability.
- No airworthiness, spaceworthiness, type-design, operational or mission approval.
- No replacement for applicable safety, development-assurance, hardware, software, security or mission-assurance processes.
- No direct mapping from GAISSF-SEC-050 criticality to DAL, SIL or other formal assurance level.
- No approved means of compliance or regulator acceptance.
- No assertion that military, defence, national-security, classified or export-controlled applications are covered without separate legal and assurance scoping.
- No mature public aerospace-AI actuarial dataset supporting generalised financial-loss estimates.
- No guarantee that controls eliminate residual risk, prevent incidents or predict future operational performance.
- No universal applicability of prompt-injection, agentic or generative-AI controls to embedded flight or spacecraft systems.
- No assertion that the conceptual architectures are approved designs.

1. Scope and Application

The guide applies according to actual intended function, architecture, authority and consequence. All 52 D1-D8 controls belong to the Applicable-Control Baseline. D9 is additional/conditional and becomes mandatory where Physical AI or cyber-physical actuation is in scope. Where D9 applies, use GAISSF D9 alongside PAI-SF™ v1.0 as applicable; CSX-PAISF-004 is mapping only and does not establish equivalence, evidence sufficiency or automatic conformance.

2. Aerospace AI System Classes

Class	Applicability principle
Embedded aviation ML	Certification-relevant or operationally consequential software/model component; prompts and GenAI controls apply only if present.
Aviation enterprise and ground AI	Maintenance, engineering, airline, airport and ATM support; data, GenAI, identity and workflow controls often dominate.
Onboard space autonomy	Delayed intervention, bounded action, resource and mission constraints dominate.
Space ground segment	Command separation, privileged access, telemetry integrity and provider dependencies dominate.
Physical AI	Sensor, actuator, environment, timing and independent safety constraints dominate.

3. Criticality and Non-equivalence

GAISSF-SEC-050 criticality is an implementation prioritisation aid. It shall not be converted directly into a DAL, SIL, software level, hardware level, failure-condition classification or mission class. Those determinations arise from the applicable safety and system-development process.

4. General Aerospace Implementation Principles

- Treat AI as an integrated system, not a model-only asset.
- Bind evidence to exact hardware, software, model, data, prompt, retrieval and threshold baselines.
- Separate recommendation, approval and command authority.
- Make degraded mode, communications loss, sensor failure and rollback explicit.
- Assess common-mode effects across fleet, constellation and shared providers.
- Record data and simulation limitations, especially rare-event scarcity and simulator-to-reality gaps.
- Coordinate security containment with safety, mission and operational authority.
- Do not use financial estimates without an explicit formula, assumptions, range and confidence.

D7 boundary and evidence taxonomy. D7 remains the canonical GAISSF domain Human & Societal Harms and retains its social-engineering, phishing, deepfake, out-of-band-authentication and external-attack control set. Aerospace safety, human factors and bias concerns may be addressed elsewhere but do not redefine D7. This guide does not create a competing T1-T4 evidence-tier model.

5. Control-by-Control Aerospace Interpretation

D1-CTL-01 - DATASET PROVENANCE & POISONING PREVENTION

Illustrative implementation starting point (non-prescriptive)	Hash verification + source allowlist + poisoning detection.
Validation status	ODA3 technical sector interpretation; organisation-specific engineering and independent specialist review remain required.
System class	data
Aviation applicability	For flight, maintenance and airport use, distinguish certified or controlled engineering data from operational, passenger, weather, navigation and third-party feeds. Record aircraft type, tail/fleet coverage, geography, season and sensor configuration.
Space applicability	For spacecraft and ground systems, record mission phase, orbit or trajectory regime, payload mode, ground-station source, time correlation and communications gaps. Preserve provenance across onboard filtering and delayed downlink.
Aerospace interpretation	Establish traceable lineage from acquisition through cleaning, labelling, simulation, transformation, training, evaluation and operational use. Treat rare-event scarcity and simulator-to-reality differences as explicit evidence limitations.
Minimum evidence	Control owner and applicability decision; system/configuration baseline; data-specific design evidence; dated operating or test record; open exceptions and residual-risk approval.
Enhanced evidence	Independent review; representative degraded/adversarial test; corroborating telemetry; supplier evidence; change-impact and common-mode analysis where material.
Assessment consideration	Select a material dataset and reproduce its lineage,

Illustrative implementation starting point (non-prescriptive)	Hash verification + source allowlist + poisoning detection.
	legal/contractual rights, transformations, integrity checks, coverage limits and linkage to the model release that consumed it.
Failure indicators	Missing source or licence; mixed controlled/uncontrolled data; unrecorded synthetic augmentation; absent rare-event coverage analysis; stale navigation or maintenance data; dataset hash does not match release evidence.
Notably Absent	No airworthiness, spaceworthiness, certification, legal-compliance or regulator-acceptance conclusion; no universal threshold or approved design is asserted.

D1-CTL-02 - MODEL EXTRACTION RESISTANCE

Illustrative implementation starting point (non-prescriptive)	Rate limiting + diversity detection + extraction monitoring.
Validation status	ODA3 technical sector interpretation; organisation-specific engineering and independent specialist review remain required.
System class	agentic
Aviation applicability	Relevant when AI can initiate, sequence or recommend actions affecting dispatch, maintenance, ground movement, UAS, AAM or operational systems. Human oversight must be feasible within the time available; “human in the loop” alone is not evidence of effective control.
Space applicability	Relevant to onboard autonomy, fault management, planning, rendezvous/proximity operations and ground automation. Intervention may be delayed or unavailable, requiring pre-authorised action envelopes and independent constraints.
Aerospace interpretation	Define permitted actions, authority boundaries, preconditions, rate and resource limits, independent monitors, inhibit/abort paths and recovery behaviour. Separate recommendation from command authority.
Minimum evidence	Control owner and applicability decision; system/configuration baseline; agentic-specific design evidence; dated operating or test record; open exceptions and residual-risk approval.
Enhanced evidence	Independent review; representative degraded/adversarial test; corroborating telemetry; supplier evidence; change-impact and common-mode analysis where material.
Assessment consideration	Attempt an out-of-scope action, dependency failure and communications-loss scenario. Verify that the system remains inside its authorised action envelope and produces complete decision and command evidence.
Failure indicators	Agent obtains broader credentials than intended; action chain is not bounded; operator cannot distinguish recommendation from command; no safe state under link loss; autonomous recovery conflicts with mission constraints.
Notably Absent	No airworthiness, spaceworthiness, certification, legal-compliance or regulator-acceptance conclusion; no universal threshold or approved design is asserted.

D1-CTL-03 - BEHAVIORAL DRIFT DETECTION

Illustrative implementation starting point (non-prescriptive)	Baseline profiling + KL divergence monitoring + accuracy tracking.
Validation status	ODA3 technical sector interpretation; organisation-specific engineering and independent specialist review remain required.
System class	operations
Aviation applicability	Monitor by aircraft type, fleet, route/environment, software/model baseline and operational context. Ensure alerts integrate with operations, maintenance, safety and security processes without creating unsafe workload.
Space applicability	Monitor by spacecraft, constellation, mission phase, payload mode, ground station and command baseline. Design for intermittent downlink and delayed forensic access.
Aerospace interpretation	Define baseline behaviour, telemetry coverage, time synchronisation, thresholds, alert owners, operational response and evidence retention. Detect model/data drift, provider changes, policy bypass and control degradation.
Minimum evidence	Control owner and applicability decision; system/configuration baseline; operations-specific design evidence; dated operating or test record; open exceptions and residual-risk approval.
Enhanced evidence	Independent review; representative degraded/adversarial test; corroborating telemetry; supplier evidence; change-impact and common-mode analysis where material.
Assessment consideration	Trace a production alert or simulated threshold crossing from telemetry through triage, operational decision, corrective action and closure. Confirm gaps in downlink or logging are visible rather than silently ignored.
Failure indicators	No per-baseline monitoring; time stamps cannot be correlated; fleet averages hide one configuration; alert has no operational owner; provider model changes without detection; log retention is shorter than investigation need.
Notably Absent	No airworthiness, spaceworthiness, certification, legal-compliance or regulator-acceptance conclusion; no universal threshold or approved design is asserted.

D1-CTL-04 - FEDERATED LEARNING POISONING PREVENTION

Illustrative implementation starting point (non-prescriptive)	Gradient anomaly detection + robust aggregation.
Validation status	ODA3 technical sector interpretation; organisation-specific engineering and independent specialist review remain required.
System class	governance
Aviation applicability	Applies to aircraft programmes, airlines, airports, ATM support, UAS and MRO. The accountable decision path must align with the organisation holding design, production, operating or maintenance authority; GAISSF-SEC-050 does not assign that legal authority.
Space applicability	Applies to launch, spacecraft, payload and ground-segment organisations. Mission authority, range

Illustrative implementation starting point (non-prescriptive)	Gradient anomaly detection + robust aggregation.
	authority, payload ownership and ground-operations accountability must be separated where responsibilities differ.
Aerospace interpretation	Translate the control into named aerospace decision rights, approval gates, escalation thresholds and retained records. Tie scope to a specific product, fleet, mission, site, software/model baseline and operating configuration.
Minimum evidence	Control owner and applicability decision; system/configuration baseline; governance-specific design evidence; dated operating or test record; open exceptions and residual-risk approval.
Enhanced evidence	Independent review; representative degraded/adversarial test; corroborating telemetry; supplier evidence; change-impact and common-mode analysis where material.
Assessment consideration	Sample a consequential decision and trace who approved it, which baseline and evidence were considered, which dissent or uncertainty was recorded, and how the decision would be revisited after change.
Failure indicators	Unowned AI risk; approval by a role without authority; fleet or mission scope omitted; temporary exception without expiry; decision record cannot be linked to the deployed baseline.
Notably Absent	No airworthiness, spaceworthiness, certification, legal-compliance or regulator-acceptance conclusion; no universal threshold or approved design is asserted.

D1-CTL-05 - EMBEDDING SPACE ROBUSTNESS

Illustrative implementation starting point (non-prescriptive)	Adversarial evaluation and robustness measurement under defined threat models and test conditions; select methods appropriate to the model, data and operational context.
Validation status	ODA3 technical sector interpretation; organisation-specific engineering and independent specialist review remain required.
System class	governance
Aviation applicability	Applies to aircraft programmes, airlines, airports, ATM support, UAS and MRO. The accountable decision path must align with the organisation holding design, production, operating or maintenance authority; GAISSF-SEC-050 does not assign that legal authority.
Space applicability	Applies to launch, spacecraft, payload and ground-segment organisations. Mission authority, range authority, payload ownership and ground-operations accountability must be separated where responsibilities differ.
Aerospace interpretation	Translate the control into named aerospace decision rights, approval gates, escalation thresholds and retained records. Tie scope to a specific product, fleet, mission, site, software/model baseline and operating configuration.
Minimum evidence	Control owner and applicability decision; system/configuration baseline; governance-specific design evidence; dated operating or test record; open exceptions and residual-risk approval.

Illustrative implementation starting point (non-prescriptive)	Adversarial evaluation and robustness measurement under defined threat models and test conditions; select methods appropriate to the model, data and operational context.
Enhanced evidence	Independent review; representative degraded/adversarial test; corroborating telemetry; supplier evidence; change-impact and common-mode analysis where material.
Assessment consideration	Sample a consequential decision and trace who approved it, which baseline and evidence were considered, which dissent or uncertainty was recorded, and how the decision would be revisited after change.
Failure indicators	Unowned AI risk; approval by a role without authority; fleet or mission scope omitted; temporary exception without expiry; decision record cannot be linked to the deployed baseline.
Notably Absent	No airworthiness, spaceworthiness, certification, legal-compliance or regulator-acceptance conclusion; no universal threshold or approved design is asserted.

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

Illustrative implementation starting point (non-prescriptive)	PQC signing (ML-DSA/SLH-DSA) + PQC key exchange (ML-KEM).
Validation status	ODA3 technical sector interpretation; organisation-specific engineering and independent specialist review remain required.
System class	governance
Aviation applicability	Applies to aircraft programmes, airlines, airports, ATM support, UAS and MRO. The accountable decision path must align with the organisation holding design, production, operating or maintenance authority; GAISSF-SEC-050 does not assign that legal authority.
Space applicability	Applies to launch, spacecraft, payload and ground-segment organisations. Mission authority, range authority, payload ownership and ground-operations accountability must be separated where responsibilities differ.
Aerospace interpretation	Translate the control into named aerospace decision rights, approval gates, escalation thresholds and retained records. Tie scope to a specific product, fleet, mission, site, software/model baseline and operating configuration.
Minimum evidence	Control owner and applicability decision; system/configuration baseline; governance-specific design evidence; dated operating or test record; open exceptions and residual-risk approval.
Enhanced evidence	Independent review; representative degraded/adversarial test; corroborating telemetry; supplier evidence; change-impact and common-mode analysis where material.
Assessment consideration	Sample a consequential decision and trace who approved it, which baseline and evidence were considered, which dissent or uncertainty was recorded, and how the decision would be revisited after change.
Failure indicators	Unowned AI risk; approval by a role without authority; fleet or mission scope omitted; temporary exception without expiry; decision record cannot be linked to the

Illustrative implementation starting point (non-prescriptive)	PQC signing (ML-DSA/SLH-DSA) + PQC key exchange (ML-KEM).
	deployed baseline.
Notably Absent	No airworthiness, spaceworthiness, certification, legal-compliance or regulator-acceptance conclusion; no universal threshold or approved design is asserted.

D1-CTL-07 - LORA/ADAPTER INTEGRITY VERIFICATION

Illustrative implementation starting point (non-prescriptive)	Adapter scanning + provenance verification + registry allowlist.
Validation status	ODA3 technical sector interpretation; organisation-specific engineering and independent specialist review remain required.
System class	governance
Aviation applicability	Applies to aircraft programmes, airlines, airports, ATM support, UAS and MRO. The accountable decision path must align with the organisation holding design, production, operating or maintenance authority; GAISSF-SEC-050 does not assign that legal authority.
Space applicability	Applies to launch, spacecraft, payload and ground-segment organisations. Mission authority, range authority, payload ownership and ground-operations accountability must be separated where responsibilities differ.
Aerospace interpretation	Translate the control into named aerospace decision rights, approval gates, escalation thresholds and retained records. Tie scope to a specific product, fleet, mission, site, software/model baseline and operating configuration.
Minimum evidence	Control owner and applicability decision; system/configuration baseline; governance-specific design evidence; dated operating or test record; open exceptions and residual-risk approval.
Enhanced evidence	Independent review; representative degraded/adversarial test; corroborating telemetry; supplier evidence; change-impact and common-mode analysis where material.
Assessment consideration	Sample a consequential decision and trace who approved it, which baseline and evidence were considered, which dissent or uncertainty was recorded, and how the decision would be revisited after change.
Failure indicators	Unowned AI risk; approval by a role without authority; fleet or mission scope omitted; temporary exception without expiry; decision record cannot be linked to the deployed baseline.
Notably Absent	No airworthiness, spaceworthiness, certification, legal-compliance or regulator-acceptance conclusion; no universal threshold or approved design is asserted.

D1-CTL-08 - MODEL MERGE ATTACK DETECTION

Illustrative implementation starting point (non-prescriptive)	Pre-registration behavioural evaluation + regression testing.
Validation status	ODA3 technical sector interpretation; organisation-specific engineering and independent specialist review remain required.

Illustrative implementation starting point (non-prescriptive)	Pre-registration behavioural evaluation + regression testing.
System class	governance
Aviation applicability	Applies to aircraft programmes, airlines, airports, ATM support, UAS and MRO. The accountable decision path must align with the organisation holding design, production, operating or maintenance authority; GAISSF-SEC-050 does not assign that legal authority.
Space applicability	Applies to launch, spacecraft, payload and ground-segment organisations. Mission authority, range authority, payload ownership and ground-operations accountability must be separated where responsibilities differ.
Aerospace interpretation	Translate the control into named aerospace decision rights, approval gates, escalation thresholds and retained records. Tie scope to a specific product, fleet, mission, site, software/model baseline and operating configuration.
Minimum evidence	Control owner and applicability decision; system/configuration baseline; governance-specific design evidence; dated operating or test record; open exceptions and residual-risk approval.
Enhanced evidence	Independent review; representative degraded/adversarial test; corroborating telemetry; supplier evidence; change-impact and common-mode analysis where material.
Assessment consideration	Sample a consequential decision and trace who approved it, which baseline and evidence were considered, which dissent or uncertainty was recorded, and how the decision would be revisited after change.
Failure indicators	Unowned AI risk; approval by a role without authority; fleet or mission scope omitted; temporary exception without expiry; decision record cannot be linked to the deployed baseline.
Notably Absent	No airworthiness, spaceworthiness, certification, legal-compliance or regulator-acceptance conclusion; no universal threshold or approved design is asserted.

D1-CTL-09 - QUANTIZATION BACKDOOR SCREENING

Illustrative implementation starting point (non-prescriptive)	Cross-precision behavioural comparison + delta threshold monitoring.
Validation status	ODA3 technical sector interpretation; organisation-specific engineering and independent specialist review remain required.
System class	governance
Aviation applicability	Applies to aircraft programmes, airlines, airports, ATM support, UAS and MRO. The accountable decision path must align with the organisation holding design, production, operating or maintenance authority; GAISSF-SEC-050 does not assign that legal authority.
Space applicability	Applies to launch, spacecraft, payload and ground-segment organisations. Mission authority, range authority, payload ownership and ground-operations accountability must be separated where responsibilities differ.
Aerospace interpretation	Translate the control into named aerospace decision

Illustrative implementation starting point (non-prescriptive)	Cross-precision behavioural comparison + delta threshold monitoring.
	rights, approval gates, escalation thresholds and retained records. Tie scope to a specific product, fleet, mission, site, software/model baseline and operating configuration.
Minimum evidence	Control owner and applicability decision; system/configuration baseline; governance-specific design evidence; dated operating or test record; open exceptions and residual-risk approval.
Enhanced evidence	Independent review; representative degraded/adversarial test; corroborating telemetry; supplier evidence; change-impact and common-mode analysis where material.
Assessment consideration	Sample a consequential decision and trace who approved it, which baseline and evidence were considered, which dissent or uncertainty was recorded, and how the decision would be revisited after change.
Failure indicators	Unowned AI risk; approval by a role without authority; fleet or mission scope omitted; temporary exception without expiry; decision record cannot be linked to the deployed baseline.
Notably Absent	No airworthiness, spaceworthiness, certification, legal-compliance or regulator-acceptance conclusion; no universal threshold or approved design is asserted.

D2-CTL-01 - DIRECT PROMPT INJECTION PREVENTION

Illustrative implementation starting point (non-prescriptive)	Input validation + adversarial pattern matching + system prompt isolation + guardrail sidecar.
Validation status	ODA3 technical sector interpretation; organisation-specific engineering and independent specialist review remain required.
System class	genai
Aviation applicability	Primarily relevant to maintenance copilots, technical-publication support, engineering assistants, airport/airline operations and ground systems. It is not presumed applicable to deterministic embedded flight control unless such technology is actually used.
Space applicability	Primarily relevant to engineering, mission planning, analyst support, ground operations and knowledge retrieval. Onboard applicability requires explicit confirmation of model type, authority and resource allocation.
Aerospace interpretation	Control prompts, retrieval sources, context, tool access, generated content and human approval. Technical instructions, code and operational recommendations require traceable sources and qualified verification before use.
Minimum evidence	Control owner and applicability decision; system/configuration baseline; genai-specific design evidence; dated operating or test record; open exceptions and residual-risk approval.
Enhanced evidence	Independent review; representative degraded/adversarial test; corroborating telemetry; supplier evidence; change-impact and common-mode analysis where material.

Illustrative implementation starting point (non-prescriptive)	Input validation + adversarial pattern matching + system prompt isolation + guardrail sidecar.
Assessment consideration	Run representative direct and indirect injection tests; inspect retrieval provenance and tool permissions; sample outputs for source traceability; verify that consequential use requires documented human approval.
Failure indicators	Generated maintenance step used without controlled-source verification; hidden instruction in retrieved document changes behaviour; proprietary or export-controlled data sent to an unauthorised service; tool action exceeds approved scope.
Notably Absent	No airworthiness, spaceworthiness, certification, legal-compliance or regulator-acceptance conclusion; no universal threshold or approved design is asserted.

D2-CTL-02 - INDIRECT PROMPT INJECTION PREVENTION

Illustrative implementation starting point (non-prescriptive)	Contextual separation + source allowlisting + output validation + RAG sanitization pipeline.
Validation status	ODA3 technical sector interpretation; organisation-specific engineering and independent specialist review remain required.
System class	genai
Aviation applicability	Primarily relevant to maintenance copilots, technical-publication support, engineering assistants, airport/airline operations and ground systems. It is not presumed applicable to deterministic embedded flight control unless such technology is actually used.
Space applicability	Primarily relevant to engineering, mission planning, analyst support, ground operations and knowledge retrieval. Onboard applicability requires explicit confirmation of model type, authority and resource allocation.
Aerospace interpretation	Control prompts, retrieval sources, context, tool access, generated content and human approval. Technical instructions, code and operational recommendations require traceable sources and qualified verification before use.
Minimum evidence	Control owner and applicability decision; system/configuration baseline; genai-specific design evidence; dated operating or test record; open exceptions and residual-risk approval.
Enhanced evidence	Independent review; representative degraded/adversarial test; corroborating telemetry; supplier evidence; change-impact and common-mode analysis where material.
Assessment consideration	Run representative direct and indirect injection tests; inspect retrieval provenance and tool permissions; sample outputs for source traceability; verify that consequential use requires documented human approval.
Failure indicators	Generated maintenance step used without controlled-source verification; hidden instruction in retrieved document changes behaviour; proprietary or export-controlled data sent to an unauthorised service; tool action exceeds approved scope.

