ODA3 INSTITUTE

BUILDING AN AI ASSURANCE EVIDENCE PACK

A Practitioner's Field Guide to Defensible Evidence Preparation

Doc ID: ODA3-2026-08-WHP-COM-015
Framework Context: GAISSF Ecosystem
Classification: Public — GEL v1.0
Status: **FINAL**
Date: 15 August 2026

ODA3 INSTITUTE

BUILDING AN AI ASSURANCE EVIDENCE PACK

A Practitioner's Field Guide to Defensible Evidence Preparation

Doc ID: ODA3-2026-08-WHP-COM-015
Document Type: WHP — Adoption / Practitioner Field Guide
Framework Context: GAISSF Ecosystem
Classification: Public — GEL v1.0
Status: FINAL
Date: 13 August 2026
Audience: CISOs, Security Architects, AI Governance Leads, AI Security Teams, Risk and Compliance Teams, Control Owners, Evidence Custodians, Internal Assurance Teams, GRC Teams, Evidence-Preparation Teams

---

Publication Boundary

This is an informative practitioner field guide. It explains how an organization can organize, preserve, evaluate, and prepare evidence for AI assurance activities. It does not create, modify, replace, or supersede controlled GAISSF Ecosystem requirements, assessment methodology, scoring rules, auditor procedures, certification decision logic, or scheme rules.

An evidence pack is **not** an assessment result. The existence, quantity, or apparent completeness of artifacts does not by itself establish control effectiveness, conformity, certification, regulatory compliance, system safety, or operational assurance.

Where this guide conflicts with a current controlled ODA3 source, the controlled source governs.

Methodology Note

This guide uses the current GAISSF™ adoption baseline for evidence classes and evidence-index concepts and the controlled ODA3 evidence methodology only for public-safe evidence-quality concepts. Across the WHP practitioner suite, the recurring discipline is the same: **assurance is demonstrated, not asserted; artifact existence alone does not establish evidence sufficiency, control effectiveness, system safety, conformity, or regulatory compliance.**

COM-015 is an **evidence-preparation guide for later assurance review and assessment use**. It does not determine framework applicability, select a conformance profile, perform a control assessment, assign confidence, determine conformity, authorize certification, or replace sector-specific safety/legal processes.

The public guide may explain evidence preparation, indexing, provenance, coverage, limitations, contradiction and handoff-preparation questions. It does not publish controlled confidence scales or caps, assessor workpaper logic, certification gates, proprietary testing logic, universal sampling rules, or assessor decision rules.

ODA3 Institute was founded in March 2026. This guide therefore does not claim proprietary deployment history, independently measured audit-reduction or incident-reduction outcomes, assessor-consistency results, or field validation across all sectors. Examples are illustrative unless explicitly identified otherwise.

Suite Position and Handoffs

COM-015 sits between implementation activity and formal assessment use:

scope / applicability → implementation activity → evidence generation → COM-015 evidence preparation → COM-016 assessment use

The adjacent practitioner documents remain distinct:

- **COM-009 — The First 30 Days with GAISSF™:** implementation planning and an internal management checkpoint; it does not make an assurance conclusion.
- **COM-010 — Applying GAISSF™: Worked Scenarios:** sector/application reasoning with illustrative evidence; examples do not establish real-system applicability or evidence sufficiency.
- **COM-011 — When AI Breaks:** short incident-response orientation.
- **COM-012 — Classifying AI Incidents:** UAIF™ classification and uncertainty discipline.
- **COM-013 — AI Incident Response Field Guide:** full UAIF™ / AI-IRF™ incident-response integration.
- **COM-014 — Securing Physical AI:** PAI-SF™ practitioner guidance and the physical-AI / functional-safety boundary.
- **COM-016 — The GAISSF™ Assessment Field Guide:** planned next-layer guide that may consume an evidence pack but must not be pre-empted by COM-015.
- **COM-017 — Implementing GAISSF™: A Practitioner's Guide:** planned deeper implementation companion; implementation artifacts may become evidence only when they satisfy the evidence conditions explained here.

Acronym Reference

Acronym	Expansion
AI-IRF™	AI Incident Response Framework
GAISSF™	Global AI Safety and Security Framework
GEL	GAISSF Ecosystem License
GRC	Governance, Risk and Compliance
IAM	Identity and Access Management
ODA3	ODA3 Institute
PAI-SF™	Physical AI Security Framework
SIEM	Security Information and Event Management
UAIF™	Unified AI Incident Framework
WHP	White Paper / Practitioner Guide (ODA3 source family)

How to Use This Guide

Use the guide in the way that matches the current practitioner problem:

Situation	Start here	Then use
Building a first evidence pack	Section 2 — Scope	Sections 3, 9, 14, 41
Repairing a document dump	Section 1.1 — Artifact is not conclusion	Sections 9–20 and 33
Heavy provider dependency	Section 26 — Third-Party /	Sections 9, 17–20, 32

Situation	Start here	Then use
	Hosted AI	
Agentic / tool-using AI	Section 25	Sections 24, 28 and 41
Physical AI	Section 29	COM-014 plus Sections 28, 32 and 40
Incident-driven review	Section 28	COM-012 / COM-013 plus Sections 17–20
Preparing a handoff	Sections 30–33	Sections 41–46

The document does not need to be read linearly. The evidence index and requirement–evidence matrix are the two main navigation artifacts: use them to link the detailed material back to a defined scope and evidence objective.

---

1. Why AI Assurance Needs an Evidence Architecture

AI assurance fails surprisingly early when organizations begin with the question:

“What documents do we need?”

The better first question is:

“What do we need to establish, for which system, over what period, and what evidence could support that conclusion?”

Organizations commonly possess large volumes of potentially relevant material: policies, architecture diagrams, model cards, evaluation reports, logs, tickets, approvals, risk registers, supplier attestations, dashboards, incident records, training records and meeting minutes. These artifacts become assurance evidence only when their relationship to a bounded claim can be demonstrated.

A useful evidence architecture therefore connects:

requirement → evidence objective → source → artifact → provenance → scope and period → preparation-quality review → limitation or contradiction → later assurance / assessment use

This matters particularly for AI because the assessed behavior may depend on more than a single application or model. Relevant evidence can span:

- enterprise governance;
- model or service configuration;
- training, fine-tuning or retrieval data;
- prompts and system instructions;
- identity and access controls;
- tools and agent permissions;
- model-provider changes;
- monitoring and evaluation systems;
- human review;
- incident response;
- supplier controls; and
- physical interfaces where AI can cause or influence physical action.

An evidence pack should make these relationships inspectable rather than forcing a reviewer to reconstruct them from an unstructured document repository.

Key principle: Assurance is demonstrated, not asserted. The purpose of COM-015 is to make the evidentiary basis inspectable without converting the existence of artifacts into an assurance conclusion.

1.1 Artifact is not conclusion

A policy proves, at most, something about documented intent and approval. It does not by itself prove operation.

A configuration export can demonstrate a technical state at a particular time. It does not necessarily establish that the state existed throughout the relevant period.

A test report can establish that a defined test was performed under defined conditions. It does not automatically establish production effectiveness outside those conditions.

A dashboard can summarize operations. It does not automatically prove completeness or lineage of the underlying telemetry.

A supplier certificate or attestation can be relevant. It does not automatically establish the customer's end-to-end control.

