

GenAI Assure Framework

Version 2.1

Ethical • Legal • Safe • Accountable AI Governance

An assurance framework for organisations deploying third-party AI systems

Publication edition • Version 2.1 • 30 September 2026. Approved for publication by ELSA AI LTD. © 2026 ELSA AI LTD. CC BY-ND 4.0

Author: Faisal Ali, AAISM, CISM, CRISC

Owner: ELSA AI LTD

© 2026 ELSA AI LTD

Contents

Contents	2
1. Document control	4
1.1 Version history	4
1.2 Claims discipline	4
1.3 Owner decisions and publication terms	Error! Bookmark not defined.
2. Executive preface	5
3. Normative conventions, definitions and abbreviations	6
3.1 Normative language	6
3.2 Definitions	6
3.3 Abbreviations	7
4. Scope and architecture taxonomy	8
4.1 Purpose and audience	8
4.2 Technology scope	8
4.3 Out of scope	8
5. Actor classification	9
5.1 Why actor status is assessed per use case	9
5.2 Roles recognised by the framework	9
5.3 Modification decision tree	9
6. Governance model	10
6.1 Roles, authority and accountability	10
6.2 The advisory boundary	10
6.3 Framework and document change control	10
7. Risk classification method	11
7.1 Design intent	11
7.2 The seven stages	11
7.3 Mandatory overrides	11
7.4 Inherent-risk dimensions	12
7.5 Tiers and what they buy	12
8. Control architecture and catalogue	13
8.1 Atomic control schema	13
8.2 Reading the catalogue tables	13
8.3 Domain GV: Governance and accountability	14
8.4 Domain LR: Legal and regulatory screening	14
8.5 Domain DP: Data protection and privacy	15
8.6 Domain TS: Technical security	16
8.7 Domain AI: AI application security	16
8.8 Domain RG: RAG and knowledge systems	17
8.9 Domain AA: Agentic AI	17

8.10 Domain SP: Supplier and foundation-model assurance	18
8.11 Domain CM: Change and configuration management.....	19
8.12 Domain LG: Logging, monitoring and detection	19
8.13 Domain IR: Incident management and notification.....	20
8.14 Domain RS: Resilience and continuity	20
8.15 Domain HO: Human oversight and competence.....	21
8.16 Domain SQ: Safety, quality and human impact	21
8.17 Domain TE: Testing, evaluation, verification and validation	22
8.18 Domain EV: Evidence and assurance	22
8.19 Domain EX: Exceptions and risk acceptance	23
9. AI security threat model	24
10. Assurance model	25
10.1 The distinction v1.1 collapsed	25
10.2 Conclusion rules	25
11. TEVV methodology.....	26
11.1 Test specifications, not test labels.....	26
11.2 Acceptance criteria rules.....	26
12. Legal and regulatory implementation	27
12.1 EU AI Act deployer obligations (high-risk systems)	27
12.2 FRIA and the UK overlay	27
12.3 Incident notification decision matrix.....	28
13. Operational playbooks and implementation	29
13.1 Shadow AI operations playbook.....	29
13.2 Evidence architecture	29
13.3 Measurement and reporting	29
13.4 Implementation profiles	29
14. Standards and regulatory crosswalk	31
14.1 Source register	31
15. Continuous improvement	32
16. Publication Gate, licence and attribution	32
16.1 Publication Gate.....	32
16.2 Attribution.....	32
16.3 Licence	32
Annex A: Assessment finding traceability	34

1. Document control

Item	Value
Document title	GenAI Assure Framework
Version	2.1
Status	Publication edition. Approved for publication by ELSA AI LTD.
Document type	Assurance framework, controlled publication
Author	Faisal Ali, AAISM, CISM, CRISC
Owner	ELSA AI LTD
Copyright	© 2026 ELSA AI LTD
Slogan	Ethical • Legal • Safe • Accountable AI Governance
Publication date	30 September 2026
Review basis	Full architectural redesign in response to the Phase 1 assessment of v1.1 (66 findings, GA-001 to GA-066)
Licence	Creative Commons Attribution-NoDerivatives 4.0 International (CC BY-ND 4.0); Section 16.3

1.1 Version history

Version	Date	Change summary	Status
1.0	2025	Initial framework: six compound control groupings, tiering rubric, TEVV-Lite playbooks, evidence-pack concept.	Superseded
1.1	31 January 2026	Extended tool coverage, additional detections and documentation requirements.	Superseded. Not published.
2.0	July 2026	Architecture baseline approved by the owner in response to the Phase 1 assessment: atomic control catalogue, seven-stage risk method, actor classification, AI security threat model, agentic AI domain, corrected legal mapping. Design baseline only.	Internal architecture baseline
2.1	30 September 2026	Publication edition. Owner decisions closed; licence and trademark terms adopted; assurance boundary established; totals corrected to 115 controls across 17 domains. Traceability at Annex A.	Approved for publication

1.2 Claims discipline

This framework supports assurance; it does not establish production readiness, certification, accreditation, regulatory endorsement or legal compliance. References to laws and standards are subject to the crosswalk limitations in Section 14 and the assurance boundary in Section 10.2.

2. Executive preface

Version 1.1 of this framework packed six broad control groupings, sample configurations and selected legal statements into a condensed implementation guide. The Phase 1 assessment of July 2026 rated it Developing overall, with Initial ratings for control design, risk classification, AI security, testing and regulatory accuracy, and an Absent rating for agentic AI readiness. It found 66 defects, 21 of them Critical. Version 2.1 is the redesign that assessment demanded, not an editorial pass over the old text.