Illustrative implementation starting point (non-prescriptive)	Contextual separation + source allowlisting + output validation + RAG sanitization pipeline.
Notably Absent	No airworthiness, spaceworthiness, certification, legal-compliance or regulator-acceptance conclusion; no universal threshold or approved design is asserted.

D2-CTL-03 - JAILBREAK RESISTANCE TESTING

Illustrative implementation starting point (non-prescriptive)	Quarterly red-team prompt library + adversarial training + automated refusal monitoring.
Validation status	ODA3 technical sector interpretation; organisation-specific engineering and independent specialist review remain required.
System class	assurance
Aviation applicability	Use representative aircraft, simulator, environmental, operational and human-in-the-loop conditions. Evidence may support—but does not replace—applicable certification and safety processes.
Space applicability	Use mission-phase, orbital/trajectory, communications, radiation/resource and ground-segment scenarios. Include delayed detection and inaccessible-asset conditions.
Aerospace interpretation	Predefine acceptance criteria, configurations, datasets, scenarios and uncertainty treatment. Test expected, edge, degraded, adversarial and recovery conditions and preserve enough detail for independent repetition.
Minimum evidence	Control owner and applicability decision; system/configuration baseline; assurance-specific design evidence; dated operating or test record; open exceptions and residual-risk approval.
Enhanced evidence	Independent review; representative degraded/adversarial test; corroborating telemetry; supplier evidence; change-impact and common-mode analysis where material.
Assessment consideration	Reperform a risk-based sample, confirm the tested configuration matches the release, inspect failed cases and verify that limitations were carried into operational restrictions and monitoring.
Failure indicators	Only average-case metrics; test data overlaps training data; no degraded-mode or rare-event scenarios; unapproved threshold changes; assessor lacks independence; failures are excluded from reported results.
Notably Absent	No airworthiness, spaceworthiness, certification, legal-compliance or regulator-acceptance conclusion; no universal threshold or approved design is asserted.

D2-CTL-04 - MULTI-MODAL INJECTION DEFENSE

Illustrative implementation starting point (non-prescriptive)	Multi-modal content scanning + steganography detection + modality-specific guardrails.
Validation status	ODA3 technical sector interpretation; organisation-specific engineering and independent specialist review remain required.
System class	governance
Aviation applicability	Applies to aircraft programmes, airlines, airports, ATM support, UAS and MRO. The accountable decision path

Illustrative implementation starting point (non-prescriptive)	Multi-modal content scanning + steganography detection + modality-specific guardrails.
	must align with the organisation holding design, production, operating or maintenance authority; GAISSF-SEC-050 does not assign that legal authority.
Space applicability	Applies to launch, spacecraft, payload and ground-segment organisations. Mission authority, range authority, payload ownership and ground-operations accountability must be separated where responsibilities differ.
Aerospace interpretation	Translate the control into named aerospace decision rights, approval gates, escalation thresholds and retained records. Tie scope to a specific product, fleet, mission, site, software/model baseline and operating configuration.
Minimum evidence	Control owner and applicability decision; system/configuration baseline; governance-specific design evidence; dated operating or test record; open exceptions and residual-risk approval.
Enhanced evidence	Independent review; representative degraded/adversarial test; corroborating telemetry; supplier evidence; change-impact and common-mode analysis where material.
Assessment consideration	Sample a consequential decision and trace who approved it, which baseline and evidence were considered, which dissent or uncertainty was recorded, and how the decision would be revisited after change.
Failure indicators	Unowned AI risk; approval by a role without authority; fleet or mission scope omitted; temporary exception without expiry; decision record cannot be linked to the deployed baseline.
Notably Absent	No airworthiness, spaceworthiness, certification, legal-compliance or regulator-acceptance conclusion; no universal threshold or approved design is asserted.

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

Illustrative implementation starting point (non-prescriptive)	Parameter schema validation + allowlist enforcement + sandboxed execution.
Validation status	ODA3 technical sector interpretation; organisation-specific engineering and independent specialist review remain required.
System class	agentic
Aviation applicability	Relevant when AI can initiate, sequence or recommend actions affecting dispatch, maintenance, ground movement, UAS, AAM or operational systems. Human oversight must be feasible within the time available; “human in the loop” alone is not evidence of effective control.
Space applicability	Relevant to onboard autonomy, fault management, planning, rendezvous/proximity operations and ground automation. Intervention may be delayed or unavailable, requiring pre-authorised action envelopes and independent constraints.
Aerospace interpretation	Define permitted actions, authority boundaries, preconditions, rate and resource limits, independent monitors, inhibit/abort paths and recovery behaviour.

Illustrative implementation starting point (non-prescriptive)	Parameter schema validation + allowlist enforcement + sandboxed execution.
	Separate recommendation from command authority.
Minimum evidence	Control owner and applicability decision; system/configuration baseline; agentic-specific design evidence; dated operating or test record; open exceptions and residual-risk approval.
Enhanced evidence	Independent review; representative degraded/adversarial test; corroborating telemetry; supplier evidence; change-impact and common-mode analysis where material.
Assessment consideration	Attempt an out-of-scope action, dependency failure and communications-loss scenario. Verify that the system remains inside its authorised action envelope and produces complete decision and command evidence.
Failure indicators	Agent obtains broader credentials than intended; action chain is not bounded; operator cannot distinguish recommendation from command; no safe state under link loss; autonomous recovery conflicts with mission constraints.
Notably Absent	No airworthiness, spaceworthiness, certification, legal-compliance or regulator-acceptance conclusion; no universal threshold or approved design is asserted.

D2-CTL-06 - CROSS-CONTEXT HIJACKING MITIGATION

Illustrative implementation starting point (non-prescriptive)	Context window segmentation + prompt anchoring + attention boundary enforcement.
Validation status	ODA3 technical sector interpretation; organisation-specific engineering and independent specialist review remain required.
System class	governance
Aviation applicability	Applies to aircraft programmes, airlines, airports, ATM support, UAS and MRO. The accountable decision path must align with the organisation holding design, production, operating or maintenance authority; GAISSF-SEC-050 does not assign that legal authority.
Space applicability	Applies to launch, spacecraft, payload and ground-segment organisations. Mission authority, range authority, payload ownership and ground-operations accountability must be separated where responsibilities differ.
Aerospace interpretation	Translate the control into named aerospace decision rights, approval gates, escalation thresholds and retained records. Tie scope to a specific product, fleet, mission, site, software/model baseline and operating configuration.
Minimum evidence	Control owner and applicability decision; system/configuration baseline; governance-specific design evidence; dated operating or test record; open exceptions and residual-risk approval.
Enhanced evidence	Independent review; representative degraded/adversarial test; corroborating telemetry; supplier evidence; change-impact and common-mode analysis where material.
Assessment consideration	Sample a consequential decision and trace who approved it, which baseline and evidence were

Illustrative implementation starting point (non-prescriptive)	Context window segmentation + prompt anchoring + attention boundary enforcement.
	considered, which dissent or uncertainty was recorded, and how the decision would be revisited after change.
Failure indicators	Unowned AI risk; approval by a role without authority; fleet or mission scope omitted; temporary exception without expiry; decision record cannot be linked to the deployed baseline.
Notably Absent	No airworthiness, spaceworthiness, certification, legal-compliance or regulator-acceptance conclusion; no universal threshold or approved design is asserted.

D3-CTL-01 - LEAST AGENCY ENFORCEMENT

Illustrative implementation starting point (non-prescriptive)	Role-based tool scoping + policy-as-code + dynamic permission revocation.
Validation status	ODA3 technical sector interpretation; organisation-specific engineering and independent specialist review remain required.
System class	governance
Aviation applicability	Applies to aircraft programmes, airlines, airports, ATM support, UAS and MRO. The accountable decision path must align with the organisation holding design, production, operating or maintenance authority; GAISSF-SEC-050 does not assign that legal authority.
Space applicability	Applies to launch, spacecraft, payload and ground-segment organisations. Mission authority, range authority, payload ownership and ground-operations accountability must be separated where responsibilities differ.
Aerospace interpretation	Translate the control into named aerospace decision rights, approval gates, escalation thresholds and retained records. Tie scope to a specific product, fleet, mission, site, software/model baseline and operating configuration.
Minimum evidence	Control owner and applicability decision; system/configuration baseline; governance-specific design evidence; dated operating or test record; open exceptions and residual-risk approval.
Enhanced evidence	Independent review; representative degraded/adversarial test; corroborating telemetry; supplier evidence; change-impact and common-mode analysis where material.
Assessment consideration	Sample a consequential decision and trace who approved it, which baseline and evidence were considered, which dissent or uncertainty was recorded, and how the decision would be revisited after change.
Failure indicators	Unowned AI risk; approval by a role without authority; fleet or mission scope omitted; temporary exception without expiry; decision record cannot be linked to the deployed baseline.
Notably Absent	No airworthiness, spaceworthiness, certification, legal-compliance or regulator-acceptance conclusion; no universal threshold or approved design is asserted.

D3-CTL-02 - INTER-AGENT COMMUNICATION SECURITY

Illustrative implementation starting point (non-prescriptive)	mTLS for agent mesh + message signing + payload validation.
Validation status	ODA3 technical sector interpretation; organisation-specific engineering and independent specialist review remain required.
System class	agentic
Aviation applicability	Relevant when AI can initiate, sequence or recommend actions affecting dispatch, maintenance, ground movement, UAS, AAM or operational systems. Human oversight must be feasible within the time available; “human in the loop” alone is not evidence of effective control.
Space applicability	Relevant to onboard autonomy, fault management, planning, rendezvous/proximity operations and ground automation. Intervention may be delayed or unavailable, requiring pre-authorised action envelopes and independent constraints.
Aerospace interpretation	Define permitted actions, authority boundaries, preconditions, rate and resource limits, independent monitors, inhibit/abort paths and recovery behaviour. Separate recommendation from command authority.
Minimum evidence	Control owner and applicability decision; system/configuration baseline; agentic-specific design evidence; dated operating or test record; open exceptions and residual-risk approval.
Enhanced evidence	Independent review; representative degraded/adversarial test; corroborating telemetry; supplier evidence; change-impact and common-mode analysis where material.
Assessment consideration	Attempt an out-of-scope action, dependency failure and communications-loss scenario. Verify that the system remains inside its authorised action envelope and produces complete decision and command evidence.
Failure indicators	Agent obtains broader credentials than intended; action chain is not bounded; operator cannot distinguish recommendation from command; no safe state under link loss; autonomous recovery conflicts with mission constraints.
Notably Absent	No airworthiness, spaceworthiness, certification, legal-compliance or regulator-acceptance conclusion; no universal threshold or approved design is asserted.

D3-CTL-03 - AGENTIC PROMPT CHAINING DETECTION

Illustrative implementation starting point (non-prescriptive)	Cross-session behavioural correlation + chain pattern detection + anomaly scoring.
Validation status	ODA3 technical sector interpretation; organisation-specific engineering and independent specialist review remain required.
System class	genai
Aviation applicability	Primarily relevant to maintenance copilots, technical-publication support, engineering assistants, airport/airline operations and ground systems. It is not presumed applicable to deterministic embedded flight control unless such technology is actually used.

Illustrative implementation starting point (non-prescriptive)	Cross-session behavioural correlation + chain pattern detection + anomaly scoring.
Space applicability	Primarily relevant to engineering, mission planning, analyst support, ground operations and knowledge retrieval. Onboard applicability requires explicit confirmation of model type, authority and resource allocation.
Aerospace interpretation	Control prompts, retrieval sources, context, tool access, generated content and human approval. Technical instructions, code and operational recommendations require traceable sources and qualified verification before use.
Minimum evidence	Control owner and applicability decision; system/configuration baseline; genai-specific design evidence; dated operating or test record; open exceptions and residual-risk approval.
Enhanced evidence	Independent review; representative degraded/adversarial test; corroborating telemetry; supplier evidence; change-impact and common-mode analysis where material.
Assessment consideration	Run representative direct and indirect injection tests; inspect retrieval provenance and tool permissions; sample outputs for source traceability; verify that consequential use requires documented human approval.
Failure indicators	Generated maintenance step used without controlled-source verification; hidden instruction in retrieved document changes behaviour; proprietary or export-controlled data sent to an unauthorised service; tool action exceeds approved scope.
Notably Absent	No airworthiness, spaceworthiness, certification, legal-compliance or regulator-acceptance conclusion; no universal threshold or approved design is asserted.

D3-CTL-04 - EMBODIED AI SAFETY CONTROLS

Illustrative implementation starting point (non-prescriptive)	Sensor integrity verification + safety interlocks + fail-safe state enforcement.
Validation status	ODA3 technical sector interpretation; organisation-specific engineering and independent specialist review remain required.
System class	governance
Aviation applicability	Applies to aircraft programmes, airlines, airports, ATM support, UAS and MRO. The accountable decision path must align with the organisation holding design, production, operating or maintenance authority; GAISSF-SEC-050 does not assign that legal authority.
Space applicability	Applies to launch, spacecraft, payload and ground-segment organisations. Mission authority, range authority, payload ownership and ground-operations accountability must be separated where responsibilities differ.
Aerospace interpretation	Translate the control into named aerospace decision rights, approval gates, escalation thresholds and retained records. Tie scope to a specific product, fleet, mission, site, software/model baseline and operating configuration.

Illustrative implementation starting point (non-prescriptive)	Sensor integrity verification + safety interlocks + fail-safe state enforcement.
Minimum evidence	Control owner and applicability decision; system/configuration baseline; governance-specific design evidence; dated operating or test record; open exceptions and residual-risk approval.
Enhanced evidence	Independent review; representative degraded/adversarial test; corroborating telemetry; supplier evidence; change-impact and common-mode analysis where material.
Assessment consideration	Sample a consequential decision and trace who approved it, which baseline and evidence were considered, which dissent or uncertainty was recorded, and how the decision would be revisited after change.
Failure indicators	Unowned AI risk; approval by a role without authority; fleet or mission scope omitted; temporary exception without expiry; decision record cannot be linked to the deployed baseline.
Notably Absent	No airworthiness, spaceworthiness, certification, legal-compliance or regulator-acceptance conclusion; no universal threshold or approved design is asserted.

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

Illustrative implementation starting point (non-prescriptive)	SPIFFE/SPIRE workload identity + short-lived certificates + continuous attestation.
Validation status	ODA3 technical sector interpretation; organisation-specific engineering and independent specialist review remain required.
System class	agentic
Aviation applicability	Relevant when AI can initiate, sequence or recommend actions affecting dispatch, maintenance, ground movement, UAS, AAM or operational systems. Human oversight must be feasible within the time available; “human in the loop” alone is not evidence of effective control.
Space applicability	Relevant to onboard autonomy, fault management, planning, rendezvous/proximity operations and ground automation. Intervention may be delayed or unavailable, requiring pre-authorised action envelopes and independent constraints.
Aerospace interpretation	Define permitted actions, authority boundaries, preconditions, rate and resource limits, independent monitors, inhibit/abort paths and recovery behaviour. Separate recommendation from command authority.
Minimum evidence	Control owner and applicability decision; system/configuration baseline; agentic-specific design evidence; dated operating or test record; open exceptions and residual-risk approval.
Enhanced evidence	Independent review; representative degraded/adversarial test; corroborating telemetry; supplier evidence; change-impact and common-mode analysis where material.
Assessment consideration	Attempt an out-of-scope action, dependency failure and communications-loss scenario. Verify that the system remains inside its authorised action envelope and produces complete decision and command evidence.

Illustrative implementation starting point (non-prescriptive)	SPIFFE/SPIRE workload identity + short-lived certificates + continuous attestation.
Failure indicators	Agent obtains broader credentials than intended; action chain is not bounded; operator cannot distinguish recommendation from command; no safe state under link loss; autonomous recovery conflicts with mission constraints.
Notably Absent	No airworthiness, spaceworthiness, certification, legal-compliance or regulator-acceptance conclusion; no universal threshold or approved design is asserted.

D3-CTL-06 - PERSISTENT MEMORY EXFILTRATION PREVENTION

Illustrative implementation starting point (non-prescriptive)	User-scoped memory isolation + encryption at rest + query-level access controls.
Validation status	ODA3 technical sector interpretation; organisation-specific engineering and independent specialist review remain required.
System class	governance
Aviation applicability	Applies to aircraft programmes, airlines, airports, ATM support, UAS and MRO. The accountable decision path must align with the organisation holding design, production, operating or maintenance authority; GAISSF-SEC-050 does not assign that legal authority.
Space applicability	Applies to launch, spacecraft, payload and ground-segment organisations. Mission authority, range authority, payload ownership and ground-operations accountability must be separated where responsibilities differ.
Aerospace interpretation	Translate the control into named aerospace decision rights, approval gates, escalation thresholds and retained records. Tie scope to a specific product, fleet, mission, site, software/model baseline and operating configuration.
Minimum evidence	Control owner and applicability decision; system/configuration baseline; governance-specific design evidence; dated operating or test record; open exceptions and residual-risk approval.
Enhanced evidence	Independent review; representative degraded/adversarial test; corroborating telemetry; supplier evidence; change-impact and common-mode analysis where material.
Assessment consideration	Sample a consequential decision and trace who approved it, which baseline and evidence were considered, which dissent or uncertainty was recorded, and how the decision would be revisited after change.
Failure indicators	Unowned AI risk; approval by a role without authority; fleet or mission scope omitted; temporary exception without expiry; decision record cannot be linked to the deployed baseline.
Notably Absent	No airworthiness, spaceworthiness, certification, legal-compliance or regulator-acceptance conclusion; no universal threshold or approved design is asserted.

D3-CTL-07 - SECURE MEMORY LIFECYCLE MANAGEMENT

Illustrative implementation starting point (non-prescriptive)	Cryptographic deletion + lifecycle policy enforcement + retention auditing.
Validation status	ODA3 technical sector interpretation; organisation-specific engineering and independent specialist review remain required.
System class	governance
Aviation applicability	Applies to aircraft programmes, airlines, airports, ATM support, UAS and MRO. The accountable decision path must align with the organisation holding design, production, operating or maintenance authority; GAISSF-SEC-050 does not assign that legal authority.
Space applicability	Applies to launch, spacecraft, payload and ground-segment organisations. Mission authority, range authority, payload ownership and ground-operations accountability must be separated where responsibilities differ.
Aerospace interpretation	Translate the control into named aerospace decision rights, approval gates, escalation thresholds and retained records. Tie scope to a specific product, fleet, mission, site, software/model baseline and operating configuration.
Minimum evidence	Control owner and applicability decision; system/configuration baseline; governance-specific design evidence; dated operating or test record; open exceptions and residual-risk approval.
Enhanced evidence	Independent review; representative degraded/adversarial test; corroborating telemetry; supplier evidence; change-impact and common-mode analysis where material.
Assessment consideration	Sample a consequential decision and trace who approved it, which baseline and evidence were considered, which dissent or uncertainty was recorded, and how the decision would be revisited after change.
Failure indicators	Unowned AI risk; approval by a role without authority; fleet or mission scope omitted; temporary exception without expiry; decision record cannot be linked to the deployed baseline.
Notably Absent	No airworthiness, spaceworthiness, certification, legal-compliance or regulator-acceptance conclusion; no universal threshold or approved design is asserted.

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

Illustrative implementation starting point (non-prescriptive)	Automated BOM generation + version tracking + registry synchronization.
Validation status	ODA3 technical sector interpretation; organisation-specific engineering and independent specialist review remain required.
System class	governance
Aviation applicability	Applies to aircraft programmes, airlines, airports, ATM support, UAS and MRO. The accountable decision path must align with the organisation holding design, production, operating or maintenance authority; GAISSF-SEC-050 does not assign that legal authority.
Space applicability	Applies to launch, spacecraft, payload and ground-

Illustrative implementation starting point (non-prescriptive)	Automated BOM generation + version tracking + registry synchronization.
	segment organisations. Mission authority, range authority, payload ownership and ground-operations accountability must be separated where responsibilities differ.
Aerospace interpretation	Translate the control into named aerospace decision rights, approval gates, escalation thresholds and retained records. Tie scope to a specific product, fleet, mission, site, software/model baseline and operating configuration.
Minimum evidence	Control owner and applicability decision; system/configuration baseline; governance-specific design evidence; dated operating or test record; open exceptions and residual-risk approval.
Enhanced evidence	Independent review; representative degraded/adversarial test; corroborating telemetry; supplier evidence; change-impact and common-mode analysis where material.
Assessment consideration	Sample a consequential decision and trace who approved it, which baseline and evidence were considered, which dissent or uncertainty was recorded, and how the decision would be revisited after change.
Failure indicators	Unowned AI risk; approval by a role without authority; fleet or mission scope omitted; temporary exception without expiry; decision record cannot be linked to the deployed baseline.
Notably Absent	No airworthiness, spaceworthiness, certification, legal-compliance or regulator-acceptance conclusion; no universal threshold or approved design is asserted.

