ODA3 INSTITUTE

PRACTITIONER FIELD GUIDE

THE GAISSF™ ASSESSMENT FIELD GUIDE

A Practitioner's Guide to the Assessment Lifecycle, Evidence Review and Decision Boundaries

Doc ID: ODA3-2026-08-WHP-COM-016

Framework Context: GAISSF Ecosystem

Classification: Public — GEL v1.0

Status: FINAL

Date: 18 August 2026

ODA3 INSTITUTE

THE GAISSF™ ASSESSMENT FIELD GUIDE

A Practitioner's Guide to the Assessment Lifecycle, Evidence Review and Decision Boundaries

Doc ID: ODA3-2026-08-WHP-COM-016
Document Type: WHP — Adoption / Practitioner Field Guide
Framework Context: GAISSF Ecosystem
Classification: Public — GEL v1.0
Status: FINAL
Date: 13 August 2026
Publisher: ODA3 Institute

Audience: CISOs, AI Governance Leads, Security Architects, System and Service Owners, Control Owners, Evidence Custodians, Risk and Compliance Teams, Internal Assurance Teams, Assessment Coordinators, GRC Teams, Executive Sponsors

---

Publication Boundary

This is an informative practitioner field guide about the **GAISSF assessment lifecycle**. It is intended to help an organization understand what happens after evidence preparation and how to participate in an assessment without confusing assessment preparation, assessment work, technical review, reporting, certification, or regulatory approval.

This guide does **not** create, modify, replace, reproduce, or supersede:

- GAISSF™ normative requirements or control statements;
- controlled GAISSF assessment methodology;
- assessor workpaper procedures;
- evidence-confidence categories or caps;
- control-determination categories or decision rules;
- capability or maturity scales;
- sampling formulas or sample-size rules;
- aggregation or gate logic;
- finding-severity algorithms or thresholds;
- assessor calibration rules;
- certification eligibility thresholds;
- certification decision logic; or
- scheme rules.

The controlled methodology remains authoritative for authorized assessment work. Where this guide conflicts with a current controlled ODA3 source, the controlled source governs.

An assessment report is distinct from a certification decision.

The current controlled assessment-methodology family is a **design and pilot-development baseline**. The controlling methodology register states that live assessment is not approved, and the

current certification-interface baseline does not authorize certificate issuance. Accordingly, COM-016 must not be read as announcing a live ODA3 assessment service, an issued GAISSF certification scheme, or current certification availability.

Where a later exact methodology version receives bounded operational approval, or an issued scheme is subsequently approved, the applicable approval record and scheme rules—not this guide—will govern that activity.

GEL and Commercial-Delivery Boundary

Public availability of this guide does not grant commercial assessment, audit, certification, implementation, paid training, software-embedding or client-delivery rights.

Where a third party performs work using GAISSF Ecosystem licensed material as a commercial methodology, applicable GEL v1.0 terms and current ODA3 partner/service authorization rules govern.

This guide does not authorize any person or organization to represent itself as an ODA3 assessor, auditor, certification body, accredited body, approved provider, endorsed partner or equivalent.

Methodology Note

COM-016 is written from the **practitioner / assessed-organization perspective** rather than as a public assessor manual.

It explains:

- what an assessment needs from the organization;
- how scope and evidence handoff interact;
- why evidence may be challenged or supplemented;
- what kinds of assessment activity an organization may encounter;
- how contrary evidence, exceptions and missing evidence affect the process;
- why findings and remediation remain traceable;
- why independent technical review exists;
- what an assessment report can and cannot establish; and
- why certification, where applicable under an issued scheme, is a separate decision.

It does **not** tell an assessor exactly how to assign controlled results.

The public practitioner journey used in this guide is deliberately higher-level:

prepare → participate → clarify → remediate where needed → understand the bounded outcome

This is **not** the controlled assessor workflow and should not be used to infer one. The exact ordering, phase gates, procedure-selection logic, sampling, decision points, review states, workpaper fields and authorized outcomes remain controlled.

The guide may describe activities an organization can encounter—such as evidence requests, interviews, technical testing, factual-accuracy review, remediation, independent review and reporting—without specifying the controlled sequence in which an assessor must perform or disposition them.

Current-state note: the controlled methodology register distinguishes design-baseline status, pilot authorization and bounded operational approval. Approval is version-specific. In the source baseline verified for this publication, the methodology design is controlled and pilot materials are prepared, but live conformity assessment and certification issuance are not approved.

How to Use This Guide

You do not need to read all 63 sections before an assessment.

Use COM-016 in three ways:

A. Prepare for participation

Start with:

- **Section 3** — purpose and intended reliance;
- **Section 4** — scope, object, units and period;
- **Section 6** — roles;
- **Section 7** — evidence handoff; and
- **Section 46** — participation checklist.

B. Respond during assessment activity

Use:

- **Sections 8–15** when evidence is challenged;
- **Sections 17–21** when you need to understand a request, interview, walkthrough or test;
- **Section 24** when an exception appears;
- **Sections 26–30** for factual correction, response, remediation and later evidence; and
- **Sections 48–54** for organization-side trackers.

C. Understand the bounded outcome

Use:

- **Section 31** — independent technical review;
- **Section 32** — closing discussion;
- **Sections 33–36** — report, limitations, certification separation and claims; and
- **Section 57** — reassessment triggers.

For agentic or physical AI, also use **Sections 22–23**.
For incident-related evidence, use **Section 21.3** together with COM-012 / COM-013.

Practitioner Participation Map

The following is an **organization-side orientation**, not an assessor workflow or controlled methodology:

Before participation	While responding	Before relying on the outcome
confirm purpose and scope	answer the exact question asked	reconcile factual accuracy
nominate one coordinator	link responses to evidence IDs	confirm unresolved limitations
freeze/version the evidence handoff	distinguish fact, representation and unknown	understand what technical review can still change
disclose material changes	preserve contrary evidence	read scope, period and limitations with the conclusion
identify provider/access constraints	log new or late evidence	keep assessment and certification separate
prepare secure access routes	escalate material change or	use only authorized external

Before participation	While responding	Before relying on the outcome
	safety issues	claims

Five Practitioner Rules

1. **Answer the evidence question, not the hoped-for outcome.**
2. **Keep the assessed period separate from later remediation.**
3. **Do not hide contradictions, provider unknowns or inaccessible evidence.**
4. **Route material scope change early rather than patching the pack silently.**
5. **Never detach a conclusion from its scope, period and limitations.**

The practitioner-side **scope record, evidence handoff, query tracker, remediation tracker and subsequent-evidence log** are navigation aids for the assessed organization. They are not ODA3 assessor workpapers and do not reproduce the controlled assessment record set.

PART I — BEFORE ASSESSMENT

1. Why Assessment Is Not Document Review

A weak assessment starts with a document room and asks whether enough files exist.

A stronger assessment starts with a bounded question:

What was assessed, against which criteria, for which system and period, using what evidence and procedures, and what conclusion can that record support?

The difference matters because the same artifact can be:

- relevant to one control but not another;
- current for one system version but stale for another;
- strong evidence of design but weak evidence of operation;
- useful context but insufficient for a material conclusion;
- provider evidence that does not cover the customer's integration;
- valid point-in-time evidence but not representative of an operating period; or
- contradicted by another artifact.

An assessment therefore cannot be reduced to checking whether policies, screenshots, logs or certificates exist.

The practitioner-visible relationship is simpler:

defined criteria + bounded scope + evidence claims + evidence items + known limitations → controlled assessment use

COM-015 organizes and prepares those evidence inputs. COM-016 explains why they may then be challenged, supplemented, tested or limited during assessment participation.

The exact assessor procedure/result chain remains controlled and is not reproduced here.

1.1 Three misconceptions to remove early

Misconception 1 — "We submitted the evidence, so the control should pass."

Submission is not determination. An artifact can be authentic and relevant yet still not cover the required period, population or operating condition.

Misconception 2 — "The assessor is asking for more evidence because our first pack was poor."

Sometimes. But additional evidence can also be a normal consequence of:

- sampling;
- contradiction;
- material change;
- third-party reliance;
- a new dependency;
- a test exception;
- insufficient operating opportunities;
- a scope question; or
- the need to distinguish design from operation.

Misconception 3 — "A favorable assessment means certification."

Assessment and certification are separate. A report may be an input to a certification function only where an issued scheme applies. A favorable assessment is not self-executing certification.

1.2 Five Common Participation Failures

1. Outcome-first answering

The organization answers what it wants the assessor to conclude instead of the exact factual or evidentiary question.

2. Evidence flooding

Teams upload large volumes of loosely related material instead of linking a bounded claim to the strongest available source.

3. Silent scope drift

A model, provider, tool, environment or permission changes, but the assessment team continues using the original scope record.

4. Historical rewriting

Later remediation is presented as though it operated throughout the assessed period.

5. Informal side-channel resolution

Material questions, disagreements or evidence are handled in chat, meetings or email without being linked back to the governed assessment-response record.

These are practitioner failure patterns, not formal assessment findings.

---

2. Suite Position and Handoff

The practitioner sequence is:

scope / applicability → implementation activity → evidence generation → COM-015 evidence preparation → COM-016 assessment participation / use → bounded reporting, remediation or reassessment as applicable

COM-015 and COM-016 solve different problems.

Guide	Primary question	Output
COM-015	How should we organize defensible evidence?	Versioned evidence handoff / preparation pack
COM-016	What happens when that evidence is assessed?	Practitioner understanding of assessment participation and boundaries

COM-016 should not be used to reverse-engineer controlled assessor logic. It should be used to make the organization a more effective, accurate and disciplined participant.