The evidence pack exists to preserve these distinctions.

---

2. Start with Scope, Not Documents

Evidence collection should not begin until the system or review object is sufficiently identifiable.

As a practitioner starter, record:

Scope element	Practitioner question
System / service	What exactly is being examined?
Business process	What business activity does it support or influence?
Model / service / version	Which model, provider service or version is relevant?
Environment	Production, staging, region, tenant, site or other boundary?
Data / retrieval	Which material data and retrieval sources affect operation?
Tools / integrations	What can the AI system call, change, retrieve or trigger?
Suppliers	Which controls or evidence depend on third parties?
Human roles	Who approves, reviews, overrides or responds?
Physical interface	Can outputs cause or influence physical action?
Evidence period	What period must the evidence represent?
Exclusions	What is explicitly outside the review?

2.1 Why the period matters

A configuration captured today may not demonstrate the configuration that existed three months ago. A supplier report may cover a different period. A model evaluation may pre-date a material model update. A policy may have been approved after the period under review.

Record both:

- **collection date** — when the evidence was obtained; and
- **evidence period/state** — what period or system state it actually represents.

They are not interchangeable.

2.2 Material change

"Material" defined for this guide: a change, condition or evidence gap is material if it could reasonably affect a conclusion about the assessed system's evidence relevance, usefulness for later review, or the validity of prior evidence for the applicable scope and period. Materiality is contextual; this guide does not establish a universal quantitative threshold. Organizations should apply their own governed risk and change process to determine materiality for a given system and claim.

Evidence should be reconsidered when a material change affects its relevance. Examples include:

- model or provider version change;
- material prompt/system-instruction change;
- new tool or agent permission;
- new data source;
- new deployment region;
- monitoring redesign;
- significant incident;
- safety-control change;
- supplier change; or
- material business-process change.

"Latest available" does not necessarily mean "relevant to the assessed state."

2.3 Scoping and Applicability Handoff

COM-015 does not decide whether Appendix P.1, the full GAISSF™ domain set, PAI-SF™, UAIF™, AI-IRF™, or another framework/profile is applicable to a specific system. Those determinations belong to the applicable framework sources and the organization's governed scoping process.

Where COM-009 or another approved scoping/applicability process has already created a system boundary, profile decision or applicable-control baseline, COM-015 should **consume and preserve that record** rather than silently re-perform or broaden the decision. Where the evidence process exposes an unrecorded dependency, material change or contradiction, route the issue back for scope/applicability reassessment.

---

3. Requirement → Evidence Objective → Claim

An evidence request should be derived from what needs to be established.

Example:

Requirement: access to a material AI capability is controlled.

A weak request is:

"Provide the access-control policy."

A stronger evidence objective is:

Establish how access is authorized, technically enforced, reviewed, changed and revoked for the assessed system during the relevant period.

Potential evidence may then include:

- policy/procedure;
- role and entitlement design;
- configuration export;
- access-review records;
- joiner/mover/leaver tickets;
- privileged-access logs;
- exceptions; and
- evidence of remediation where inappropriate access was identified.

The policy remains useful, but it is no longer mistaken for the entire control.

3.1 Bounded claims

For each evidence objective, ask:

1. What could this evidence support?
2. What could it not support?
3. Which system/unit does it apply to?
4. Which period does it represent?
5. What additional evidence would be needed for a broader conclusion?

This prevents a common assurance failure: making a system-wide or period-wide claim from a narrow artifact.

---

4. Six GAISSF Adoption Evidence Classes

The evidence architecture uses six practical classes.

Cross-framework note: these are the six evidence classes used by the current GAISSF™ adoption guidance: governance, design, technical, operational, supplier and corrective action. They are not a universal taxonomy that every GAISSF Ecosystem framework must reproduce. For example, COM-014 / PAI-SF™ uses a physical-AI practitioner grouping of design, test, runtime, incident, third-party and governance evidence. Those groupings serve different implementation purposes and should not be forced into false one-to-one equivalence.

Evidence class	Typical examples	Primary use
Governance	charters, approvals, minutes, risk decisions	accountability and authorized decisions
Design	policies, procedures, architectures, threat models	intended control design
Technical	configurations, test outputs, scans, model evaluations	technical implementation and tested state
Operational	logs, tickets, review records, monitoring outputs	evidence of operation over time
Supplier	contracts, attestations, due diligence, service reports	third-party reliance and responsibility boundary

Evidence class	Typical examples	Primary use
Corrective action	findings, fixes, validation, closure approval	issue-to-remediation traceability

No class is inherently “better” in all cases. Its value depends on the claim.

4.1 Governance evidence

Governance evidence can establish who owns a decision, what was approved, which risk was accepted, or whether a review forum operated.

Useful characteristics include:

- authorized approver;
- date;
- defined scope;
- decision or outcome;
- supporting rationale;
- follow-up obligations.

A governance document without evidence that required actions occurred may remain design/governance evidence rather than operating evidence.

4.2 Design evidence

Design evidence explains how a control is intended to work.

Examples:

- security policy;
- AI use policy;
- architecture;
- threat model;
- procedure;
- role model;
- monitoring design;
- incident playbook.

A well-written procedure does not prove execution. Design evidence should be connected to technical or operational evidence where operation is the claim.

4.3 Technical evidence

Technical evidence may include:

- configuration exports;
- model evaluation results;
- adversarial test results;
- IAM configuration;
- deployment records;
- security scan results;
- signed artifacts;
- system-generated configuration state.

Record the tested environment, version, date, methodology and known limitations. A result from a laboratory environment should not silently become evidence about production.

4.4 Operational evidence

Operational evidence demonstrates what happened during operation.

Examples:

- logs;
- alert and investigation records;
- periodic review records;
- monitoring outputs;
- change tickets;
- approval records;
- incident records;
- exception handling.

Period coverage is particularly important. A single clean week should not be presented as evidence for a year without a justified basis.

4.5 Supplier evidence

Supplier evidence may include:

- contract/SLA;
- security schedules;
- service descriptions;
- audit/attestation reports;
- provider evaluations;
- change notifications;
- incident notifications;
- due-diligence records.

Always ask: **what responsibility remains with the customer?**

Provider evidence should be mapped to the inherited control or service boundary rather than treated as evidence for the whole AI system.

4.6 Corrective-action evidence

A finding is not closed because a team says it is closed.

A defensible corrective-action chain is:

finding → assigned action → implemented change → validation → closure decision

Retain evidence of the validation, not only the remediation ticket.

---

5. Design Evidence vs Operating Evidence

One of the most important distinctions in an evidence pack is between **control design** and **control operation**.

Design statement	Operating evidence question
Access must be reviewed quarterly.	Were the reviews performed, with what population, exceptions and outcomes?
Models must be evaluated before release.	Which version was evaluated, under what

Design statement	Operating evidence question
	criteria, and what happened to failures?
Human review is required.	Did authorized reviewers actually review material cases?
Provider changes must be assessed.	Were change notices received, triaged and acted on?
Incidents must be escalated.	Were incidents identified and escalated under the defined process?

A mature evidence pack contains both where both are needed.

6. Direct, Corroborating and Contextual Support

Practitioners benefit from distinguishing how evidence contributes without assigning public numeric weights.

Direct evidence addresses the evidence objective closely. A system-generated access export for the assessed production tenant may directly support a question about configured access.