D4-CTL-02 - MODEL FILE & ARTIFACT SCANNING

Illustrative implementation starting point (non-prescriptive)	Static analysis + deserialization sandboxing + signature verification.
Validation status	ODA3 technical sector interpretation; organisation-specific engineering and independent specialist review remain required.
System class	governance
Aviation applicability	Applies to aircraft programmes, airlines, airports, ATM support, UAS and MRO. The accountable decision path must align with the organisation holding design, production, operating or maintenance authority; GAISSF-SEC-050 does not assign that legal authority.
Space applicability	Applies to launch, spacecraft, payload and ground-segment organisations. Mission authority, range authority, payload ownership and ground-operations accountability must be separated where responsibilities differ.
Aerospace interpretation	Translate the control into named aerospace decision rights, approval gates, escalation thresholds and retained records. Tie scope to a specific product, fleet, mission, site, software/model baseline and operating configuration.
Minimum evidence	Control owner and applicability decision; system/configuration baseline; governance-specific design evidence; dated operating or test record; open

Illustrative implementation starting point (non-prescriptive)	Static analysis + deserialization sandboxing + signature verification.
	exceptions and residual-risk approval.
Enhanced evidence	Independent review; representative degraded/adversarial test; corroborating telemetry; supplier evidence; change-impact and common-mode analysis where material.
Assessment consideration	Sample a consequential decision and trace who approved it, which baseline and evidence were considered, which dissent or uncertainty was recorded, and how the decision would be revisited after change.
Failure indicators	Unowned AI risk; approval by a role without authority; fleet or mission scope omitted; temporary exception without expiry; decision record cannot be linked to the deployed baseline.
Notably Absent	No airworthiness, spaceworthiness, certification, legal-compliance or regulator-acceptance conclusion; no universal threshold or approved design is asserted.

D4-CTL-03 - MODEL HUB & REGISTRY VETTING

Illustrative implementation starting point (non-prescriptive)	Provenance verification + license compliance + security scorecard.
Validation status	ODA3 technical sector interpretation; organisation-specific engineering and independent specialist review remain required.
System class	governance
Aviation applicability	Applies to aircraft programmes, airlines, airports, ATM support, UAS and MRO. The accountable decision path must align with the organisation holding design, production, operating or maintenance authority; GAISSF-SEC-050 does not assign that legal authority.
Space applicability	Applies to launch, spacecraft, payload and ground-segment organisations. Mission authority, range authority, payload ownership and ground-operations accountability must be separated where responsibilities differ.
Aerospace interpretation	Translate the control into named aerospace decision rights, approval gates, escalation thresholds and retained records. Tie scope to a specific product, fleet, mission, site, software/model baseline and operating configuration.
Minimum evidence	Control owner and applicability decision; system/configuration baseline; governance-specific design evidence; dated operating or test record; open exceptions and residual-risk approval.
Enhanced evidence	Independent review; representative degraded/adversarial test; corroborating telemetry; supplier evidence; change-impact and common-mode analysis where material.
Assessment consideration	Sample a consequential decision and trace who approved it, which baseline and evidence were considered, which dissent or uncertainty was recorded, and how the decision would be revisited after change.
Failure indicators	Unowned AI risk; approval by a role without authority; fleet or mission scope omitted; temporary exception without expiry; decision record cannot be linked to the

Illustrative implementation starting point (non-prescriptive)	Provenance verification + license compliance + security scorecard.
	deployed baseline.
Notably Absent	No airworthiness, spaceworthiness, certification, legal-compliance or regulator-acceptance conclusion; no universal threshold or approved design is asserted.

D4-CTL-04 - MCP SERVER BEHAVIORAL MONITORING

Illustrative implementation starting point (non-prescriptive)	Tool-call logging + anomaly detection + access control enforcement.
Validation status	ODA3 technical sector interpretation; organisation-specific engineering and independent specialist review remain required.
System class	operations
Aviation applicability	Monitor by aircraft type, fleet, route/environment, software/model baseline and operational context. Ensure alerts integrate with operations, maintenance, safety and security processes without creating unsafe workload.
Space applicability	Monitor by spacecraft, constellation, mission phase, payload mode, ground station and command baseline. Design for intermittent downlink and delayed forensic access.
Aerospace interpretation	Define baseline behaviour, telemetry coverage, time synchronisation, thresholds, alert owners, operational response and evidence retention. Detect model/data drift, provider changes, policy bypass and control degradation.
Minimum evidence	Control owner and applicability decision; system/configuration baseline; operations-specific design evidence; dated operating or test record; open exceptions and residual-risk approval.
Enhanced evidence	Independent review; representative degraded/adversarial test; corroborating telemetry; supplier evidence; change-impact and common-mode analysis where material.
Assessment consideration	Trace a production alert or simulated threshold crossing from telemetry through triage, operational decision, corrective action and closure. Confirm gaps in downlink or logging are visible rather than silently ignored.
Failure indicators	No per-baseline monitoring; time stamps cannot be correlated; fleet averages hide one configuration; alert has no operational owner; provider model changes without detection; log retention is shorter than investigation need.
Notably Absent	No airworthiness, spaceworthiness, certification, legal-compliance or regulator-acceptance conclusion; no universal threshold or approved design is asserted.

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

Illustrative implementation starting point (non-prescriptive)	Contractual security requirements + penetration testing + data flow mapping.
Validation status	ODA3 technical sector interpretation; organisation-specific engineering and independent specialist review remain required.

Illustrative implementation starting point (non-prescriptive)	Contractual security requirements + penetration testing + data flow mapping.
System class	supply
Aviation applicability	Applies across OEM, avionics, software, model, cloud, data, MRO and engineering suppliers. Requirements should flow down according to design and operational authority, safety relevance and access to controlled technical data.
Space applicability	Applies across spacecraft, launch, ground-station, payload, cloud, analytics and communications suppliers. Consider long support periods, component substitution, export restrictions and remote support jurisdictions.
Aerospace interpretation	Require component and model provenance, subcontractor visibility, signed releases, vulnerability and incident cooperation, change notification, support/EOL commitments and evidence delivery. Avoid relying on generic attestations where technical evidence is available.
Minimum evidence	Control owner and applicability decision; system/configuration baseline; supply-specific design evidence; dated operating or test record; open exceptions and residual-risk approval.
Enhanced evidence	Independent review; representative degraded/adversarial test; corroborating telemetry; supplier evidence; change-impact and common-mode analysis where material.
Assessment consideration	Trace one critical supplier component from contract and due diligence through build manifest, release approval, deployment, vulnerability notification and support obligation.
Failure indicators	Unknown fourth party; missing model or adapter hash; supplier can change hosted model without notice; no support-end plan; remote privileged access is not recorded; contractual evidence rights are absent.
Notably Absent	No airworthiness, spaceworthiness, certification, legal-compliance or regulator-acceptance conclusion; no universal threshold or approved design is asserted.

D4-CTL-06 - SHADOW AI DISCOVERY & GOVERNANCE

Illustrative implementation starting point (non-prescriptive)	Network traffic analysis + SaaS discovery + policy enforcement.
Validation status	ODA3 technical sector interpretation; organisation-specific engineering and independent specialist review remain required.
System class	governance
Aviation applicability	Applies to aircraft programmes, airlines, airports, ATM support, UAS and MRO. The accountable decision path must align with the organisation holding design, production, operating or maintenance authority; GAISSF-SEC-050 does not assign that legal authority.
Space applicability	Applies to launch, spacecraft, payload and ground-segment organisations. Mission authority, range authority, payload ownership and ground-operations accountability must be separated where responsibilities differ.
Aerospace interpretation	Translate the control into named aerospace decision

Illustrative implementation starting point (non-prescriptive)	Network traffic analysis + SaaS discovery + policy enforcement.
	rights, approval gates, escalation thresholds and retained records. Tie scope to a specific product, fleet, mission, site, software/model baseline and operating configuration.
Minimum evidence	Control owner and applicability decision; system/configuration baseline; governance-specific design evidence; dated operating or test record; open exceptions and residual-risk approval.
Enhanced evidence	Independent review; representative degraded/adversarial test; corroborating telemetry; supplier evidence; change-impact and common-mode analysis where material.
Assessment consideration	Sample a consequential decision and trace who approved it, which baseline and evidence were considered, which dissent or uncertainty was recorded, and how the decision would be revisited after change.
Failure indicators	Unowned AI risk; approval by a role without authority; fleet or mission scope omitted; temporary exception without expiry; decision record cannot be linked to the deployed baseline.
Notably Absent	No airworthiness, spaceworthiness, certification, legal-compliance or regulator-acceptance conclusion; no universal threshold or approved design is asserted.

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

Illustrative implementation starting point (non-prescriptive)	Dependency scanning + CVE matching + automated patching.
Validation status	ODA3 technical sector interpretation; organisation-specific engineering and independent specialist review remain required.
System class	governance
Aviation applicability	Applies to aircraft programmes, airlines, airports, ATM support, UAS and MRO. The accountable decision path must align with the organisation holding design, production, operating or maintenance authority; GAISSF-SEC-050 does not assign that legal authority.
Space applicability	Applies to launch, spacecraft, payload and ground-segment organisations. Mission authority, range authority, payload ownership and ground-operations accountability must be separated where responsibilities differ.
Aerospace interpretation	Translate the control into named aerospace decision rights, approval gates, escalation thresholds and retained records. Tie scope to a specific product, fleet, mission, site, software/model baseline and operating configuration.
Minimum evidence	Control owner and applicability decision; system/configuration baseline; governance-specific design evidence; dated operating or test record; open exceptions and residual-risk approval.
Enhanced evidence	Independent review; representative degraded/adversarial test; corroborating telemetry; supplier evidence; change-impact and common-mode analysis where material.

Illustrative implementation starting point (non-prescriptive)	Dependency scanning + CVE matching + automated patching.
Assessment consideration	Sample a consequential decision and trace who approved it, which baseline and evidence were considered, which dissent or uncertainty was recorded, and how the decision would be revisited after change.
Failure indicators	Unowned AI risk; approval by a role without authority; fleet or mission scope omitted; temporary exception without expiry; decision record cannot be linked to the deployed baseline.
Notably Absent	No airworthiness, spaceworthiness, certification, legal-compliance or regulator-acceptance conclusion; no universal threshold or approved design is asserted.

D5-CTL-01 - HARMFUL CONTENT BLOCKING

Illustrative implementation starting point (non-prescriptive)	Content safety classifier + refusal engine.
Validation status	ODA3 technical sector interpretation; organisation-specific engineering and independent specialist review remain required.
System class	genai
Aviation applicability	Primarily relevant to maintenance copilots, technical-publication support, engineering assistants, airport/airline operations and ground systems. It is not presumed applicable to deterministic embedded flight control unless such technology is actually used.
Space applicability	Primarily relevant to engineering, mission planning, analyst support, ground operations and knowledge retrieval. Onboard applicability requires explicit confirmation of model type, authority and resource allocation.
Aerospace interpretation	Control prompts, retrieval sources, context, tool access, generated content and human approval. Technical instructions, code and operational recommendations require traceable sources and qualified verification before use.
Minimum evidence	Control owner and applicability decision; system/configuration baseline; genai-specific design evidence; dated operating or test record; open exceptions and residual-risk approval.
Enhanced evidence	Independent review; representative degraded/adversarial test; corroborating telemetry; supplier evidence; change-impact and common-mode analysis where material.
Assessment consideration	Run representative direct and indirect injection tests; inspect retrieval provenance and tool permissions; sample outputs for source traceability; verify that consequential use requires documented human approval.
Failure indicators	Generated maintenance step used without controlled-source verification; hidden instruction in retrieved document changes behaviour; proprietary or export-controlled data sent to an unauthorised service; tool action exceeds approved scope.
Notably Absent	No airworthiness, spaceworthiness, certification, legal-compliance or regulator-acceptance conclusion; no universal threshold or approved design is asserted.

D5-CTL-02 - PII LEAKAGE PREVENTION

Illustrative implementation starting point (non-prescriptive)	PII detection + masking + access controls.
Validation status	ODA3 technical sector interpretation; organisation-specific engineering and independent specialist review remain required.
System class	governance
Aviation applicability	Applies to aircraft programmes, airlines, airports, ATM support, UAS and MRO. The accountable decision path must align with the organisation holding design, production, operating or maintenance authority; GAISSF-SEC-050 does not assign that legal authority.
Space applicability	Applies to launch, spacecraft, payload and ground-segment organisations. Mission authority, range authority, payload ownership and ground-operations accountability must be separated where responsibilities differ.
Aerospace interpretation	Translate the control into named aerospace decision rights, approval gates, escalation thresholds and retained records. Tie scope to a specific product, fleet, mission, site, software/model baseline and operating configuration.
Minimum evidence	Control owner and applicability decision; system/configuration baseline; governance-specific design evidence; dated operating or test record; open exceptions and residual-risk approval.
Enhanced evidence	Independent review; representative degraded/adversarial test; corroborating telemetry; supplier evidence; change-impact and common-mode analysis where material.
Assessment consideration	Sample a consequential decision and trace who approved it, which baseline and evidence were considered, which dissent or uncertainty was recorded, and how the decision would be revisited after change.
Failure indicators	Unowned AI risk; approval by a role without authority; fleet or mission scope omitted; temporary exception without expiry; decision record cannot be linked to the deployed baseline.
Notably Absent	No airworthiness, spaceworthiness, certification, legal-compliance or regulator-acceptance conclusion; no universal threshold or approved design is asserted.

D5-CTL-03 - COPYRIGHT DETECTION

Illustrative implementation starting point (non-prescriptive)	n-gram overlap detection + refusal.
Validation status	ODA3 technical sector interpretation; organisation-specific engineering and independent specialist review remain required.
System class	governance
Aviation applicability	Applies to aircraft programmes, airlines, airports, ATM support, UAS and MRO. The accountable decision path must align with the organisation holding design, production, operating or maintenance authority; GAISSF-SEC-050 does not assign that legal authority.
Space applicability	Applies to launch, spacecraft, payload and ground-segment organisations. Mission authority, range authority, payload ownership and ground-operations

Illustrative implementation starting point (non-prescriptive)	n-gram overlap detection + refusal.
	accountability must be separated where responsibilities differ.
Aerospace interpretation	Translate the control into named aerospace decision rights, approval gates, escalation thresholds and retained records. Tie scope to a specific product, fleet, mission, site, software/model baseline and operating configuration.
Minimum evidence	Control owner and applicability decision; system/configuration baseline; governance-specific design evidence; dated operating or test record; open exceptions and residual-risk approval.
Enhanced evidence	Independent review; representative degraded/adversarial test; corroborating telemetry; supplier evidence; change-impact and common-mode analysis where material.
Assessment consideration	Sample a consequential decision and trace who approved it, which baseline and evidence were considered, which dissent or uncertainty was recorded, and how the decision would be revisited after change.
Failure indicators	Unowned AI risk; approval by a role without authority; fleet or mission scope omitted; temporary exception without expiry; decision record cannot be linked to the deployed baseline.
Notably Absent	No airworthiness, spaceworthiness, certification, legal-compliance or regulator-acceptance conclusion; no universal threshold or approved design is asserted.

D5-CTL-04 - AI WATERMARKING ROBUSTNESS

Illustrative implementation starting point (non-prescriptive)	C2PA-compliant watermarking + tamper resistance testing.
Validation status	ODA3 technical sector interpretation; organisation-specific engineering and independent specialist review remain required.
System class	governance
Aviation applicability	Applies to aircraft programmes, airlines, airports, ATM support, UAS and MRO. The accountable decision path must align with the organisation holding design, production, operating or maintenance authority; GAISSF-SEC-050 does not assign that legal authority.
Space applicability	Applies to launch, spacecraft, payload and ground-segment organisations. Mission authority, range authority, payload ownership and ground-operations accountability must be separated where responsibilities differ.
Aerospace interpretation	Translate the control into named aerospace decision rights, approval gates, escalation thresholds and retained records. Tie scope to a specific product, fleet, mission, site, software/model baseline and operating configuration.
Minimum evidence	Control owner and applicability decision; system/configuration baseline; governance-specific design evidence; dated operating or test record; open exceptions and residual-risk approval.
Enhanced evidence	Independent review; representative degraded/adversarial test; corroborating telemetry;

Illustrative implementation starting point (non-prescriptive)	C2PA-compliant watermarking + tamper resistance testing.
	supplier evidence; change-impact and common-mode analysis where material.
Assessment consideration	Sample a consequential decision and trace who approved it, which baseline and evidence were considered, which dissent or uncertainty was recorded, and how the decision would be revisited after change.
Failure indicators	Unowned AI risk; approval by a role without authority; fleet or mission scope omitted; temporary exception without expiry; decision record cannot be linked to the deployed baseline.
Notably Absent	No airworthiness, spaceworthiness, certification, legal-compliance or regulator-acceptance conclusion; no universal threshold or approved design is asserted.

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

Illustrative implementation starting point (non-prescriptive)	Data minimization + purpose limitation + machine unlearning.
Validation status	ODA3 technical sector interpretation; organisation-specific engineering and independent specialist review remain required.
System class	governance
Aviation applicability	Applies to aircraft programmes, airlines, airports, ATM support, UAS and MRO. The accountable decision path must align with the organisation holding design, production, operating or maintenance authority; GAISSF-SEC-050 does not assign that legal authority.
Space applicability	Applies to launch, spacecraft, payload and ground-segment organisations. Mission authority, range authority, payload ownership and ground-operations accountability must be separated where responsibilities differ.
Aerospace interpretation	Translate the control into named aerospace decision rights, approval gates, escalation thresholds and retained records. Tie scope to a specific product, fleet, mission, site, software/model baseline and operating configuration.
Minimum evidence	Control owner and applicability decision; system/configuration baseline; governance-specific design evidence; dated operating or test record; open exceptions and residual-risk approval.
Enhanced evidence	Independent review; representative degraded/adversarial test; corroborating telemetry; supplier evidence; change-impact and common-mode analysis where material.
Assessment consideration	Sample a consequential decision and trace who approved it, which baseline and evidence were considered, which dissent or uncertainty was recorded, and how the decision would be revisited after change.
Failure indicators	Unowned AI risk; approval by a role without authority; fleet or mission scope omitted; temporary exception without expiry; decision record cannot be linked to the deployed baseline.
Notably Absent	No airworthiness, spaceworthiness, certification, legal-compliance or regulator-acceptance conclusion; no

Illustrative implementation starting point (non-prescriptive)	Data minimization + purpose limitation + machine unlearning.
	universal threshold or approved design is asserted.

D5-CTL-06 - PRIVACY-PRESERVING ML VALIDATION

Illustrative implementation starting point (non-prescriptive)	Differential privacy + membership inference testing.
Validation status	ODA3 technical sector interpretation; organisation-specific engineering and independent specialist review remain required.
System class	assurance
Aviation applicability	Use representative aircraft, simulator, environmental, operational and human-in-the-loop conditions. Evidence may support—but does not replace—applicable certification and safety processes.
Space applicability	Use mission-phase, orbital/trajectory, communications, radiation/resource and ground-segment scenarios. Include delayed detection and inaccessible-asset conditions.
Aerospace interpretation	Predefine acceptance criteria, configurations, datasets, scenarios and uncertainty treatment. Test expected, edge, degraded, adversarial and recovery conditions and preserve enough detail for independent repetition.
Minimum evidence	Control owner and applicability decision; system/configuration baseline; assurance-specific design evidence; dated operating or test record; open exceptions and residual-risk approval.
Enhanced evidence	Independent review; representative degraded/adversarial test; corroborating telemetry; supplier evidence; change-impact and common-mode analysis where material.
Assessment consideration	Reperform a risk-based sample, confirm the tested configuration matches the release, inspect failed cases and verify that limitations were carried into operational restrictions and monitoring.
Failure indicators	Only average-case metrics; test data overlaps training data; no degraded-mode or rare-event scenarios; unapproved threshold changes; assessor lacks independence; failures are excluded from reported results.
Notably Absent	No airworthiness, spaceworthiness, certification, legal-compliance or regulator-acceptance conclusion; no universal threshold or approved design is asserted.

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

Illustrative implementation starting point (non-prescriptive)	Approval workflow + policy enforcement + audit log.
Validation status	ODA3 technical sector interpretation; organisation-specific engineering and independent specialist review remain required.
System class	agentic
Aviation applicability	Relevant when AI can initiate, sequence or recommend actions affecting dispatch, maintenance, ground movement, UAS, AAM or operational systems. Human oversight must be feasible within the time available; “human in the loop” alone is not evidence of effective control.

Illustrative implementation starting point (non-prescriptive)	Approval workflow + policy enforcement + audit log.
Space applicability	Relevant to onboard autonomy, fault management, planning, rendezvous/proximity operations and ground automation. Intervention may be delayed or unavailable, requiring pre-authorised action envelopes and independent constraints.
Aerospace interpretation	Define permitted actions, authority boundaries, preconditions, rate and resource limits, independent monitors, inhibit/abort paths and recovery behaviour. Separate recommendation from command authority.
Minimum evidence	Control owner and applicability decision; system/configuration baseline; agentic-specific design evidence; dated operating or test record; open exceptions and residual-risk approval.
Enhanced evidence	Independent review; representative degraded/adversarial test; corroborating telemetry; supplier evidence; change-impact and common-mode analysis where material.
Assessment consideration	Attempt an out-of-scope action, dependency failure and communications-loss scenario. Verify that the system remains inside its authorised action envelope and produces complete decision and command evidence.
Failure indicators	Agent obtains broader credentials than intended; action chain is not bounded; operator cannot distinguish recommendation from command; no safe state under link loss; autonomous recovery conflicts with mission constraints.
Notably Absent	No airworthiness, spaceworthiness, certification, legal-compliance or regulator-acceptance conclusion; no universal threshold or approved design is asserted.