Five architectural changes define this version:

1. **Atomic controls.** The six compound controls are replaced by a catalogue of 115 single-requirement controls across 17 domains. Each control carries one objective, one owner, defined applicability, an evidence specification and a test reference, so an assessor can conclude pass or fail on each requirement individually (GA-024, GA-025).
2. **A sequenced risk method.** Additive scoring is withdrawn. Classification now runs through seven ordered stages, beginning with prohibited-use screening and mandatory legal overrides that no numerical score can offset. Budget, timeline and desired trust signal no longer influence the assessed tier (GA-020, GA-022).
3. **A current threat model.** The framework now maps controls to the OWASP Top 10 for LLM and GenAI Applications 2025, the OWASP agentic AI threat guidance and MITRE ATLAS, and adds a dedicated agentic AI control domain covering autonomy tiers, tool authorisation, containment, interruption and reversibility (GA-026 to GA-029).
4. **Corrected legal mapping.** Every legal statement flagged by the assessment is rewritten: DPIA and FRIA gates are now 100 per cent preconditions rather than 95 per cent targets, lawful basis is assessed per use case, the Article 17 erasure error is removed, EU AI Act Articles 26, 27 and 73 are stated at obligation level, and the UK overlay reflects the Data (Use and Access) Act 2025 (GA-015, GA-016, GA-033, GA-034, GA-039 to GA-041).
5. **An honest assurance model.** Control presence is no longer treated as effectiveness. The framework separates design, implementation and operating effectiveness, and defines evidence sufficiency, sampling and conclusion rules (GA-052).

What v1.1 got right survives. The deployer orientation, the Shadow AI operations playbook, the evidence-pack architecture, the lifecycle gates and the operational security content are retained and strengthened rather than discarded, as the assessment recommended.

Reading order. Sections 3 to 5 establish language, scope and actors. Sections 6 and 7 set governance and risk classification. Section 8 is the control catalogue. Sections 9 to 13 cover the threat model, assurance, testing, legal implementation and operations. Sections 14 to 16 govern the framework itself.

3. Normative conventions, definitions and abbreviations

3.1 Normative language

The framework uses controlled normative language. These words carry the following force wherever they appear in capitals:

Keyword	Force
MUST	A mandatory requirement for every use case within the control's stated applicability. Non-conformity requires a formal exception under domain EX, unless the requirement is listed as non-exceptionable.
MUST NOT	An absolute prohibition within the control's stated applicability.
SHOULD	A recommended practice. Departure is permitted where the organisation documents the rationale and any compensating measure.
MAY	A permitted option. No justification is needed either way.

No control applies universally. Every control states its applicability by risk tier, architecture class and actor role (Section 8.1). An organisation implements the controls its classified use cases attract, records permitted exclusions with rationale, and treats everything else as out of scope. This replaces the v1.1 instruction that organisations “must implement and operate all controls”, which the assessment found unworkable (GA-002).

3.2 Definitions

Term	Definition in this framework
AI system	A machine-based system, as defined in Article 3(1) of Regulation (EU) 2024/1689, that infers from inputs how to generate outputs such as predictions, content, recommendations or decisions that can influence physical or virtual environments.
AI model	The trained artefact (for example a foundation model, fine-tuned model or task-specific model) that produces inferences within an AI system.
AI tool	A third-party product or service through which the organisation consumes AI capability, including SaaS assistants, hosted model APIs, embedded AI features and no-code AI platforms.
Agent / agentic system	An AI system that plans and executes multi-step actions toward a goal, which may include invoking tools, using credentials, writing to systems of record or delegating to other agents, with reduced human involvement per action.
Use case	A defined business application of an AI tool, identified by a unique use_case_id, with a named accountable owner, a risk tier and an approved data scope.
Deployer	A natural or legal person using an AI system under its authority in the course of a professional activity, per Regulation (EU) 2024/1689.
Provider	A person or body that develops an AI system or has one developed and places it on the market or puts it into service under its own name or trademark, per the same Regulation.
Control	A single testable requirement with one objective, one accountable owner role, a defined evidence specification and a defined test.
Evidence	A record produced or collected to demonstrate that a control is designed, implemented or operating, meeting the specification in the control's schema.
Assurance	A supported conclusion about control design, implementation or operating effectiveness, reached through the methods in Section 10. Presence of a control artefact is not assurance.
Operating effectiveness	Demonstrated performance of a control over a defined period, evidenced through testing of a defined sample, not through configuration screenshots alone.
Exception	A time-limited, approved and compensated departure from a MUST requirement, recorded under domain EX.

Term	Definition in this framework
Shadow AI	AI tools or use cases in operation without sanction under this framework.
Zero-tolerance event	A test or production outcome defined in advance as unacceptable at any frequency, triggering automatic failure or incident handling regardless of aggregate pass rates.

3.3 Abbreviations

ADM automated decision-making. **AIMS** AI management system. **AICM** CSA AI Controls Matrix. **DLP** data loss prevention. **DPIA** data protection impact assessment. **DUAA** Data (Use and Access) Act 2025. **FRIA** fundamental rights impact assessment. **GPAI** general-purpose AI. **HRIA** human rights impact assessment. **IAM** identity and access management. **MCP** Model Context Protocol. **PAN** primary account number. **RAG** retrieval-augmented generation. **SIEM** security information and event management. **TEVV** testing, evaluation, verification and validation. **WORM** write once, read many.