Corroborating evidence supports or challenges another item. A ticket showing approval may corroborate a configuration change.

Contextual evidence explains the environment but may not establish the control. An architecture presentation can help orient a reviewer but may not establish deployed state.

Unverified leads point to potentially useful evidence but should not be treated as established evidence until validated.

The categories are contextual, not universal labels. The same artifact may play different roles for different claims.

7. Human Statements and Representations

Interviews and representations are useful, particularly where processes involve human judgment or provider-side information is inaccessible. They should remain identifiable as representations.

Record:

- speaker/source;
- role;
- date;
- subject;
- whether the statement is first-hand;
- supporting evidence;
- unresolved contradictions.

Do not silently convert:

“The provider told us no material change occurred”

into:

“No material change occurred.”

The first is a provider representation. The second is a factual conclusion requiring an adequate evidence basis.

8. Machine-Generated and AI-Generated Evidence

Automated evidence can be highly valuable because it may be repeatable and high-volume. It can also fail invisibly if its source, query, transformation or coverage is unknown.

For material automated evidence, preserve where appropriate:

- source system;
- query/export method;
- time zone;
- filters;
- population;
- collection time;
- transformation steps;
- tool/version;
- operator;
- integrity controls.

8.1 AI-generated summaries

AI may help summarize a large evidence corpus. The summary should not replace the underlying source.

Where an AI-generated summary is used:

1. retain authoritative source data or a sufficient access path;
2. record the generating model/tool version and material prompt, query, rule, parameters or transformations where relevant and permissible;
3. validate important extracted facts and proportionately consider hallucination, omission, drift, aggregation and reproducibility risk;
4. preserve contradictory source evidence;
5. maintain accountable human/system ownership; and
6. do not treat generated interpretation as independent merely because software produced it or as a substitute for the underlying records.

9. Build the Evidence Index

Every material artifact should have a stable evidence ID.

The current GAISSF™ Adoption Guide defines the public minimum evidence-index fields as **Evidence ID, Control ID, Scope element, Owner, Period, Source, Integrity, Status, Location, and Limitations**. COM-015 retains those fields and recommends additional preparation metadata where useful.

Field	Status / purpose
Evidence ID	**GAISSF public minimum** — unique identifier

Field	Status / purpose
Control ID / requirement linkage	**GAISSF public minimum** — mapped control; COM-015 may add sub-requirement/evidence-objective context
Scope element	**GAISSF public minimum** — system, business unit, supplier or location
Owner / custodian	**GAISSF public minimum** — evidence custodian
Period	**GAISSF public minimum** — date range/state represented
Source	**GAISSF public minimum** — system or process of origin
Integrity	**GAISSF public minimum** — hash, signature or repository control where appropriate
Status	**GAISSF public minimum** — draft, approved, expired or superseded
Location	**GAISSF public minimum** — controlled repository path or secure reference
Limitations	**GAISSF public minimum** — known gaps, sampling limits or assumptions
Evidence title	COM-015 preparation enrichment
Evidence class	COM-015 preparation enrichment
Version	COM-015 preparation enrichment
Collection date	COM-015 preparation enrichment
Provenance detail	COM-015 preparation enrichment
Transformation / redaction	COM-015 preparation enrichment
Related evidence IDs	COM-015 preparation enrichment

9.1 Stable IDs

An evidence ID should not depend entirely on a mutable filename. Example: EV-ACCESS-0042.

9.2 Superseded records

Do not delete an artifact merely because a newer version exists when the older artifact is relevant to the assessed period. Mark it superseded and preserve its period relevance.

---

10. Provenance

Provenance answers:

- Where did this evidence originate?
- Who or what created it?
- Who collected it?
- How was it extracted?
- Was it transformed?
- Can the transformation be reconstructed?
- Which environment/version does it describe?

Weak provenance can cap the usefulness of otherwise impressive evidence.

A screenshot copied into a presentation with no source, timestamp or environment may be visually persuasive but evidentially weak.

11. Integrity and Authenticity

Not every artifact needs a cryptographic hash. Integrity measures should be proportionate to sensitivity, materiality and the way evidence is handled. No single mechanism is universally required by COM-015; the appropriate control depends on the evidence type, source, sensitivity, system context, and intended assurance use. The practitioner question is whether origin, completeness, provenance and resistance to undetected alteration can be inspected — not which specific tool is used.

Possible measures include:

- controlled repositories;
- access controls;
- immutable or append-only logs;
- source-system exports;
- signatures;
- hashes;
- approval workflows;
- version history;
- chain-of-custody records.

11.1 Screenshots

Screenshots are sometimes necessary, but record:

- source system;
- environment;
- date/time;
- relevant account/context;
- what the screenshot excludes;
- whether an export or underlying record exists.

A screenshot should not be used to imply period-wide operation unless supported by period evidence.

11.2 Redaction

Redaction should preserve evidentiary meaning.

Record:

- what category of information was removed;
- why;
- who performed the transformation where appropriate;
- whether the original remains available under controlled access.

12. Currency and Evidence Period

Evidence becomes stale when the system changes materially or when the evidence no longer represents the claim period.

Examples:

- evaluation report refers to model v3; production uses v4;
- supplier report covers the previous service architecture;
- architecture diagram predates an agent integration;
- access review predates a privilege redesign;
- monitoring report covers only a pre-production pilot.

Do not “refresh” evidence by changing the filename or cover date. Obtain or generate evidence that actually represents the current state.

---

13. Coverage and Representativeness

Coverage asks whether the evidence represents enough of the relevant population, period, environment or cases for the intended use.

Practitioner questions include:

- What population exists?
- What portion is represented?
- Are high-risk cases included?
- Are failures and exceptions visible?
- Are all material environments covered?
- Are maintenance periods excluded?
- Are low-frequency but high-consequence events represented?
- Did the system materially change during the period?

This guide does not prescribe a universal sample-size formula. Sampling decisions belong to the applicable controlled methodology, scheme, engagement or legal requirement.

---

14. Evidence Repository Design

A practical evidence repository should be navigable by someone other than the person who assembled it.

Recommended structure:

```
00_README_AND_SCOPE/
01_REQUIREMENT_AND_EVIDENCE_MATRIX/
02_EVIDENCE_INDEX/
03_GOVERNANCE/
04_DESIGN/
05_TECHNICAL/
06_OPERATIONAL/
07_SUPPLIER/
08_INCIDENTS_EXCEPTIONS_CORRECTIVE_ACTION/
```


09_LIMITATIONS_AND_GAPS/
10_CHANGE_LOG/
11_SECURE_REFERENCES/

Do not duplicate highly sensitive data merely to fit the folder structure. A secure reference may be preferable.

The index is the authoritative navigation layer.

14.1 Practical Pack States

For internal document control, a pack may move through:

Working Collection → Internal Quality Review → Handoff Version → Addendum / Change Set → Archived Version

These are document-management states only. They are **not** ODA3 assessment, audit, conformity, certification or readiness designations.

At handoff, freeze or snapshot the pack sufficiently to identify what was actually supplied. Record the pack version, handoff date/time, open gaps, known exclusions and any secure references that remain external to the pack.

Subsequent evidence should be added through a controlled change/addendum record rather than silently replacing the handed-off evidence set.