D6-CTL-02 - AUDIT TRAIL COMPLETENESS

Illustrative implementation starting point (non-prescriptive)	Structured logging + SIEM integration + retention enforcement.
Validation status	ODA3 technical sector interpretation; organisation-specific engineering and independent specialist review remain required.
System class	governance
Aviation applicability	Applies to aircraft programmes, airlines, airports, ATM support, UAS and MRO. The accountable decision path must align with the organisation holding design, production, operating or maintenance authority; GAISSF-SEC-050 does not assign that legal authority.
Space applicability	Applies to launch, spacecraft, payload and ground-segment organisations. Mission authority, range authority, payload ownership and ground-operations accountability must be separated where responsibilities differ.
Aerospace interpretation	Translate the control into named aerospace decision rights, approval gates, escalation thresholds and retained records. Tie scope to a specific product, fleet, mission, site, software/model baseline and operating configuration.
Minimum evidence	Control owner and applicability decision; system/configuration baseline; governance-specific design evidence; dated operating or test record; open

Illustrative implementation starting point (non-prescriptive)	Structured logging + SIEM integration + retention enforcement.
	exceptions and residual-risk approval.
Enhanced evidence	Independent review; representative degraded/adversarial test; corroborating telemetry; supplier evidence; change-impact and common-mode analysis where material.
Assessment consideration	Sample a consequential decision and trace who approved it, which baseline and evidence were considered, which dissent or uncertainty was recorded, and how the decision would be revisited after change.
Failure indicators	Unowned AI risk; approval by a role without authority; fleet or mission scope omitted; temporary exception without expiry; decision record cannot be linked to the deployed baseline.
Notably Absent	No airworthiness, spaceworthiness, certification, legal-compliance or regulator-acceptance conclusion; no universal threshold or approved design is asserted.

D6-CTL-03 - AI MODEL CARD COMPLETENESS

Illustrative implementation starting point (non-prescriptive)	Standardized template + version control + public accessibility.
Validation status	ODA3 technical sector interpretation; organisation-specific engineering and independent specialist review remain required.
System class	governance
Aviation applicability	Applies to aircraft programmes, airlines, airports, ATM support, UAS and MRO. The accountable decision path must align with the organisation holding design, production, operating or maintenance authority; GAISSF-SEC-050 does not assign that legal authority.
Space applicability	Applies to launch, spacecraft, payload and ground-segment organisations. Mission authority, range authority, payload ownership and ground-operations accountability must be separated where responsibilities differ.
Aerospace interpretation	Translate the control into named aerospace decision rights, approval gates, escalation thresholds and retained records. Tie scope to a specific product, fleet, mission, site, software/model baseline and operating configuration.
Minimum evidence	Control owner and applicability decision; system/configuration baseline; governance-specific design evidence; dated operating or test record; open exceptions and residual-risk approval.
Enhanced evidence	Independent review; representative degraded/adversarial test; corroborating telemetry; supplier evidence; change-impact and common-mode analysis where material.
Assessment consideration	Sample a consequential decision and trace who approved it, which baseline and evidence were considered, which dissent or uncertainty was recorded, and how the decision would be revisited after change.
Failure indicators	Unowned AI risk; approval by a role without authority; fleet or mission scope omitted; temporary exception without expiry; decision record cannot be linked to the

Illustrative implementation starting point (non-prescriptive)	Standardized template + version control + public accessibility.
	deployed baseline.
Notably Absent	No airworthiness, spaceworthiness, certification, legal-compliance or regulator-acceptance conclusion; no universal threshold or approved design is asserted.

D6-CTL-04 - AI INCIDENT RESPONSE READINESS

Illustrative implementation starting point (non-prescriptive)	AI-IR runbook + tabletop exercises + containment automation.
Validation status	ODA3 technical sector interpretation; organisation-specific engineering and independent specialist review remain required.
System class	incident
Aviation applicability	Coordinate cyber incident response with safety, maintenance, operational control, continued airworthiness and regulator/customer notification processes as applicable. Containment must not create a more hazardous condition.
Space applicability	Coordinate mission anomaly, cyber incident, range/launch safety and ground-segment response. Preserve command history and account for limited access to the affected asset.
Aerospace interpretation	Define AI-specific detection, classification, containment, rollback, operational restriction, evidence preservation, stakeholder communication and return-to-service criteria. Include near misses and common-mode fleet/constellation risk.
Minimum evidence	Control owner and applicability decision; system/configuration baseline; incident-specific design evidence; dated operating or test record; open exceptions and residual-risk approval.
Enhanced evidence	Independent review; representative degraded/adversarial test; corroborating telemetry; supplier evidence; change-impact and common-mode analysis where material.
Assessment consideration	Exercise a scenario involving unsafe output or compromised update. Verify authority to inhibit or roll back, evidence preservation, downstream notification and objective recovery criteria.
Failure indicators	Security team disables a function without safety/mission assessment; no model/data quarantine path; copied artefacts remain in use; return to service is based on elapsed time rather than verified criteria.
Notably Absent	No airworthiness, spaceworthiness, certification, legal-compliance or regulator-acceptance conclusion; no universal threshold or approved design is asserted.

D6-CTL-05 - MODEL DEPRECATION & DECOMMISSIONING

Illustrative implementation starting point (non-prescriptive)	Access revocation + decommission audit + scheduled lifecycle.
Validation status	ODA3 technical sector interpretation; organisation-specific engineering and independent specialist review remain required.
System class	governance

Illustrative implementation starting point (non-prescriptive)	Access revocation + decommission audit + scheduled lifecycle.
Aviation applicability	Applies to aircraft programmes, airlines, airports, ATM support, UAS and MRO. The accountable decision path must align with the organisation holding design, production, operating or maintenance authority; GAISSF-SEC-050 does not assign that legal authority.
Space applicability	Applies to launch, spacecraft, payload and ground-segment organisations. Mission authority, range authority, payload ownership and ground-operations accountability must be separated where responsibilities differ.
Aerospace interpretation	Translate the control into named aerospace decision rights, approval gates, escalation thresholds and retained records. Tie scope to a specific product, fleet, mission, site, software/model baseline and operating configuration.
Minimum evidence	Control owner and applicability decision; system/configuration baseline; governance-specific design evidence; dated operating or test record; open exceptions and residual-risk approval.
Enhanced evidence	Independent review; representative degraded/adversarial test; corroborating telemetry; supplier evidence; change-impact and common-mode analysis where material.
Assessment consideration	Sample a consequential decision and trace who approved it, which baseline and evidence were considered, which dissent or uncertainty was recorded, and how the decision would be revisited after change.
Failure indicators	Unowned AI risk; approval by a role without authority; fleet or mission scope omitted; temporary exception without expiry; decision record cannot be linked to the deployed baseline.
Notably Absent	No airworthiness, spaceworthiness, certification, legal-compliance or regulator-acceptance conclusion; no universal threshold or approved design is asserted.

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

Illustrative implementation starting point (non-prescriptive)	Contractual security requirements + annual assessment + audit rights.
Validation status	ODA3 technical sector interpretation; organisation-specific engineering and independent specialist review remain required.
System class	supply
Aviation applicability	Applies across OEM, avionics, software, model, cloud, data, MRO and engineering suppliers. Requirements should flow down according to design and operational authority, safety relevance and access to controlled technical data.
Space applicability	Applies across spacecraft, launch, ground-station, payload, cloud, analytics and communications suppliers. Consider long support periods, component substitution, export restrictions and remote support jurisdictions.
Aerospace interpretation	Require component and model provenance, subcontractor visibility, signed releases, vulnerability and incident cooperation, change notification,

Illustrative implementation starting point (non-prescriptive)	Contractual security requirements + annual assessment + audit rights.
	support/EOL commitments and evidence delivery. Avoid relying on generic attestations where technical evidence is available.
Minimum evidence	Control owner and applicability decision; system/configuration baseline; supply-specific design evidence; dated operating or test record; open exceptions and residual-risk approval.
Enhanced evidence	Independent review; representative degraded/adversarial test; corroborating telemetry; supplier evidence; change-impact and common-mode analysis where material.
Assessment consideration	Trace one critical supplier component from contract and due diligence through build manifest, release approval, deployment, vulnerability notification and support obligation.
Failure indicators	Unknown fourth party; missing model or adapter hash; supplier can change hosted model without notice; no support-end plan; remote privileged access is not recorded; contractual evidence rights are absent.
Notably Absent	No airworthiness, spaceworthiness, certification, legal-compliance or regulator-acceptance conclusion; no universal threshold or approved design is asserted.

D6-CTL-07 - AI RESILIENCE & BUSINESS CONTINUITY

Illustrative implementation starting point (non-prescriptive)	Failover systems + degraded mode + RTO/RPO definition.
Validation status	ODA3 technical sector interpretation; organisation-specific engineering and independent specialist review remain required.
System class	incident
Aviation applicability	Coordinate cyber incident response with safety, maintenance, operational control, continued airworthiness and regulator/customer notification processes as applicable. Containment must not create a more hazardous condition.
Space applicability	Coordinate mission anomaly, cyber incident, range/launch safety and ground-segment response. Preserve command history and account for limited access to the affected asset.
Aerospace interpretation	Define AI-specific detection, classification, containment, rollback, operational restriction, evidence preservation, stakeholder communication and return-to-service criteria. Include near misses and common-mode fleet/constellation risk.
Minimum evidence	Control owner and applicability decision; system/configuration baseline; incident-specific design evidence; dated operating or test record; open exceptions and residual-risk approval.
Enhanced evidence	Independent review; representative degraded/adversarial test; corroborating telemetry; supplier evidence; change-impact and common-mode analysis where material.
Assessment consideration	Exercise a scenario involving unsafe output or compromised update. Verify authority to inhibit or roll back, evidence preservation, downstream notification

Illustrative implementation starting point (non-prescriptive)	Failover systems + degraded mode + RTO/RPO definition.
	and objective recovery criteria.
Failure indicators	Security team disables a function without safety/mission assessment; no model/data quarantine path; copied artefacts remain in use; return to service is based on elapsed time rather than verified criteria.
Notably Absent	No airworthiness, spaceworthiness, certification, legal-compliance or regulator-acceptance conclusion; no universal threshold or approved design is asserted.

D7-CTL-H01 - AI-GENERATED PHISHING SIMULATION

Illustrative implementation starting point (non-prescriptive)	Simulation campaigns + click tracking + remedial training.
Validation status	ODA3 technical sector interpretation; organisation-specific engineering and independent specialist review remain required.
System class	governance
Aviation applicability	Applies to aircraft programmes, airlines, airports, ATM support, UAS and MRO. The accountable decision path must align with the organisation holding design, production, operating or maintenance authority; GAISSF-SEC-050 does not assign that legal authority.
Space applicability	Applies to launch, spacecraft, payload and ground-segment organisations. Mission authority, range authority, payload ownership and ground-operations accountability must be separated where responsibilities differ.
Aerospace interpretation	Translate the control into named aerospace decision rights, approval gates, escalation thresholds and retained records. Tie scope to a specific product, fleet, mission, site, software/model baseline and operating configuration.
Minimum evidence	Control owner and applicability decision; system/configuration baseline; governance-specific design evidence; dated operating or test record; open exceptions and residual-risk approval.
Enhanced evidence	Independent review; representative degraded/adversarial test; corroborating telemetry; supplier evidence; change-impact and common-mode analysis where material.
Assessment consideration	Sample a consequential decision and trace who approved it, which baseline and evidence were considered, which dissent or uncertainty was recorded, and how the decision would be revisited after change.
Failure indicators	Unowned AI risk; approval by a role without authority; fleet or mission scope omitted; temporary exception without expiry; decision record cannot be linked to the deployed baseline.
Notably Absent	No airworthiness, spaceworthiness, certification, legal-compliance or regulator-acceptance conclusion; no universal threshold or approved design is asserted.

D7-CTL-H02 - DEEPPFAKE DETECTION TRAINING

Illustrative implementation starting point (non-prescriptive)	Training modules + quiz + simulated attacks.
Validation status	ODA3 technical sector interpretation; organisation-

Illustrative implementation starting point (non-prescriptive)	Training modules + quiz + simulated attacks.
	specific engineering and independent specialist review remain required.
System class	genai
Aviation applicability	Primarily relevant to maintenance copilots, technical-publication support, engineering assistants, airport/airline operations and ground systems. It is not presumed applicable to deterministic embedded flight control unless such technology is actually used.
Space applicability	Primarily relevant to engineering, mission planning, analyst support, ground operations and knowledge retrieval. Onboard applicability requires explicit confirmation of model type, authority and resource allocation.
Aerospace interpretation	Control prompts, retrieval sources, context, tool access, generated content and human approval. Technical instructions, code and operational recommendations require traceable sources and qualified verification before use.
Minimum evidence	Control owner and applicability decision; system/configuration baseline; genai-specific design evidence; dated operating or test record; open exceptions and residual-risk approval.
Enhanced evidence	Independent review; representative degraded/adversarial test; corroborating telemetry; supplier evidence; change-impact and common-mode analysis where material.
Assessment consideration	Run representative direct and indirect injection tests; inspect retrieval provenance and tool permissions; sample outputs for source traceability; verify that consequential use requires documented human approval.
Failure indicators	Generated maintenance step used without controlled-source verification; hidden instruction in retrieved document changes behaviour; proprietary or export-controlled data sent to an unauthorised service; tool action exceeds approved scope.
Notably Absent	No airworthiness, spaceworthiness, certification, legal-compliance or regulator-acceptance conclusion; no universal threshold or approved design is asserted.

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

Illustrative implementation starting point (non-prescriptive)	Independent channel verification + policy enforcement.
Validation status	ODA3 technical sector interpretation; organisation-specific engineering and independent specialist review remain required.
System class	governance
Aviation applicability	Applies to aircraft programmes, airlines, airports, ATM support, UAS and MRO. The accountable decision path must align with the organisation holding design, production, operating or maintenance authority; GAISSF-SEC-050 does not assign that legal authority.
Space applicability	Applies to launch, spacecraft, payload and ground-segment organisations. Mission authority, range authority, payload ownership and ground-operations accountability must be separated where responsibilities

Illustrative implementation starting point (non-prescriptive)	Independent channel verification + policy enforcement.
	differ.
Aerospace interpretation	Translate the control into named aerospace decision rights, approval gates, escalation thresholds and retained records. Tie scope to a specific product, fleet, mission, site, software/model baseline and operating configuration.
Minimum evidence	Control owner and applicability decision; system/configuration baseline; governance-specific design evidence; dated operating or test record; open exceptions and residual-risk approval.
Enhanced evidence	Independent review; representative degraded/adversarial test; corroborating telemetry; supplier evidence; change-impact and common-mode analysis where material.
Assessment consideration	Sample a consequential decision and trace who approved it, which baseline and evidence were considered, which dissent or uncertainty was recorded, and how the decision would be revisited after change.
Failure indicators	Unowned AI risk; approval by a role without authority; fleet or mission scope omitted; temporary exception without expiry; decision record cannot be linked to the deployed baseline.
Notably Absent	No airworthiness, spaceworthiness, certification, legal-compliance or regulator-acceptance conclusion; no universal threshold or approved design is asserted.

D7-CTL-H04 - AI SOCIAL ENGINEERING IR

Illustrative implementation starting point (non-prescriptive)	Tabletop exercises + IR plan + verification triggers.
Validation status	ODA3 technical sector interpretation; organisation-specific engineering and independent specialist review remain required.
System class	governance
Aviation applicability	Applies to aircraft programmes, airlines, airports, ATM support, UAS and MRO. The accountable decision path must align with the organisation holding design, production, operating or maintenance authority; GAISSF-SEC-050 does not assign that legal authority.
Space applicability	Applies to launch, spacecraft, payload and ground-segment organisations. Mission authority, range authority, payload ownership and ground-operations accountability must be separated where responsibilities differ.
Aerospace interpretation	Translate the control into named aerospace decision rights, approval gates, escalation thresholds and retained records. Tie scope to a specific product, fleet, mission, site, software/model baseline and operating configuration.
Minimum evidence	Control owner and applicability decision; system/configuration baseline; governance-specific design evidence; dated operating or test record; open exceptions and residual-risk approval.
Enhanced evidence	Independent review; representative degraded/adversarial test; corroborating telemetry; supplier evidence; change-impact and common-mode analysis where material.

Illustrative implementation starting point (non-prescriptive)	Tabletop exercises + IR plan + verification triggers.
Assessment consideration	Sample a consequential decision and trace who approved it, which baseline and evidence were considered, which dissent or uncertainty was recorded, and how the decision would be revisited after change.
Failure indicators	Unowned AI risk; approval by a role without authority; fleet or mission scope omitted; temporary exception without expiry; decision record cannot be linked to the deployed baseline.
Notably Absent	No airworthiness, spaceworthiness, certification, legal-compliance or regulator-acceptance conclusion; no universal threshold or approved design is asserted.

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

Illustrative implementation starting point (non-prescriptive)	AI-generated phishing detection + SOC tuning + response automation.
Validation status	ODA3 technical sector interpretation; organisation-specific engineering and independent specialist review remain required.
System class	governance
Aviation applicability	Applies to aircraft programmes, airlines, airports, ATM support, UAS and MRO. The accountable decision path must align with the organisation holding design, production, operating or maintenance authority; GAISSF-SEC-050 does not assign that legal authority.
Space applicability	Applies to launch, spacecraft, payload and ground-segment organisations. Mission authority, range authority, payload ownership and ground-operations accountability must be separated where responsibilities differ.
Aerospace interpretation	Translate the control into named aerospace decision rights, approval gates, escalation thresholds and retained records. Tie scope to a specific product, fleet, mission, site, software/model baseline and operating configuration.
Minimum evidence	Control owner and applicability decision; system/configuration baseline; governance-specific design evidence; dated operating or test record; open exceptions and residual-risk approval.
Enhanced evidence	Independent review; representative degraded/adversarial test; corroborating telemetry; supplier evidence; change-impact and common-mode analysis where material.
Assessment consideration	Sample a consequential decision and trace who approved it, which baseline and evidence were considered, which dissent or uncertainty was recorded, and how the decision would be revisited after change.
Failure indicators	Unowned AI risk; approval by a role without authority; fleet or mission scope omitted; temporary exception without expiry; decision record cannot be linked to the deployed baseline.
Notably Absent	No airworthiness, spaceworthiness, certification, legal-compliance or regulator-acceptance conclusion; no universal threshold or approved design is asserted.

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

Illustrative implementation starting point (non-prescriptive)	Risk classification framework + conformity assessment.
Validation status	ODA3 technical sector interpretation; organisation-specific engineering and independent specialist review remain required.
System class	governance
Aviation applicability	Applies to aircraft programmes, airlines, airports, ATM support, UAS and MRO. The accountable decision path must align with the organisation holding design, production, operating or maintenance authority; GAISSF-SEC-050 does not assign that legal authority.
Space applicability	Applies to launch, spacecraft, payload and ground-segment organisations. Mission authority, range authority, payload ownership and ground-operations accountability must be separated where responsibilities differ.
Aerospace interpretation	Translate the control into named aerospace decision rights, approval gates, escalation thresholds and retained records. Tie scope to a specific product, fleet, mission, site, software/model baseline and operating configuration.
Minimum evidence	Control owner and applicability decision; system/configuration baseline; governance-specific design evidence; dated operating or test record; open exceptions and residual-risk approval.
Enhanced evidence	Independent review; representative degraded/adversarial test; corroborating telemetry; supplier evidence; change-impact and common-mode analysis where material.
Assessment consideration	Sample a consequential decision and trace who approved it, which baseline and evidence were considered, which dissent or uncertainty was recorded, and how the decision would be revisited after change.
Failure indicators	Unowned AI risk; approval by a role without authority; fleet or mission scope omitted; temporary exception without expiry; decision record cannot be linked to the deployed baseline.
Notably Absent	No airworthiness, spaceworthiness, certification, legal-compliance or regulator-acceptance conclusion; no universal threshold or approved design is asserted.

D8-CTL-02 - ISO 42001 GAP ANALYSIS

Illustrative implementation starting point (non-prescriptive)	Gap analysis methodology + remediation tracking.
Validation status	ODA3 technical sector interpretation; organisation-specific engineering and independent specialist review remain required.
System class	governance
Aviation applicability	Applies to aircraft programmes, airlines, airports, ATM support, UAS and MRO. The accountable decision path must align with the organisation holding design, production, operating or maintenance authority; GAISSF-SEC-050 does not assign that legal authority.
Space applicability	Applies to launch, spacecraft, payload and ground-segment organisations. Mission authority, range authority, payload ownership and ground-operations

Illustrative implementation starting point (non-prescriptive)	Gap analysis methodology + remediation tracking.
	accountability must be separated where responsibilities differ.
Aerospace interpretation	Translate the control into named aerospace decision rights, approval gates, escalation thresholds and retained records. Tie scope to a specific product, fleet, mission, site, software/model baseline and operating configuration.
Minimum evidence	Control owner and applicability decision; system/configuration baseline; governance-specific design evidence; dated operating or test record; open exceptions and residual-risk approval.
Enhanced evidence	Independent review; representative degraded/adversarial test; corroborating telemetry; supplier evidence; change-impact and common-mode analysis where material.
Assessment consideration	Sample a consequential decision and trace who approved it, which baseline and evidence were considered, which dissent or uncertainty was recorded, and how the decision would be revisited after change.
Failure indicators	Unowned AI risk; approval by a role without authority; fleet or mission scope omitted; temporary exception without expiry; decision record cannot be linked to the deployed baseline.
Notably Absent	No airworthiness, spaceworthiness, certification, legal-compliance or regulator-acceptance conclusion; no universal threshold or approved design is asserted.