4. Scope and architecture taxonomy

4.1 Purpose and audience

The framework governs organisations that procure, configure and operate third-party AI systems. It is written for the people who carry that work: accountable executives, risk and compliance functions, security architects, data protection officers, procurement leads and the business owners of individual AI use cases. Model development from scratch remains out of scope, but the boundary between deploying and providing is now determined by the actor classification in Section 5, not assumed (GA-003).

4.2 Technology scope

The framework applies to third-party AI in all of the following architecture classes. Each use case records which classes it involves; controls state their applicability against these classes (GA-004).

Class	Architecture	Typical examples
AC-01	SaaS AI assistants and copilots	Enterprise chat assistants, productivity copilots, custom copilot builders
AC-02	Hosted model APIs	Foundation-model APIs consumed directly or through gateways and model routers
AC-03	Retrieval-augmented generation	Vector stores, embedding pipelines, enterprise search over private corpora
AC-04	Fine-tuned or adapted models	Hosted fine-tuning, adapters, instruction-tuned variants of third-party models
AC-05	Multimodal systems	Image, audio, video and document inputs or outputs, including speech and vision
AC-06	Workflow orchestration	No-code and low-code AI platforms, automation chains, AI steps in business workflows
AC-07	AI agents	Single agents that plan and execute actions, including tool and function calling
AC-08	Multi-agent systems	Delegating, cooperating or supervising agent arrangements
AC-09	Browser and computer-use agents	Agents that operate a browser, desktop or terminal on the organisation's estate
AC-10	Plugin and MCP-style integrations	Tool ecosystems, connectors, function registries exposed to models
AC-11	Embedded AI	AI features inside procured products, devices or vertical applications
AC-12	Traditional predictive ML	Third-party scoring, classification and forecasting services deployed by the organisation
AC-13	Meeting and communication AI	Transcription, summarisation and recording assistants processing internal communications

4.3 Out of scope

Ground-up model development, training-infrastructure engineering and the provider-side obligations of organisations that place AI systems on the market under their own name are out of scope, **except** where the actor classification in Section 5 determines that the organisation has become a provider for a specific system. In that event the framework flags the provider obligations that arise and requires specialist legal support; it does not claim to satisfy them (GA-005).

5. Actor classification

5.1 Why actor status is assessed per use case

Regulation (EU) 2024/1689 attaches different obligations to providers, deployers, importers, distributors, authorised representatives and product manufacturers. v1.1 assumed the organisation is always a deployer. That assumption fails in practice: under Article 25, an organisation that puts its name or trademark on a high-risk system, substantially modifies one, or modifies the intended purpose of a system so that it becomes high-risk is treated as a provider, with provider obligations. Actor status is therefore assessed and recorded for every use case, and re-assessed on material change (GA-003, GA-005).

5.2 Roles recognised by the framework

Role	Working description	Framework consequence
Deployer	Uses an AI system under its authority in a professional activity.	Full framework applies. EU AI Act deployer obligations per Section 12.1 where the system is high-risk.
Provider (including by conversion)	Develops or has developed a system and supplies it under its own name, or is treated as provider under Article 25.	Escalation to legal counsel MUST occur. Provider obligations sit outside this framework's assurance scope and are flagged, not satisfied.
Importer / distributor	Places on the EU market, or makes available, a system produced by another.	Verification and cooperation duties flagged for legal assessment where the organisation resells or bundles AI.
Product manufacturer	Integrates a high-risk AI system into a product placed on the market under its name.	Escalation to legal counsel MUST occur before launch.
Authorised representative	Acts on written mandate of a non-EU provider.	Out of framework scope; flagged for legal assessment.

5.3 Modification decision tree

Before approval, and again on any material change, each use case answers the following in order. The first “yes” determines the route.

1. **Own-name supply.** Will the organisation market or supply the system, or a product containing it, under its own name or trademark? Yes: treat as provider route; escalate to counsel.
2. **Substantial modification.** Does fine-tuning, adaptation, system-prompt engineering, tool grafting or pipeline change materially alter the system's behaviour, safety characteristics or intended purpose beyond the provider's instructions for use? Yes: obtain a documented legal determination of actor status before go-live. Hosted fine-tuning and adapters (class AC-04) always trigger this question.
3. **Purpose change into high risk.** Is the organisation using a system for a purpose that brings it within a high-risk category the provider did not intend? Yes: provider-conversion risk; escalate to counsel.
4. **None of the above.** Record deployer status with the assessment date and assessor, and proceed under this framework.

Control anchor. This decision tree is enforced by controls GA-LR-02 and GA-LR-03 and is a precondition of Stage 2 of the risk method in Section 7.

6. Governance model

6.1 Roles, authority and accountability

The framework assigns each governance duty to a named role. One person can hold several roles in a small organisation, but the duties, and the segregations marked below, do not merge (GA-024 root cause: ownerless compound controls).

Role	Authority and accountability
Accountable executive	Approves the AI governance policy, owns aggregate AI risk, receives exception and incident reporting, chairs or sponsors the AI governance body.
AI governance body	Cross-functional board with defined membership, quorum and decision rights. Approves use cases at Tiers 3 and 4, approves acceptance criteria, adjudicates exceptions within its authority, reviews measurement reporting.
Use-case owner	Business owner of a specific use case. Owns the risk decision, the harm model, the data scope and the operational performance of that use case.
Security function	Advisory and assurance role. Identifies and quantifies risk, designs and tests security controls, and reports findings. It does not accept business risk and does not sign off business deployment decisions.
Data protection officer / privacy lead	Owns DPIA methodology, lawful basis assessment quality, rights procedures and privacy review of logging and evidence stores.
Legal counsel	Owns actor-status determinations, statutory interpretation, FRIA triggers, notification decisions and licence matters.
Control owners	One accountable role per control, named in the control record. Responsible for design, operation and evidence of that single control.
Independent tester / internal audit	Tests controls it does not own. For Tier 3 and Tier 4 use cases the tester MUST be independent of the control owner.