---

15. Five Public Evidence Review Dimensions

Before handoff, review each material evidence chain qualitatively across five dimensions.

15.1 Relevance

Does the evidence address the specific evidence objective?

15.2 Independence

How free are the source and the relevant evidence-generation or review process from responsibility, interest or self-review? Is corroboration appropriate?

Independence is not binary, and no universal source hierarchy overrides relevance. A system-owner configuration or operational record may be the most probative evidence for one claim and insufficient alone for another.

15.3 Currency

Does the evidence represent the relevant version, state and period?

15.4 Coverage

Does it represent the necessary population, time span, environments and material cases?

15.5 Integrity

Can source, provenance and material transformations be assessed?

These are **review dimensions, not a public scoring formula**.

---

16. From Artifact to Defensible Evidence

COM-015 does not determine whether evidence is **sufficient** for a formal assessment conclusion. Sufficiency is a reasoned determination within the applicable controlled assessment process.

Before handoff, a practitioner can still test whether the evidence basis is well prepared:

1. Is the item identifiable?
2. Is it attributable?
3. Is it linked to the correct scope and period?
4. Is its integrity assessable?
5. Is it lawfully available for the intended review/use?
6. Does it address a defined requirement/evidence objective?
7. Is the source and collection basis appropriate to the claim?
8. Is coverage explicit?
9. Has likely corroborating evidence been identified where relevant?
10. Is contrary evidence visible?
11. Are limitations explicit?
12. Is the evidence chain organized for handoff to the applicable assurance or assessment process?

Passing this preparation check does **not** create an ODA3 or third-party status such as *assessment-ready*, *audit-ready*, *certification-ready*, *conforming* or *certifiable*. Those are not outcomes of COM-015.

---

17. Contrary Evidence

A defensible pack includes material evidence that challenges the desired conclusion.

Examples:

- a failed red-team test;
- an unauthorized tool invocation;
- an incident during an otherwise clean monitoring period;
- an access-review exception;
- an evaluation regression;
- a supplier notification inconsistent with an internal assumption.

Do not remove a material adverse artifact because remediation occurred later. Instead preserve:

adverse evidence → response → remediation → validation → current status

The history can itself become useful assurance evidence.

---

18. Missing Evidence

“Missing” can mean different things:

- requested but not provided;
- never created;

- no longer retained;
- inaccessible because of supplier boundaries;
- corrupted;
- outside the evidence period;
- not reproducible;
- unknown whether it exists.

Record the reason.

Missing evidence does not automatically prove control failure. It can, however, limit what can be concluded.

---

19. Notably Absent Discipline

A strong evidence pack documents important absences carefully.

Weak:

No incidents occurred.

Better:

No qualifying incident record was identified in the reviewed incident register for the stated period.

Stronger still:

No qualifying incident record was identified in the reviewed incident register and linked escalation queue for the stated period. This statement is limited to those records and does not establish that no unrecorded event occurred.

The required discipline is:

search basis + period + coverage + limitation

“No evidence observed” is not equivalent to “did not occur.”

---

20. Evidence Gaps and Remediation

Create a gap record when an evidence objective cannot be adequately supported.

Recommended fields:

- Gap ID
- Requirement/control
- Missing or limited evidence
- Reason
- Review consequence
- Owner
- Action
- Target date where used operationally
- Replacement evidence

- Validation
- Status
- Residual limitation

Do not close the gap simply because an artifact has been uploaded. Confirm that the new evidence actually addresses the gap.

---

21. AI-Specific Evidence Pattern: Model and System Identity

A reviewer should be able to identify the system that the evidence describes.

Useful records may include:

- AI system inventory entry;
- model/provider identity;
- model/service version;
- deployment ID;
- environment;
- system owner;
- business owner;
- material integrations;
- change history.

Where provider version visibility is limited, record that limitation rather than inventing precision.

---

22. AI-Specific Evidence Pattern: Data, Retrieval and Provenance

Depending on the system, evidence may include:

- dataset inventory;
- source and lineage;
- ingestion approvals;
- retrieval-source inventory;
- access controls;
- quality checks;
- change records;
- retention/deletion records;
- provenance metadata.

Do not imply that a dataset record establishes legality, representativeness, safety or fitness for purpose unless the relevant evidence supports those separate conclusions.

---

23. AI-Specific Evidence Pattern: Model Integrity and Evaluation

Potential evidence:

- evaluation plan;
- acceptance criteria;
- test dataset description;
- model/service version;
- evaluation output;
- adversarial test report;
- regression result;
- remediation;
- rollback decision;
- residual limitation.

Preserve failed results and retests. A final passing report without the relevant failure/remediation history can obscure the actual control story.

---

24. AI-Specific Evidence Pattern: Runtime and Monitoring

Potential evidence:

- telemetry inventory;
- logging configuration;
- alert rules;
- monitoring output;
- drift/performance reports;
- investigation tickets;
- alert dispositions;
- retention configuration;
- coverage-gap records.

A dashboard is not enough if the underlying telemetry population, filters and retention cannot be explained.

---

25. AI-Specific Evidence Pattern: Agentic and Tool-Using Systems

For systems able to call tools or take actions, evidence may need to establish:

- declared purpose;
- allowed tools;
- delegated permissions;
- identity used for actions;
- approval boundaries;
- transaction/action limits;
- tool-call logs;
- escalation;
- intervention;
- kill/disable mechanism;

- rollback or recovery.

A content filter does not establish that an agent lacked authority to perform a high-impact action. Authority and action evidence must be examined separately.

26. AI-Specific Evidence Pattern: Third-Party / Hosted AI

Build a responsibility boundary.

Question	Evidence
What service is inherited?	contract/service description
What criteria were applied, and by whom?	report/test criteria and competence information where available
What does provider evidence cover?	entities, services, regions, versions, report/attestation scope
Does it cover the relevant period and material changes?	report period/change notices
What exceptions or carve-outs exist?	qualifications, subservice exclusions, incidents
What remains customer-controlled?	responsibility/interface matrix
How are provider changes learned?	notification/change process
What can customer telemetry observe?	logs/monitoring
Can sufficient underlying evidence lawfully be accessed if needed?	contractual/access-rights record
What remains unknown?	residual limitation record

Do not turn a provider assertion into an enterprise-observed fact. A third-party certification or assurance report is scoped evidence; its conclusion does not transfer into GAISSF™ conformity.

26.1 Provider-Evidence Triage

For each material evidence objective, assign the operational evidence source:

Source state	Meaning	Practitioner action
Customer-observed	Enterprise can inspect the relevant evidence directly	Index the enterprise source and provenance
Provider-provided	Evidence is supplied by the provider	Record provider scope, period, criteria and limitations
Shared / interface	Both parties hold part of the evidence chain	Link both sides and record the interface dependency
Provider-only / inaccessible	Relevant evidence exists or may exist but cannot be accessed	Record contractual/access limitation and residual unknown
Unknown	It is not yet known who can evidence the objective	Open an evidence gap and assign an owner

This table is an evidence-collection aid, not a control-rating mechanism.

27. AI-Specific Evidence Pattern: Human Oversight

Human oversight should be evidenced as an operating mechanism where operation matters.

Potential evidence:

- role definition;
- decision-rights matrix;
- competency/training;
- review records;
- override records;
- escalation;
- workload/capacity evidence where material;
- sampled decisions;
- exception handling.