2.1 What a good COM-015 handoff gives the assessment

A strong handoff normally makes visible:

- exact pack version;
- system/service identity;
- environment and deployment context;
- evidence period;
- prior scope/applicability decisions;
- requirement–evidence mapping;
- evidence index;
- provenance;
- secure references;
- provider/customer responsibility boundaries;
- known limitations;
- contrary evidence;
- open gaps;
- unresolved provider access;
- change history; and
- the frozen handoff state.

These inputs reduce avoidable discovery work. They do not eliminate assessment procedures.

Scope handoff rule: COM-016 consumes prior scope/applicability records; it does not silently re-perform the COM-009/COM-010 reasoning process. If assessment evidence reveals a material dependency, exclusion, change or contradiction that makes the recorded scope unreliable, that issue is routed back for governed scope/applicability reconsideration.

2.2 The handoff does not freeze reality

The evidence pack may be frozen for version control, but the AI system may continue to change.

If a material model, provider, configuration, retrieval, tool, permission, architecture or operating-environment change occurs after handoff, the organization must surface it. Silence creates a scope and temporal integrity problem.

2.3 Adjacent Practitioner-Guide Boundaries

Document	Role relative to COM-016	Consistency rule
COM-009 — The First 30 Days with GAISSF™	Initial implementation planning / management checkpoint	Day 30 is a management checkpoint, not an assurance or conformance conclusion. COM-016 does not convert elapsed time into assessment readiness.
COM-010 — Applying GAISSF™: Worked Scenarios	Hypothetical scoping and application reasoning	Scenarios are teaching aids. COM-016 does not infer applicability, sufficiency or assessment outcome from similarity to a scenario.
COM-012 — Classifying AI Incidents	UAIF™ classification and uncertainty discipline	COM-016 may consume incident records as evidence but does not invent UAIF classifications or legal-reportability determinations.
COM-013 — AI Incident Response Field Guide	UAIF™ / AI-IRF™ incident-response integration	Assessment does not replace incident classification or response. Facts, hypotheses, provider representations, unknowns and later reclassification remain traceable.
COM-014 — Securing Physical AI	PAI-SF™ physical-AI assurance and safety boundary	COM-016 preserves the D9/PAI-SF boundary and keeps assurance closure separate from sector safety / operational authorization.
COM-015 — Building an AI Assurance Evidence Pack	Evidence preparation and handoff	COM-016 consumes the evidence pack but does not treat pack completeness as sufficiency, conformity or readiness.
COM-017 — Implementing GAISSF™: A Practitioner's Guide	Forthcoming deeper implementation companion	Implementation guidance belongs in COM-017. Implementation outputs may become evidence; COM-016 does not prescribe implementation design merely to obtain a favorable assessment.

3. Engagement Purpose and Intended Reliance

Before fieldwork starts, the organization should understand **why** the assessment is being performed.

A practical engagement-purpose record should answer:

- What decision will the assessment inform?
- Who expects to rely on the result?
- What assessment purpose or route is intended, and is that route actually authorized under the current ODA3 methodology / scheme state?
- What system or service is in scope?
- What time period is relevant?
- What external claims are expected or prohibited?
- Are there regulators, customers, boards, insurers, investors or contracting parties who may receive the result?
- Are there confidentiality or distribution restrictions?
- Is a specific business deadline driving the work?
- What happens if the assessment cannot reach a defensible conclusion?

3.1 Intended reliance shapes rigor, not truth

A higher-stakes use can require deeper evidence and stronger assurance activity, but intended reliance does not change the underlying facts.

An assessment performed for board decision support, a major customer, a regulated deployment or a future certification interface may require greater rigor than a narrow internal diagnostic review. That difference must be governed by the authorized assessment process—not improvised by the assessed organization.

3.2 Avoid outcome-driven engagement language

Problematic objectives include:

- “get us certified quickly”;
- “confirm we comply”;
- “produce a clean report for the board”;
- “avoid findings before the customer review”; or
- “show that our controls are effective.”

A better formulation is:

Perform the authorized assessment for the declared scope and period and report the supported result, findings, limitations and unresolved uncertainty.

The purpose of an assessment is not to manufacture a favorable outcome. It is to create a defensible one.

---

4. Assessment Object, Units, Period and Scope

An assessment conclusion is meaningful only when attached to an identifiable object.

The organization should be prepared to describe the same major scope dimensions used by the controlled scope architecture:

- **legal and organizational** — legal entity, accountable executive, business units, functions, owners and shared services;
- **technical** — AI systems, models, agents, tools, datasets, APIs, infrastructure, interfaces, safety mechanisms and security tooling;

- **lifecycle** — relevant design, procurement, development, validation, deployment, operation, change, incident-response and retirement stages;
- **environment** — development, test, staging, production, recovery and relevant external or field environments;
- **geographic and jurisdictional** — locations, data regions, operating territories, user populations and materially relevant legal constraints;
- **temporal** — assessment period, operating period, evidence cutoff, known change windows and point-in-time elements;
- **third party** — suppliers, hosted services, foundation-model providers, integrators, data providers and outsourced control operators; and
- **use and exposure** — intended and reasonably foreseeable use, prohibited use, threat exposure, impact and materially affected stakeholders.

Also record explicit exclusions and dependencies that cross the proposed boundary.

These dimensions describe a factual assessment boundary. They do **not** make COM-016 a legal-applicability or regulatory-compliance determination.

4.1 Assessment object vs assessment unit

The **assessment object** is the bounded thing the conclusion is about.

An **assessment unit** is a specific implementation or operating component within that object against which evidence may need to be examined.

A single object can therefore include several units:

- central governance;
- one production application;
- one cloud-hosted model service;
- a retrieval layer;
- an identity service;
- a provider-operated control;
- a regional deployment; or
- a local operational team.

This distinction matters because evidence for a central policy may not establish implementation at every unit.

4.2 Period is part of scope

A current screenshot cannot automatically prove what operated six months ago.

A historical report cannot automatically prove the current deployment.

Record:

- assessment period;
- evidence cutoff;
- important deployment changes;
- major incidents;
- material provider changes;
- migration dates;
- policy/control changes; and
- periods where the control had no opportunity to operate.

4.3 Scope exclusions require discipline

An exclusion is not automatically a favorable assessment fact.

If an excluded system, geography, provider or process is materially connected to the assessed object, the assessment may need to consider whether the exclusion makes the remaining claim misleading.

Practitioners should not respond to adverse evidence by narrowing scope after the fact without transparent governance.

---

5. Scope Change and Material Change

Scope is not a one-time form.

Assessment integrity depends on detecting when the declared boundary no longer represents the system being assessed.

Potential material-change triggers include:

- model replacement;
- major model-version change;
- provider migration;
- new fine-tuning;
- retrieval-source change;
- new external tool or plugin;
- agent authority expansion;
- privileged-access change;
- new data class;
- new geography;
- new customer or business process;
- new physical operating environment;
- firmware/controller change;
- autonomy-mode change;
- incident exposing previously unknown behavior;
- major remediation;
- architecture redesign; or
- acquisition or integration of another service.

5.1 What the organization should do

When a material change occurs:

1. record the change;
2. identify the effective date;
3. identify affected systems, controls and evidence;
4. tell the assessment coordinator promptly;
5. identify whether prior evidence is now stale or only partially representative;
6. identify new provider or system dependencies;
7. preserve pre-change evidence where relevant; and

8. allow the authorized assessment process to determine whether scope, testing or timing must change.

Do not silently update the evidence pack and erase the prior state.

5.2 Change can improve a control and still complicate the assessment

A remediation or system upgrade may genuinely improve security. It can also make it impossible to use pre-change evidence to prove post-change operation—or vice versa.

The correct response is temporal precision, not optimistic blending.

---

6. Roles and Participation

Assessment is easier when every participant knows what they own.

Participant	Practitioner responsibility
Executive Sponsor	Confirms authority, purpose, escalation path and organizational commitment
Assessment Coordinator	Coordinates access, evidence, meetings, questions, changes and responses
System / Service Owner	Confirms architecture, identity, operating context and material changes
Control Owner	Explains control design, operation, exceptions and remediation
Evidence Custodian	Provides authoritative records or controlled access
Provider Liaison	Coordinates provider evidence, responsibilities and limitations
AI / Security SME	Explains technical behavior, controls, evaluations and attack/abuse context
Privacy / Legal / Compliance	Advises on lawful access, confidentiality and external obligations
Internal Assurance	Supports factual traceability and challenge without substituting for independent assessment
Authorized Assessment Personnel	Perform the controlled assessment work
Independent Technical Reviewer	Challenges the file and issuance basis under controlled rules
Certification Function	Separate decision role only where an issued scheme applies

6.1 One person can hold several internal roles

Small organizations may not have a separate evidence custodian, control owner and assessment coordinator.

That is acceptable as an operating reality, provided the responsibilities are still clear.

The organization should not, however, assume that combining internal roles removes the independence requirements applied to external assessment, technical review or certification decision-making.

6.2 Keep response ownership clear

Every evidence request, query, factual correction and remediation action should have one accountable response owner.

A useful rule is:

One owner may coordinate the answer; several people may contribute the evidence.

This prevents contradictory responses from different teams.

6.3 Minimum Participation Packet

Before substantive assessment activity begins, the organization should be able to produce or point to:

- one current scope statement;
- one named assessment coordinator;
- one current system/service architecture reference;
- one evidence-handoff version;
- one list of system/control owners;
- one provider/dependency list;
- one secure-access route for sensitive evidence;
- one material-change log;
- one known-gap / limitation summary; and
- one incident / major-event reference list where relevant.

This is a coordination aid, not an assessment-readiness gate.

---

7. Evidence Handoff from COM-015

The evidence handoff should be treated as a controlled snapshot.