6.2 The advisory boundary

A recurring failure in AI programmes is the quiet conversion of security review into business sign-off. This framework keeps the boundary explicit: security identifies, quantifies and reports risk; the use-case owner and, at higher tiers, the governance body decide whether to accept it. Where a stakeholder requests security “approval”, the request is reframed as a risk advisory with a documented business acceptance. Every acceptance is recorded with the accepting role, date and residual-risk statement.

6.3 Framework and document change control

Changes to this framework are themselves change-controlled: a change log records each revision, revision markers such as “UPDATED REQUIREMENT” never appear in normative text (GA-063), and duplicated normative requirements are prohibited. Each requirement lives once, in the control catalogue, and other sections cross-reference it (GA-064, GA-065).

7. Risk classification method

7.1 Design intent

v1.1 added five dimension scores into a tier. Under that arithmetic, a recruitment screening tool with strong scores on convenience dimensions could land in a fast lane despite a decisive legal characteristic. The assessment rated the method structurally unsafe (GA-020). The replacement is sequential: legal screens and overrides come first and cannot be offset, scoring comes later and only within the space the overrides leave open, and available budget never touches the classification (GA-022). ISO/IEC 23894:2023 supplies the risk-management discipline the stages implement.

7.2 The seven stages

Stage	Name	What happens	Possible outcomes
1	Prohibited-use screening	The use case is checked against the prohibited AI practices in Article 5 of Regulation (EU) 2024/1689 and against the organisation's own prohibited-use list.	Blocked outright, or proceed.
2	Actor classification	The Section 5 decision tree runs. Provider-conversion routes leave this framework's scope.	Deployer route confirmed, or legal escalation.
3	System classification	The system is classified: high-risk (Article 6 and Annex III), transparency-obligated (Article 50), GPAI-based, or other. Jurisdictional applicability (EU, UK, other) is recorded.	Classification recorded with rationale.
4	Mandatory overrides	Defined characteristics force a minimum tier regardless of any score (Section 7.3).	Minimum tier set, or no override.
5	Inherent-risk evaluation	Nine dimensions are rated against anchored descriptors (Section 7.4). The highest-band dimension anchors the inherent rating; dimensions cannot dilute one another.	Inherent risk band.
6	Control-strength evaluation	Existing control design and operating evidence for this use case's applicable controls are assessed using Section 10 methods.	Control strength rating.
7	Residual-risk acceptance	The residual position is accepted, treated or rejected by the authority the tier demands. Non-exceptionable requirements (GA-EX-02) cannot be waived here.	Approval, conditions, or rejection.

7.3 Mandatory overrides

Any of the following sets the floor shown, before and regardless of Stage 5 scoring:

Trigger	Minimum tier
System classified high-risk under Article 6 / Annex III	Tier 4
Solely automated decisions with legal or similarly significant effect on individuals (UK GDPR Articles 22A to 22D or GDPR Article 22 territory)	Tier 4
Special category or criminal-offence data processed at scale	Tier 4
Safety-of-life, clinical or critical-infrastructure function	Tier 4
Children or vulnerable persons materially affected	Tier 4
Agentic capability with external, irreversible action (payment, record deletion, legal commitment)	Tier 4
Agentic capability with external but reversible action	Tier 3 minimum
Employment, credit, housing, insurance or eligibility decisions supported (not solely automated)	Tier 3 minimum
Confidential or trade-secret corpus exposed to a third-party service	Tier 3 minimum

7.4 Inherent-risk dimensions

Stage 5 rates each dimension Low, Medium, High or Very High against anchored descriptors maintained in the method annex. The dimensions replace the v1.1 five-factor set and add the factors the assessment found missing (GA-021):

- **Data sensitivity** of inputs, retrieval corpora and outputs.
- **Decision impact** on individuals, customers, patients or the organisation.
- **Autonomy level** using the A0 to A4 tiers in Section 8, domain AA.
- **Affected population and vulnerability**, including scale of exposure.
- **Reversibility** of the actions or decisions the system produces or executes.
- **Model capability and generality**, including GPAI-based systems and frontier-capability models.
- **External action capability**: tool calling, browsing, code execution, transactions.
- **Dependency criticality**: what fails in the business if this system fails or is withdrawn.
- **Novelty and change velocity** of the architecture and provider.

7.5 Tiers and what they buy

Tier 1 (Minimal), Tier 2 (Standard), Tier 3 (Enhanced) and Tier 4 (Maximum) determine which controls apply, the assurance level demanded (Section 10), approval authority and TEVV depth. Resourcing, budget and delivery timeline are planned through the implementation profiles in Section 13.4 *after* classification. They are never inputs to it. A use case the organisation cannot afford to govern at its assessed tier is deferred or rejected, not reclassified (GA-022).

8. Control architecture and catalogue

8.1 Atomic control schema

v1.1 carried six catalogue entries; GA-TP-001 alone bundled IAM, secrets, DLP, network security, cryptography and configuration management. One passing component could conceal five failing ones, and no owner, test or conclusion could attach to the whole (GA-024). Version 2.1 replaces them with 115 atomic controls in 17 domains. Every control record holds the full schema below; the catalogue tables in this section print the five fields an assessor reaches for first, and the complete records live in the control register that accompanies the framework (GA-025).

