ODA3 INSTITUTE

ADOPTION GUIDE

Applying GAISSF™

Worked Scenarios

Doc ID: ODA3-2026-08-WHP-COM-010

Classification: Public / Freely Distributable

Executive Summary

This guide illustrates how GAISSF™ Ecosystem scoping and evidence discipline can be applied across diverse, regulated operating environments. The scenarios are hypothetical. They are designed to teach a reasoning process, not to determine framework applicability, regulatory obligations, control sufficiency, conformance, certification, or system safety for another system.

Framework applicability must be determined from the specific system and context. It must not be inferred from a sector label, a similar-sounding use case, organizational size, or convenience.

METHODOLOGY NOTE

Each scenario follows one canonical reasoning pattern:
SCENARIO REASONING PATTERN — System characteristics → scope/applicability → relevant requirements → implementation claims → evidence → unresolved questions → reassessment.
Sector regulation may materially affect scope, evidence, and assurance expectations. Applicable obligations must be established separately for the jurisdiction, entity, regulated activity, and use case.

Scenario Selection Guide

Use this table only to find a useful starting scenario. Sector selection is a reading aid, not an applicability determination. If a system spans multiple contexts, review every materially relevant scenario.

Context	Start with	Context	Start with
General enterprise / internal AI	1	Banking / financial services	2
Capital markets	3	Insurance	4
Healthcare / clinical AI	5	Medical-device AI	6
Automotive	7	Drones / aviation	8
Manufacturing / industrial robotics	9	Energy / critical infrastructure	10
Telecommunications	11	Defence	12
Government / public sector	13	Employment / HR	14
Third-party / foundation-model integration	15	Potential AI event / incident	16

EVIDENCE QUALITY

Illustrative artifacts are not automatically sufficient evidence. Evidence should be evaluated for relevance, authenticity, traceability, completeness, currency, and linkage to the specific system and determined scope. Artifact existence alone does not establish evidence sufficiency, control effectiveness, system safety, conformance, certification, or regulatory compliance.

1. Enterprise: Internal RAG Assistant

A company deploys an internal AI assistant with read access to a bounded document repository. Outputs are recommendations reviewed by staff; the assistant has no autonomous external action.

Scoping factors to consider

- Capability/exposure: internal-facing, limited autonomy, indirect business consequences.
- Architecture/data: enterprise documents, retrieval layer, identity/access controls, and a third-party model provider.
- Assurance focus: retrieval and data boundaries, output handling, human review, model-provider dependency and change.

Frameworks that may become relevant

- GAISSF™: applicable scope/profile must be determined for the specific system. Appendix P.1 may be considered only where its authoritative applicability criteria are satisfied.
- UAIF™/AI-IRF™: may become relevant if anomalous or potentially harmful behaviour requires event evaluation or response.

Illustrative evidence

- documented retrieval, identity and data-access boundaries;
- records of human review and escalation;
- vendor/model-change records and evidence distinguishing provider representations from independently supportable evidence.

> **NOTABLY ABSENT**
> Does not establish P.1 eligibility for internal assistants, evidence sufficiency, or the adequacy of any particular access-control implementation.

2. Banking and Financial Services: AI-Assisted Credit/Fraud Decision Support

A regulated financial institution uses AI to support fraud alerts and credit-risk review. Human decision-makers retain final authority, but AI outputs can materially influence customer treatment.

Scoping factors to consider

- Consequential decisions affecting individuals and financial outcomes.
- Sensitive financial and identity data; model/provider dependencies.
- Need to establish whether human review is substantive, informed, timely and capable of changing the outcome.

Frameworks that may become relevant

- GAISSF™: scope should reflect consequential decision support, data exposure, system architecture and third-party dependencies.
- UAIF™/AI-IRF™: may become relevant where anomalous outputs or control failures require incident determination and proportionate response.

Illustrative evidence

- decision, review and override records;
- model/version and input-data lineage;
- provider-change records and evidence supporting monitoring, escalation and correction of erroneous recommendations.

> **NOTABLY ABSENT**
> Does not determine banking, model-risk, consumer-protection, data-protection or AI-law obligations in any jurisdiction.

3. Capital Markets: AI-Assisted Trading Surveillance

A securities firm uses AI to prioritise potentially suspicious trading patterns for analyst review. The system does not itself block trades or make enforcement decisions.