A checkbox labeled “human in the loop” is not sufficient evidence of meaningful oversight.

Where human authority is material, align the evidence question with the five-dimension practitioner pattern used in COM-014:

1. **Awareness** — did the person know intervention/review was required?
2. **Comprehension** — could the person understand the relevant state?
3. **Time** — was there enough time to act meaningfully?
4. **Authority** — was the person actually permitted to intervene or decide?
5. **Effectiveness** — did the intervention/review mechanism materially constrain or change the system outcome?

COM-015 does not define universal pass/fail thresholds for those dimensions.

---

28. AI-Specific Evidence Pattern: Incidents and Exceptions

Relevant evidence may include:

- detection or intake record;
- preserved input/output and other volatile evidence;
- system/model/service/version;
- timestamped timeline;
- confirmed facts separated from working hypotheses;
- classification record **where classification is supportable**;
- classification-change history and the evidence that prompted change;
- provider communications clearly identified as provider representations;
- enterprise-observed evidence and provider-side unknowns kept distinguishable;
- investigation;
- containment/recovery;
- exception/risk acceptance;
- corrective action;
- closure validation.

UAIF™ classification boundary: a detection, signal or investigative observation is not itself a UAIF™ record_type. Under UAIF™ v1.0, record_type is an L0 — Record Identity & Workflow field with two normative values: Incident and Vulnerability. Observation is not a normative record_type value. Response and evidence preservation may proceed while classification remains provisional. Do not create additional record_type values or unsupported UAIF fields.

Where UAIF™ or AI-IRF™ is used, follow the current normative technical specification and machine-readable schema. COM-015 does not redefine UAIF classification, severity, legal-reporting status, AI-IRF lifecycle phases, or response controls.

29. AI-Specific Evidence Pattern: Physical AI

Physical-AI evidence may cross digital, cyber-physical and physical domains.

Where material, aim to reconstruct the same perception-to-effect chain used in COM-014:

environment → sensing/perception → inference → autonomy state → command authorization → command → actuator state → safety-controller state → operator action → physical outcome → recovery/revalidation

Potential sources therefore include:

- perception/sensor data and health/calibration state;
- model/software/firmware version;
- autonomy mode and operating-boundary state;
- command authorization and command sequence;
- actuator feedback;
- safety-controller events;
- emergency-stop or degraded-state records;
- environmental conditions;
- maintenance state;
- operator actions and observations;
- physical outcome/consequence records;
- recovery and revalidation records.

Preserve the boundary between AI assurance and sector-specific safety, engineering and regulatory regimes. PAI-SF™ v1.0 is the dedicated physical-AI framework; GAISSF™ Domain 9 remains retained for citation continuity while PAI-SF supersedes and expands that physical-AI security-assurance layer.

An AI assurance evidence pack does not establish functional-safety certification, airworthiness, roadworthiness, medical-device approval, machinery conformity, sector authorization or safe return to service. Security closure and evidence completeness do not substitute for the applicable safety or operational authority.

Physical-AI Evidence Coordination Note: for systems where AI outputs can create or influence physical action, the Evidence Custodian should explicitly coordinate with Engineering, Safety, OT/Operations, and IT to obtain: perception-to-effect chain logs and their integrity state; safe-state or degraded-mode entry records; human override/intervention effectiveness evidence under real operating conditions; and return-to-service authority and revalidation records. COM-015 evidence does not constitute functional-safety certification, product approval, sector authorization, or return-to-

service authority. That authority remains with the applicable Safety, Engineering, Regulatory or Operational owner.

30. The Requirement–Evidence Matrix

The matrix prevents the evidence pack from becoming a disconnected archive.

Recommended fields:

Requirement/control Evidence Evidence Scope Period Owner Limitation/gap Action

objective IDs

Do not place an assessment score or certification decision in the public preparation matrix.

31. Evidence-Pack README

The README should allow an authorized reviewer to understand the pack before opening individual evidence.

Include:

- pack ID/version;
- assessed system/service;
- business context;
- scope;
- exclusions;
- evidence period;
- model/provider/version information where available;
- repository owner;
- access instructions;
- sensitive-material handling;
- secure-reference process;
- known limitations;
- change history.

32. Sensitive Evidence

Collect only what is necessary.

Prefer a secure reference or controlled extract when full sensitive datasets are unnecessary.

Consider:

- personal data;
- secrets/credentials;

- source code;
- confidential supplier material;
- customer data;
- security-sensitive architecture;
- incident evidence;
- regulated records.

Redaction is acceptable only where the evidentiary meaning remains intact and the transformation is documented.

Retention should follow applicable legal, contractual, organizational and scheme requirements. Legal holds and contractual duties may override routine disposal; where disposal is permitted, preserve an appropriate destruction record. This guide does not establish a universal retention period.

32.1 Secure-Reference Manifest

When sensitive evidence remains in a controlled source system instead of being copied into the pack, record enough metadata for an authorized reviewer to locate and understand it without exposing the material unnecessarily.

Recommended fields:

- Evidence ID;
- source system / repository;
- secure path, query or record reference;
- evidence custodian;
- access owner / approval route;
- scope and period;
- sensitivity classification;
- whether redaction or controlled-viewing is required;
- collection or viewing method;
- known access limitations;
- expiry / retention dependency where applicable.

Do not include passwords, tokens, API keys, private keys or other secrets in the manifest.

---

33. Evidence-Pack Quality Review

Boundary: this is an internal preparation-quality check, not a control assessment, conformity determination, confidence assignment, audit opinion, certification-readiness decision or ODA3 readiness designation.

Before handoff, check:

- scope is explicit;
- evidence period is explicit;
- system/model/service/version is identifiable;
- requirement mappings resolve;
- every artifact has an evidence ID;
- repository links resolve;

- superseded evidence is marked;
- transformations/redactions are recorded;
- design evidence is not mislabeled as operating evidence;
- provider representations are distinguishable;
- stale evidence is visible;
- material contradictions are included;
- missing evidence is recorded;
- open exceptions are traceable;
- sensitive evidence is appropriately controlled;
- unsupported “not applicable” decisions are challenged;
- unsupported positive claims are removed.

33.1 Five Common Pack Failures

A pack should normally be repaired before handoff when it exhibits one or more of these patterns:

1. **Unbounded** — the system, version, environment or period is unclear.
2. **Untraceable** — artifacts cannot be tied to source, owner or provenance.
3. **Policy-only** — documented intent is presented as operating evidence.
4. **One-sided** — contrary evidence, exceptions or gaps are omitted.
5. **Uncontrolled** — sensitive evidence, duplicates, redactions or versions cannot be reliably managed.

These are practitioner warning patterns, not controlled assessment findings.

---

34. Worked Example A — Policy-Rich, Operation-Poor

Scenario: An enterprise has an approved AI access-control policy, a role matrix and a quarterly review procedure.

Available evidence: policy, procedure and role design.

Missing: completed access reviews, exception records and evidence of revocation.

Correct interpretation: design appears documented. The pack does not yet demonstrate that the quarterly review operated during the relevant period.

Gap: obtain period-relevant review records or record the limitation.

Notably Absent: no conclusion that access was ineffective merely because operating evidence is missing; equally, no positive operating-effectiveness conclusion is justified from policy volume alone.