Recommended handoff components:

1. README / scope statement;
2. pack ID and version;
3. assessment-object reference;
4. evidence period;
5. evidence index;
6. requirement–evidence matrix;
7. secure-reference manifest;
8. provider-evidence boundary;
9. gap register;
10. limitation summary;
11. contrary-evidence summary;
12. change log;
13. sensitive-evidence handling note; and
14. pack handoff manifest.

7.1 The assessor should not have to infer the evidence architecture

A folder named audit_docs_final_v4 is not an evidence architecture.

A reviewer should be able to answer:

- What is this artifact?
- Who owns it?
- What period does it cover?
- What requirement or evidence claim is it intended to support?
- What system/version does it apply to?
- What is its source?
- Has it been transformed or redacted?
- What are its limitations?
- Is it provider-provided or enterprise-observed?
- Is there related contradictory evidence?

7.2 New evidence after handoff

Later evidence should arrive through an addendum, request-response, subsequent-evidence log or other governed mechanism.

Avoid replacing the original handoff pack silently.

7.3 COM-015 Pack States Are Not Assessment States

COM-015 may use document-management states such as working collection, internal quality review, handoff version, addendum/change set and archived version.

Those labels manage the evidence package only. They do **not** mean:

- assessment-ready;
- evidence sufficient;
- control effective;
- conforming;
- audit passed;
- certification-ready; or
- ODA3 approved.

COM-016 preserves that boundary.

PART II — HOW EVIDENCE IS INTERROGATED

8. Evidence Admission and Traceability — Public View

Evidence can be useful only if its identity and context are sufficiently clear.

Practitioner questions include:

- Is the item uniquely identifiable?
- Is the source known?
- Is the owner or custodian known?

- Is the system/version clear?
- Is the relevant time period clear?
- Is the location retrievable?
- Is provenance assessable?
- Are material transformations known?
- Is the evidence lawfully available for review?
- Can the item be linked to a defined evidence claim?

The public-safe evidence-admission concepts used here are consistent with the controlled evidence baseline: an item should be sufficiently **identifiable, attributable, scope-linked, time-linked, integrity-assessable and lawfully available** before it is relied upon as assessment evidence.

Admission into an evidence record is not the same as sufficiency, reliability or a positive control conclusion.

An item may be important enough to preserve even when it is not yet strong enough to support a conclusion.

For example:

- an anonymous screenshot may identify a lead;
- an unverified vendor statement may identify where evidence should exist;
- a ticket may identify a possible incident;
- a dashboard may identify a metric worth tracing.

The assessment process may then seek the authoritative underlying record.

---

9. Evidence Claims

A common failure is to map an artifact directly to a control without stating what it is meant to establish.

Use an evidence claim.

Example:

Requirement: privileged access to an AI production system is controlled.

Potential evidence claims:

- privileged roles are defined;
- privileged users are approved;
- privileged access is periodically reviewed;
- revoked users lose access;
- emergency access is controlled;
- relevant actions are logged;
- exceptions are identified and managed.

The same policy may support only the first claim.

A role export may support the current configured population.

Review records may support historical review activity.

Tickets may support exception handling.

Logs may support evidence of actual use.

Assessment becomes more precise when artifacts are attached to bounded claims rather than broad control labels.

---

10. Design, Implementation and Operation

Three questions recur across assessment work.

10.1 Design

What is supposed to happen?

Typical evidence:

- policy;
- standard;
- architecture;
- procedure;
- control design;
- threat model;
- approval record;
- responsibility matrix.

10.2 Implementation

Has the designed mechanism actually been put into use in the assessed environment?

Typical evidence:

- configuration;
- deployed architecture;
- enabled monitoring;
- access assignment;
- integration record;
- installed safeguard;
- activated workflow.

10.3 Operation

Did the mechanism operate as intended over the relevant period or opportunities?

Typical evidence:

- period logs;
- completed reviews;
- incident records;
- monitoring outputs;
- exceptions;
- sampled cases;
- change records;
- alerts and responses.

A mature evidence conversation asks all three questions separately.

10.4 Why one successful walkthrough is limited

A walkthrough can be useful.

It can show:

- the control exists;
- people know how it works;
- a process can be demonstrated;
- a configuration is currently present.

It does not necessarily establish:

- period-wide consistency;
- every relevant population;
- absence of bypass;
- response to rare conditions;
- historical operation; or
- operation before a recent remediation.

---

11. Relevance, Independence, Provenance, Currency, Coverage and Integrity

Assessment evidence should be challenged through the five public-safe evidence-quality lenses used across the GAISSF methodology:

relevance, independence, currency, coverage and integrity.

These dimensions are not a public score and are not averaged in COM-016. **Provenance** supports the review by showing where evidence came from, how it was handled and what transformations occurred.

11.1 Relevance

Does the evidence address the correct requirement, object, unit, environment and period?

A highly authoritative document can still be irrelevant to a particular claim.

11.2 Independence

How does the source or evidence-generation process relate to the person or mechanism responsible for the asserted control?

Independence is claim-specific. A directly generated system configuration can be more probative than an external certificate with mismatched scope.

No universal source hierarchy overrides relevance.

11.3 Provenance

Where did the evidence come from?

Record:

- system/process of origin;
- owner/custodian;
- extraction or collection method;

- date/time;
- transformations;
- redactions; and
- chain of custody where material.

11.4 Currency

Does the evidence represent the assessed state?

Ask:

- Was the system materially changed after the record?
- Did a provider change the model or service?
- Was the policy superseded?
- Is a certificate/report outside its useful period?
- Did permissions change after the screenshot?
- Was the evidence collected before a major remediation?

11.5 Coverage

What does the evidence actually represent?

Coverage can refer to:

- time;
- transactions;
- users;
- systems;
- geographies;
- models;
- environments;
- cases;
- providers; or
- operating conditions.

Do not use the word **complete** when the represented population has not been defined.

11.6 Integrity

Can the reviewer assess whether the item is authentic, complete and unchanged enough for the intended use?

Possible integrity supports include:

- controlled repository;
- signed export;
- audit trail;
- system-of-record reference;
- version control;
- hash where appropriate;
- immutable log; and
- documented transformation.

Not every evidence item requires a hash. The method should fit the evidence and risk.

A cryptographic hash can help show that a file has not changed after hashing. It does **not** prove that the source was complete, truthful, correctly generated or representative of the assessed population.

---

12. Corroboration

One source can be enough for some claims.

Other claims require corroboration.

A practitioner should expect additional evidence where a material assertion depends heavily on:

- management judgement;
- one mutable source;
- one person;
- a provider outside direct access;
- an unusual exception;
- an automated summary;
- a high-impact control;
- a current screenshot for a historical claim; or
- evidence with integrity limitations.

Examples:

Claim	Useful corroboration
Policy was implemented	implementation records, communication, configuration
Control operated	system record plus cases or independent monitoring
No relevant incident found	search criteria plus logs/tickets and inquiry
Provider control can be relied on	provider assurance plus scope/period mapping and customer-interface evidence
Automated metric is accurate	data lineage plus calculation logic and exception tracing
Feature was disabled	inventory/configuration plus contrary-indicator search

Corroboration is not about maximizing file count. It is about reducing unsupported inference.

---

13. Contrary Evidence

Contrary evidence is information that conflicts with the proposed conclusion.

Examples:

- a control owner says all access reviews were completed, but one quarter has no review record;
- a dashboard reports full coverage, but one system is absent from the telemetry inventory;
- a provider states a feature is disabled, but logs show API calls to that feature;
- policy requires approval before model promotion, but one deployment record shows no approval;

- monitoring appears healthy, but an incident shows a high-impact alert failed.

13.1 Do not bury the contradiction

The correct response is to:

1. preserve it;
2. assess authenticity;
3. identify affected scope;
4. determine whether it is isolated or systemic;
5. seek additional evidence where appropriate;
6. test whether the proposed conclusion must be narrowed; and
7. preserve unresolved uncertainty.

Favorable evidence volume does not erase a material contradiction.

13.2 Do not manufacture symmetry

Not every contradictory artifact has equal weight.

A stale email and an authoritative system log are not automatically equivalent.

The assessment process evaluates the evidence in context.

---

14. Missing and Inaccessible Evidence

Missing evidence can mean different things.

It can indicate:

- the control did not operate;
- the record was not retained;
- the record exists but is inaccessible;
- a provider controls the evidence;
- the population was never defined;
- the organization looked in the wrong place;
- the event did not occur;
- the evidence was destroyed;
- the evidence is legally restricted; or
- the system cannot generate the required record.

These are not interchangeable.

14.1 What practitioners should record

For a material evidence gap, record:

- evidence objective;
- expected source;
- reason it is missing/inaccessible;
- who controls access;
- period affected;
- whether an alternative source exists;

- what limitation remains; and
- action owner.

14.2 “No evidence observed” is not “did not occur”

Where an absence matters, record:

- what records were searched;
- which period;
- what systems;
- what query or method;
- what coverage;
- what limitations remain.

---

15. Provider Evidence and Inheritance

Hosted AI introduces a common assurance problem: important controls may be operated partly or entirely by another organization.

A provider may supply:

- service description;
- contract;
- security documentation;
- independent assurance report;
- test report;
- attestation;
- architecture statement;
- incident notice;
- change notice;
- service telemetry; or
- customer-facing audit evidence.

The enterprise still needs to understand what that evidence covers.

15.1 Provider evidence review questions

Ask:

- What exact provider service/entity does the evidence cover?
- What region or deployment?
- What period?
- What version or service state?
- What criteria were used?
- Were material exceptions reported?
- What customer responsibilities are excluded?
- Are subcontractors/subservices relevant?
- What changed after the provider report period?
- What can the customer observe independently?
- What remains inaccessible?