Scoping factors to consider

- High-volume monitoring with false-positive and false-negative consequences.
- Market/client data, third-party dependencies and frequent model/rule changes.

- Human escalation and a defensible audit trail are material considerations.

Frameworks that may become relevant

- GAISSF™: scope/profile should be determined from system capability, data, exposure, architecture and assurance objectives.
- UAIF™: may help structure evaluation of material anomalous system behaviour; AI-IRF™ may support response where warranted.

Illustrative evidence

- change records for models/rules;
- analyst review and escalation records;
- evidence distinguishing model alerts, provider dependencies and final human conclusions.

> **NOTABLY ABSENT**
> Does not establish market-abuse surveillance adequacy or securities-regulatory compliance.

4. Insurance: AI-Assisted Underwriting and Claims Triage

An insurer uses AI to rank underwriting referrals and prioritise claims for human review. The system influences workflow but does not autonomously issue or deny policies or claims.

Scoping factors to consider

- Consequential decisions affecting individuals and financial outcomes.
- Personal, financial and possibly sensitive data; potential for differential outcomes.
- Human authority, meaningful review, override and rationale preservation.

Frameworks that may become relevant

- GAISSF™: scope should reflect consequential influence, data handling, architecture and third-party dependencies.
- UAIF™/AI-IRF™: may become relevant for anomalous scoring patterns or harmful behaviour requiring evaluation and response.

Illustrative evidence

- records of human review, overrides and rationale;
- input/data lineage and model/version records;
- provider-change records and outcome-review evidence where performed.

> **NOTABLY ABSENT**
> Does not establish fairness, underwriting, claims, insurance-law or regulatory requirements.

5. Healthcare: Clinical Decision-Support AI

A hospital uses AI to prioritise radiology cases and surface possible findings to clinicians. Licensed clinicians remain responsible for diagnosis and treatment decisions.

Scoping factors to consider

- Potential patient consequence despite human decision authority.
- Clinical data, model/provider updates and workflow integration.
- Risk of automation bias, over-reliance and ineffective human review.

Frameworks that may become relevant

- GAISSF™: applicable scope should reflect patient consequence, data, architecture, deployment context and third-party dependencies.
- UAIF™/AI-IRF™: may become relevant where anomalous outputs or failures require structured event evaluation and response.

Illustrative evidence

- records of clinician review and overrides;
- model/version, provider and deployment-change records;
- monitoring, escalation and correction records.

> **NOTABLY ABSENT**
> Does not establish clinical effectiveness, medical standard of care, healthcare-law compliance or regulatory authorisation.

6. Medical Devices / Physical AI: AI-Enabled Therapeutic Device

An AI-enabled medical device uses sensor inputs to adjust a bounded therapy parameter while enforcing independent device safety limits and clinician-configured constraints.

Scoping factors to consider

- Direct physical consequence and patient safety exposure.
- Sensor → model → decision → actuator chain.
- Software, hardware, communications, update and human-override dependencies.

Frameworks that may become relevant

- GAISSF™: remains applicable according to determined scope for the system's non-physical-AI requirements.
- PAI-SF™: should be considered where the system's physical or cyber-physical capabilities, autonomy, architecture and potential physical consequences place physical-AI security requirements within scope. Where PAI-SF™ applies, it supersedes and expands upon GAISSF™ Domain D9 (Physical AI Safety), which is retained in the current GAISSF™ release for continuity but is not the authoritative source for detailed physical-AI security requirements.

Illustrative evidence

- actuator/therapy boundary configuration records;
- test and evaluation records for safety-related logic and fail-safe behaviour;
- override, fault, near-miss and relevant change records.

> **NOTABLY ABSENT**
> Does not replace medical-device safety engineering, clinical validation, sector certification or regulatory authorisation.

7. Automotive: AI-Enabled ADAS Function

A vehicle uses AI perception to support lane positioning and hazard detection. The system can influence steering within its defined operational design domain while a human driver remains responsible for supervision.

Scoping factors to consider

- Perception → decision → action → physical consequence.
- Sensor integrity, environmental limits, software updates and fallback behaviour.

- Human-machine interaction and takeover assumptions.

Frameworks that may become relevant

- GAISSF™: remains applicable according to determined scope for the system's non-physical-AI requirements.
- PAI-SF™: should be considered where the system's physical or cyber-physical capabilities, autonomy, architecture and potential physical consequences place physical-AI security requirements within scope. Where PAI-SF™ applies, it supersedes and expands upon GAISSF™ Domain D9 (Physical AI Safety), which is retained in the current GAISSF™ release for continuity but is not the authoritative source for detailed physical-AI security requirements.