34.1 Filled Evidence-Pack Fragment

Illustrative mapping:

Evidence ID	Evidence	Class	Period / scope	Limitation
EV-ACCESS-001	Approved AI access-control policy	Design	Enterprise / current version	Does not show operation
EV-ACCESS-002	Production role export	Technical	Production tenant / 30 Jun	Point-in-time only
EV-ACCESS-003	Quarterly access-	Operational	Q2 / production	Two exceptions

Evidence ID	Evidence	Class	Period / scope	Limitation
	review record		tenant	remain open
EV-ACCESS-004	Revocation ticket for one exception	Corrective action	3 Jul / named account	Does not close second exception

Illustrative requirement–evidence mapping:

Evidence objective	Evidence IDs	Preparation note
Establish documented access-review requirement	EV-ACCESS-001	Design intent visible
Establish configured access population at a point in time	EV-ACCESS-002	Period-wide operation not established
Establish Q2 review activity	EV-ACCESS-003	Review performed; open exceptions visible
Establish remediation of identified exception	EV-ACCESS-004	One exception remediated; one remains unresolved

The example demonstrates traceability. It does **not** assign a confidence level, control determination, conformity result or readiness status.

---

35. Worked Example B — Strong Test, Weak Scope Linkage

Scenario: A red-team report shows extensive adversarial testing.

Problem: the report does not identify whether the tested model configuration matches production.

Correct treatment: retain the report as potentially relevant technical evidence but record the scope/version uncertainty.

Next evidence: deployment/version records, test configuration, production configuration and change history.

Boundary: technical depth does not cure a missing production linkage.

---

36. Worked Example C — Supplier Certificate with Scope Mismatch

Scenario: A hosted-model provider supplies a third-party assurance report.

Problem: the report covers provider infrastructure controls but not the customer's prompt layer, retrieval system, agent permissions or business workflow.

Correct treatment: map the report only to the inherited provider controls it supports.

Customer evidence still needed: integration controls, IAM, retrieval, monitoring, human oversight, application-layer testing and other applicable controls.

Notably Absent: possession of the supplier report does not establish end-to-end assurance.

---

37. Worked Example D — Dashboard with Unverifiable Lineage

Scenario: A polished dashboard reports "100% monitoring coverage."

Problem: the evidence pack contains no telemetry inventory, query logic, data-source mapping or coverage definition.

Correct treatment: the dashboard is a lead or summary until lineage and coverage can be validated.

Next evidence: source inventory, query/export logic, filters, retention, excluded systems and reconciliation.

---

38. Worked Example E — Material Contrary Incident

Scenario: Twelve months of monitoring reports are favorable, but one incident shows that a high-impact alert failed to trigger.

Incorrect approach: exclude the incident as an anomaly.

Correct approach: preserve it as material contrary evidence; document investigation, cause, remediation, retest and residual limitation.

Favorable volume does not erase a material contradiction.

---

39. Worked Example F — AI-Generated Evidence Summary

Scenario: An internal AI assistant summarizes thousands of tickets and reports that all required reviews occurred.

Problem: the summary does not identify source completeness, extraction logic or unsupported inferences.

Correct treatment: use the summary for navigation or triage. Validate the underlying source set and material conclusions.

The generated summary alone is insufficient.

---

40. Worked Example G — Physical-AI Evidence Pack

Scenario: An autonomous inspection robot deviates from its expected route and triggers an emergency stop.

Relevant evidence may include:

- sensor/perception records;
- localization data;
- actuator commands;
- safety-controller record;
- emergency-stop activation;
- model/software version;
- environmental conditions;
- maintenance state;
- operator observation;
- incident/investigation record.

The evidence pack should not infer AI causality merely because the system contains AI. Confirmed facts, working hypotheses and unresolved causes should remain distinguishable. It should also not determine that the robot is safe to return to service. Those conclusions require the applicable technical, security, safety and sector processes.

41. Practitioner Tool — Build Flow and Quick Start

41.1 Eight-Step Build Flow

1. **Bound the pack** — system/service, environment, version, period, exclusions and prior applicability decisions.
2. **Define evidence objectives** — state what each requirement needs the evidence to help establish.
3. **Identify sources** — enterprise systems, people, providers, repositories and secure references.
4. **Index before copying** — assign stable Evidence IDs and metadata.
5. **Collect and preserve context** — provenance, period, version, transformations and sensitivity.
6. **Challenge the pack** — look for stale material, policy-only evidence, contradictions, exceptions, missing evidence and provider unknowns.
7. **Package the handoff** — README, evidence index, requirement–evidence matrix, gap register, secure-reference manifest and change log.
8. **Freeze the handoff version** — record exactly what was supplied; handle later evidence through an addendum/change set.

41.2 Illustrative 60-Minute First Pass

This is an orientation exercise for a **small pilot slice**, not a mandatory response time, assessment SLA or complete-pack target.

Time	Action	Output
0–10 min	Choose one in-scope system and one evidence period	Draft pack identity
10–20 min	Choose one requirement/control and define one evidence objective	One matrix row
20–35 min	Locate 3–5 candidate artifacts or secure references	Initial Evidence IDs
35–45 min	Check source, period, scope, provenance and limitations	Evidence notes
45–55 min	Record contrary/missing evidence and provider unknowns	Gap entries
55–60 min	Write a short README and freeze the pilot snapshot	Mini handoff pack

The exercise is successful if it exposes what evidence exists, what it can support, and what remains unresolved. It is not successful merely because all boxes are filled.

41.3 Operating Roles

A single person may hold several roles in a small organization. The roles are separated here so accountability remains visible.

Role	Primary responsibility in the evidence pack
Evidence Pack Owner	Coordinates scope, index, repository, version and handoff
System / Service Owner	Confirms system identity, environment, changes and dependencies
Control Owner	Explains the control objective, implementation and exceptions
Evidence Custodian	Provides or securely references source records
AI / Security SME	Explains technical meaning and limitations
Provider Liaison	Obtains provider evidence, scope, change notices and access information
Privacy / Legal / Compliance	Advises on lawful access, minimization, redaction and obligations where applicable
Internal Assurance / Review	Challenges traceability, contradiction, gaps and claims without becoming the controlled ODA3 assessment process

Where an external party is engaged, GEL v1.0 and applicable ODA3 authorization/partner rules govern whether that activity is permitted as Commercial Use.

41.4 Evidence-Pack Build Checklist

- ☐ System/service and scope identified.
- ☐ Evidence period recorded.
- ☐ Prior applicability/profile decisions referenced rather than silently re-created.
- ☐ Evidence objectives defined.
- ☐ Likely evidence sources and custodians identified.
- ☐ Provider/customer/shared/unknown evidence boundaries recorded.
- ☐ Stable Evidence IDs assigned before uncontrolled copying.
- ☐ Provenance, period, version and transformations recorded.
- ☐ Sensitive material minimized or securely referenced.
- ☐ Design evidence separated from operating evidence.
- ☐ Contrary evidence and exceptions retained.
- ☐ Missing evidence and access limitations recorded.
- ☐ Requirement–evidence matrix completed for the selected scope.
- ☐ Open gaps assigned to owners.
- ☐ Evidence-Pack Quality Review performed.
- ☐ README, index, gap register and change log included.
- ☐ Handoff version frozen or snapshotted.
- ☐ Subsequent changes handled by addendum/change set rather than silent replacement.