- Is the relied-upon control actually at the provider/customer interface?

15.2 Four provider evidence states

State	Meaning
Customer-observed	Enterprise can inspect the evidence directly
Provider-provided	Provider supplies evidence or assurance
Shared/interface	Each side holds part of the evidence chain
Provider-only / inaccessible	Material evidence cannot be directly accessed

A fifth state—**unknown**—should be treated as a gap until responsibility is clarified.

For inherited or shared controls, the assessment needs a defensible basis for the operator, service, scope, responsibility split, interface, evidence access, exceptions and material change notification. Provider reputation, a certification badge or a contract clause alone does not establish effective inheritance.

15.3 Certificate does not transfer conformity

A provider certificate or assurance report has its own scope, criteria and period.

It is not an automatic conclusion about the customer's end-to-end AI system.

PART III — ASSESSMENT PROCEDURES

16. Assessment Planning

Assessment planning converts the declared scope into an authorized work programme.

The practitioner does not need to know the controlled procedure-selection algorithm, but should expect planning to consider:

- object and units;
- period;
- framework/control baseline;
- material dependencies;
- provider reliance;
- novelty;
- impact;
- architecture complexity;
- evidence availability;
- access;
- populations;
- prior incidents;
- prior findings;
- system change;
- physical-AI consequence;
- agentic authority;
- specialist requirements;
- confidentiality; and

- intended reliance.

16.1 What the organization should prepare before planning discussions

Prepare:

- architecture/current-state overview;
- list of system owners;
- provider list;
- evidence pack version;
- known material changes;
- prior major incidents;
- open evidence gaps;
- unavailable provider evidence;
- operating periods;
- major populations;
- data-access constraints;
- safety/physical interfaces;
- key business deadlines.

16.2 Duration is not universal

Assessment effort can vary materially with:

- scope;
- risk;
- novelty;
- evidence quality;
- number of units;
- provider opacity;
- physical access;
- specialist needs;
- change during fieldwork;
- contradictions; and
- remediation.

This guide does not define a universal duration, site-day minimum, interview set or fee.

16.3 Escalate Early When Any of These Appear

Tell the assessment coordinator promptly when:

- the declared system/version no longer matches production;
- a provider or tool changed materially;
- a high-impact incident or safety event occurs;
- a critical evidence source is inaccessible;
- legal/confidentiality constraints block requested access;
- contrary evidence changes the understanding of a control;
- a population cannot be reconciled;
- remediation materially changes the assessed state; or
- a requested procedure could create operational or safety risk.

Early escalation preserves scope and evidence integrity. It is not an adverse finding by itself.

---

17. Procedure Categories

A practitioner may encounter several categories of assessment activity.

17.1 Inspection

Review of documents, configurations, records, logs, tickets, reports or system settings.

17.2 Observation

Watching a control or process operate.

Examples:

- approval workflow;
- escalation process;
- human override;
- model deployment;
- incident response;
- emergency-stop procedure.

17.3 Inquiry

Structured questions to people with relevant responsibility or knowledge.

Inquiry can clarify facts, but unsupported statements do not automatically replace underlying evidence.

17.4 Reperformance

Repeating a calculation, query, review, control step or other procedure to assess whether the observed result can be reproduced.

17.5 Recalculation

Recomputing a metric or result from underlying data.

17.6 Configuration inspection

Checking actual settings, permissions, policy enforcement, routing, tool access, monitoring or deployment parameters.

17.7 Population reconciliation

Determining what the complete set of relevant items/events should be before selecting cases.

17.8 Selection / sampling

Examining a subset or targeted set of items under the controlled assessment method.

This guide does not define the selection formula or sample size.

17.9 Trace reconstruction

Reconstructing a sequence such as:

input → model/service → tool call → authorization → action → monitoring → response

17.10 Technical testing

Security, model, runtime, integration, agentic or other technical procedures appropriate to the control objective.

Detailed controlled scripts remain outside this guide.

Public Boundary on Procedure Selection

The categories above explain what a practitioner may encounter. They do not define which procedure, depth, selection method, sample size, test sequence or expansion rule an authorized assessor must use. Those decisions remain within ASM-004 and the applicable approved assessment plan.

---

18. Population and Selection — Public View

A population is the complete set of events, items, configurations, users, transactions, deployments or other units to which a conclusion is intended to apply.

Examples:

- all privileged-user assignments during the period;
- all production model deployments;
- all quarterly access reviews;
- all model-change approvals;
- all high-impact agent transactions;
- all relevant provider change notices;
- all physical-AI emergency-stop tests.

18.1 Why population definition comes first

If the organization cannot define the population, it becomes difficult to know whether selected evidence is representative.

A hand-picked set of successful cases can prove those cases occurred. It does not automatically establish the broader population.

Where selected items are intended to support a positive population-level conclusion, the assessed organization should expect the authorized assessment process to control the final selection rather than relying on a subject-curated favorable set.

18.2 What practitioners should expect

The assessment team may ask for:

- source of population;
- total count;
- date range;
- reconciliation;
- exclusion logic;
- missing items;
- duplicate handling;
- strata or categories;
- selected case IDs; and

- evidence that the organization did not remove adverse items.

18.3 No universal sample-size rule in this guide

COM-016 does not publish a universal sample-size table.

Selection depth depends on the controlled methodology, the intended inference and the actual assessment circumstances.

---

19. Walkthroughs and Reperformance

Walkthroughs are valuable because they make a process visible.

They are also easy to overinterpret.

A walkthrough can help answer:

- Who performs the control?
- Where is it performed?
- What systems are used?
- What approval is required?
- What evidence is generated?
- What happens on exception?
- What is the escalation route?

Reperformance can add stronger support by independently repeating the process or calculation.

19.1 Walkthrough trap

Scenario: a team demonstrates an access review perfectly during a meeting.

That proves the team can demonstrate the process now.

It does not automatically prove:

- every required review occurred;
- historical reviews were complete;
- exceptions were resolved;
- the same process operated before recent changes.

The assessment may therefore request period evidence.

---

20. Technical Testing

Technical testing should remain linked to an assessment question.

Do not begin with:

“Run every test we have.”

Begin with:

“What proposition about this control needs to be evaluated?”

Possible public-safe areas include:

- access/configuration testing;
- model/package identity verification;
- adversarial evaluation;
- runtime monitoring;
- logging;
- tool permission testing;
- retrieval/data-flow checks;
- update/change testing;
- rollback validation;
- boundary testing;
- human-intervention demonstration;
- physical-AI safe/degraded-state evidence.

The detailed test procedure remains controlled.

20.1 Test failure is evidence

A failed test is not something to hide because it “hurts the score.”

It may be the most important evidence in the assessment.

The organization should preserve:

- test identity;
- date;
- system/version;
- procedure context;
- expected behavior;
- observed behavior;
- relevant logs;
- affected scope;
- immediate action;
- retest evidence where applicable.

---

21. Interviews and Representations

People are essential evidence sources because many control facts are not obvious from system records alone.

Interviews can explain:

- ownership;
- intended process;
- architecture;
- exception handling;
- provider responsibilities;
- rationale;
- undocumented constraints;
- recent changes.

But an interview statement should not be treated as self-verifying merely because it came from a senior executive or system owner.

21.1 Useful interview discipline

Separate:

- confirmed fact;
- representation;
- interpretation;
- hypothesis;
- unknown;
- requested follow-up evidence.

21.2 Avoid “management said so”

For a material positive claim, ask what record, system state, test, approval or operational evidence can corroborate the representation.

21.3 Incident Evidence as an Assessment Input

COM-012 and COM-013 govern the practitioner classification/response boundary for AI incidents.

When incident-related material enters an assessment, COM-016 should preserve, where available:

- the existing incident or event identifier;
- confirmed facts;
- working hypotheses;
- provider representations;
- enterprise-observed evidence;
- unresolved unknowns;
- classification state and change history;
- response actions and their timing;
- affected system/version;
- relevant evidence preserved during response; and
- corrective-action / recovery evidence.

An anomalous output, alert, tool call, safety event or vulnerability signal is **not automatically a UAIF incident classification** merely because it appears in an assessment file.

COM-016 does not:

- invent an Observation record type;
- redefine UAIF fields or severity;
- determine legal incident status or regulatory reportability;
- replace AI-IRF™ response activities; or
- infer provider causality from provider involvement alone.

If classification remains uncertain, preserve that uncertainty. Assessment may examine what the event says about controls and evidence, but the incident-classification and response records remain governed by their applicable UAIF™ / AI-IRF™ sources.

22. Agentic-AI Assessment Patterns

Agentic systems require assessment attention beyond output quality.

A system may produce safe-looking text while holding significant authority.

Relevant dimensions include:

- agent identity;
- delegated authority;
- tools;
- credentials;
- allowed actions;
- transaction limits;
- approval requirements;
- execution boundaries;
- memory/state;
- monitoring;
- disablement;
- rollback;
- provider change;
- tool-definition change.

22.1 Separate three questions

Authority

What is the agent allowed to do?

Execution

What did the agent actually do?

Outcome

What effect followed?

These can require different evidence.

22.2 Tool metadata and permission changes

If a tool's metadata, schema, endpoint, permission or behavior can change after approval, the assessment may need to examine:

- change control;
- reauthorization;
- monitoring;
- versioning;
- detection of altered tool definitions;
- credential boundaries;
- logs of tool calls;
- evidence of human approval where required.

Do not assume that an approved integration remains equivalent after a material tool change.

23. Physical-AI Assessment Patterns

Physical AI adds a kinetic consequence layer.