Illustrative evidence

- operational-design-domain documentation;
- sensor/model change and test/evaluation records;
- fallback, takeover and fault-event records.

> **NOTABLY ABSENT**
> Does not establish roadworthiness, functional safety, vehicle type approval or autonomous-driving authorisation.

8. Aviation / Drones: Autonomous Industrial Inspection Drone

A manufacturing operator uses a drone to inspect equipment and, within tightly bounded conditions, trigger a pre-authorised facility safety interlock when a hazard is detected.

Scoping factors to consider

- Autonomous physical action and cyber-physical consequence.
- Onboard sensors/models, navigation, communications and facility-control integration.
- Confidence thresholds, command boundaries and human override.

Frameworks that may become relevant

- GAISSF™: remains applicable according to determined scope for the system's non-physical-AI requirements.
- PAI-SF™: should be considered where the system's physical or cyber-physical capabilities, autonomy, architecture and potential physical consequences place physical-AI security requirements within scope. Where PAI-SF™ applies, it supersedes and expands upon GAISSF™ Domain D9 (Physical AI Safety), which is retained in the current GAISSF™ release for continuity but is not the authoritative source for detailed physical-AI security requirements.

Illustrative evidence

- trigger-logic and test/evaluation records;
- navigation/communications failure-handling records;
- override, incident and near-miss records.

> **NOTABLY ABSENT**
> Does not establish aviation approval, airworthiness, operational authorisation or facility-safety compliance.

9. Manufacturing: AI-Controlled Robotic Production Cell

A factory uses AI vision and planning to coordinate a robotic production cell. Robots operate inside safeguarded zones but can alter movement paths within configured limits.

Scoping factors to consider

- Physical motion, human proximity and OT integration.
- Vision/model integrity, actuator control, safety interlocks and update mechanisms.
- Production availability and physical consequence.

Frameworks that may become relevant

- GAISSF™: remains applicable according to determined scope for the system's non-physical-AI requirements.
- PAI-SF™: should be considered where the system's physical or cyber-physical capabilities, autonomy, architecture and potential physical consequences place physical-AI security requirements within scope. Where PAI-SF™ applies, it supersedes and expands upon GAISSF™ Domain D9 (Physical AI Safety), which is retained in the current GAISSF™ release for continuity but is not the authoritative source for detailed physical-AI security requirements.

Illustrative evidence

- robot/AI configuration and change records;
- interlock, stop and recovery test/evaluation records;
- incident, near-miss and maintenance records relevant to AI-driven behaviour.

> **NOTABLY ABSENT**
> Does not replace machinery safety, occupational safety, functional safety or industrial certification.

10. Energy / Critical Infrastructure: AI-Assisted Plant Operations

An energy operator uses AI to recommend equipment loading and maintenance actions. Operators approve actions, but recommendations can influence critical-service availability.

Scoping factors to consider

- Critical-service consequence and potential cascading impact.
- OT/IT data flows, sensor integrity and provider dependencies.
- Need for bounded authority, rollback and resilience.

Frameworks that may become relevant

- GAISSF™: scope should reflect critical-service consequence, architecture, data, human authority and third-party dependencies.
- PAI-SF™: should be considered where the system's physical or cyber-physical capabilities, autonomy, architecture and potential physical consequences place physical-AI security requirements within scope.
- UAIF™/AI-IRF™: may support event determination and response for AI-specific failures.

Illustrative evidence

- records separating recommendation from operator action;
- configuration/change, provider and rollback records;
- incident and resilience-test evidence relevant to AI dependencies.

> **NOTABLY ABSENT**
> Does not establish energy-sector reliability, safety, cybersecurity or regulatory compliance.

11. Telecommunications: AI-Driven Network Optimisation

A telecom operator uses AI to recommend and, for changes within pre-authorised bounds, automatically adjust network parameters to manage congestion and service quality.

Scoping factors to consider

- Large-scale automated change and availability consequence.
- Network telemetry, tool access, provider dependency, change boundaries and rollback.
- Potential customer/service impact across many users.

Frameworks that may become relevant

- GAISSF™: scope should reflect autonomous change capability, scale, architecture and connected systems.
- UAIF™/AI-IRF™: may become relevant when abnormal changes or service impacts require evaluation and response.