42. Practitioner Tool — Evidence Index Template

Field	Classification	Entry
Evidence ID*	GAISSF public minimum	
Control ID / requirement linkage*	GAISSF public minimum	
Scope element*	GAISSF public minimum	
Owner / custodian*	GAISSF public minimum	
Evidence period*	GAISSF public minimum	
Source system / process*	GAISSF public minimum	
Integrity note*	GAISSF public minimum — hash, signature, governed repository control, or other applicable integrity mechanism	
Status*	GAISSF public minimum	
Location / secure reference*	GAISSF public minimum	
Known limitations*	GAISSF public minimum	
Evidence title	COM-015 preparation enrichment	
Evidence class	COM-015 preparation enrichment	
Version	COM-015 preparation enrichment	
Collection date	COM-015 preparation enrichment	
Provenance detail	COM-015 preparation enrichment	
Transformation / redaction note	COM-015 preparation enrichment	
Related evidence IDs	COM-015 preparation enrichment	
Sensitivity / access note	COM-015 preparation enrichment	
Source-system record ID	COM-015 preparation enrichment	
Provider/customer/shared state	COM-015 preparation enrichment	

* identifies fields carried from the current GAISSF™ Adoption Guide public minimum Evidence Index. All other fields are COM-015 preparation enrichments and do not modify that source requirement.

43. Practitioner Tool — Requirement–Evidence Matrix

Requirement / control	Evidence objective	Evidence IDs	Scope	Period	Owner	Limitation / gap	Action

Use rule: one requirement may have several evidence objectives, and one Evidence ID may support several objectives. Keep the mapping rationale bounded to what the artifact can actually support.

44. Practitioner Tool — Evidence Gap Register

Gap ID	Requirement / objective	Missing or limited evidence	Reason	Consequence for handoff	Owner	Action	Status

Suggested internal status values may be simple document-management terms such as **Open, Actioned, Awaiting Evidence, Accepted as Limitation, Closed**. They are not control-assessment determinations or ODA3 status values.

45. Practitioner Tool — Evidence Item Review Card

For each material item ask:

- **Identity:** Is it uniquely identifiable?
- **Attribution:** Is its source and owner known?
- **Scope:** Does it apply to the assessed unit/environment?
- **Period:** Does it represent the relevant state/period?
- **Relevance:** Does it address the evidence objective?
- **Independence:** What reliance is placed on its source?
- **Currency:** Has material change made it stale?
- **Coverage:** What population, period, cases or environments are represented?
- **Integrity:** Can provenance and transformation be assessed?
- **Contrary evidence:** Is there conflicting evidence?
- **Limitations:** What does this item not establish?
- **Related evidence:** What corroboration is needed?

No numeric score is assigned by this public card.

A practical note can be recorded as:

- **Use as contextual support**
- **Use with stated limitation**
- **Seek additional corroboration**
- **Replace / refresh source**
- **Leave unresolved and record the gap**

These are evidence-preparation notes only. They are not assessor dispositions, confidence categories or control determinations.

46. Practitioner Tool — Evidence-Pack Handoff

46.1 Handoff Manifest

Record:

- Pack ID and version;
- system/service and environment;
- evidence period;
- handoff date/time;
- pack owner and contact;
- evidence count or index version;
- open gaps and accepted limitations;
- sensitive / secure-reference items;
- known exclusions;
- provider-access constraints;
- addenda/change sets issued after handoff.

Where practical, preserve an integrity reference for the handed-off package or its manifest. This identifies what was supplied; it does not prove that every underlying source was truthful or complete.

46.2 Handoff Checklist

Grouped for scannability at handoff. All conditions carried unchanged from the original flat checklist; grouping and headers are new organizational structure, not new requirements. The evidence-architecture chain restated under Group A is drawn verbatim from Section 1 (new placement, not new wording).

A. Scope & Traceability

- requirement → evidence objective → source → artifact → provenance → scope and period → preparation-quality review → limitation or contradiction → later assurance / assessment use (Sec 1)
- ☐ Scope and evidence period explicit.
- ☐ System/model/service/version identifiable.
- ☐ Applicable requirements or prior applicability decisions referenced.
- ☐ Evidence objectives defined.
- ☐ Every submitted artifact has an Evidence ID.
- ☐ Evidence index complete for the selected scope.
- ☐ Evidence locations and secure references resolve for authorized users.

B. Evidence Quality

- ☐ Design evidence not presented as operating evidence.

C. Third-Party & Sensitivity

- ☐ Provider representations distinguished from enterprise observations.
- ☐ Provider/customer/shared/unknown responsibility boundary visible.
- ☐ Sensitive evidence minimized and access-controlled.

D. Transparency & Gaps

- ☐ Material contrary evidence included.
- ☐ Missing evidence/access limitations visible.

- ☐ Exceptions and corrective actions traceable.
- ☐ Open gaps and limitations summarized in the README/manifest.
- ☐ “Notably Absent” statements identify search basis, period, coverage and limitations.

E. Integrity & Version Control

- ☐ Superseded evidence identified.
- ☐ Material transformations/redactions recorded.
- ☐ Pack version frozen/snapshotted; subsequent change process defined.

F. Handoff Boundary

- ☐ No unsupported certification/compliance/safety/readiness claim.

---

47. What an Evidence Pack Does Not Prove

The evidence pack itself does not prove:

- conformity;
- certification;
- regulatory compliance;
- legal compliance;
- system safety;
- absence of incidents;
- absence of hidden operation;
- absence of vulnerabilities;
- operating effectiveness;
- maturity;
- optimization;
- regulator recognition;
- suitability for a particular certification route;
- an official status of *assessment-ready*, *audit-ready*, *certification-ready*, *certifiable*, *approved* or *endorsed*; or
- authority for a third party to perform ODA3-branded assessment, audit, certification or accreditation services.

A completed evidence pack or completed checklist is therefore a preparation artifact, not a status designation. Those conclusions or authorities require the applicable assessment, legal, regulatory, safety, licensing or scheme process.

---

48. Controlled-Methodology Firewall

COM-015 may explain:

- evidence classes;
- evidence indexing;
- provenance;
- scope and period linkage;

- qualitative evidence review dimensions;
- contrary evidence;
- limitations;
- gaps;
- secure handling;
- evidence preparation and handoff.

COM-015 must **not** publish or invent:

- the controlled ASM-003 confidence scale, confidence caps or confidence-assignment rules;
- final certification-route confidence thresholds;
- mandatory-control confidence floors;
- numeric weighting or averaging of evidence dimensions;
- universal sample-size formulas;
- universal minimum operating periods;
- controlled assessor decision rules;
- certification decision logic;
- proprietary test protocols;
- unapproved evidence-retention durations;
- official *assessment-ready*, *audit-ready* or *certification-ready* designations;
- authority to offer ODA3-branded assessment, audit, certification or accreditation services.

Claims Firewall

COM-015 may support a factual statement such as:

“The organization has assembled an evidence pack mapped to identified GAISSF™ controls for internal review.”

COM-015 does **not** by itself support statements such as:

- “ODA3 assessment-ready”;
- “GAISSF audit-ready”;
- “certification-ready”;
- “GAISSF compliant”;
- “GAISSF certified”;
- “ODA3 approved”;
- “ODA3 endorsed”; or
- “independently assured”.