Schema field	Content
Control ID	GA-<domain>-<nn>, stable across versions.
Title and objective	One objective per control. No compound requirements.
Normative requirement	The MUST / MUST NOT / SHOULD statement.
Applicability	Risk tiers, architecture classes (AC-01 to AC-13) and actor roles the control binds.
Owner	The single accountable role.
Evidence specification	What records demonstrate the control, their source and correlation keys.
Test procedure and frequency	Reference to the TEVV specification (Section 11), the assurance level it supports, and how often it runs.
Failure and exception route	What happens on failure; whether the control is exceptionable under domain EX.

Worked example: one control in full schema

Field	GA-AA-04 (Action limits for agents)
Objective	Agents cannot take unbounded or irreversible external actions.
Requirement	Each agent use case MUST enforce configured action limits covering spend ceilings, record-modification scope and a deny-list of irreversible action classes that require human approval under GA-AA-05.
Applicability	Tiers 2 to 4; classes AC-07, AC-08, AC-09; deployer role.
Owner	Use-case owner, with the platform team operating the enforcement point.
Evidence	Limit configuration export, enforcement-point logs keyed by use_case_id and agent identity, denied-action records, quarterly limit review minutes.
Test	TEVV spec TS-AA-04: attempt actions beyond each configured limit in a controlled environment; expected result is denial and alert. Pre-deployment, then quarterly, and on any tool-scope change.
Failure route	Deployment blocked or agent suspended. Exceptionable only for spend ceilings, never for the irreversible-action deny-list.

8.2 Reading the catalogue tables

Applicability shorthand: “T2+” means Tiers 2, 3 and 4; “All” means every tier or every architecture class; specific classes are listed by AC number. All controls bind the deployer role unless stated. **Owner shorthand:** Exec (accountable executive), Gov (AI governance body), UCO (use-case owner), Sec (security function), DPO (data protection lead), Legal, Ops (IT operations or platform team), Audit (independent tester), P&C (people and capability function).

8.3 Domain GV: Governance and accountability

Establishes who decides, who advises and who is accountable, before any tool is switched on.

ID	Requirement	Owner	Applicability	Evidence and test
GA-GV-01	An AI governance policy MUST be approved by the accountable executive and reviewed at least annually and on material regulatory change.	Exec	All	Approved policy, review record
GA-GV-02	Every sanctioned use case MUST have one named accountable owner recorded in the inventory.	Gov	All	Inventory extract with owner field populated
GA-GV-03	A complete AI use-case inventory MUST be maintained, with a unique use_case_id, architecture classes, actor status, tier and data scope per entry.	Gov	All	Inventory export, reconciliation against discovery (GA-LG-05)
GA-GV-04	An AI governance body MUST operate with defined membership, quorum and decision rights, and MUST approve all Tier 3 and Tier 4 use cases before go-live.	Exec	T3+	Terms of reference, approval minutes keyed to use_case_id
GA-GV-05	Security MUST hold an advisory and assurance role; business risk acceptance MUST be recorded against the use-case owner or governance body, never the security function.	Exec	All	Risk acceptance records showing accepting role
GA-GV-06	For Tier 3 and Tier 4 use cases, the person testing a control MUST NOT be its owner.	Audit	T3+	Test records showing tester identity vs control owner
GA-GV-07	A named owner MUST monitor regulatory and standards change on a defined cadence and maintain a change log feeding framework updates.	Legal	All	Watch log, resulting change records
GA-GV-08	Framework and policy documents MUST be version-controlled with a change log; revision markers MUST NOT appear in normative text.	Gov	All	Change log, document QA record

Assessment findings resolved: GA-002, GA-022 (governance separation), GA-024, GA-061, GA-063.

8.4 Domain LR: Legal and regulatory screening

Runs the legal gates that no score can offset. Statements of law in this domain are summaries to drive workflow; counsel owns interpretation.

Assurance of LR controls evaluates the decision process and retained evidence under the Section 10.2 ceiling. It does not substitute for an organisation's legal determination.

ID	Requirement	Owner	Applicability	Evidence and test
GA-LR-01	Every proposed use case MUST be screened against the prohibited AI practices in Article 5 of Regulation (EU) 2024/1689 and the organisation's prohibited-use list before any other assessment.	Legal	All	Screening record at Stage 1
GA-LR-02	Actor status MUST be determined per use case using the Section 5 decision tree, recorded, and re-assessed on material change including fine-tuning, purpose change or own-name supply.	Legal	All	Actor determination record with date and assessor
GA-LR-03	Each use case MUST record whether the system is high-risk (Article 6 and Annex III), transparency-obligated (Article 50), GPAI-based or other, with rationale.	Legal	All	Classification record
GA-LR-04	For each high-risk system deployed, the Article 26 obligation map in Section 12.1 MUST be completed with evidence per obligation before go-live.	UCO	T4, high-risk	Completed obligation map
GA-LR-05	Where the Article 27 statutory FRIA trigger applies, a FRIA MUST be completed and required notifications made before deployment. Voluntary human-rights impact assessments MUST be labelled as voluntary and kept distinct.	Legal	High-risk, in-scope deployers	FRIA record, notification evidence
GA-LR-06	Where Article 50 applies, affected persons MUST be informed they are interacting with AI, and synthetic or manipulated content MUST be disclosed as such.	UCO	AC-01, 05, 13; T2+	Interface notices, labelling configuration

ID	Requirement	Owner	Applicability	Evidence and test
GA-LR-07	Registration obligations applying to the organisation's deployments (including public-authority registration duties) MUST be identified and completed before go-live.	Legal	High-risk, in-scope	Registration evidence
GA-LR-08	Jurisdictional applicability (EU, UK, other) MUST be assessed per deployment, and jurisdiction-specific overlays applied; requirements from one jurisdiction MUST NOT be presented as universal.	Legal	All	Applicability assessment

Assessment findings resolved: GA-003, GA-005, GA-010, GA-016, GA-039, GA-047 scope aspects, GA-050 (Article 50 duties).