D8-CTL-03 - GPAI TECHNICAL DOCUMENTATION VERIFICATION

Illustrative implementation starting point (non-prescriptive)	Technical documentation + training data summary + copyright attestation.
Validation status	ODA3 technical sector interpretation; organisation-specific engineering and independent specialist review remain required.
System class	governance
Aviation applicability	Applies to aircraft programmes, airlines, airports, ATM support, UAS and MRO. The accountable decision path must align with the organisation holding design, production, operating or maintenance authority; GAISSF-SEC-050 does not assign that legal authority.
Space applicability	Applies to launch, spacecraft, payload and ground-segment organisations. Mission authority, range authority, payload ownership and ground-operations accountability must be separated where responsibilities differ.
Aerospace interpretation	Translate the control into named aerospace decision rights, approval gates, escalation thresholds and retained records. Tie scope to a specific product, fleet, mission, site, software/model baseline and operating configuration.
Minimum evidence	Control owner and applicability decision; system/configuration baseline; governance-specific design evidence; dated operating or test record; open exceptions and residual-risk approval.
Enhanced evidence	Independent review; representative degraded/adversarial test; corroborating telemetry;

Illustrative implementation starting point (non-prescriptive)	Technical documentation + training data summary + copyright attestation.
	supplier evidence; change-impact and common-mode analysis where material.
Assessment consideration	Sample a consequential decision and trace who approved it, which baseline and evidence were considered, which dissent or uncertainty was recorded, and how the decision would be revisited after change.
Failure indicators	Unowned AI risk; approval by a role without authority; fleet or mission scope omitted; temporary exception without expiry; decision record cannot be linked to the deployed baseline.
Notably Absent	No airworthiness, spaceworthiness, certification, legal-compliance or regulator-acceptance conclusion; no universal threshold or approved design is asserted.

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

Illustrative implementation starting point (non-prescriptive)	Incident classification + notification workflow + SLA monitoring.
Validation status	ODA3 technical sector interpretation; organisation-specific engineering and independent specialist review remain required.
System class	incident
Aviation applicability	Coordinate cyber incident response with safety, maintenance, operational control, continued airworthiness and regulator/customer notification processes as applicable. Containment must not create a more hazardous condition.
Space applicability	Coordinate mission anomaly, cyber incident, range/launch safety and ground-segment response. Preserve command history and account for limited access to the affected asset.
Aerospace interpretation	Define AI-specific detection, classification, containment, rollback, operational restriction, evidence preservation, stakeholder communication and return-to-service criteria. Include near misses and common-mode fleet/constellation risk.
Minimum evidence	Control owner and applicability decision; system/configuration baseline; incident-specific design evidence; dated operating or test record; open exceptions and residual-risk approval.
Enhanced evidence	Independent review; representative degraded/adversarial test; corroborating telemetry; supplier evidence; change-impact and common-mode analysis where material.
Assessment consideration	Exercise a scenario involving unsafe output or compromised update. Verify authority to inhibit or roll back, evidence preservation, downstream notification and objective recovery criteria.
Failure indicators	Security team disables a function without safety/mission assessment; no model/data quarantine path; copied artefacts remain in use; return to service is based on elapsed time rather than verified criteria.
Notably Absent	No airworthiness, spaceworthiness, certification, legal-compliance or regulator-acceptance conclusion; no universal threshold or approved design is asserted.

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

Illustrative implementation starting point (non-prescriptive)	Secure development practices + attestation.
Validation status	ODA3 technical sector interpretation; organisation-specific engineering and independent specialist review remain required.
System class	governance
Aviation applicability	Applies to aircraft programmes, airlines, airports, ATM support, UAS and MRO. The accountable decision path must align with the organisation holding design, production, operating or maintenance authority; GAISSF-SEC-050 does not assign that legal authority.
Space applicability	Applies to launch, spacecraft, payload and ground-segment organisations. Mission authority, range authority, payload ownership and ground-operations accountability must be separated where responsibilities differ.
Aerospace interpretation	Translate the control into named aerospace decision rights, approval gates, escalation thresholds and retained records. Tie scope to a specific product, fleet, mission, site, software/model baseline and operating configuration.
Minimum evidence	Control owner and applicability decision; system/configuration baseline; governance-specific design evidence; dated operating or test record; open exceptions and residual-risk approval.
Enhanced evidence	Independent review; representative degraded/adversarial test; corroborating telemetry; supplier evidence; change-impact and common-mode analysis where material.
Assessment consideration	Sample a consequential decision and trace who approved it, which baseline and evidence were considered, which dissent or uncertainty was recorded, and how the decision would be revisited after change.
Failure indicators	Unowned AI risk; approval by a role without authority; fleet or mission scope omitted; temporary exception without expiry; decision record cannot be linked to the deployed baseline.
Notably Absent	No airworthiness, spaceworthiness, certification, legal-compliance or regulator-acceptance conclusion; no universal threshold or approved design is asserted.

D9-CTL-01 - PHYSICAL HARM BOUNDARY ENFORCEMENT

Illustrative implementation starting point (non-prescriptive)	Define independent physical safety constraints, operating envelopes and command limits appropriate to the platform and applicable safety/design-assurance process. Use architecture-specific monitors or interlocks where justified; do not infer a universal DAL/SIL or certification level from GAISSF.
Validation status	ODA3 technical sector interpretation; organisation-specific engineering and independent specialist review remain required.
System class	physical
Aviation applicability	Relevant to perception, UAS/AAM, autonomous taxi, inspection, robotics and any AI linked to aircraft or ground equipment. Evaluate sensor diversity, environment, timing and independent physical

Illustrative implementation starting point (non-prescriptive)	Define independent physical safety constraints, operating envelopes and command limits appropriate to the platform and applicable safety/design-assurance process. Use architecture-specific monitors or interlocks where justified; do not infer a universal DAL/SIL or certification level from GAISSF.
	safeguards.
Space applicability	Relevant to spacecraft autonomy, launch, robotics, rendezvous/proximity operations and physical ground equipment. Consider communications delay, inaccessible assets and environmental extremes.
Aerospace interpretation	Define physical operating boundaries, sensor plausibility checks, actuator/command limits, fail-safe or fail-operational behaviour, human intervention and independent safety constraints.
Minimum evidence	Control owner and applicability decision; system/configuration baseline; physical-specific design evidence; dated operating or test record; open exceptions and residual-risk approval.
Enhanced evidence	Independent review; representative degraded/adversarial test; corroborating telemetry; supplier evidence; change-impact and common-mode analysis where material.
Assessment consideration	Test representative environmental, sensor-degradation, spoofing, timing, resource and actuator-boundary scenarios on suitable simulation, hardware-in-the-loop or physical test assets.
Failure indicators	Single sensor drives consequential action without plausibility check; model is validated only in laboratory conditions; override path is unavailable under fault; unsafe command persists after confidence collapses.
Notably Absent	No airworthiness, spaceworthiness, certification, legal-compliance or regulator-acceptance conclusion; no universal threshold or approved design is asserted.

D9-CTL-02 - SAFE STATE AND GRACEFUL DEGRADATION

Illustrative implementation starting point (non-prescriptive)	Define safe and degraded states, transition criteria, recovery conditions and authority for the specific aircraft, spacecraft, ground system or mission context. Timing and fail-safe/fail-operational behaviour must be derived from applicable hazard, mission and system analyses rather than a universal ladder.
Validation status	ODA3 technical sector interpretation; organisation-specific engineering and independent specialist review remain required.
System class	governance
Aviation applicability	Applies to aircraft programmes, airlines, airports, ATM support, UAS and MRO. The accountable decision path must align with the organisation holding design, production, operating or maintenance authority; GAISSF-SEC-050 does not assign that legal authority.
Space applicability	Applies to launch, spacecraft, payload and ground-segment organisations. Mission authority, range authority, payload ownership and ground-operations accountability must be separated where responsibilities differ.
Aerospace interpretation	Translate the control into named aerospace decision

Illustrative implementation starting point (non-prescriptive)	Define safe and degraded states, transition criteria, recovery conditions and authority for the specific aircraft, spacecraft, ground system or mission context. Timing and fail-safe/fail-operational behaviour must be derived from applicable hazard, mission and system analyses rather than a universal ladder.
	rights, approval gates, escalation thresholds and retained records. Tie scope to a specific product, fleet, mission, site, software/model baseline and operating configuration.
Minimum evidence	Control owner and applicability decision; system/configuration baseline; governance-specific design evidence; dated operating or test record; open exceptions and residual-risk approval.
Enhanced evidence	Independent review; representative degraded/adversarial test; corroborating telemetry; supplier evidence; change-impact and common-mode analysis where material.
Assessment consideration	Sample a consequential decision and trace who approved it, which baseline and evidence were considered, which dissent or uncertainty was recorded, and how the decision would be revisited after change.
Failure indicators	Unowned AI risk; approval by a role without authority; fleet or mission scope omitted; temporary exception without expiry; decision record cannot be linked to the deployed baseline.
Notably Absent	No airworthiness, spaceworthiness, certification, legal-compliance or regulator-acceptance conclusion; no universal threshold or approved design is asserted.

D9-CTL-03 - HUMAN OVERRIDE AND EMERGENCY STOP

Illustrative implementation starting point (non-prescriptive)	Provide human or independent override/inhibit mechanisms appropriate to the architecture, operational concept and applicable safety/mission process. Demonstrate that the AI function cannot improperly defeat the approved override path; do not assume a universal hardware emergency-stop design.
Validation status	ODA3 technical sector interpretation; organisation-specific engineering and independent specialist review remain required.
System class	governance
Aviation applicability	Applies to aircraft programmes, airlines, airports, ATM support, UAS and MRO. The accountable decision path must align with the organisation holding design, production, operating or maintenance authority; GAISSF-SEC-050 does not assign that legal authority.
Space applicability	Applies to launch, spacecraft, payload and ground-segment organisations. Mission authority, range authority, payload ownership and ground-operations accountability must be separated where responsibilities differ.
Aerospace interpretation	Translate the control into named aerospace decision rights, approval gates, escalation thresholds and retained records. Tie scope to a specific product, fleet, mission, site, software/model baseline and operating configuration.
Minimum evidence	Control owner and applicability decision;

Illustrative implementation starting point (non-prescriptive)	Provide human or independent override/inhibit mechanisms appropriate to the architecture, operational concept and applicable safety/mission process. Demonstrate that the AI function cannot improperly defeat the approved override path; do not assume a universal hardware emergency-stop design.
	system/configuration baseline; governance-specific design evidence; dated operating or test record; open exceptions and residual-risk approval.
Enhanced evidence	Independent review; representative degraded/adversarial test; corroborating telemetry; supplier evidence; change-impact and common-mode analysis where material.
Assessment consideration	Sample a consequential decision and trace who approved it, which baseline and evidence were considered, which dissent or uncertainty was recorded, and how the decision would be revisited after change.
Failure indicators	Unowned AI risk; approval by a role without authority; fleet or mission scope omitted; temporary exception without expiry; decision record cannot be linked to the deployed baseline.
Notably Absent	No airworthiness, spaceworthiness, certification, legal-compliance or regulator-acceptance conclusion; no universal threshold or approved design is asserted.

D9-CTL-04 - CYBER-PHYSICAL ATTACK DETECTION

Illustrative implementation starting point (non-prescriptive)	Use layered cyber-physical anomaly detection across sensors, commands and AI behaviour where material. Detection thresholds and plausibility criteria must be derived, justified and validated for the relevant platform and operating environment; no universal statistical threshold is prescribed.
Validation status	ODA3 technical sector interpretation; organisation-specific engineering and independent specialist review remain required.
System class	physical
Aviation applicability	Relevant to perception, UAS/AAM, autonomous taxi, inspection, robotics and any AI linked to aircraft or ground equipment. Evaluate sensor diversity, environment, timing and independent physical safeguards.
Space applicability	Relevant to spacecraft autonomy, launch, robotics, rendezvous/proximity operations and physical ground equipment. Consider communications delay, inaccessible assets and environmental extremes.
Aerospace interpretation	Define physical operating boundaries, sensor plausibility checks, actuator/command limits, fail-safe or fail-operational behaviour, human intervention and independent safety constraints.
Minimum evidence	Control owner and applicability decision; system/configuration baseline; physical-specific design evidence; dated operating or test record; open exceptions and residual-risk approval.
Enhanced evidence	Independent review; representative degraded/adversarial test; corroborating telemetry; supplier evidence; change-impact and common-mode analysis where material.
Assessment consideration	Test representative environmental, sensor-degradation,

Illustrative implementation starting point (non-prescriptive)	Use layered cyber-physical anomaly detection across sensors, commands and AI behaviour where material. Detection thresholds and plausibility criteria must be derived, justified and validated for the relevant platform and operating environment; no universal statistical threshold is prescribed.
	spoofing, timing, resource and actuator-boundary scenarios on suitable simulation, hardware-in-the-loop or physical test assets.
Failure indicators	Single sensor drives consequential action without plausibility check; model is validated only in laboratory conditions; override path is unavailable under fault; unsafe command persists after confidence collapses.
Notably Absent	No airworthiness, spaceworthiness, certification, legal-compliance or regulator-acceptance conclusion; no universal threshold or approved design is asserted.

D9-CTL-05 - PHYSICAL ENVIRONMENT INTEGRITY MONITORING

Illustrative implementation starting point (non-prescriptive)	Monitor sensor health, calibration, plausibility, environmental limits and redundancy as applicable. Define degraded-sensor behaviour and authority reduction from system/hazard analysis, with thresholds justified for the specific platform rather than inherited as universal values.
Validation status	ODA3 technical sector interpretation; organisation-specific engineering and independent specialist review remain required.
System class	operations
Aviation applicability	Monitor by aircraft type, fleet, route/environment, software/model baseline and operational context. Ensure alerts integrate with operations, maintenance, safety and security processes without creating unsafe workload.
Space applicability	Monitor by spacecraft, constellation, mission phase, payload mode, ground station and command baseline. Design for intermittent downlink and delayed forensic access.
Aerospace interpretation	Define baseline behaviour, telemetry coverage, time synchronisation, thresholds, alert owners, operational response and evidence retention. Detect model/data drift, provider changes, policy bypass and control degradation.
Minimum evidence	Control owner and applicability decision; system/configuration baseline; operations-specific design evidence; dated operating or test record; open exceptions and residual-risk approval.
Enhanced evidence	Independent review; representative degraded/adversarial test; corroborating telemetry; supplier evidence; change-impact and common-mode analysis where material.
Assessment consideration	Trace a production alert or simulated threshold crossing from telemetry through triage, operational decision, corrective action and closure. Confirm gaps in downlink or logging are visible rather than silently ignored.
Failure indicators	No per-baseline monitoring; time stamps cannot be correlated; fleet averages hide one configuration; alert has no operational owner; provider model changes without detection; log retention is shorter than investigation need.

Illustrative implementation starting point (non-prescriptive)	Monitor sensor health, calibration, plausibility, environmental limits and redundancy as applicable. Define degraded-sensor behaviour and authority reduction from system/hazard analysis, with thresholds justified for the specific platform rather than inherited as universal values.
Notably Absent	No airworthiness, spaceworthiness, certification, legal-compliance or regulator-acceptance conclusion; no universal threshold or approved design is asserted.

D9-CTL-06 - ACTUATOR COMMAND VERIFICATION

Illustrative implementation starting point (non-prescriptive)	Verify consequential actuator or command outputs against approved bounds, state/sequence plausibility, rate/resource limits and authority constraints. Independence and design-assurance rigor must be proportionate to the applicable safety/mission process; GAISSF does not prescribe a universal SIL, DAL or certification level.
Validation status	ODA3 technical sector interpretation; organisation-specific engineering and independent specialist review remain required.
System class	physical
Aviation applicability	Relevant to perception, UAS/AAM, autonomous taxi, inspection, robotics and any AI linked to aircraft or ground equipment. Evaluate sensor diversity, environment, timing and independent physical safeguards.
Space applicability	Relevant to spacecraft autonomy, launch, robotics, rendezvous/proximity operations and physical ground equipment. Consider communications delay, inaccessible assets and environmental extremes.
Aerospace interpretation	Define physical operating boundaries, sensor plausibility checks, actuator/command limits, fail-safe or fail-operational behaviour, human intervention and independent safety constraints.
Minimum evidence	Control owner and applicability decision; system/configuration baseline; physical-specific design evidence; dated operating or test record; open exceptions and residual-risk approval.
Enhanced evidence	Independent review; representative degraded/adversarial test; corroborating telemetry; supplier evidence; change-impact and common-mode analysis where material.
Assessment consideration	Test representative environmental, sensor-degradation, spoofing, timing, resource and actuator-boundary scenarios on suitable simulation, hardware-in-the-loop or physical test assets.
Failure indicators	Single sensor drives consequential action without plausibility check; model is validated only in laboratory conditions; override path is unavailable under fault; unsafe command persists after confidence collapses.
Notably Absent	No airworthiness, spaceworthiness, certification, legal-compliance or regulator-acceptance conclusion; no universal threshold or approved design is asserted.

D9-CTL-07 - PHYSICAL INCIDENT EVIDENCE PRESERVATION

Illustrative implementation starting point (non-prescriptive)	Preserve the sensor, model, command, safety-monitor, override and health records needed to reconstruct relevant physical-AI events. Buffer duration, retention, integrity controls and event-freeze behavior must be set from safety, mission, legal, investigation and architecture needs; no universal retention duration is prescribed.
Validation status	ODA3 technical sector interpretation; organisation-specific engineering and independent specialist review remain required.
System class	incident
Aviation applicability	Coordinate cyber incident response with safety, maintenance, operational control, continued airworthiness and regulator/customer notification processes as applicable. Containment must not create a more hazardous condition.
Space applicability	Coordinate mission anomaly, cyber incident, range/launch safety and ground-segment response. Preserve command history and account for limited access to the affected asset.
Aerospace interpretation	Define AI-specific detection, classification, containment, rollback, operational restriction, evidence preservation, stakeholder communication and return-to-service criteria. Include near misses and common-mode fleet/constellation risk.
Minimum evidence	Control owner and applicability decision; system/configuration baseline; incident-specific design evidence; dated operating or test record; open exceptions and residual-risk approval.
Enhanced evidence	Independent review; representative degraded/adversarial test; corroborating telemetry; supplier evidence; change-impact and common-mode analysis where material.
Assessment consideration	Exercise a scenario involving unsafe output or compromised update. Verify authority to inhibit or roll back, evidence preservation, downstream notification and objective recovery criteria.
Failure indicators	Security team disables a function without safety/mission assessment; no model/data quarantine path; copied artefacts remain in use; return to service is based on elapsed time rather than verified criteria.
Notably Absent	No airworthiness, spaceworthiness, certification, legal-compliance or regulator-acceptance conclusion; no universal threshold or approved design is asserted.

6. Conceptual Reference Architectures

RA-01 - Onboard aviation inference

Conceptual flow	Sensors -> deterministic pre-processing -> inference partition -> independent plausibility/safety monitor -> advisory or bounded actuator interface. Separate maintenance/loading interface and signed configuration store.
Trust boundaries	Aircraft data bus, maintenance interface, update loader, crew display and actuator boundary.

Conceptual flow	Sensors -> deterministic pre-processing -> inference partition -> independent plausibility/safety monitor -> advisory or bounded actuator interface. Separate maintenance/loading interface and signed configuration store.
Principal safeguards	Independent monitor, hard command limits, crew inhibit/override, safe degradation and configuration verification.
Evidence points	Resource saturation, sensor disagreement, monitor intervention, command rejection and exact loaded baseline.
Status	Informative conceptual pattern; not an approved design.

RA-02 - Airline/airport ground AI

Conceptual flow	Operational feeds and enterprise data -> controlled data platform -> model service -> workflow application -> dispatcher/airport/maintenance approval. Aircraft connectivity remains brokered and least-privileged.
Trust boundaries	Cloud/provider, airline network, airport partner, mobile endpoint and aircraft/ground interface.
Principal safeguards	Human approval for consequential action, source validation, network segmentation, provider change detection and rollback.
Evidence points	Recommendation provenance, user decision, provider/model version, data freshness and downstream action.
Status	Informative conceptual pattern; not an approved design.

RA-03 - Maintenance and engineering GenAI

Conceptual flow	Controlled technical-publication repository -> retrieval/index pipeline -> isolated model service -> answer with citations -> licensed engineer/maintainer verification -> controlled work record.
Trust boundaries	External documents, model provider, retrieval index, engineering workspace and maintenance execution system.
Principal safeguards	No direct work-order authority; controlled-source allowlist; prompt/retrieval injection defence; export-data boundary; citation and revision checks.
Evidence points	Prompt, retrieved passages, model/version, cited revision, reviewer decision and rejected answer.
Status	Informative conceptual pattern; not an approved design.

RA-04 - Autonomous UAS/AAM

Conceptual flow	Multi-modal sensors and navigation -> perception/fusion -> planner -> independent constraints/geofence -> flight-control interface; remote operator and lost-link manager remain separate.
Trust boundaries	Air vehicle, command link, remote station, navigation service, payload and fleet service.
Principal safeguards	Action envelope, lost-link state machine, geofence, independent collision safeguard, operator inhibit and safe landing/hold logic.
Evidence points	Sensor health, OOD/confidence, planner commands, constraint rejections, link state and operator

Conceptual flow	Multi-modal sensors and navigation -> perception/fusion -> planner -> independent constraints/geofence -> flight-control interface; remote operator and lost-link manager remain separate.
	interventions.
Status	Informative conceptual pattern; not an approved design.