For current market-facing practitioner use, COM-014 establishes the public framework boundary: GAISSF™ Domain 9 remains retained for citation continuity, while PAI-SF™ v1.0 is the dedicated framework that supersedes and expands the physical-AI security-assurance layer represented by D9. This does not create two independent public assurance obligations.

Some controlled assessment-design sources predate that finalized practitioner wording and describe PAI-SF as supplementary to D9. COM-016 does not use that older wording to override the current COM-014 public boundary; exact assessment-methodology reconciliation remains a controlled source-management matter before live use.

Where PAI-SF™ applies, assessment evidence may need to reconstruct:

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

Potential evidence includes:

- sensor/perception state;
- calibration/health;
- model/software/firmware version;
- autonomy mode;
- operating boundary;
- command path;
- authorization/interlocks;
- actuator feedback;
- safety-controller events;
- emergency stop;
- degraded mode;
- operator intervention;
- environmental conditions;
- maintenance;
- recovery;
- revalidation.

23.1 Security and safety authorities remain distinct

A security/assurance assessment can identify evidence, findings, remediation and verification.

It does not itself establish:

- functional-safety certification;
- roadworthiness;
- airworthiness;
- medical-device approval;
- machinery conformity;
- regulatory approval;
- safe return to service.

Where a physical system is involved, expect explicit coordination across:

engineering + safety + operations + security + IT/OT + governance

Two-Authority Return-to-Service Boundary

Keep two decision tracks separate:

1. **security / assurance evidence closure** — whether the relevant security or assurance issue has been addressed and evidenced; and
2. **safety / operational authorization** — whether the competent sector, engineering, safety or operating authority permits return to service.

Completion of Track 1 does not grant Track 2.

COM-016 also does not collapse PAI-SF's physical-AI practitioner evidence groupings into the GAISSF adoption evidence classes used by COM-015. Those are different public organizing lenses and should be mapped only where useful, not treated as identical taxonomies.

PART IV — FROM EXCEPTIONS TO FINDINGS

24. Test Exceptions

A test exception is an observed condition that differs from the expected result.

Examples:

- one selected access review is incomplete;
- one deployment lacks an approval record;
- one provider change notice arrived after production change;
- one agent action exceeds the expected permission boundary;
- one emergency-stop test fails;
- one model evaluation cannot be reproduced.

An exception is an input to assessment analysis.

It is not automatically a formal finding. The controlled finding methodology requires validation, requirement mapping, population/extent analysis and consideration of contrary evidence before an exception becomes a finding.

24.1 What should be preserved

Record:

- test/case identity;
- expected condition;
- observed condition;
- source evidence;
- date/period;
- affected unit;
- population relationship;
- immediate explanation;
- related exceptions;
- potential cause;

- follow-up action.

The controlled assessment process determines how the exception affects the assessment conclusion.

---

25. Finding Anatomy

A finding should be understandable without relying on vague language such as “control weak” or “needs improvement.”

A useful public anatomy includes:

- **Condition** — what was observed;
- **Criteria** — what should have existed or occurred;
- **Evidence** — what supports the statement;
- **Extent** — what systems/units/periods are affected;
- **Consequence** — why the condition matters;
- **Cause status** — known, provisional or not established;
- **Response** — what the organization says;
- **Action** — what correction/remediation is planned;
- **Verification** — what evidence later supports correction;
- **History** — opening, changes, closure/reopening where applicable.

25.1 Cause must not be invented

It is acceptable to state:

- cause not yet established;
- evidence suggests a possible cause;
- provider investigation pending.

Do not infer cause simply because it is plausible.

25.2 Improvement opportunity vs failure

Not every useful recommendation is a failure.

A good assessment preserves the difference between:

- established requirement failure;
- evidence limitation;
- improvement opportunity;
- emerging risk;
- urgent factual escalation.

The controlled methodology owns the formal classification.

---

26. Factual Accuracy

The assessed organization should have a fair opportunity to correct factual errors.

Examples of legitimate factual correction:

- wrong system version;
- incorrect date;
- evidence attributed to the wrong owner;
- provider scope misunderstood;
- an artifact omitted from the record;
- architecture component misidentified.

26.1 Factual correction is not outcome bargaining

The organization may:

- provide correcting evidence;
- point out misinterpretation;
- challenge scope;
- document disagreement;
- explain context.

It should not:

- pressure the assessor to remove supported adverse evidence;
- demand a different classification for commercial reasons;
- substitute a management promise for evidence;
- rewrite the assessed period using later remediation.

26.2 Maintain an Assessment Query / Response Tracker

Recommended fields:

- Query ID;
- date;
- subject;
- question;
- owner;
- response;
- evidence IDs;
- factual correction vs new evidence vs disagreement;
- status;
- follow-up.

---

27. Management Response

A management response is valuable when it is precise.

A useful response may state:

- agreement/disagreement;
- factual clarification;
- business context;
- immediate containment;
- planned correction;

- accountable owner;
- target date;
- residual limitation;
- dependency;
- requested retest.

Avoid statements such as:

- “risk accepted, therefore resolved”;
- “we plan to fix it, so close the finding”;
- “no customer complained, so impact is low”;
- “the provider says it is fine”;
- “this is an industry-wide issue.”

Those statements may be relevant context. They do not erase evidence.

Risk-acceptance boundary: management may accept business risk within its authority, but that acceptance does not convert an unmet GAISSF requirement into conformity and does not, by itself, create assessor closure.

---

28. Corrective Action

Corrective action should address the condition and, where appropriate, the underlying cause and affected population.

A useful corrective-action record includes:

- finding reference;
- containment;
- root-cause status;
- correction;
- systemic remediation;
- affected units/population;
- owner;
- completion evidence;
- retest/verification requirement;
- residual risk/limitation;
- date.

28.1 Plan is not implementation

A plan can establish intent.

It does not prove the correction has been implemented.

28.2 Implementation is not operating effectiveness

A newly deployed safeguard can prove implementation.

It may need time or operating opportunities before sustained operation can be evaluated.

---

29. Verification and Closure

Verification asks whether the claimed corrective action is supported by evidence.

COM-016 explains the practitioner's evidentiary role only. Formal closure type, closure authority, required verification depth and any resulting control/domain/overall reassessment remain governed by the controlled methodology.

Possible public-safe verification activity:

- inspect new configuration;
- review revised process;
- test affected population;
- reperformance;
- review post-change logs;
- repeat scenario;
- inspect provider evidence;
- confirm issue-to-action traceability.

Closure does not rewrite history.

If the original assessed-period condition existed, the historical record remains.

A later status can say the issue was corrected or verified after the period.

It should not imply the original condition never existed.

---

30. Assessment Cutoff and Subsequent Evidence

An assessment needs a point after which evidence handling becomes controlled.

Why?

Because otherwise:

- evidence can arrive indefinitely;
- the assessed period becomes unclear;
- later remediation can blur historical truth;
- report issuance never stabilizes.

30.1 Practical timing categories

Before cutoff

Evidence is generally handled within normal fieldwork.

After cutoff but before closing

New evidence may require a documented decision on whether and how it can be considered.

After closing but before issuance

The item may be a factual correction, new evidence, remediation evidence or a material event requiring reopened work.

After issuance

Use correction, follow-up, surveillance or reassessment mechanisms as applicable.

30.2 Subsequent Evidence

Use **Section 53 — Subsequent Evidence Log** for the organization-side chronology.

For each late item, preserve at minimum:

- when it was received;
- what period it represents;
- why it was not in the original handoff;
- whether it is a factual correction, new evidence, remediation evidence or a material event;
- related evidence / finding references; and
- any potential scope effect the organization needs to flag.

The organization records the chronology. The controlled methodology determines how, or whether, the item affects assessment work or reporting.

PART V — REVIEW, REPORTING AND CLAIMS

31. Technical Review

Technical review is an independent challenge of whether the assessment file supports the proposed report.

It is a quality gate.

It is **not a second assessment** and it is not certification. The reviewer evaluates whether the recorded assessment work supports issuance and may require rework, additional procedures or escalation; the reviewer does not manufacture support through report wording.

31.1 What technical review challenges

At a high level, expect review of:

- scope and object;
- evidence chain;
- testing;
- exceptions;
- findings;
- limitations;
- material judgement;
- report wording;
- claim boundary;
- version/release integrity.

31.2 Why technical review matters to the assessed organization

A finding or conclusion discussed at the closing meeting may still be subject to technical review.

That is not necessarily inconsistency.

It reflects the fact that the assessment file must withstand independent challenge before issuance.

31.3 No report-only repair

If the report says more than the evidence supports, the correct answer is not merely softer wording.

The underlying assessment record may need correction, clarification, additional evidence or additional procedure.

---

32. Closing Meeting

The closing meeting is not a negotiation over the desired result.

Its purpose is to ensure that the organization understands:

- scope;
- period;
- work performed at a suitable level;
- material findings;
- limitations;
- factual corrections;
- disagreements;
- outstanding items;
- next steps;
- technical review;
- reporting;
- corrective action; and
- certification separation where relevant.

32.1 “No surprise, no promise”

Material concerns should normally be raised once sufficiently validated so the organization has an opportunity to preserve evidence and respond.

But early communication is not a promise that:

- no further findings will arise;
- the report will be favorable;
- a preliminary conclusion cannot change after technical review.

Any conclusions discussed before technical review and authorization remain provisional for communication purposes.

32.2 Prepare for the Closing Discussion

Before the closing discussion, reconcile five things:

1. **scope** — exact object, period and material changes;
2. **evidence** — open requests, late evidence and unresolved access constraints;
3. **facts** — submitted corrections and remaining factual disagreements;
4. **remediation** — what existed during the period versus what changed later; and
5. **claims** — what can and cannot be said before technical review / issuance.