8.5 Domain DP: Data protection and privacy

Replaces the v1.1 shortcuts with per-case legal method. Completion percentages below 100 are not tolerated where a statutory precondition applies.

ID	Requirement	Owner	Applicability	Evidence and test
GA-DP-01	Lawful basis MUST be assessed and documented per use case against purpose, necessity, data types, relationship and reasonable expectations. Generic lawful bases assigned to whole processing classes MUST NOT be used.	DPO	All	Per-case lawful basis assessment
GA-DP-02	Where the GDPR Article 35 threshold is met, a DPIA MUST be completed before go-live. This is a blocking gate: mandatory-gate compliance is 100 per cent, and unmet prerequisites block deployment.	DPO	All	DPIA record linked to go-live approval
GA-DP-03	Where UK GDPR Articles 22A to 22D apply to significant automated decisions, the required safeguards (information, human intervention, representations, contestation) MUST be implemented, with the stricter conditions for special category data observed.	DPO	UK ADM cases	Safeguard design and operation evidence
GA-DP-04	Each use case MUST define approved data classes for prompts, uploads and retrieval corpora; data outside the approved classes MUST NOT be submitted.	UCO	All	Data scope record, DLP alignment (GA-TS-06)
GA-DP-05	A subject access procedure MUST cover AI records: discovery across prompts, outputs and logs, identity matching, third-party redaction and rights-conflict handling under Article 15.	DPO	T2+	Procedure, executed request samples
GA-DP-06	Erasure and restriction requests MUST receive individual legal analysis. Technical difficulty of removing data from a trained model MUST NOT be treated as an Article 17(3) exemption, and restriction MUST NOT be applied as an automatic substitute for erasure.	DPO	All	Request handling records with legal analysis
GA-DP-07	International transfers involved in AI services MUST be assessed, safeguarded and recorded in a transfer register, with region pinning applied where required (GA-SP-04).	DPO	All	Transfer register entries
GA-DP-08	De-identification measures MUST be selected against a documented threat model, distinguishing masking, keyed hashing, tokenisation and encryption. Unsalted deterministic hashing MUST NOT be relied on for low-entropy identifiers, and format-preserving encryption MUST NOT be described as tokenisation.	Sec	T2+	Technique selection record per data flow
GA-DP-09	Log and evidence stores holding personal data MUST themselves carry data protection controls: classification, access restriction, retention limits and secure deletion (see GA-EV-02, GA-EV-03).	DPO	All	Store-level privacy review

Assessment findings resolved: GA-009, GA-012, GA-013, GA-015, GA-033, GA-034, GA-035, GA-041, GA-043, GA-065.

8.6 Domain TS: Technical security

The conventional security spine, restated one requirement at a time and stripped of misattributed citations. Where a threshold below is an internal design choice, it says so.

ID	Requirement	Owner	Applicability	Evidence and test
GA-TS-01	Access to sanctioned AI tools MUST use federated identity with multi-factor authentication.	Ops	All	IdP configuration, coverage report
GA-TS-02	Entitlements to AI tools and their data scopes MUST follow least privilege and be recertified on a defined cycle.	Ops	T2+	Entitlement reviews
GA-TS-03	API keys, tokens and service credentials for AI services MUST be vaulted and rotated; credentials MUST NOT be embedded in prompts, notebooks or workflow definitions.	Sec	AC-02, 06, 07, 10	Vault inventory, secret-scanning results
GA-TS-04	Transport and storage cryptography MUST meet an internal, risk-based cryptographic baseline set by the security function. The baseline is organisational policy; it is not attributed to NIST SP 800-52 or any external source as a universal mandate, and supplier deviations are handled by risk assessment rather than a fixed retirement deadline.	Sec	All	Baseline document, configuration evidence
GA-TS-05	Network egress to AI services MUST be restricted to sanctioned endpoints via allow-listing at the egress control point.	Sec	T2+	Egress policy export, blocked-attempt logs
GA-TS-06	DLP for AI channels MUST implement jurisdiction-specific identifier patterns defined by the privacy and security functions. PAN detection MUST be implemented where, and only where, cardholder data scope exists, aligned to PCI DSS v4.0.1; Luhn-based detection is internal implementation guidance, not a PCI citation.	Sec	T2+	Pattern library with jurisdiction mapping, detection logs
GA-TS-07	Tenant, workspace and data-boundary isolation settings offered by each AI service MUST be configured and verified.	Ops	All	Configuration evidence per service
GA-TS-08	A configuration baseline MUST exist per sanctioned AI tool, with drift detection and remediation.	Ops	T2+	Baseline, drift reports

Assessment findings resolved: GA-007, GA-008, GA-009 (pattern sourcing), GA-010, GA-011, GA-064.

8.7 Domain AI: AI application security

Controls for the attack surface v1.1 did not model: the model interface itself. Mapped to the OWASP Top 10 for LLM and GenAI Applications 2025 in Section 9.