RA-05 - Satellite onboard autonomy

Conceptual flow	Telemetry/state estimator -> fault/perception/planning model -> bounded command sequencer -> independent spacecraft constraints -> subsystem commands; ground uploads signed policy and baseline.
Trust boundaries	Onboard partitions, payload, TT&C, ground mission control and update path.
Principal safeguards	Pre-authorised action set, resource/thermal constraints, safe-mode dominance, command authentication, delayed intervention assumptions.
Evidence points	Decision rationale, state estimate, command sequence, constraint result, resource margin and downlink gaps.
Status	Informative conceptual pattern; not an approved design.

RA-06 - Satellite ground-segment AI

Conceptual flow	Ground stations and telemetry ingest -> time correlation/data quality -> analytics/model services -> mission/SOC workflow -> authorised command-generation process with independent approval.
Trust boundaries	Ground station, mission network, cloud/provider, partner data, command system and remote support.
Principal safeguards	Separation of analytics from command authority, dual approval for material commands, privileged access monitoring and immutable command history.
Evidence points	Input provenance, clock health, analyst/model recommendation, approvals, command disposition and provider changes.
Status	Informative conceptual pattern; not an approved design.

RA-07 - Launch operations AI

Conceptual flow	Vehicle/ground telemetry -> synchronised stream processing -> anomaly model -> launch-team display; launch authority and deterministic redline/abort systems remain independent.
Trust boundaries	Vehicle, pad/ground systems, range data, vendor feeds and countdown network.
Principal safeguards	AI cannot override redlines or launch authority; time-sync monitoring; one-way or controlled interfaces; tested hold/continue presentation.
Evidence points	Telemetry timestamps, feature values, anomaly score, displayed alert, human disposition and countdown state.
Status	Informative conceptual pattern; not an approved design.

RA-08 - Federated aerospace AI platform

Conceptual flow	Business units retain local data/model environments; central governance, identity, registry, policy and monitoring services enforce shared minimum controls while sector authorities approve local releases.
Trust boundaries	Organisational, cloud, supplier, jurisdiction and mission/programme boundaries.
Principal safeguards	Tenant isolation, scoped identities, signed registry, local approval, data residency and cross-fleet common-dependency analysis.
Evidence points	Tenant/model inventory, policy decisions, registry events, shared-service changes and cross-domain incidents.
Status	Informative conceptual pattern; not an approved design.

RA-09 - Disconnected/intermittently connected system

Conceptual flow	Local sensor/data store -> onboard/edge model -> local monitor and bounded action -> durable event buffer -> authenticated delayed synchronisation and ground reassessment.
Trust boundaries	Connected/disconnected state, local/ground authority, delayed update and evidence transfer boundary.
Principal safeguards	Offline-safe policy, expiring authority, local rollback, buffer integrity, conflict resolution and no dependence on unavailable cloud safeguards.
Evidence points	Connectivity state, local decisions, buffered logs, authority expiry, sync conflicts and post-link reconciliation.
Status	Informative conceptual pattern; not an approved design.

7. Aerospace Use Cases

UC-01 - Predictive aircraft maintenance

Environment	Aviation/MRO
Mode	Advisory
Data Inputs	Sensor histories, removal records, operating cycles and maintenance findings
Primary Threats	Training-data coverage gaps; maintenance-record manipulation; fleet/configuration mismatch
Failure Modes	Missed impending failure; excessive removals; maintenance action on wrong configuration
Critical Controls	D1.2 D1.3 D2.1 D2.2 D6.1 D6.3 D7.3
Evidence	System definition; data/model provenance; risk assessment; control evidence; validation report; human-oversight record; deployment baseline; monitoring and incident records
Validation	Prognostic calibration by component/fleet; back-testing on unseen removals; false-negative review
Constraints	Long component lives; censored failure data; fleet modifications
Monitoring	Calibration, missed-event rate, alert-to-disposition time, override reason
Incident Scenario	A low-confidence alert is treated as mandatory and

Environment	Aviation/MRO
	causes removal of a serviceable component.
Residual Risk	Organisation-specific acceptance required; this record does not establish approval.
Notably Absent	No assertion that the use case is approved, certified or suitable for a particular platform.

UC-02 - Engine-health monitoring

Environment	Aviation
Mode	Decision support
Data Inputs	Engine sensor streams, flight phase, environment and maintenance configuration
Primary Threats	Sensor spoofing; time-alignment error; configuration drift; common-mode model defect
Failure Modes	Undetected degradation or false shutdown/maintenance recommendation
Critical Controls	D2.1 D3.1 D6.3 D7.2 D7.3 D8.2 D9.1
Evidence	System definition; data/model provenance; risk assessment; control evidence; validation report; human-oversight record; deployment baseline; monitoring and incident records
Validation	Cross-engine and physics-informed plausibility tests; phase-specific thresholds
Constraints	High-frequency data; sensor replacement; environmental extremes
Monitoring	Residual trend, sensor disagreement, confidence collapse, per-tail anomalies
Incident Scenario	A sensor bias propagates across the fleet model and suppresses a genuine deterioration trend.
Residual Risk	Organisation-specific acceptance required; this record does not establish approval.
Notably Absent	No assertion that the use case is approved, certified or suitable for a particular platform.

UC-03 - Avionics anomaly detection

Environment	Aviation
Mode	Monitoring
Data Inputs	Bus messages, equipment status, fault logs and configuration
Primary Threats	Adversarial event sequence; logging gaps; benign novelty; compromised maintenance interface
Failure Modes	Missed cyber/safety anomaly or disruptive false alarm
Critical Controls	D3.1 D4.1 D6.3 D7.2 D7.3 D8.1
Evidence	System definition; data/model provenance; risk assessment; control evidence; validation report; human-oversight record; deployment baseline; monitoring and incident records
Validation	Replay known faults and novel sequences; validate timing and false-alarm burden
Constraints	Deterministic buses; certification baseline; limited onboard compute
Monitoring	Coverage by bus/message; alert precision; unclassified anomaly backlog
Incident Scenario	An unusual but valid maintenance state is classified as attack and triggers an unsafe operational restriction.

Environment	Aviation
Residual Risk	Organisation-specific acceptance required; this record does not establish approval.
Notably Absent	No assertion that the use case is approved, certified or suitable for a particular platform.

UC-04 - Autonomous taxiing

Environment	Aviation/Physical AI
Mode	Autonomous
Data Inputs	Cameras, lidar/radar, airport map, route clearance and vehicle state
Primary Threats	Visual adversarial conditions; map corruption; object occlusion; command-link compromise
Failure Modes	Collision, runway incursion or loss of control
Critical Controls	D3.1 D3.5 D6.3 D7.4 D8.2 D9.1 D9.2 D9.3
Evidence	System definition; data/model provenance; risk assessment; control evidence; validation report; human-oversight record; deployment baseline; monitoring and incident records
Validation	Closed-course, HIL and edge-weather testing; stop-distance and override verification
Constraints	Lighting/weather; airport markings; braking and latency limits
Monitoring	Obstacle detection margin, override latency, boundary violations, safe-stop success
Incident Scenario	A temporary work vehicle is not represented in training data and the system fails to stop.
Residual Risk	Organisation-specific acceptance required; this record does not establish approval.
Notably Absent	No assertion that the use case is approved, certified or suitable for a particular platform.

UC-05 - Flight-path optimisation

Environment	Aviation
Mode	Decision support
Data Inputs	Weather, NOTAMs, aircraft performance, airspace and fuel state
Primary Threats	Stale weather/NOTAM; objective manipulation; constraint omission
Failure Modes	Unsafe or non-compliant route; fuel or delay penalty
Critical Controls	D1.4 D2.1 D5.1 D6.1 D6.3 D7.3
Evidence	System definition; data/model provenance; risk assessment; control evidence; validation report; human-oversight record; deployment baseline; monitoring and incident records
Validation	Constraint satisfaction and conservative fallback across disrupted scenarios
Constraints	Rapidly changing conditions; dispatcher/crew authority
Monitoring	Rejected recommendations, stale-source detection, constraint violations
Incident Scenario	Optimisation omits a temporary restriction and recommends an unavailable route.
Residual Risk	Organisation-specific acceptance required; this record does not establish approval.
Notably Absent	No assertion that the use case is approved, certified or

Environment	Aviation
	suitable for a particular platform.

UC-06 - Air-traffic conflict decision support

Environment	ATM
Mode	Decision support
Data Inputs	Surveillance tracks, flight plans, sector state and controller inputs
Primary Threats	Track corruption; latency; automation bias; overload-induced alert suppression
Failure Modes	Loss of separation or unnecessary intervention
Critical Controls	D2.2 D5.2 D6.3 D7.3 D7.4 D9.2
Evidence	System definition; data/model provenance; risk assessment; control evidence; validation report; human-oversight record; deployment baseline; monitoring and incident records
Validation	Rare-conflict simulation, latency injection and controller-in-the-loop evaluation
Constraints	Strict timing; human workload; national ATM rules
Monitoring	Alert timeliness, controller acceptance, nuisance alert rate, missed conflict
Incident Scenario	A late but confident recommendation causes the controller to defer an earlier safe resolution.
Residual Risk	Organisation-specific acceptance required; this record does not establish approval.
Notably Absent	No assertion that the use case is approved, certified or suitable for a particular platform.

UC-07 - UAS perception and obstacle avoidance

Environment	UAS
Mode	Autonomous
Data Inputs	EO/IR, radar/lidar, navigation and terrain data
Primary Threats	Adversarial patterns; GNSS spoofing; sensor occlusion; edge-compute saturation
Failure Modes	Collision, airspace excursion or mission loss
Critical Controls	D3.5 D6.3 D7.4 D8.2 D9.1 D9.2 D9.4
Evidence	System definition; data/model provenance; risk assessment; control evidence; validation report; human-oversight record; deployment baseline; monitoring and incident records
Validation	Flight test, HIL, sensor-degradation and lost-link scenarios
Constraints	Size/weight/power; weather; remote intervention latency
Monitoring	OOD detection, sensor disagreement, geofence events, lost-link response
Incident Scenario	Sun glare blinds the primary camera while compute saturation delays the fallback sensor.
Residual Risk	Organisation-specific acceptance required; this record does not establish approval.
Notably Absent	No assertion that the use case is approved, certified or suitable for a particular platform.

UC-08 - Advanced-air-mobility fleet operations

Environment	AAM
Mode	Automation support
Data Inputs	Vehicle status, vertiport capacity, weather, crew and maintenance state
Primary Threats	Fleet-wide model defect; scheduling objective conflict; degraded communications
Failure Modes	Unsafe dispatch, congestion or stranded aircraft
Critical Controls	D1.3 D2.4 D6.3 D7.3 D8.4 D9.5
Evidence	System definition; data/model provenance; risk assessment; control evidence; validation report; human-oversight record; deployment baseline; monitoring and incident records
Validation	Fleet simulation with weather, maintenance and vertiport disruptions
Constraints	Emerging operations; high automation; heterogeneous suppliers
Monitoring	Dispatch overrides, common-mode alerts, recovery time, constraint breaches
Incident Scenario	A shared dispatch model underestimates battery reserve during an unexpected weather shift.
Residual Risk	Organisation-specific acceptance required; this record does not establish approval.
Notably Absent	No assertion that the use case is approved, certified or suitable for a particular platform.

UC-09 - Maintenance-manual assistant

Environment	MRO/GenAI
Mode	Generative assistant
Data Inputs	Controlled manuals, service bulletins, task cards and approved engineering data
Primary Threats	Hallucination; indirect prompt injection; obsolete revision; controlled-data leakage
Failure Modes	Incorrect maintenance instruction or unauthorised disclosure
Critical Controls	D2.1 D3.2 D5.1 D5.2 D6.2 D7.2 D8.1
Evidence	System definition; data/model provenance; risk assessment; control evidence; validation report; human-oversight record; deployment baseline; monitoring and incident records
Validation	Source-grounded answer tests; revision control; adversarial document injection
Constraints	Controlled publications; licensed personnel; offline hangar use
Monitoring	Citation coverage, obsolete-source use, human rejection, leakage events
Incident Scenario	A malicious note embedded in a PDF instructs the assistant to omit a mandatory inspection step.
Residual Risk	Organisation-specific acceptance required; this record does not establish approval.
Notably Absent	No assertion that the use case is approved, certified or suitable for a particular platform.

UC-10 - Engineering design copilot

Environment	Engineering/GenAI
Mode	Generative assistant
Data Inputs	Requirements, design models, standards, prior analyses and code
Primary Threats	IP leakage; unsupported design claim; retrieval poisoning; model-provider change
Failure Modes	Design defect, rework or export-control breach
Critical Controls	D1.4 D2.1 D2.4 D3.4 D5.2 D6.4
Evidence	System definition; data/model provenance; risk assessment; control evidence; validation report; human-oversight record; deployment baseline; monitoring and incident records
Validation	Benchmark on traceable engineering tasks; independent review; data-boundary tests
Constraints	Proprietary/controlled data; tool qualification; reproducibility
Monitoring	Unsupported citation rate, review rejection, provider-version drift
Incident Scenario	A generated design rationale cites a non-existent constraint and is copied into a review package.
Residual Risk	Organisation-specific acceptance required; this record does not establish approval.
Notably Absent	No assertion that the use case is approved, certified or suitable for a particular platform.

UC-11 - Safety-adjacent code generation

Environment	Engineering/GenAI
Mode	Generative assistant
Data Inputs	Requirements, APIs, coding standards and test harnesses
Primary Threats	Insecure code; hidden dependency; licence contamination; false test confidence
Failure Modes	Regression, vulnerability or invalid assurance evidence
Critical Controls	D2.4 D3.4 D3.5 D6.2 D6.4
Evidence	System definition; data/model provenance; risk assessment; control evidence; validation report; human-oversight record; deployment baseline; monitoring and incident records
Validation	Static/dynamic analysis, independent review, traceability and negative testing
Constraints	Assurance objectives; language/tool constraints; reproducible build
Monitoring	Generated-code defect density, review escapes, dependency policy violations
Incident Scenario	Generated error handling passes nominal tests but fails under resource exhaustion.
Residual Risk	Organisation-specific acceptance required; this record does not establish approval.
Notably Absent	No assertion that the use case is approved, certified or suitable for a particular platform.

UC-12 - Non-destructive inspection

Environment	Manufacturing/MRO
Mode	Perception

Environment	Manufacturing/MRO
Data Inputs	Radiography, ultrasound, thermography or visual imagery plus part metadata
Primary Threats	Image manipulation; calibration drift; rare defect underrepresentation
Failure Modes	Defect escape or unnecessary scrappage
Critical Controls	D2.2 D6.1 D6.3 D7.3 D9.2
Evidence	System definition; data/model provenance; risk assessment; control evidence; validation report; human-oversight record; deployment baseline; monitoring and incident records
Validation	Blind defect sets, equipment/calibration variation and human-reader comparison
Constraints	Part geometry; equipment variation; low defect prevalence
Monitoring	Probability of detection, false call rate, calibration drift, inspector overrides
Incident Scenario	A new composite lay-up creates image features outside the validated defect set.
Residual Risk	Organisation-specific acceptance required; this record does not establish approval.
Notably Absent	No assertion that the use case is approved, certified or suitable for a particular platform.

UC-13 - Manufacturing quality inspection

Environment	Manufacturing
Mode	Perception
Data Inputs	Line images, process parameters, part serial and configuration
Primary Threats	Camera shift; lighting drift; label error; bypass of rejected-part workflow
Failure Modes	Common-mode manufacturing escape or production disruption
Critical Controls	D2.2 D3.1 D6.3 D7.3 D8.2
Evidence	System definition; data/model provenance; risk assessment; control evidence; validation report; human-oversight record; deployment baseline; monitoring and incident records
Validation	Golden-part, defect challenge, line-change and fail-safe tests
Constraints	Production speed; rework; configuration changes
Monitoring	Escape rate, false reject, camera health, manual bypass
Incident Scenario	A camera replacement changes colour response and increases undetected surface defects.
Residual Risk	Organisation-specific acceptance required; this record does not establish approval.
Notably Absent	No assertion that the use case is approved, certified or suitable for a particular platform.

UC-14 - Digital-twin anomaly analysis

Environment	Engineering/Operations
Mode	Decision support
Data Inputs	Simulation state, telemetry, configuration and environment
Primary Threats	Model-form error; stale configuration; simulator-to-

Environment	Engineering/Operations
	reality gap
Failure Modes	Incorrect maintenance or mission decision
Critical Controls	D1.3 D2.1 D6.1 D6.3 D7.3
Evidence	System definition; data/model provenance; risk assessment; control evidence; validation report; human-oversight record; deployment baseline; monitoring and incident records
Validation	Validate twin fidelity by regime; quantify residuals and extrapolation limits
Constraints	Multiple fidelity levels; configuration synchronisation
Monitoring	Twin residuals, out-of-domain use, baseline mismatch
Incident Scenario	The twin remains on a pre-modification mass model and misattributes a performance anomaly.
Residual Risk	Organisation-specific acceptance required; this record does not establish approval.
Notably Absent	No assertion that the use case is approved, certified or suitable for a particular platform.

UC-15 - Satellite collision-risk decision support

Environment	Space
Mode	Decision support
Data Inputs	Tracking observations, conjunction messages, covariance, manoeuvre constraints
Primary Threats	Tracking-data poisoning; covariance error; latency; space-weather uncertainty
Failure Modes	Collision, unnecessary manoeuvre or mission-life loss
Critical Controls	D2.1 D6.3 D7.3 D8.2 D9.2
Evidence	System definition; data/model provenance; risk assessment; control evidence; validation report; human-oversight record; deployment baseline; monitoring and incident records
Validation	Historical conjunction replay, uncertainty propagation and delayed-data tests
Constraints	Sparse observations; operator coordination; fuel limits
Monitoring	Data age, covariance quality, recommendation reversals, decision latency
Incident Scenario	A delayed update narrows uncertainty after the manoeuvre decision window has closed.
Residual Risk	Organisation-specific acceptance required; this record does not establish approval.
Notably Absent	No assertion that the use case is approved, certified or suitable for a particular platform.

UC-16 - Spacecraft fault detection, isolation and recovery

Environment	Space/Physical AI
Mode	Autonomous
Data Inputs	Telemetry, subsystem state, fault rules and onboard model
Primary Threats	Novel fault; correlated sensor failure; unsafe autonomous recovery
Failure Modes	Loss of spacecraft, payload or safe mode
Critical Controls	D3.5 D6.3 D7.4 D8.2 D9.1 D9.3
Evidence	System definition; data/model provenance; risk assessment; control evidence; validation report; human-

Environment	Space/Physical AI
	oversight record; deployment baseline; monitoring and incident records
Validation	Fault injection, HIL, link-loss and safe-mode transition tests
Constraints	Limited compute; delayed ground intervention; radiation effects
Monitoring	Isolation accuracy, safe-mode success, recovery oscillation, command conflicts
Incident Scenario	A correlated sensor fault causes the model to isolate the healthy subsystem and repeatedly reset.
Residual Risk	Organisation-specific acceptance required; this record does not establish approval.
Notably Absent	No assertion that the use case is approved, certified or suitable for a particular platform.

UC-17 - Autonomous satellite operations

Environment	Space
Mode	Autonomous
Data Inputs	Mission plan, spacecraft state, constraints, ephemeris and resources
Primary Threats	Goal conflict; command-sequence error; compromised planning input
Failure Modes	Mission degradation or unsafe proximity/attitude state
Critical Controls	D3.1 D3.5 D7.4 D8.2 D9.1 D9.3
Evidence	System definition; data/model provenance; risk assessment; control evidence; validation report; human-oversight record; deployment baseline; monitoring and incident records
Validation	Bounded-planning, resource, communications-loss and independent-monitor tests
Constraints	Long latency; limited energy; infrequent updates
Monitoring	Constraint violations, plan aborts, monitor interventions, resource margin
Incident Scenario	An optimiser consumes thermal margin to meet a science objective and blocks a later safe-mode action.
Residual Risk	Organisation-specific acceptance required; this record does not establish approval.
Notably Absent	No assertion that the use case is approved, certified or suitable for a particular platform.

UC-18 - Onboard image analysis

Environment	Space
Mode	Perception/analytics
Data Inputs	Payload imagery, calibration and mission priorities
Primary Threats	Radiation-induced corruption; domain shift; compression artefact; priority manipulation
Failure Modes	Missed event or wasted downlink capacity
Critical Controls	D2.2 D6.3 D7.3 D9.2
Evidence	System definition; data/model provenance; risk assessment; control evidence; validation report; human-oversight record; deployment baseline; monitoring and incident records
Validation	Orbital/seasonal variation, corruption and compression tests

Environment	Space
Constraints	Power, compute and downlink limits
Monitoring	Confidence by geography, corruption detection, downlink selection quality
Incident Scenario	Cloud patterns outside the training distribution are misclassified as priority targets.
Residual Risk	Organisation-specific acceptance required; this record does not establish approval.
Notably Absent	No assertion that the use case is approved, certified or suitable for a particular platform.

UC-19 - Launch-countdown anomaly detection

Environment	Launch
Mode	Decision support
Data Inputs	Vehicle telemetry, ground systems, procedures and environmental data
Primary Threats	Timing misalignment; false alarm; compromised ground feed; rare-event scarcity
Failure Modes	Unsafe continuation or unnecessary hold/abort
Critical Controls	D3.1 D6.3 D7.2 D8.1 D9.2
Evidence	System definition; data/model provenance; risk assessment; control evidence; validation report; human-oversight record; deployment baseline; monitoring and incident records
Validation	Countdown replay, fault injection, time-sync failure and operator-in-loop tests
Constraints	One-time mission; strict timing; authority hierarchy
Monitoring	Alert lead time, nuisance holds, time-sync health, unclassified anomalies
Incident Scenario	A clock offset combines unrelated signals into a false high-confidence anomaly.
Residual Risk	Organisation-specific acceptance required; this record does not establish approval.
Notably Absent	No assertion that the use case is approved, certified or suitable for a particular platform.