Use **Section 55 — Closing Meeting Preparation Checklist** as the detailed organization-side checklist.

33. Assessment Report

Where an assessment report is actually authorized and issued under the applicable ODA3 process, it should remain inseparable from the context that gives the conclusion meaning.

At minimum, a practitioner should expect clear identification of:

- report identity/version;
- assessed object;
- criteria/framework version;
- period;
- boundary;
- material exclusions;
- relevant work basis;
- findings;
- limitations;
- conclusion;
- subsequent events where material;
- intended users/distribution;
- claim restrictions.

33.1 Scope inseparability

Do not quote:

“The assessment found…”

without the scope.

A system-level result should not be presented as enterprise-wide unless the assessment actually covered that enterprise-wide scope.

33.2 Temporal accuracy

Later remediation should not be written as if it operated throughout the assessed period.

A report may distinguish:

- assessed-period condition;
- later corrective action;
- verified post-period status.

34. Limitations and Notably Absent

A strong report explains both what was established and what was not.

Potential limitations include:

- inaccessible provider evidence;
- short operating period;
- partial population;

- material system change;
- unresolved contradiction;
- unavailable logs;
- physical access limitation;
- untested operating mode;
- model/provider opacity;
- legal/confidentiality restriction.

34.1 Notably Absent discipline

A statement such as **“no relevant incident record or qualifying event was identified in the reviewed sources”** should be supported by:

- records searched;
- period;
- systems;
- search method;
- coverage;
- limitations.

It does not prove that no incident occurred.

It is an evidence-search statement, not a new UAIF classification, legal-reportability determination or proof of event absence.

34.2 Silence is not a favorable result

An untested condition should not be presented as implicitly satisfactory merely because the report does not mention a problem.

---

35. Assessment vs Certification

This is one of the most important boundaries in the guide.

Public model:

assessment work → independent technical review → authorized assessment report → separate certification decision interface where an issued scheme applies

35.1 A favorable assessment is not self-executing certification

A later certification decision, where applicable, may require:

- an eligible assessment package;
- independent decision authority;
- applicable scheme rules;
- unresolved-findings treatment;
- validity/surveillance rules;
- contractual/mark rules;
- other scheme conditions.

Those rules are not published in COM-016.

35.2 Current Certification Status

The current controlled source state is explicit:

- ASM-012 is a **Controlled Design Baseline**;
- final tier thresholds, mandatory-control designations, validity periods, surveillance cycles, marks and public-registry design are not approved;
- external certification issuance is not authorized under that baseline; and
- the methodology master register records live assessment and certification as not approved.

Therefore **GAISSF certification is not currently available under the controlled design baseline used for this guide**.

A later scheme launch would require separate approved scheme instruments and version-specific authority. COM-016 must not be used to imply that approval has already occurred.

35.3 Assessor and certification decision-maker are not the same role

The separation protects impartiality.

The person/team that performed or technically reviewed the assessment should not simply convert its own result into a certification grant.

---

36. Report Use and Claim Boundaries

A report recipient should use the complete, authorized report or approved statement.

Avoid:

- extracting one favorable sentence;
- omitting material limitations;
- removing the period;
- generalizing from one system to the enterprise;
- calling an assessment a certification;
- calling readiness work an audit pass;
- turning a provider report into customer-system conformity;
- using ODA3 names/marks beyond authorization.

36.1 Claim ladder

The strength of external wording should never exceed the strongest authorized evidence and report basis.

Examples of increasingly risky wording:

Lower-risk factual statement:
“System X participated in a GAISSF assessment for the stated scope and period.”

Requires an actually issued, authorized assessment report under an approved ODA3 assessment process:
“The authorized assessment report concluded [exact authorized conclusion] for System X, subject to the stated scope and limitations.”

Not supported by COM-016 alone:
“GAISSF compliant.”
“GAISSF certified.”

“ODA3 approved.”
“Regulator accepted.”
“Safe to deploy.”

---

37. Reassessment and Surveillance Concepts

An assessment conclusion can become stale.

Potential triggers for reassessment or follow-up include:

- material model change;
- provider change;
- agent authority expansion;
- new data source;
- architecture change;
- significant incident;
- major control redesign;
- critical provider issue;
- new geography;
- physical operating-boundary change;
- safety-interface change;
- repeated finding;
- evidence no longer representative;
- major regulatory/contractual change.

COM-016 does not define a universal surveillance interval. The current controlled design baseline does not approve a live surveillance cycle. Where a later issued scheme exists, its applicable scheme rules govern.

PART VI — WORKED EXAMPLES

Use each example with the same three questions:

1. **What can the current evidence establish?**
2. **What remains unresolved or contradicted?**
3. **What should the organization do next without overstating the outcome?**

The examples are hypothetical. They do not create sampling, finding, scoring or certification rules.

38. Worked Example — Complete Pack, Incomplete Operating Basis

Scenario: An organization submits an excellent COM-015 pack for privileged access.

The pack includes:

- approved policy;
- role design;
- current user export;

- current configuration;
- one completed access review;
- tickets for recent revocations.

Assessment question: Did the access-review control operate over the intended period?

Issue: The evidence period is twelve months, but only the most recent review is available.

Practitioner mistake: argue that the policy plus current configuration demonstrates the control.

Better response:

1. acknowledge what the evidence establishes;
2. identify missing period records;
3. search for review evidence in the authoritative system;
4. document retention/access limitations;
5. identify alternative corroboration;
6. avoid claiming period-wide operation if it cannot be supported.

Key lesson: a complete pack can still contain an incomplete basis for a specific conclusion.

---

39. Worked Example — Provider Assurance, Customer Integration Gap

Scenario: A hosted-model provider supplies a high-quality independent assurance report.

The enterprise uses the provider through:

- its own retrieval layer;
- custom tools;
- privileged API keys;
- an internal agent;
- customer-specific access controls.

Mistake: treat the provider report as evidence that the whole AI solution is assured.

Assessment response:

- identify provider-report scope;
- identify customer responsibilities;
- identify integration controls;
- examine customer telemetry;
- examine tool/credential boundaries;
- document provider-only evidence;
- preserve unknowns.

Key lesson: provider assurance is scoped evidence, not transferred enterprise conformity.

---

40. Worked Example — One Material Contradiction

Scenario: Monitoring dashboards show a year of favorable results.

One existing incident record shows that a high-impact alert failed during a material event.

Mistake: describe the incident as an anomaly and exclude it.

Better treatment:

- preserve the incident;
- validate the evidence;
- determine affected system and period;
- inspect monitoring lineage;
- determine whether similar cases exist;
- expand evidence/testing where appropriate;
- keep unresolved uncertainty visible.

Key lesson: favorable volume cannot erase material contrary evidence.

COM-016 does not need to reclassify the incident to use the preserved incident evidence in assessment. If classification is disputed or incomplete, preserve that state and use COM-012/COM-013 for the applicable classification/response discipline.

---

41. Worked Example — Material Change After Handoff

Scenario: The evidence pack is frozen on 1 August.

On 5 August, the provider silently changes the deployed model version.

Fieldwork is still underway.

Mistake: leave the change out because it happened after handoff.

Better response:

- record the change;
- identify effective date;
- identify affected controls/evidence;
- notify assessment coordinator;
- assess whether earlier evidence remains representative;
- capture new model/service identity;
- identify whether additional procedures are needed.

Key lesson: frozen evidence does not freeze the system.

---

42. Worked Example — Policy and Walkthrough vs Sustained Operation

Scenario: A team has a strong model-change policy and demonstrates the approval workflow flawlessly.

Historical deployment records show several releases but the evidence does not include approval records for two.

Mistake: rely on the walkthrough as proof that historical releases were approved.

Better response:

- preserve the walkthrough as design/implementation evidence;
- identify historical deployment population;
- reconcile approval records;
- investigate missing cases;
- preserve exceptions.

Key lesson: current demonstration does not automatically prove historical operating effectiveness.

---

43. Worked Example — Agent Authority Expansion

Scenario: An enterprise agent initially has read-only tools.

Later, one tool is updated to support write actions, but the enterprise approval record is not renewed.

The content filter continues to perform well.

Mistake: conclude the agent remains low-risk because output safety is strong.

Better assessment questions:

- When did authority change?
- Who approved the tool?
- Did tool metadata/schema change?
- Were permissions reauthorized?
- Are write actions logged?
- Are high-impact actions gated?
- Can credentials be accessed?
- What monitoring detects authority expansion?

Key lesson: output safety and action authority are different assurance questions.

---

44. Worked Example — Physical-AI Finding and Return-to-Service

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

Assessment evidence identifies:

- localization disagreement;
- autonomy-state change;
- command history;
- emergency-stop event;
- operator intervention;
- maintenance state.

A security remediation is implemented.

Mistake: state that the robot is safe to return to service because the security issue was corrected.

Better response:

- verify the security remediation as applicable;
- preserve the event and assessed-period finding;

- provide evidence to the appropriate safety/operations function;
- require applicable revalidation;
- keep return-to-service authority separate.

Key lesson: assurance closure does not substitute for safety authorization.

---

45. Worked Example — Post-Assessment Remediation

Scenario: A finding is established for the assessed period.

The organization corrects the issue two weeks later before the report is finalized.

Mistake: ask that the original finding be removed.

Better treatment:

- preserve the original assessed-period condition;
- record remediation date;
- verify the correction where applicable;
- state post-period status separately;
- preserve traceability from finding to closure.

Key lesson: remediation changes current status, not historical truth.