ID	Requirement	Owner	Applicability	Evidence and test
GA-AI-01	Each use case MUST define input trust boundaries and an instruction hierarchy distinguishing system instructions, user input and retrieved or third-party content, and configure available controls to enforce it.	Sec	T2+, AC-01 to 03, 06, 07, 10	Boundary design record, configuration
GA-AI-02	Direct and indirect prompt injection testing MUST be performed before deployment and on material change, using the TEVV specification for the use case.	Sec	T2+	Test results against TS-AI-02
GA-AI-03	System prompts MUST be treated as confidential configuration; leakage testing MUST be included in pre-deployment TEVV for Tier 3 and above.	Sec	T3+	Leakage test results
GA-AI-04	Model outputs MUST be validated and safely handled before any downstream execution: no direct execution of generated code or commands, output encoding for downstream interpreters, and schema validation for structured outputs.	Ops	T2+, AC-02, 06, 07, 10	Output-handling design, validation config
GA-AI-05	Sensitive-information disclosure testing (personal data, secrets, confidential corpus content) MUST run pre-deployment and on corpus or model change.	Sec	T2+	Test results

ID	Requirement	Owner	Applicability	Evidence and test
GA-AI-06	Rate limits, spend ceilings and quota alerts MUST be configured per use case to contain unbounded consumption and denial-of-wallet abuse.	Ops	AC-02, 06, 07	Limit configuration, alert records
GA-AI-07	Jailbreak-resistance evaluation MUST be performed at a depth proportionate to tier, using documented attack corpora.	Sec	T3+	Evaluation report
GA-AI-08	Use cases presenting generated content to decision-makers or the public MUST implement overreliance mitigations appropriate to the case: AI labelling, source citation where retrieval is used, and stated limitations.	UCO	T2+	Interface evidence, user guidance

Assessment findings resolved: GA-027, GA-028, GA-036 (overreliance and synthetic content aspects).

8.8 Domain RG: RAG and knowledge systems

Retrieval pipelines carry instructions as well as facts. These controls treat the corpus as an input trust problem, not a storage problem.

ID	Requirement	Owner	Applicability	Evidence and test
GA-RG-01	Knowledge sources feeding retrieval MUST be approved, with provenance recorded per source and ingestion restricted to approved connectors.	UCO	AC-03	Source register, connector config
GA-RG-02	Content MUST be screened before indexing for embedded instructions, malware indicators and data-classification violations.	Sec	AC-03, T2+	Screening pipeline evidence, quarantine records
GA-RG-03	Retrieval MUST enforce the querying user's entitlements at query time; the vector store MUST NOT become a bypass of source-system access control.	Ops	AC-03	Entitlement enforcement test results
GA-RG-04	Vector stores MUST enforce tenant and collection isolation between use cases and data domains.	Ops	AC-03	Isolation configuration and test
GA-RG-05	Poisoning and indirect-injection testing of the retrieval corpus MUST be included in TEVV for Tier 3 and above.	Sec	AC-03, T3+	Test results against TS-RG-05
GA-RG-06	Corrections, deletions and data-subject rights outcomes MUST propagate to derived artefacts including embeddings and caches within a defined period.	DPO	AC-03	Propagation procedure, executed samples

Assessment findings resolved: GA-029, and the retrieval aspects of GA-004.

8.9 Domain AA: Agentic AI

The domain v1.1 lacked entirely. Autonomy is tiered, bounded, identified, contained, interruptible and reversible, in that order.

ID	Requirement	Owner	Applicability	Evidence and test
GA-AA-01	Every agent use case MUST be assigned an autonomy tier: A0 human executes; A1 suggest only; A2 act with per-action approval; A3 act within bounded scope with post-hoc review; A4 autonomous within a sandbox with no irreversible external action. The tier feeds Stage 4 and Stage 5 of the risk method.	UCO	AC-07 to 09	Autonomy tier record
GA-AA-02	Tools and functions available to an agent MUST be allow-listed per use case with per-tool authorisation scopes; unlisted tools MUST be unreachable.	Sec	AC-07 to 10	Tool registry, scope configuration
GA-AA-03	Agents MUST run under distinct non-human identities with their own credentials and entitlements; shared or borrowed human credentials MUST NOT be used.	Ops	AC-07 to 09	Identity inventory, credential audit
GA-AA-04	Configured action limits MUST cover spend ceilings, record-modification scope and a deny-list of irreversible action classes (full schema in Section 8.1).	UCO	T2+, AC-07 to 09	Limit config, denied-action logs

ID	Requirement	Owner	Applicability	Evidence and test
GA-AA-05	Human approval gates MUST intercept defined action classes before execution, and an interruption capability MUST allow a human to halt an agent immediately at any point.	UCO	AC-07 to 09	Gate configuration, interruption test results
GA-AA-06	Browsing, code execution and computer-use capabilities MUST run in contained environments with no default reach into production systems or unrestricted networks.	Sec	AC-09, T2+	Sandbox architecture, containment test
GA-AA-07	Agent memory persistence MUST be approved per use case, with content controls on what may be stored, and a purge capability covering rights requests and incident response.	DPO	AC-07 to 09	Memory policy, purge test
GA-AA-08	Delegation between agents MUST be bounded by explicit authority limits, loop and recursion protection MUST be configured, and the full action chain across agents MUST be logged with correlation keys.	Sec	AC-08	Delegation config, chain logs

Assessment findings resolved: GA-026, GA-021 (autonomy dimension), GA-004 (classes AC-07 to AC-10).

8.10 Domain SP: Supplier and foundation-model assurance

Certificates prove an audit happened; they do not prove a model suits the use. This domain assesses the model, the service and the exit, separately.