UC-20 - Mission planning

Environment	Space/Aviation
Mode	Decision support
Data Inputs	Objectives, constraints, resources, weather/orbit and asset state
Primary Threats	Incorrect constraint; objective gaming; stale state; planner non-determinism
Failure Modes	Mission failure, unsafe allocation or resource exhaustion
Critical Controls	D1.3 D3.5 D6.1 D7.3
Evidence	System definition; data/model provenance; risk assessment; control evidence; validation report; human-oversight record; deployment baseline; monitoring and incident records
Validation	Constraint-based test oracle, worst-case resource and contingency scenarios
Constraints	Complex dependencies; changing mission priorities
Monitoring	Constraint violations, plan churn, human modifications, reserve margin
Incident Scenario	The planner satisfies mission score by scheduling an

Environment	Space/Aviation
	activity inside a communications blackout.
Residual Risk	Organisation-specific acceptance required; this record does not establish approval.
Notably Absent	No assertion that the use case is approved, certified or suitable for a particular platform.

UC-21 - Ground-station cybersecurity monitoring

Environment	Space ground segment
Mode	Security monitoring
Data Inputs	Network, identity, command, telemetry and provider logs
Primary Threats	Log poisoning; alert evasion; compromised model; excessive privilege
Failure Modes	Undetected command-path compromise or denial of mission service
Critical Controls	D3.2 D4.1 D6.2 D7.2 D8.1 D8.2
Evidence	System definition; data/model provenance; risk assessment; control evidence; validation report; human-oversight record; deployment baseline; monitoring and incident records
Validation	Attack replay, missing-log, identity and command-correlation tests
Constraints	Mixed legacy/cloud systems; 24x7 mission operations
Monitoring	Command anomaly coverage, log loss, triage latency, false positives
Incident Scenario	A compromised support account issues valid-looking commands outside its normal mission window.
Residual Risk	Organisation-specific acceptance required; this record does not establish approval.
Notably Absent	No assertion that the use case is approved, certified or suitable for a particular platform.

UC-22 - Space-weather prediction

Environment	Space
Mode	Forecasting
Data Inputs	Solar observations, models and historical impacts
Primary Threats	Data outage; calibration shift; overconfidence in rare events
Failure Modes	Incorrect mission protection or scheduling decision
Critical Controls	D2.2 D6.1 D6.3 D7.3
Evidence	System definition; data/model provenance; risk assessment; control evidence; validation report; human-oversight record; deployment baseline; monitoring and incident records
Validation	Hindcast extreme events, uncertainty calibration and data-outage tests
Constraints	Rare severe events; evolving instruments
Monitoring	Forecast calibration, missed severe event, data-source health
Incident Scenario	A model trained on one solar cycle understates uncertainty during a different regime.
Residual Risk	Organisation-specific acceptance required; this record does not establish approval.
Notably Absent	No assertion that the use case is approved, certified or

Environment	Space
	suitable for a particular platform.

UC-23 - Remote-sensing analytics

Environment	Space/Analytics
Mode	Analytics
Data Inputs	Imagery, geolocation, labels and customer tasking
Primary Threats	Dataset bias; geolocation error; misuse; adversarial camouflage
Failure Modes	Incorrect analytical conclusion or rights/geopolitical harm
Critical Controls	D1.4 D2.1 D5.1 D5.3 D6.3
Evidence	System definition; data/model provenance; risk assessment; control evidence; validation report; human-oversight record; deployment baseline; monitoring and incident records
Validation	Geography/season/sensor stratification, geolocation and misuse evaluation
Constraints	Dual-use context; licensing; changing sensors
Monitoring	Performance by geography, geolocation error, misuse flags
Incident Scenario	A model performs poorly in one region but global aggregate metrics conceal the gap.
Residual Risk	Organisation-specific acceptance required; this record does not establish approval.
Notably Absent	No assertion that the use case is approved, certified or suitable for a particular platform.

UC-24 - Crew decision support

Environment	Aviation/Human spaceflight
Mode	Decision support
Data Inputs	Vehicle state, procedures, environment and mission context
Primary Threats	Automation bias; alert overload; stale procedure; misleading explanation
Failure Modes	Delayed or incorrect crew action
Critical Controls	D5.2 D6.3 D7.3 D9.3
Evidence	System definition; data/model provenance; risk assessment; control evidence; validation report; human-oversight record; deployment baseline; monitoring and incident records
Validation	Human-in-the-loop abnormal/emergency scenarios and explanation tests
Constraints	High workload; time pressure; variable crew expertise
Monitoring	Response time, reliance, disagreement, alert burden
Incident Scenario	A confident but poorly sourced recommendation causes the crew to delay a checklist action.
Residual Risk	Organisation-specific acceptance required; this record does not establish approval.
Notably Absent	No assertion that the use case is approved, certified or suitable for a particular platform.

8. Threat and Failure Register

T-01 - Training-data poisoning

Category	Data/Development
Affected Asset	AI-enabled aerospace system, dependent platform and operational process
Initiating Condition	Threat-specific initiating condition described in the main guide
Likely Consequence	Safety, mission, availability, confidentiality, integrity or assurance impact depending on context
Detectability	Usually low before targeted integrity analysis; improves with trusted acquisition, signed manifests, statistical outlier review and independent source reconciliation.
Control Objectives	Reduce attack opportunity; detect deviations; constrain physical/operational effect; support safe containment and recovery.
Evidence Expectations	Source manifests and hashes; acquisition approvals; poisoning challenge set; label-review record; clean-room rebuild evidence.
Residual Risk	Record undetectable or low-observability cases and the authority accepting them.
Limitations	Detectability and consequence depend on architecture, environment, data access and operational authority.

T-02 - Sensor deception and adversarial input

Category	Physical/Operations
Affected Asset	AI-enabled aerospace system, dependent platform and operational process
Initiating Condition	Threat-specific initiating condition described in the main guide
Likely Consequence	Safety, mission, availability, confidentiality, integrity or assurance impact depending on context
Detectability	High when redundant modalities and physical plausibility checks disagree; low for a standalone sensor or attack that remains physically plausible.
Control Objectives	Reduce attack opportunity; detect deviations; constrain physical/operational effect; support safe containment and recovery.
Evidence Expectations	Sensor-fusion design; spoofing/occlusion tests; plausibility thresholds; raw-sensor retention; monitor and override logs.
Residual Risk	Record undetectable or low-observability cases and the authority accepting them.
Limitations	Detectability and consequence depend on architecture, environment, data access and operational authority.

T-03 - GNSS/PNT spoofing

Category	Physical/Operations
Affected Asset	AI-enabled aerospace system, dependent platform and operational process
Initiating Condition	Threat-specific initiating condition described in the main guide
Likely Consequence	Safety, mission, availability, confidentiality, integrity or assurance impact depending on context

Category	Physical/Operations
Detectability	Moderate to high with multi-source navigation, inertial comparison, signal-quality monitoring and time-consistency checks; otherwise low.
Control Objectives	Reduce attack opportunity; detect deviations; constrain physical/operational effect; support safe containment and recovery.
Evidence Expectations	PNT architecture; spoofing/jamming tests; cross-check logs; degraded-navigation procedure; event reconstruction data.
Residual Risk	Record undetectable or low-observability cases and the authority accepting them.
Limitations	Detectability and consequence depend on architecture, environment, data access and operational authority.

T-04 - Model backdoor or trojan

Category	Model/Supply chain
Affected Asset	AI-enabled aerospace system, dependent platform and operational process
Initiating Condition	Threat-specific initiating condition described in the main guide
Likely Consequence	Safety, mission, availability, confidentiality, integrity or assurance impact depending on context
Detectability	Low through normal functional testing; improves with provenance, weight scanning, trigger search, behavioural differential testing and reproducible build comparison.
Control Objectives	Reduce attack opportunity; detect deviations; constrain physical/operational effect; support safe containment and recovery.
Evidence Expectations	Model hash and origin; training/build records; backdoor test suite; differential results; supplier attestation and independent review.
Residual Risk	Record undetectable or low-observability cases and the authority accepting them.
Limitations	Detectability and consequence depend on architecture, environment, data access and operational authority.

T-05 - Compromised model or software update

Category	Platform/Change
Affected Asset	AI-enabled aerospace system, dependent platform and operational process
Initiating Condition	Threat-specific initiating condition described in the main guide
Likely Consequence	Safety, mission, availability, confidentiality, integrity or assurance impact depending on context
Detectability	High where signature, provenance and baseline comparison fail; low where a valid signing key or release process is compromised without independent checks.
Control Objectives	Reduce attack opportunity; detect deviations; constrain physical/operational effect; support safe containment and recovery.
Evidence Expectations	Signed manifest; key custody; release approval; reproducible build; staged deployment; rollback and post-install verification.

Category	Platform/Change
Residual Risk	Record undetectable or low-observability cases and the authority accepting them.
Limitations	Detectability and consequence depend on architecture, environment, data access and operational authority.

T-06 - Model or concept drift

Category	Operations
Affected Asset	AI-enabled aerospace system, dependent platform and operational process
Initiating Condition	Threat-specific initiating condition described in the main guide
Likely Consequence	Safety, mission, availability, confidentiality, integrity or assurance impact depending on context
Detectability	Moderate where representative labels or proxy indicators arrive promptly; low where outcomes are delayed, rare or unobservable.
Control Objectives	Reduce attack opportunity; detect deviations; constrain physical/operational effect; support safe containment and recovery.
Evidence Expectations	Monitoring baseline; stratified drift metrics; outcome reconciliation; threshold rationale; retraining/restriction decision records.
Residual Risk	Record undetectable or low-observability cases and the authority accepting them.
Limitations	Detectability and consequence depend on architecture, environment, data access and operational authority.

T-07 - Ground-segment compromise

Category	Platform/Operations
Affected Asset	AI-enabled aerospace system, dependent platform and operational process
Initiating Condition	Threat-specific initiating condition described in the main guide
Likely Consequence	Safety, mission, availability, confidentiality, integrity or assurance impact depending on context
Detectability	High for identity, network or command anomalies with complete correlated logs; reduced by legacy gaps, shared accounts or missing time synchronisation.
Control Objectives	Reduce attack opportunity; detect deviations; constrain physical/operational effect; support safe containment and recovery.
Evidence Expectations	Architecture and identity map; privileged-session logs; command history; network telemetry; incident exercise and containment evidence.
Residual Risk	Record undetectable or low-observability cases and the authority accepting them.
Limitations	Detectability and consequence depend on architecture, environment, data access and operational authority.

T-08 - Automation bias and over-reliance

Category	Human factors
Affected Asset	AI-enabled aerospace system, dependent platform and operational process
Initiating Condition	Threat-specific initiating condition described in the main

Category	Human factors
	guide
Likely Consequence	Safety, mission, availability, confidentiality, integrity or assurance impact depending on context
Detectability	Indirect; detected through disagreement, override, response-time and scenario-study evidence rather than technical telemetry alone.
Control Objectives	Reduce attack opportunity; detect deviations; constrain physical/operational effect; support safe containment and recovery.
Evidence Expectations	Human-factors study; training records; recommendation/decision logs; disagreement analysis; workload and reliance observations.
Residual Risk	Record undetectable or low-observability cases and the authority accepting them.
Limitations	Detectability and consequence depend on architecture, environment, data access and operational authority.

T-09 - Unsafe handover or intervention

Category	Human factors/Physical
Affected Asset	AI-enabled aerospace system, dependent platform and operational process
Initiating Condition	Threat-specific initiating condition described in the main guide
Likely Consequence	Safety, mission, availability, confidentiality, integrity or assurance impact depending on context
Detectability	High in controlled tests measuring detection, comprehension and intervention time; uncertain in untested rare operational states.
Control Objectives	Reduce attack opportunity; detect deviations; constrain physical/operational effect; support safe containment and recovery.
Evidence Expectations	Handover design; timing budget; simulator/HIL results; operator comprehension; override availability and post-event logs.
Residual Risk	Record undetectable or low-observability cases and the authority accepting them.
Limitations	Detectability and consequence depend on architecture, environment, data access and operational authority.

T-10 - Supplier provenance or assurance opacity

Category	Supply chain
Affected Asset	AI-enabled aerospace system, dependent platform and operational process
Initiating Condition	Threat-specific initiating condition described in the main guide
Likely Consequence	Safety, mission, availability, confidentiality, integrity or assurance impact depending on context
Detectability	High when required artefacts are absent, inconsistent or cannot be tied to the delivered version; technical compromise may remain undetected.
Control Objectives	Reduce attack opportunity; detect deviations; constrain physical/operational effect; support safe containment and recovery.
Evidence Expectations	AI/SBOM; model and adapter hashes; dataset

Category	Supply chain
	provenance; subcontractor list; change history; evidence-rights and audit records.
Residual Risk	Record undetectable or low-observability cases and the authority accepting them.
Limitations	Detectability and consequence depend on architecture, environment, data access and operational authority.

T-11 - Resource exhaustion

Category	Platform/Operations
Affected Asset	AI-enabled aerospace system, dependent platform and operational process
Initiating Condition	Threat-specific initiating condition described in the main guide
Likely Consequence	Safety, mission, availability, confidentiality, integrity or assurance impact depending on context
Detectability	High with resource telemetry, quotas and stress tests; low where saturation causes telemetry or safeguards to fail simultaneously.
Control Objectives	Reduce attack opportunity; detect deviations; constrain physical/operational effect; support safe containment and recovery.
Evidence Expectations	Capacity model; load and stress tests; CPU/GPU/memory/power telemetry; degradation logic; safe-failure and recovery results.
Residual Risk	Record undetectable or low-observability cases and the authority accepting them.
Limitations	Detectability and consequence depend on architecture, environment, data access and operational authority.

T-12 - Telemetry and forensic gap

Category	Operations/Incident
Affected Asset	AI-enabled aerospace system, dependent platform and operational process
Initiating Condition	Threat-specific initiating condition described in the main guide
Likely Consequence	Safety, mission, availability, confidentiality, integrity or assurance impact depending on context
Detectability	Directly detectable through coverage and integrity health checks; the underlying incident may remain unknowable once evidence is lost.
Control Objectives	Reduce attack opportunity; detect deviations; constrain physical/operational effect; support safe containment and recovery.
Evidence Expectations	Logging specification; coverage map; clock tests; loss alerts; retention settings; immutable storage and investigation rehearsal.
Residual Risk	Record undetectable or low-observability cases and the authority accepting them.
Limitations	Detectability and consequence depend on architecture, environment, data access and operational authority.

T-13 - Generative hallucination

Category	Output/GenAI
Affected Asset	AI-enabled aerospace system, dependent platform and

Category	Output/GenAI
	operational process
Initiating Condition	Threat-specific initiating condition described in the main guide
Likely Consequence	Safety, mission, availability, confidentiality, integrity or assurance impact depending on context
Detectability	Moderate with source-grounding checks, unsupported-claim detection and qualified review; low for novel technical synthesis without a reliable reference set.
Control Objectives	Reduce attack opportunity; detect deviations; constrain physical/operational effect; support safe containment and recovery.
Evidence Expectations	Citation and retrieval logs; claim-to-source sampling; semantic consistency tests; human verification records; rejected-output analysis.
Residual Risk	Record undetectable or low-observability cases and the authority accepting them.
Limitations	Detectability and consequence depend on architecture, environment, data access and operational authority.

T-14 - Prompt or retrieval injection

Category	GenAI/Platform
Affected Asset	AI-enabled aerospace system, dependent platform and operational process
Initiating Condition	Threat-specific initiating condition described in the main guide
Likely Consequence	Safety, mission, availability, confidentiality, integrity or assurance impact depending on context
Detectability	Moderate to high using adversarial corpora, instruction-boundary monitoring and tool-call review; novel indirect channels remain possible.
Control Objectives	Reduce attack opportunity; detect deviations; constrain physical/operational effect; support safe containment and recovery.
Evidence Expectations	Injection test corpus; retrieval provenance; tool permission logs; policy-enforcement tests; blocked and successful attack traces.
Residual Risk	Record undetectable or low-observability cases and the authority accepting them.
Limitations	Detectability and consequence depend on architecture, environment, data access and operational authority.

T-15 - Fleet or constellation common-mode defect

Category	Systemic/Operations
Affected Asset	AI-enabled aerospace system, dependent platform and operational process
Initiating Condition	Threat-specific initiating condition described in the main guide
Likely Consequence	Safety, mission, availability, confidentiality, integrity or assurance impact depending on context
Detectability	Low before deployment if diversity and common-dependency analysis are weak; high after correlated anomalies appear across identical baselines.
Control Objectives	Reduce attack opportunity; detect deviations; constrain physical/operational effect; support safe containment

Category	Systemic/Operations
	and recovery.
Evidence Expectations	Baseline inventory; common-cause analysis; staged rollout; canary results; cross-fleet telemetry; rollback exercise and impact assessment.
Residual Risk	Record undetectable or low-observability cases and the authority accepting them.
Limitations	Detectability and consequence depend on architecture, environment, data access and operational authority.

9. Evidence Catalogue

E-01 - Governance mandate

Purpose	Provide attributable, current evidence of control design or operation.
Owner	Assigned according to system and lifecycle authority.
Minimum Content	Scope, authority, membership, decision rights, escalation, meeting cadence and executive approval.
Lifecycle Stage	Defined by evidence type and system lifecycle.
Retention	Set from safety, mission, legal, contractual, investigation, evidence-preservation and approved records requirements. GAISSF conformance tiers do not define universal retention periods.
Integrity	Versioned, attributable, access-controlled and protected from unauthorised alteration.
Independent Assessor or Reviewer Use	Corroborate design, implementation and operating effectiveness for the stated scope and period.
Common Weaknesses	Generic committee charter; no authority over engineering/operations; decisions not linked to systems.

E-02 - AI system inventory record

Purpose	Provide attributable, current evidence of control design or operation.
Owner	Assigned according to system and lifecycle authority.
Minimum Content	System ID, owner, intended function, model/provider, platform, interfaces, deployment, lifecycle, criticality and prohibited uses.
Lifecycle Stage	Defined by evidence type and system lifecycle.
Retention	Set from safety, mission, legal, contractual, investigation, evidence-preservation and approved records requirements. GAISSF conformance tiers do not define universal retention periods.
Integrity	Versioned, attributable, access-controlled and protected from unauthorised alteration.
Independent Assessor or Reviewer Use	Corroborate design, implementation and operating effectiveness for the stated scope and period.
Common Weaknesses	Model-only inventory; missing ground/onboard components; stale fleet or supplier records.

E-03 - Intended-function statement

Purpose	Provide attributable, current evidence of control design or operation.
Owner	Assigned according to system and lifecycle authority.

Purpose	Provide attributable, current evidence of control design or operation.
Minimum Content	Operational purpose, users, decisions/actions, inputs/outputs, environment, authority, success/failure and prohibited use.
Lifecycle Stage	Defined by evidence type and system lifecycle.
Retention	Set from safety, mission, legal, contractual, investigation, evidence-preservation and approved records requirements. GAISSF conformance tiers do not define universal retention periods.
Integrity	Versioned, attributable, access-controlled and protected from unauthorised alteration.
Independent Assessor or Reviewer Use	Corroborate design, implementation and operating effectiveness for the stated scope and period.
Common Weaknesses	Marketing description; no action authority or failure consequence; scope changes not controlled.

E-04 - Operational-design-domain definition

Purpose	Provide attributable, current evidence of control design or operation.
Owner	Assigned according to system and lifecycle authority.
Minimum Content	Environmental, geographic, platform, sensor, communications, workload and human-availability limits plus exit conditions.
Lifecycle Stage	Defined by evidence type and system lifecycle.
Retention	Set from safety, mission, legal, contractual, investigation, evidence-preservation and approved records requirements. GAISSF conformance tiers do not define universal retention periods.
Integrity	Versioned, attributable, access-controlled and protected from unauthorised alteration.
Independent Assessor or Reviewer Use	Corroborate design, implementation and operating effectiveness for the stated scope and period.
Common Weaknesses	ODD limited to weather; no degraded sensors, link loss or maintenance configuration.

E-05 - Safety-assessment interface record

Purpose	Provide attributable, current evidence of control design or operation.
Owner	Assigned according to system and lifecycle authority.
Minimum Content	Related hazards, failure conditions, assumptions, safety requirements, monitors, independence and unresolved security interactions.
Lifecycle Stage	Defined by evidence type and system lifecycle.
Retention	Set from safety, mission, legal, contractual, investigation, evidence-preservation and approved records requirements. GAISSF conformance tiers do not define universal retention periods.
Integrity	Versioned, attributable, access-controlled and protected from unauthorised alteration.
Independent Assessor or Reviewer Use	Corroborate design, implementation and operating effectiveness for the stated scope and period.
Common Weaknesses	Claims “covered by safety case” without traceability; security-induced hazards omitted.

E-06 - Security risk assessment

Purpose	Provide attributable, current evidence of control design or operation.
Owner	Assigned according to system and lifecycle authority.
Minimum Content	Assets, threats, attack paths, likelihood/consequence criteria, existing controls, residual risk and acceptance authority.
Lifecycle Stage	Defined by evidence type and system lifecycle.
Retention	Set from safety, mission, legal, contractual, investigation, evidence-preservation and approved records requirements. GAISSF conformance tiers do not define universal retention periods.
Integrity	Versioned, attributable, access-controlled and protected from unauthorised alteration.
Independent Assessor or Reviewer Use	Corroborate design, implementation and operating effectiveness for the stated scope and period.
Common Weaknesses	Generic AI risks; no command/data path; no supplier or physical effect analysis.