PART VII — PRACTITIONER TOOLS

46. Assessment Participation Checklist

Before kickoff

- ☐ Executive purpose is clear.
- ☐ Intended users/reliance are identified.
- ☐ Assessment object is defined.
- ☐ System/service/version is identified.
- ☐ Assessment period is understood.
- ☐ COM-015 evidence pack is frozen/versioned.
- ☐ Major evidence gaps are known.
- ☐ Provider dependencies are identified.
- ☐ Sensitive-evidence access route is defined.
- ☐ Key system/control owners are assigned.
- ☐ Material recent changes are disclosed.
- ☐ Material incident records and unresolved AI-event signals relevant to scope are disclosed.

During fieldwork

- ☐ Evidence requests have owners.
- ☐ New evidence is version-controlled.
- ☐ Material changes are reported.

- ☐ Provider statements are labeled as provider representations.
- ☐ Contrary evidence is preserved.
- ☐ Interviews are supported by records where appropriate.
- ☐ Factual corrections are distinguished from new evidence.
- ☐ Late evidence is logged.
- ☐ Findings are not negotiated through undocumented channels.
- ☐ Sensitive evidence remains minimized and access-controlled.

Before closing

- ☐ Open queries are reconciled.
- ☐ Factual corrections are submitted.
- ☐ Known disagreements are documented.
- ☐ Open evidence limitations are visible.
- ☐ Remediation completed after the period is clearly dated.
- ☐ Subsequent evidence is logged.
- ☐ Claim restrictions are understood.

Stop and Escalate Before Continuing If

- ☐ the assessment object is no longer the deployed system;
- ☐ a material provider/tool/model change has not been incorporated;
- ☐ a safety or operational event requires separate authority;
- ☐ a requested evidence source cannot be lawfully or safely accessed;
- ☐ a material contradiction is being excluded for convenience; or
- ☐ someone proposes changing historical evidence to reflect later remediation.

These are practitioner escalation triggers, not assessor dispositions.

---

47. Practitioner Evidence-Handoff Check

Question	Yes / No / Note
Exact pack ID/version recorded?	
System/service/version recorded?	
Assessment period recorded?	
Scope/applicability source referenced?	
Evidence index present?	
Requirement–evidence matrix present?	
Secure references resolve for authorized users?	
Provider/customer responsibility boundary visible?	
Contrary evidence included?	
Evidence gaps visible?	
Sensitive-evidence controls documented?	
Change log current?	

Question	Yes / No / Note
Handoff date/time recorded?	

Boundary: this is an organization-side preparation check. Completion is not assessment clearance, assessment readiness, evidence sufficiency, conformity, or any controlled assessor status.

---

48. Evidence Request Tracker

Field	Entry
Request ID	
Date	
Requirement / topic	
Evidence objective	
Requested artifact/source	
Owner	
Due date	
Evidence IDs supplied	
Secure reference	
Limitation	
Status	
Follow-up	

Recommended internal status terms may include:

Open / In Progress / Supplied / Partially Supplied / Inaccessible / Clarification Required / Closed

These are request-management terms only.

48.1 Evidence-Request Triage — Five Questions

Before responding to a material evidence request, ask:

1. **What exact claim or question is being tested?**
2. **Which system, unit and period does the request concern?**
3. **What is the strongest authoritative source we can provide?**
4. **Is there contrary, missing, provider-only or later evidence that must travel with it?**
5. **Does access require redaction, secure viewing or another controlled route?**

If the request is ambiguous, clarify it before sending a document dump.

---

49. Assessment Query / Response Tracker

Field	Entry
Query ID	
Date	
Related evidence/control/topic	
Question	
Response owner	

Field	Entry
Response	
Evidence IDs	
Factual correction / new evidence / explanation / disagreement	
Follow-up	
Status	

Keep management commentary separate from factual evidence.

This tracker is maintained by the assessed organization for request management. It is not an assessor query log, review-note register, workpaper or disposition record.

49.1 What a Good Response Looks Like

A useful response is:

- **specific** — answers the actual question;
- **bounded** — names system, scope and period;
- **traceable** — cites evidence IDs or secure references;
- **candid** — distinguishes fact, representation and unknown;
- **complete enough** — includes material contrary or limiting evidence;
- **controlled** — records late evidence and material changes through the proper route.

Avoid narrative advocacy where a short factual response plus evidence will do.

---

50. Scope Change Notification Template

Change ID:
Date identified:
System/service affected:
Previous state:
New state:
Effective date:
Reason for change:
Affected architecture/components:
Affected evidence:
Affected providers:
Potentially affected control/topic (organization view):
Potential period impact:
Immediate action:
Assessment coordinator notified:
Related evidence IDs:
Open questions:

The assessed organization records and communicates the change. Any assessment consequence is determined only through the applicable controlled methodology.

The organization should not use this template to assign an assessment status, control result, finding severity or required reassessment depth.

---

51. Factual Accuracy Response Template

Finding / query reference:
Statement being corrected:
Type: factual correction / contextual clarification / disagreement
Correct fact:
Evidence ID(s):
Why the original statement is inaccurate or incomplete:
Scope affected:
Period affected:
Requested correction:
Residual disagreement, if any:

Do not use this template to request a favorable outcome without evidence.

51.1 If You Disagree With a Finding or Factual Statement

Use the narrowest applicable response:

Situation	Practitioner response
Fact is wrong	identify the incorrect fact and supply the authoritative correcting evidence
Context is incomplete	add the missing context without denying the established evidence
New evidence exists	submit it through the subsequent-evidence route and identify the represented period
Interpretation is disputed	state the alternative interpretation and the evidence supporting it
Requirement or scope application is disputed	identify the exact boundary issue and request governed reconsideration
No supporting evidence exists for your position	record the disagreement honestly rather than manufacturing support

A disagreement is not resolved merely because management is senior, risk is accepted, or remediation is planned.

---

52. Organization Remediation Tracker

Field	Entry
Finding / issue reference	
Condition summary	
Affected system/scope	
Assessed-period date/period	
Management response	
Containment	
Cause status	
Corrective action	
Owner	
Target completion	
Completion evidence IDs	

Field	Entry
Evidence supporting claimed completion	
Residual limitation	
Current status	

Internal remediation status is not the same as assessor closure.

This tracker must not reproduce controlled finding-severity, closure-type, verification-status or assessor-disposition fields.

---

53. Subsequent Evidence Log

This is an organization-side chronology. It records what arrived after the original handoff and why. It does not assign the controlled treatment of later evidence.

Field	Entry
Subsequent Evidence ID	
Date received	
Source	
Related requirement/finding	
Period represented	
Why not in original handoff	
Factual correction / new evidence / remediation / material event	
Scope impact	
Related Evidence IDs	
Assessment action	
Status	

54. Assessment Limitation Register

Field	Entry
Limitation ID	
Affected system/control/topic	
Description	
Cause	
Period	
Provider/customer boundary	
Evidence impact	
Alternative evidence available?	
Action owner	
Current status	
Potential disclosure concern to flag?	

The organization may flag a limitation for attention, but it does not decide assessment or report treatment. The applicable controlled methodology governs that decision.

55. Closing Meeting Preparation Checklist

- ☐ Exact scope and period reconciled.
- ☐ Material changes disclosed.
- ☐ Evidence cutoff understood.
- ☐ Open evidence requests listed.
- ☐ Factual corrections prepared.
- ☐ Management responses prepared.
- ☐ Disagreements documented.
- ☐ Open limitations summarized.
- ☐ Provider unknowns summarized.
- ☐ Remediation status dated accurately.
- ☐ Subsequent evidence logged.
- ☐ Claim restrictions understood.
- ☐ Post-fieldwork review and reporting boundaries understood.
- ☐ Certification separation understood.

56. Report / Claim-Boundary Checklist

Before reusing an assessment conclusion in another document, website, proposal, board paper or customer communication, ask:

- ☐ Is the exact report ID/version cited?
- ☐ Is the system/object identified?
- ☐ Is the assessed period stated where material?
- ☐ Is the scope preserved?
- ☐ Are material limitations preserved?
- ☐ Is the wording no stronger than the authorized report wording?
- ☐ Is assessment clearly distinguished from certification?
- ☐ Is the statement free from regulator-approval implications?
- ☐ Is no enterprise-wide claim inferred from a system-level assessment?
- ☐ Is no safety/return-to-service claim inferred from AI assurance?
- ☐ Are ODA3 names/marks used only as authorized?

57. Reassessment Trigger Register

Trigger	Date	Affected system	Potential assessment impact	Owner	Action
Model/ provider change					

Trigger	Date	Affected system	Potential assessment impact	Owner	Action
New tool/agent authority					
Major architecture change					
Significant incident					
Major remediation					
Provider assurance expiry/change					
Physical operating-boundary change					
Safety-interface change					
New geography/use case					
Material legal/regulatory change					

This register supports governance. It does not prescribe a universal reassessment interval.

57.1 Reassessment Triage Rule

A trigger should prompt one practical question:

Does the existing scope, evidence basis or issued conclusion still represent the current system?

If the answer is unclear, route the change for governed review rather than assuming the prior conclusion still applies.

PART VIII — BOUNDARIES AND SOURCES

58. What an Assessment Does Not Prove

A GAISSF assessment does not, by itself, prove:

- absence of vulnerabilities;
- absence of incidents;

- perfect model behavior;
- absence of hidden operation;
- legal compliance;
- regulatory approval;
- fairness;
- accuracy in every context;
- product safety;
- functional-safety certification;
- airworthiness;
- roadworthiness;
- medical-device approval;
- safe return to service;
- enterprise-wide conformity outside the assessed scope;
- future security after material change;
- provider controls outside the assessed reliance boundary;
- certification;
- accreditation;
- regulator recognition.

An assessment conclusion is evidence-bounded, scope-bounded and time-bounded.