Illustrative evidence

- change-authorisation and autonomous-action boundaries;
- telemetry linking AI recommendations/actions to network outcomes;
- provider-change, rollback, incident and post-change review records.

> **NOTABLY ABSENT**
> Does not establish telecom-service, network-security or sector-regulatory compliance.

12. Defence: AI-Assisted Intelligence and Logistics Decision Support

A defence organisation uses AI to prioritise imagery for analyst review and optimise non-weapon logistics planning. Human authorities retain decision responsibility.

Scoping factors to consider

- High-consequence environment, uncertainty and sensitive data.
- Model provenance, access control, human authorisation and adversarial manipulation risk.
- Need to preserve uncertainty and provenance in analyst workflows.

Frameworks that may become relevant

- GAISSF™: applicable scope should reflect high-consequence use, architecture, data and adversarial exposure.
- UAIF™/AI-IRF™: may support structured evaluation and response to anomalous AI behaviour.

Illustrative evidence

- provenance and model/version records;
- records distinguishing AI suggestions, analyst judgements and unresolved uncertainty;
- access, change, anomaly and response records.

> **NOTABLY ABSENT**
> Does not establish suitability for weapons, targeting, autonomous lethal action, classified-system authorisation, export controls or defence accreditation.

13. Government / Public Sector: AI-Assisted Eligibility Prioritisation

A public body uses AI to prioritise applications for human caseworker review. The system does not make the final eligibility decision.

Scoping factors to consider

- Consequential decisions affecting individuals.
- Personal data, explainability, contestability and record preservation.
- Need to distinguish prioritisation from final administrative decision.

Frameworks that may become relevant

- GAISSF™: scope should reflect individual consequence, data, architecture and human authority.
- UAIF™/AI-IRF™: may become relevant for harmful anomalous behaviour requiring structured evaluation and response.

Illustrative evidence

- caseworker review/override records;
- data/model/version lineage;
- records supporting challenge, correction and investigation where applicable.

> **NOTABLY ABSENT**
> Does not establish administrative-law, public-sector, data-protection or rights compliance.

14. Employment: AI-Assisted Hiring Screening

An employer uses AI to rank applications for recruiter review. Recruiters make final decisions and can override rankings.

Scoping factors to consider

- Consequential decisions affecting individuals.
- Applicant personal data and vendor scoring model.
- Human review, override, automation-bias risk and outcome monitoring.

Frameworks that may become relevant

- GAISSF™: applicable scope/profile should be determined from capability, exposure, architecture and consequence severity.
- UAIF™/AI-IRF™: may become relevant for anomalous or potentially harmful scoring behaviour.

Illustrative evidence

- processed-data and exclusion records;
- recruiter review/override records;
- vendor/model-change records and outcome-review evidence where performed.

> **NOTABLY ABSENT**
> Does not establish fairness, non-discrimination, employment-law or regulatory requirements.

15. Third-Party / Supply Chain: Embedded Foundation Model

A company integrates an externally provided foundation model into its product. The provider controls model training and updates; the company controls configuration, invocation and product-level safeguards.

Scoping factors to consider

- Risk depends on product exposure and enabled capabilities.
- Provider-controlled updates and limited visibility into upstream controls.

- Need to distinguish vendor representations from independently verified evidence.

Frameworks that may become relevant

- GAISSF™: scope/profile should be determined for the company's own product and use of the provider model; supply-chain considerations are relevant inputs.
- UAIF™/AI-IRF™: may become relevant where provider-side changes or failures create potentially harmful events.

Illustrative evidence

- vendor representations versus independently verified evidence;
- model-update monitoring and change-impact records;
- product configuration and invocation boundaries.

> **NOTABLY ABSENT**
> Does not establish that vendor representations satisfy GAISSF™ requirements or define sufficient independent verification.

16. Cross-Sector: From AI Event to Incident Response

A deployed AI system begins producing a pattern of potentially harmful outputs. The organisation must determine what happened, preserve facts and uncertainty, and decide whether and how to respond.

Scoping factors to consider

- Begins after deployment rather than at initial scoping.
- Existing GAISSF™ scope/control baseline should be used where available; its absence should be documented where relevant.
- Facts, hypotheses and unknowns may evolve during response.