Where factual alignment is stated, the statement should remain bounded to the identified scope, artifact, version and evidence actually established. It must not imply official ODA3 audit, certification, endorsement or approval.

---

49. Source and Traceability Boundary

This guide is an informative practitioner publication. It is subordinate to the current controlling GAISSF Ecosystem framework, control, schema, assessment-methodology, licensing and publication-state sources.

Where a statement in this guide conflicts with a controlling source, the controlling source governs.

49.1 Primary Evidence-Architecture Sources

GAISSF-IMP-009 — GAISSF™ v1.0 Adoption Guide, Final Publication Edition.
Used for the six GAISSF adoption evidence classes, public minimum Evidence Index fields, evidence-period concepts and adoption-guide limitations.

ODA3-2026-08-ASM-003 — GAISSF™ Evidence Sufficiency and Confidence Standard (controlled source).
Used only for public-safe concepts concerning evidence admissibility, provenance, qualitative evidence-quality dimensions, corroboration, contrary and missing evidence, AI-generated evidence, third-party reliance, minimization and limitations.

ASM-003 remains controlled. COM-015 does not reproduce its confidence scale or caps, control-determination pathways, assessor workpaper/disposition fields, certification decision roles or other controlled assessment logic.

49.2 GAISSF™ Framework Sources

GAISSF-NOR-001 — GAISSF™ v1.0 Framework Standard.
Primary normative framework source where current and applicable.

GAISSF-NOR-004 — GAISSF™ v1.0 Control Catalogue, Final Publication Edition.
Current control-library baseline used for control references.

GAISSF-NOR-007 — GAISSF™ v1.0 Release Notes, Final Publication Edition.
Used for publication-state and supersession context.

GAISSF-NOR-008 — GAISSF™ v1.0 Change Log.
Used to avoid treating superseded control-set artifacts as current authority.

49.3 Practitioner and Framework-Boundary Sources

- **ODA3-2026-08-WHP-COM-009 — The First 30 Days with GAISSF™**
- **ODA3-2026-08-WHP-COM-010 — Applying GAISSF™: Worked Scenarios**
- **ODA3-2026-08-WHP-COM-011 — When AI Breaks: The ODA3 Guide to AI Incident Response**
- **ODA3-2026-08-WHP-COM-012 — Classifying AI Incidents: A Practical Guide to UAIF™**
- **ODA3-2026-08-WHP-COM-013 — AI Incident Response Field Guide: Using UAIF™ and AI-IRF™ for Real-World AI Incidents**
- **ODA3-2026-08-WHP-COM-014 — Securing Physical AI: A Practitioner's Guide to PAI-SF™**

These publications support practitioner routing and boundary consistency. They do not override the controlling normative, technical, schema or assessment-methodology sources.

COM-012 confirms that the publication status of a practitioner guide and the maturity state of an underlying framework are separate controls. Accordingly, COM-015 does not infer framework maturity or certification availability from the publication state of a WHP guide.

49.4 Physical-AI Boundary

PAI-SF™ v1.0 is the dedicated physical-AI framework. Its current public structure is **12 domains and 41 controls**. GAISSF™ Domain 9 remains retained for citation continuity while PAI-SF supersedes and expands the physical-AI security-assurance layer represented by D9.

PAI-SF interfaces with functional safety and sector assurance; it does not replace them. COM-015 therefore treats physical-AI evidence as an assurance input and does not create product approval, sector authorization, functional-safety certification or return-to-service authority.

49.5 GEL v1.0 Claims and Use Boundary

GAISSF Ecosystem License (GEL) v1.0 governs use of the licensed material.

For COM-015, the relevant public boundaries are:

- factual alignment or mapping does not itself constitute ODA3 certification, endorsement, audit or approval;
- public availability or redistribution does not itself grant Commercial Use rights;
- commercial consulting, implementation, assessment, audit, compliance, certification, training, software/SaaS integration and client-delivery use remain subject to applicable GEL terms and written authorization;
- no party may present a modified or derived implementation as official, ODA3-endorsed or ODA3-certified without authorization.

49.6 Naming and Publication Convention

COM-015 uses the current market-facing names and marks:

- **GAISSF™ — Global AI Safety and Security Framework**
- **UAIF™ — Unified AI Incident Framework**
- **AI-IRF™ — AI Incident Response Framework**
- **PAI-SF™ — Physical AI Security Framework**
- **GAISSF Ecosystem** — umbrella term for the integrated ODA3 framework ecosystem.

Where a historical technical source uses earlier naming or mark conventions, the current public naming convention is used in COM-015 without changing the technical authority of the underlying source.

---

50. Notably Absent

This guide deliberately does **not** establish or publish:

- a controlled evidence-confidence scale or confidence-cap logic;
- assessor workpaper, disposition or control-determination logic;
- certification-route confidence thresholds;
- mandatory-control confidence floors;
- numeric weighting of evidence-quality dimensions;
- a universal sampling formula;
- a universal minimum operating period;
- a universal evidence-retention period;
- a rule that any evidence class alone establishes conformity;
- a rule that absence of observed evidence proves absence of an event;
- a framework applicability/profile determination for a real system;
- an ODA3 *assessment-ready*, *audit-ready* or *certification-ready* designation;
- regulator recognition, certification or accreditation status;

- commercial assessment, audit, certification or client-delivery rights;
- legal or regulatory conclusions;
- product-safety, sector-approval or physical-AI return-to-service authority.

The guide also does not treat a 60-minute exercise, role assignment, internal pack state, review-card note, checklist completion or handoff manifest as an assessment result or ODA3 status.

51. Next Reading

1. **The First 30 Days with GAISSF™** — ODA3-2026-08-WHP-COM-009
2. **Applying GAISSF™: Worked Scenarios** — ODA3-2026-08-WHP-COM-010
3. **When AI Breaks: The ODA3 Guide to AI Incident Response** — ODA3-2026-08-WHP-COM-011
4. **Classifying AI Incidents: A Practical Guide to UAIF™** — ODA3-2026-08-WHP-COM-012
5. **AI Incident Response Field Guide: Using UAIF™ and AI-IRF™ for Real-World AI Incidents** — ODA3-2026-08-WHP-COM-013
6. **Securing Physical AI: A Practitioner's Guide to PAI-SF™** — ODA3-2026-08-WHP-COM-014
7. **The GAISSF™ Assessment Field Guide** — ODA3-2026-08-WHP-COM-016 — forthcoming
8. **Implementing GAISSF™: A Practitioner's Guide** — ODA3-2026-08-WHP-COM-017 — forthcoming

GEL Attribution and Use Boundary

GAISSF™, UAIF™, AI-IRF™, and PAI-SF™ are frameworks of ODA3 Institute and are referenced under the GAISSF Ecosystem License (GEL) v1.0.

This attribution does not itself authorize unrestricted commercial delivery, assessment, audit, certification, endorsement, training, software embedding or client delivery. Applicable GEL terms and any current controlled-service or partner rules govern those uses.

Published by ODA3 Institute.

Final Publication Statement

This document is the released **FINAL** edition of **ODA3-2026-08-WHP-COM-015 — Building an AI Assurance Evidence Pack**.

FINAL status applies to this practitioner guide only. It does not create or alter framework maturity, certification availability, regulator recognition, commercial delivery rights, product approval, sector authorization or any controlled assessment determination.