E-07 - Mission-assurance record

Purpose	Provide attributable, current evidence of control design or operation.
Owner	Assigned according to system and lifecycle authority.
Minimum Content	Mission objectives, critical functions, loss criteria, resilience assumptions, mission phases, contingencies and decision authority.
Lifecycle Stage	Defined by evidence type and system lifecycle.
Retention	Set from safety, mission, legal, contractual, investigation, evidence-preservation and approved records requirements. GAISSF conformance tiers do not define universal retention periods.
Integrity	Versioned, attributable, access-controlled and protected from unauthorised alteration.
Independent Assessor or Reviewer Use	Corroborate design, implementation and operating effectiveness for the stated scope and period.
Common Weaknesses	No mission-phase differentiation; success metrics without loss/recovery criteria.

E-08 - Model card

Purpose	Provide attributable, current evidence of control design or operation.
Owner	Assigned according to system and lifecycle authority.
Minimum Content	Model identifier/hash, architecture, purpose, training/evaluation data, performance by condition, limitations, uncertainty and release owner.
Lifecycle Stage	Defined by evidence type and system lifecycle.
Retention	Set from safety, mission, legal, contractual, investigation, evidence-preservation and approved records requirements. GAISSF conformance tiers do not define universal retention periods.
Integrity	Versioned, attributable, access-controlled and protected from unauthorised alteration.
Independent Assessor or Reviewer Use	Corroborate design, implementation and operating effectiveness for the stated scope and period.
Common Weaknesses	No exact weight/adaptor version; aggregate performance

Purpose	Provide attributable, current evidence of control design or operation.
	only; hosted model can change silently.

E-09 - System card

Purpose	Provide attributable, current evidence of control design or operation.
Owner	Assigned according to system and lifecycle authority.
Minimum Content	End-to-end components, human workflow, data and command flows, dependencies, safeguards, intended use, limitations and monitoring.
Lifecycle Stage	Defined by evidence type and system lifecycle.
Retention	Set from safety, mission, legal, contractual, investigation, evidence-preservation and approved records requirements. GAISSF conformance tiers do not define universal retention periods.
Integrity	Versioned, attributable, access-controlled and protected from unauthorised alteration.
Independent Assessor or Reviewer Use	Corroborate design, implementation and operating effectiveness for the stated scope and period.
Common Weaknesses	Describes model but not integrated system; omits downstream action and provider services.

E-10 - Dataset sheet and provenance manifest

Purpose	Provide attributable, current evidence of control design or operation.
Owner	Assigned according to system and lifecycle authority.
Minimum Content	Dataset ID/hash, origin, rights, collection, labels, transformations, synthetic content, coverage, quality and limitations.
Lifecycle Stage	Defined by evidence type and system lifecycle.
Retention	Set from safety, mission, legal, contractual, investigation, evidence-preservation and approved records requirements. GAISSF conformance tiers do not define universal retention periods.
Integrity	Versioned, attributable, access-controlled and protected from unauthorised alteration.
Independent Assessor or Reviewer Use	Corroborate design, implementation and operating effectiveness for the stated scope and period.
Common Weaknesses	Missing dataset/label version; undocumented augmentation; no rare-event or geography analysis.

E-11 - AI bill of materials

Purpose	Provide attributable, current evidence of control design or operation.
Owner	Assigned according to system and lifecycle authority.
Minimum Content	Model and adapter hashes, framework/runtime, datasets or provenance references, prompts, retrieval indexes, tools, licences, supplier and support status.
Lifecycle Stage	Defined by evidence type and system lifecycle.
Retention	Set from safety, mission, legal, contractual, investigation, evidence-preservation and approved records requirements. GAISSF conformance tiers do not define universal retention periods.
Integrity	Versioned, attributable, access-controlled and protected

Purpose	Provide attributable, current evidence of control design or operation.
	from unauthorised alteration.
Independent Assessor or Reviewer Use	Corroborate design, implementation and operating effectiveness for the stated scope and period.
Common Weaknesses	Missing fine-tuning adapter or system prompt; hosted model version ambiguous; dataset hash absent.

E-12 - Software bill of materials

Purpose	Provide attributable, current evidence of control design or operation.
Owner	Assigned according to system and lifecycle authority.
Minimum Content	Component, version, supplier, licence, dependency relationship, build source, vulnerability status and support/EOL.
Lifecycle Stage	Defined by evidence type and system lifecycle.
Retention	Set from safety, mission, legal, contractual, investigation, evidence-preservation and approved records requirements. GAISSF conformance tiers do not define universal retention periods.
Integrity	Versioned, attributable, access-controlled and protected from unauthorised alteration.
Independent Assessor or Reviewer Use	Corroborate design, implementation and operating effectiveness for the stated scope and period.
Common Weaknesses	Build/runtime mismatch; transitive dependencies omitted; no linkage to deployed image.

E-13 - Architecture and data-flow diagram

Purpose	Provide attributable, current evidence of control design or operation.
Owner	Assigned according to system and lifecycle authority.
Minimum Content	Trust boundaries, identities, privileges, data/model/command flows, onboard/ground/cloud allocation, monitors and recovery paths.
Lifecycle Stage	Defined by evidence type and system lifecycle.
Retention	Set from safety, mission, legal, contractual, investigation, evidence-preservation and approved records requirements. GAISSF conformance tiers do not define universal retention periods.
Integrity	Versioned, attributable, access-controlled and protected from unauthorised alteration.
Independent Assessor or Reviewer Use	Corroborate design, implementation and operating effectiveness for the stated scope and period.
Common Weaknesses	Diagram omits maintenance or provider access; no command path or independent safeguard.

E-14 - Configuration baseline

Purpose	Provide attributable, current evidence of control design or operation.
Owner	Assigned according to system and lifecycle authority.
Minimum Content	Hardware, software, model, data, prompt, retrieval, thresholds, environment and approved deviations with immutable identifiers.
Lifecycle Stage	Defined by evidence type and system lifecycle.
Retention	Set from safety, mission, legal, contractual, investigation,

Purpose	Provide attributable, current evidence of control design or operation.
	evidence-preservation and approved records requirements. GAISSF conformance tiers do not define universal retention periods.
Integrity	Versioned, attributable, access-controlled and protected from unauthorised alteration.
Independent Assessor or Reviewer Use	Corroborate design, implementation and operating effectiveness for the stated scope and period.
Common Weaknesses	Model version recorded but threshold/retrieval index omitted; deployed tail differs from record.

E-15 - Simulation plan and results

Purpose	Provide attributable, current evidence of control design or operation.
Owner	Assigned according to system and lifecycle authority.
Minimum Content	Simulation objective, model fidelity, assumptions, scenario coverage, seed/configuration, limitations, results and correlation evidence.
Lifecycle Stage	Defined by evidence type and system lifecycle.
Retention	Set from safety, mission, legal, contractual, investigation, evidence-preservation and approved records requirements. GAISSF conformance tiers do not define universal retention periods.
Integrity	Versioned, attributable, access-controlled and protected from unauthorised alteration.
Independent Assessor or Reviewer Use	Corroborate design, implementation and operating effectiveness for the stated scope and period.
Common Weaknesses	Simulation treated as reality; no fidelity bounds; failed scenarios removed.

E-16 - Verification and validation report

Purpose	Provide attributable, current evidence of control design or operation.
Owner	Assigned according to system and lifecycle authority.
Minimum Content	Requirements/acceptance criteria, test configuration, data independence, scenarios, results, failures, uncertainty and approval.
Lifecycle Stage	Defined by evidence type and system lifecycle.
Retention	Set from safety, mission, legal, contractual, investigation, evidence-preservation and approved records requirements. GAISSF conformance tiers do not define universal retention periods.
Integrity	Versioned, attributable, access-controlled and protected from unauthorised alteration.
Independent Assessor or Reviewer Use	Corroborate design, implementation and operating effectiveness for the stated scope and period.
Common Weaknesses	Tested version differs from release; only nominal conditions; unclear acceptance decision.

E-17 - Adversarial test report

Purpose	Provide attributable, current evidence of control design or operation.
Owner	Assigned according to system and lifecycle authority.
Minimum Content	Threat model, rules, attack corpus, tester independence,

Purpose	Provide attributable, current evidence of control design or operation.
	environment, success criteria, findings, exploit traces, remediation and retest.
Lifecycle Stage	Defined by evidence type and system lifecycle.
Retention	Set from safety, mission, legal, contractual, investigation, evidence-preservation and approved records requirements. GAISSF conformance tiers do not define universal retention periods.
Integrity	Versioned, attributable, access-controlled and protected from unauthorised alteration.
Independent Assessor or Reviewer Use	Corroborate design, implementation and operating effectiveness for the stated scope and period.
Common Weaknesses	Tool output without attack rationale; no physical or indirect channels; findings closed without retest.

E-18 - Human-factors assessment

Purpose	Provide attributable, current evidence of control design or operation.
Owner	Assigned according to system and lifecycle authority.
Minimum Content	User population, task analysis, workload, reliance, comprehension, handover, intervention timing, training and observed errors.
Lifecycle Stage	Defined by evidence type and system lifecycle.
Retention	Set from safety, mission, legal, contractual, investigation, evidence-preservation and approved records requirements. GAISSF conformance tiers do not define universal retention periods.
Integrity	Versioned, attributable, access-controlled and protected from unauthorised alteration.
Independent Assessor or Reviewer Use	Corroborate design, implementation and operating effectiveness for the stated scope and period.
Common Weaknesses	Usability survey only; no abnormal scenario or authority analysis.

E-19 - Operator training record

Purpose	Provide attributable, current evidence of control design or operation.
Owner	Assigned according to system and lifecycle authority.
Minimum Content	Role, system/version, limitations, prohibited uses, alerts, override, degraded mode, scenario assessment and competence decision.
Lifecycle Stage	Defined by evidence type and system lifecycle.
Retention	Set from safety, mission, legal, contractual, investigation, evidence-preservation and approved records requirements. GAISSF conformance tiers do not define universal retention periods.
Integrity	Versioned, attributable, access-controlled and protected from unauthorised alteration.
Independent Assessor or Reviewer Use	Corroborate design, implementation and operating effectiveness for the stated scope and period.
Common Weaknesses	Attendance record without competence; training does not match deployed version.

E-20 - Supplier assurance record

Purpose	Provide attributable, current evidence of control design or operation.
Owner	Assigned according to system and lifecycle authority.
Minimum Content	Criticality, due diligence, provenance, development/testing, access, subcontractors, vulnerabilities, changes, support, evidence rights and exceptions.
Lifecycle Stage	Defined by evidence type and system lifecycle.
Retention	Set from safety, mission, legal, contractual, investigation, evidence-preservation and approved records requirements. GAISSF conformance tiers do not define universal retention periods.
Integrity	Versioned, attributable, access-controlled and protected from unauthorised alteration.
Independent Assessor or Reviewer Use	Corroborate design, implementation and operating effectiveness for the stated scope and period.
Common Weaknesses	Questionnaire response uncorroborated; no delivered-version linkage; fourth parties unknown.

E-21 - Change-impact assessment

Purpose	Provide attributable, current evidence of control design or operation.
Owner	Assigned according to system and lifecycle authority.
Minimum Content	Changed artefacts, reason, affected requirements/risks/evidence, regression scope, fleet/mission impact, rollback and approval.
Lifecycle Stage	Defined by evidence type and system lifecycle.
Retention	Set from safety, mission, legal, contractual, investigation, evidence-preservation and approved records requirements. GAISSF conformance tiers do not define universal retention periods.
Integrity	Versioned, attributable, access-controlled and protected from unauthorised alteration.
Independent Assessor or Reviewer Use	Corroborate design, implementation and operating effectiveness for the stated scope and period.
Common Weaknesses	Change classified minor based on code lines; data/provider/threshold change omitted.

E-22 - Deployment manifest

Purpose	Provide attributable, current evidence of control design or operation.
Owner	Assigned according to system and lifecycle authority.
Minimum Content	Target asset/environment, exact artefact hashes, configuration, approvals, time, installer, staged group, checks and rollback point.
Lifecycle Stage	Defined by evidence type and system lifecycle.
Retention	Set from safety, mission, legal, contractual, investigation, evidence-preservation and approved records requirements. GAISSF conformance tiers do not define universal retention periods.
Integrity	Versioned, attributable, access-controlled and protected from unauthorised alteration.
Independent Assessor or Reviewer Use	Corroborate design, implementation and operating effectiveness for the stated scope and period.

Purpose	Provide attributable, current evidence of control design or operation.
Common Weaknesses	Release ticket without actual hashes; cannot identify which tails/spacecraft received update.

E-23 - Monitoring plan and logs

Purpose	Provide attributable, current evidence of control design or operation.
Owner	Assigned according to system and lifecycle authority.
Minimum Content	Indicators, baselines, strata, thresholds, owners, response, time synchronisation, coverage gaps, retention and sampled logs.
Lifecycle Stage	Defined by evidence type and system lifecycle.
Retention	Set from safety, mission, legal, contractual, investigation, evidence-preservation and approved records requirements. GAISSF conformance tiers do not define universal retention periods.
Integrity	Versioned, attributable, access-controlled and protected from unauthorised alteration.
Independent Assessor or Reviewer Use	Corroborate design, implementation and operating effectiveness for the stated scope and period.
Common Weaknesses	Dashboard without approved thresholds; fleet average hides configuration; missing-log condition not alerted.

E-24 - Incident report

Purpose	Provide attributable, current evidence of control design or operation.
Owner	Assigned according to system and lifecycle authority.
Minimum Content	Timeline, system/baseline, detection, safety/mission effect, containment, evidence, notifications, root cause, recovery and lessons.
Lifecycle Stage	Defined by evidence type and system lifecycle.
Retention	Set from safety, mission, legal, contractual, investigation, evidence-preservation and approved records requirements. GAISSF conformance tiers do not define universal retention periods.
Integrity	Versioned, attributable, access-controlled and protected from unauthorised alteration.
Independent Assessor or Reviewer Use	Corroborate design, implementation and operating effectiveness for the stated scope and period.
Common Weaknesses	IT-only narrative; model/data version unknown; containment consequences not assessed.

E-25 - Vulnerability record

Purpose	Provide attributable, current evidence of control design or operation.
Owner	Assigned according to system and lifecycle authority.
Minimum Content	Affected component/version, exploitability, exposure, safety/mission interaction, supplier status, treatment, deadline, exception and closure.
Lifecycle Stage	Defined by evidence type and system lifecycle.
Retention	Set from safety, mission, legal, contractual, investigation, evidence-preservation and approved records requirements. GAISSF conformance tiers do not define universal retention periods.

Purpose	Provide attributable, current evidence of control design or operation.
Integrity	Versioned, attributable, access-controlled and protected from unauthorised alteration.
Independent Assessor or Reviewer Use	Corroborate design, implementation and operating effectiveness for the stated scope and period.
Common Weaknesses	CVSS copied without aerospace context; unsupported component not escalated.

E-26 - Risk acceptance

Purpose	Provide attributable, current evidence of control design or operation.
Owner	Assigned according to system and lifecycle authority.
Minimum Content	Defined risk, affected scope, evidence, alternatives, interim controls, accountable authority, expiry, review trigger and communication.
Lifecycle Stage	Defined by evidence type and system lifecycle.
Retention	Set from safety, mission, legal, contractual, investigation, evidence-preservation and approved records requirements. GAISSF conformance tiers do not define universal retention periods.
Integrity	Versioned, attributable, access-controlled and protected from unauthorised alteration.
Independent Assessor or Reviewer Use	Corroborate design, implementation and operating effectiveness for the stated scope and period.
Common Weaknesses	Permanent acceptance; approver lacks authority; no expiry or affected fleet list.

E-27 - Independent assurance report

Purpose	Provide attributable, current evidence of control design or operation.
Owner	Assigned according to system and lifecycle authority.
Minimum Content	Scope, criteria, competence, independence/conflicts, methods, samples, findings, limitations, responses and conclusion.
Lifecycle Stage	Defined by evidence type and system lifecycle.
Retention	Set from safety, mission, legal, contractual, investigation, evidence-preservation and approved records requirements. GAISSF conformance tiers do not define universal retention periods.
Integrity	Versioned, attributable, access-controlled and protected from unauthorised alteration.
Independent Assessor or Reviewer Use	Corroborate design, implementation and operating effectiveness for the stated scope and period.
Common Weaknesses	Reviewer reports to implementation owner; unsupported “compliant” conclusion; limited sample undisclosed.

E-28 - Decommissioning record

Purpose	Provide attributable, current evidence of control design or operation.
Owner	Assigned according to system and lifecycle authority.
Minimum Content	Assets/models/data/credentials disabled, retention, supplier termination, residual interfaces, archive integrity and confirmation.

Purpose	Provide attributable, current evidence of control design or operation.
Lifecycle Stage	Defined by evidence type and system lifecycle.
Retention	Set from safety, mission, legal, contractual, investigation, evidence-preservation and approved records requirements. GAISSF conformance tiers do not define universal retention periods.
Integrity	Versioned, attributable, access-controlled and protected from unauthorised alteration.
Independent Assessor or Reviewer Use	Corroborate design, implementation and operating effectiveness for the stated scope and period.
Common Weaknesses	Model endpoint remains reachable; keys or retrieval indexes retained without owner.

10. Supplier Assurance

- SQ-01 Controlled data and provenance: Identify all model, adapter, dataset, simulation and HIL artefacts. Explain segregation and authorised handling of controlled technical data, CUI, export-controlled or classified information.
- SQ-02 Development and assurance independence: Describe development, verification and approval independence. Identify whether the inference function and independent safety/constraint monitor share personnel, code, toolchain or common dependencies.
- SQ-03 Delivered configuration: Provide exact model/software hashes, AI/SBOM, prompts/retrieval indexes, runtime and configuration for each delivered release.
- SQ-04 Training and evaluation data: Provide provenance, rights, transformation, synthetic-data use, coverage, rare-event limitations and contamination controls for material datasets.
- SQ-05 Secure build and release: Describe protected source, reproducible builds, signing, key custody, release approval, staged deployment, installation verification and rollback.
- SQ-06 Subcontractors and service locations: List critical subcontractors, hosted-service locations, remote support jurisdictions and any party able to change or access the delivered system.
- SQ-07 Change notification: Identify model, data, prompt, retrieval, dependency, infrastructure and support changes that trigger customer notice and reassessment; state the proposed notice period for approval.
- SQ-08 Vulnerability and incident cooperation: Describe intake, severity, notification, evidence preservation, containment support, patch/mitigation timelines and coordination with safety/mission investigations.
- SQ-09 Monitoring and drift: Specify available telemetry, model/provider change indicators, drift and performance measures, threshold governance and evidence export.
- SQ-10 Support, obsolescence and exit: State support term, EOL notice, component substitution controls, escrow/source-access options, data/model export, transition support and secure termination.
- SQ-11 Remote and privileged access: Describe identities, MFA, JIT access, approval, session recording, emergency access, customer visibility and revocation.
- SQ-12 Physical and environmental assurance: For embedded/physical AI, provide environmental, sensor-degradation, timing, resource, HIL and fail-safe/fail-operational test evidence.
- SQ-13 Human factors and training: Provide operator/maintainer task analysis, limitations, handover/override design, training material and competence evidence.
- SQ-14 Evidence and audit rights: List assurance artefacts available, retention periods, independent reports, customer test rights, audit cooperation and restrictions.
- SQ-15 Common-mode and fleet impact: Identify shared dependencies and the method for determining whether a defect affects multiple aircraft, spacecraft, customers or missions.

11. External Reference Method

The source register identifies official or recognised materials and their publication status as rechecked for this public-draft edition. GAISSF-SEC-050 does not reproduce licensed standards or claim clause-level equivalence, means of compliance or regulatory acceptance. Any future clause-level crosswalk must be performed by competent specialists and state whether the relation is overlap, supporting evidence, analogy or verified requirement mapping. Reference editions must be revalidated before reliance.

12. Illustrative Implementation Roadmap (Non-prescriptive)

Example: first 30 days

Illustrative sequence only; timing must be tailored to programme phase, system criticality, architecture, safety/mission processes and applicable obligations. Confirm scope, owner, system class, inventory, criticality and prohibited uses.

Example: first 60 days

Complete risk/threat analysis, architecture, supplier and evidence-gap review.

Example: first 90 days

Implement priority controls, validation, monitoring, incident and change processes.

Example: by 6 months

For high-consequence systems, plan independent review or assurance where appropriate to the applicable safety, mission, contractual or regulatory process, and exercise rollback/containment.

Example: by 12 months

Measure effectiveness, resolve systemic supplier/common-mode gaps and integrate lessons learned.

13. Publication Validation

Control population	PASS - 59 controls preserved; Applicable-Control Baseline = 52 D1-D8 + applicable D9
Use cases	PASS - 24 differentiated records
Threats	PASS - 15 differentiated records
Evidence catalogue	PASS - 28 differentiated specifications
Architectures	PASS - 9 differentiated conceptual patterns
Evidence taxonomy	PASS - no competing T1-T4 sector evidence model retained
External mapping claims	PASS - no unverified clause-level compliance mappings asserted
Financial claims	PASS - no ODA3-generated monetary exposure estimates included
Publication recommendation	Draft for Publication Public External Review Open
Final Publication gates	Aerospace engineering; aviation/space safety and mission assurance; cybersecurity/human factors as applicable; jurisdiction-specific legal/regulatory; accessibility/editorial; explicit authorized final release approval.