Frameworks that may become relevant

- UAIF™: may provide a common structure for determining whether the event constitutes an AI-specific incident and, where appropriate, characterising/classifying it.
- AI-IRF™: supports preparation and proportionate response; response may begin while incident determination/classification continues.
- GAISSF™: post-event findings may feed back into governance and control baselines.

Illustrative evidence

- timestamped detection and containment records;
- a contemporaneous record distinguishing known facts, working hypotheses, unresolved questions and observations not established by evidence;
- post-review linkage to applicable control baselines where they exist.

> **NOTABLY ABSENT**
> Does not establish a universal incident threshold or represent a complete UAIF™/AI-IRF™ workflow.

From Scenario to Your System

> **APPLICATION WORKFLOW**
> Select the closest scenario → identify where your system differs → determine actual scope/applicability → determine applicable requirements → establish implementation claims → evaluate evidence → record gaps and unknowns → reassess, remediate or respond as appropriate.
> Do not inherit a scenario's scope, framework selection, or evidence conclusions by analogy.

Relationship to Sector-Specific Requirements

GAISSF™ and the wider GAISSF Ecosystem do not replace sector-specific laws, regulations, supervisory requirements, safety standards, engineering standards, licensing requirements, or authorisation regimes.

Those external requirements may affect system scope, required controls, evidence expectations, assurance objectives, incident obligations, and acceptance criteria. Evidence generated through existing regulatory or assurance activities may also be relevant to GAISSF-related implementation claims, but relevance and sufficiency must be established rather than assumed.

Applicable obligations must therefore be determined separately for the jurisdiction, entity, regulated activity, system, and use case.

Cross-Sector Characteristics

IMPORTANT

The comparison below is illustrative only. It is not a risk rating, applicability determination, control-selection mechanism, maturity score, or assurance conclusion.

Context	Consequential decisions	Autonomous action	Physical effects	Third-party dependency	Human authority	Incident evidence
Finance	Often relevant	System-specific	Usually indirect	System-specific	Material consideration	Material consideration
Healthcare	Often relevant	System-specific	May be direct or indirect	System-specific	Material consideration	Material consideration
Physical AI	System-specific	Often material	Potentially direct	System-specific	Architecture-dependent	Material consideration
Critical infrastructure	System-specific	Architecture-dependent	Potentially direct	System-specific	Architecture-dependent	Material consideration
Employment	Often relevant	System-specific	Usually indirect	System-specific	Material consideration	Material consideration
Defence	System-specific	Architecture-dependent	May be direct or indirect	System-specific	Mission/use-case dependent	Material consideration

Common Assurance Patterns

Across the scenarios, recurring assurance questions concern human authority, system boundaries, third-party dependencies, change control, evidence provenance, consequential outcomes, uncertainty, and the ability to investigate potentially harmful behaviour. Their significance varies by system.

- Framework applicability must be determined from the specific system and context, not inferred from a sector label or similar example.
- More than one framework may apply to the same system. GAISSF™, UAIF™, AI-IRF™, and PAI-SF™ are not a fixed maturity sequence.
- For physical AI, PAI-SF™ should be considered where physical or cyber-physical capabilities and consequences place physical-AI risk in scope. Where PAI-SF™ applies, it supersedes and expands upon GAISSF™ Domain D9 (Physical AI Safety), retained in the current GAISSF™ release for continuity but not authoritative for detailed physical-AI requirements.
- Existing sector regulation, safety standards, and certification regimes remain separate obligations and must be evaluated independently.

KEY PRINCIPLE

Assurance is demonstrated, not asserted. Illustrative artifacts may support an implementation claim; artifact existence alone does not establish sufficiency, control effectiveness, system safety, conformance, certification, or regulatory compliance.

NOTABLY ABSENT

This guide does not enumerate detailed framework control objectives or individual control IDs, perform a formal applicability determination for any real system, provide legal advice, establish sector-specific regulatory compliance, report client outcomes, prescribe evidence automation, define universal human-review thresholds, or provide implementation-effort ratings. Every scenario is hypothetical.

Next Reading

- The 10-Minute Introduction to GAISSF™
- GAISSF™ Appendix P.1 SMB Quick-Start Guide
- The First 30 Days with GAISSF™
- Which GAISSF Ecosystem Framework Do I Need?
- Securing Physical AI: A Practitioner's Guide to PAI-SF™
- Building an AI Assurance Evidence Pack
- Implementing GAISSF™: A Practitioner's Guide