The assessment may record jurisdictional or legal constraints relevant to scope, but COM-016 does not determine statutory applicability or provide legal advice.

---

59. Controlled-Methodology Firewall

COM-016 **may explain publicly**:

- practitioner-facing assessment journey concepts;
- practitioner roles;
- object/unit/period concepts;
- evidence handoff;
- why evidence is challenged;
- high-level categories of procedure;
- design vs implementation vs operation;
- population concept;
- corroboration;
- contrary evidence;
- limitations and uncertainty;
- finding anatomy at a non-scoring level;
- factual accuracy;
- remediation traceability;
- technical-review purpose;
- reporting scope/period/limitation discipline;
- assessment/certification separation;

- material-change/reassessment triggers.

COM-016 **must not publish or reconstruct**:

- ordered assessor workflow, phase-gate sequence or decision path;
- controlled engagement-state values or rules;
- evidence-confidence categories/caps;
- control-determination values or decision rules;
- capability/maturity scales;
- sample-size rules or route floors;
- proprietary test protocols or scenario libraries;
- aggregation precedence or gate logic;
- mandatory/critical-control registers;
- finding-severity matrix/thresholds;
- controlled remediation deadlines;
- assessor judgement/override levels;
- technical-review sampling or clearance taxonomy;
- assessor competency minima/authorization rules;
- certification eligibility thresholds;
- certification decision rules;
- controlled workbook values/validation lists;
- benchmarking formulas;
- public scoring percentages.

If a practitioner needs one of those elements, it must come from the applicable controlled ODA3 methodology and authorization—not from COM-016.

---

60. GEL / Commercial-Delivery Boundary

GAISSF Ecosystem public materials may be referenced and used subject to GEL v1.0.

Public access to COM-016 does not authorize:

- paid GAISSF assessment services;
- GAISSF-branded audit;
- ODA3 certification;
- accreditation activity;
- commercial client delivery using the guide as a methodology;
- software embedding of protected control logic beyond permitted terms;
- endorsement claims;
- ODA3 assessor status.

Applicable GEL terms, partner requirements and written authorizations govern commercial delivery.

---

61. Source and Traceability Boundary

COM-016 retains source traceability without reproducing the internal methodology map.

Where COM-016 conflicts with a controlling framework or controlled methodology source, the controlling source governs.

61.1 Controlling Framework Sources

- **GAISSF™ v1.0 Framework Standard**
- **GAISSF™ v1.0 Control Catalogue**
- current GAISSF release/change-control artifacts

The current controlled assessment methodology uses the corrected **59-control / nine-domain** GAISSF baseline. COM-016 does not amend the framework or control catalogue.

61.2 Controlled Methodology Source Family

The current internal assessment-methodology family is maintained under the **ASM** series and its controlling methodology register.

For public COM-016 purposes, only the following source-state facts are necessary:

- the methodology family is currently at **Controlled Design Baseline** state;
- that state does not authorize live assessment or certification;
- approval is exact-version and scope specific;
- pilot preparation or validation activity does not create live-use authority by implication; and
- certification requires separately approved scheme instruments.

COM-016 intentionally does **not** publish the complete internal component map, dependency graph, status taxonomy, validation-state model or operational-approval sequence.

61.3 Verified Public-Safe Source Propositions

COM-016 retains only the source propositions needed for practitioner use:

1. **Scope before conclusion.** Assessment object, unit, boundary, period, dependencies and exclusions must remain explicit.
2. **Evidence before assertion.** Management explanation can direct inquiry but does not substitute for the required evidence basis.
3. **Evidence quality is multidimensional.** Relevance, independence, currency, coverage and integrity remain distinct.
4. **No document-count logic.** Evidence sufficiency is a reasoned determination, not a volume threshold.
5. **Population before inference.** Broader claims require a defined population and defensible selection/coverage.
6. **Design, implementation and operation are different assurance questions.**
7. **Contrary evidence remains visible.** Favorable evidence volume cannot cancel a material contradiction.
8. **Exception is not automatically a finding.**
9. **Remediation does not rewrite assessed-period history.**
10. **Professional judgement cannot waive criteria, manufacture evidence, suppress adverse facts or negotiate an outcome.**
11. **Independent technical review is a quality challenge, not a second assessment and not certification.**
12. **Assessment conclusion and certification decision are separate.**
13. **Methodology validation does not create operational or certification authority by implication.**

14. **Data availability does not create permission for benchmarking or public comparison.**
15. **Assessor authorization is bounded; training attendance alone does not create authority.**

61.4 Practitioner Integration Sources

- **COM-009 — The First 30 Days with GAISSF™**
- **COM-010 — Applying GAISSF™: Worked Scenarios**
- **COM-012 — Classifying AI Incidents: A Practical Guide to UAIF™**
- **COM-013 — AI Incident Response Field Guide: Using UAIF™ and AI-IRF™ for Real-World AI Incidents**
- **COM-014 — Securing Physical AI: A Practitioner's Guide to PAI-SF™**
- **COM-015 — Building an AI Assurance Evidence Pack**

COM-015 is the primary evidence-preparation handoff into COM-016. COM-014 governs the current public physical-AI practitioner boundary used by this guide.

COM-017 — Implementing GAISSF™: A Practitioner's Guide is forthcoming. COM-016 does not pre-empt its implementation role. Implementation artifacts may later become assessment evidence, but artifact creation for assessment convenience is not a substitute for sound implementation.

61.5 Cross-Document Consistency Register

Document	Role carried into COM-016	COM-016 treatment
COM-009	implementation planning / management checkpoint	elapsed time or Day-30 completion is not treated as assurance, readiness or conformity
COM-010	hypothetical scope/application reasoning	scenario similarity is not treated as applicability or evidence sufficiency
COM-012	UAIF™ classification discipline	assessment preserves classification state/uncertainty; no invented incident taxonomy
COM-013	UAIF™ / AI-IRF™ response loop	assessment consumes incident evidence without replacing classification/response
COM-014	PAI-SF™ physical-AI assurance	D9 continuity, PAI-SF expansion and two-authority return-to-service boundary preserved
COM-015	evidence preparation	evidence-pack completeness is not sufficiency, conformity or assessment readiness
COM-017	implementation	implementation design remains outside COM-016; outputs are candidate evidence, not favorable-result guarantees

61.6 Naming and Source-State Caution

Preserved GAISSF source files may contain historical naming/mark conventions. COM-016 uses the current market-facing **GAISSF™ — Global AI Safety and Security Framework** convention without changing the technical authority of the underlying source.

The status of a practitioner publication does not activate controlled methodology, live assessment authority or certification.

62. Notably Absent

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

- a rule that Day-30 completion or scenario similarity creates assessment readiness;
- a rule that COM-015 pack completeness creates evidence sufficiency or conformity;
- a new UAIF incident classification, Observation record type or legal-reportability determination;
- a rule that assurance remediation authorizes physical return to service;
- implementation design that belongs in COM-017;
- any practitioner checklist, triage card, participation packet, escalation trigger or tracker becoming an assessor disposition, readiness state, SLA or certification gate;
- an ordered public assessor workflow or phase-gate map;
- an internal ASM component/dependency map;
- an operational-approval or validation-state taxonomy;
- a public GAISSF scoring model;
- a public confidence scale;
- a public control-rating scale;
- a public capability/maturity scale;
- universal sample-size rules;
- universal operating-period requirements;
- a mandatory-control register;
- a public aggregation formula;
- finding-severity thresholds;
- remediation deadlines;
- assessor judgement/override levels;
- technical-review clearance categories;
- assessor qualification minima;
- certification eligibility thresholds;
- certificate validity periods;
- surveillance cycles;
- public registry rules;
- an issued certification scheme;
- regulator recognition;
- legal-compliance conclusions;
- safety approval;
- assessment outcome guarantees;
- live GAISSF conformity-assessment authority created by this guide;
- an operational approval for the ASM series as a whole;
- current GAISSF certificate issuance;
- public benchmark or comparative scheme-data release authority.

These are boundaries, not drafting omissions.

63. Next Reading

1. **Building an AI Assurance Evidence Pack** — ODA3-2026-08-WHP-COM-015
2. **The First 30 Days with GAISSF™** — ODA3-2026-08-WHP-COM-009
3. **Applying GAISSF™: Worked Scenarios** — ODA3-2026-08-WHP-COM-010
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. **Implementing GAISSF™: A Practitioner's Guide** — ODA3-2026-08-WHP-COM-017 — forthcoming

Publication Metadata and Suggested Citation

Document ID: ODA3-2026-08-WHP-COM-016 **Title:** *The GAISSF™ Assessment Field Guide* **Subtitle:** *A Practitioner's Guide to the Assessment Lifecycle, Evidence Review and Decision Boundaries* **Document Type:** WHP — Adoption / Practitioner Field Guide **Framework Context:** GAISSF Ecosystem **Classification:** Public — GEL v1.0 **Publisher:** ODA3 Institute **Copyright:** © 2026 ODA3 Pvt Ltd

Suggested citation: ODA3 Institute (2026), *The GAISSF™ Assessment Field Guide: A Practitioner's Guide to the Assessment Lifecycle, Evidence Review and Decision Boundaries*, ODA3-2026-08-WHP-COM-016.

No DOI, ISBN, accreditation number, certification number or regulator-recognition identifier is assigned by this guide.

GEL Attribution

Final Publication Statement

This document is the released **FINAL** edition of **ODA3-2026-08-WHP-COM-016 — The GAISSF™ Assessment Field Guide**.

FINAL status applies to this practitioner guide only.

It does not activate controlled methodology, live assessment authority, certification availability, regulator recognition, commercial delivery rights, product approval, sector authorization, or any controlled assessment determination.