ID	Requirement	Owner	Applicability	Evidence and test
GA-SP-01	AI supplier due diligence MUST assess suitability of the specific service and model for the intended use, beyond collection of SOC reports and certificates.	UCO	All	Due diligence record per service
GA-SP-02	Foundation-model assessment MUST review model provenance, provider training-data statements, published safety evaluations and model documentation for the model version in use.	Sec	T2+, AC-02, 04, 07	Model assessment record
GA-SP-03	Contracts MUST address organisational data use in training (with opt-out where offered), sub-processing, confidentiality, incident notification duties, audit or evidence rights, and exit terms.	Legal	All	Contract clause mapping
GA-SP-04	Hosting locations and data residency MUST be verified, with region pinning configured where the transfer assessment (GA-DP-07) requires it.	Ops	All	Residency evidence, pinning config
GA-SP-05	Provider change communications MUST be monitored through a named channel and owner, feeding the change process in GA-CM-01.	Ops	All	Monitoring record
GA-SP-06	Each critical AI service MUST have a documented exit and portability plan covering data retrieval, deletion attestation and replacement options.	UCO	T3+	Exit plan, tested where feasible
GA-SP-07	Dataset and connector suppliers feeding RAG or fine-tuning MUST pass the same due diligence as model suppliers, proportionate to tier.	UCO	AC-03, 04	Supplier records

Assessment findings resolved: GA-030, and the acquisition aspects of NIST SP 800-218A coverage.

8.11 Domain CM: Change and configuration management

Hosted models change without customer releases. Approval must therefore expire when the thing approved no longer exists.

ID	Requirement	Owner	Applicability	Evidence and test
GA-CM-01	A model and provider change-classification process MUST categorise changes (model version, behaviour, terms, hosting, tooling) and route each category to a defined response.	Ops	All	Change classification records
GA-CM-02	Model or version pinning MUST be configured where the service offers it for Tier 3 and above; where unavailable, the absence is recorded as a risk factor at Stage 5.	Ops	T3+	Pinning configuration or documented absence
GA-CM-03	Material model, provider or configuration changes MUST trigger regression testing against the use case's TEVV specification before continued production use.	Sec	T2+	Regression results linked to change record
GA-CM-04	Defined change classes MUST trigger reapproval by the original approval authority.	Gov	T3+	Reapproval records
GA-CM-05	Internal configuration changes to AI tools (system prompts, tool scopes, retrieval sources, limits) MUST pass through change control with rollback capability.	Ops	T2+	Change tickets, rollback evidence

Assessment findings resolved: GA-031.

8.12 Domain LG: Logging, monitoring and detection

v1.1 ordered full prompts and outputs into SIEM and WORM storage, creating a second high-risk personal-data repository. The default inverts here: metadata always, content only by justified exception.

ID	Requirement	Owner	Applicability	Evidence and test
GA-LG-01	AI interaction events MUST be logged as metadata by default: actor, time, tool, use_case_id, action class, decision outcome and correlation keys. Prompt and output content MUST NOT be logged by default.	Sec	All	Log schema, sample records
GA-LG-02	Content-level logging MAY be enabled only per use case, with documented justification, minimisation, protection, access restriction and a time limit, approved by the privacy lead.	DPO	As approved	Content-logging approvals
GA-LG-03	The event taxonomy MUST separate raw events, alerts, incidents and regulatory case records; regulatory determinations (breach status, reportability) MUST NOT be encoded as fields at event-generation time.	Sec	All	Taxonomy document, schema evidence
GA-LG-04	Detection use cases MUST cover AI-specific abuse patterns: injection indicators, anomalous tool-call chains, exfiltration signatures, agent limit breaches and consumption anomalies.	Sec	T2+	Detection rule inventory, alert samples
GA-LG-05	Shadow AI discovery MUST run continuously across available signals (network, SaaS management, endpoint, expense data), feeding the Section 13.1 playbook.	Sec	All	Discovery reports
GA-LG-06	Logs generated by high-risk systems under the organisation's control MUST be retained for at least six months per Article 26, and longer only where a documented basis exists, respecting storage limitation.	Ops	High-risk	Retention configuration
GA-LG-07	Log stores MUST enforce integrity protection and role-restricted access.	Ops	All	Access model, integrity mechanism

Assessment findings resolved: GA-012, GA-013, GA-014, GA-023 (signal side), GA-043.

8.13 Domain IR: Incident management and notification

One AI incident can trigger security, privacy, safety and AI Act duties concurrently, on different clocks. The matrix in Section 12.3 does the routing; these controls make it operable.

ID	Requirement	Owner	Applicability	Evidence and test
GA-IR-01	An AI incident taxonomy MUST classify incidents across security, privacy, safety, rights and discrimination, quality, and availability categories, with severity levels and pre-defined zero-tolerance events.	Sec	All	Taxonomy, classification records
GA-IR-02	The notification decision matrix (Section 12.3) MUST be applied to every qualifying incident, producing a documented decision per notification route and deadline.	Legal	All	Completed matrix per incident
GA-IR-03	On identifying a serious incident involving a deployed high-risk system, the provider MUST be informed without delay, and the organisation's own reporting duties assessed immediately against the Article 73 outer limits.	Legal	High-risk	Provider notification evidence, assessment record
GA-IR-04	AI-specific response runbooks MUST exist for at minimum: prompt injection compromise, retrieval poisoning, agent malfunction or runaway, model output causing harm, and provider outage or withdrawal.	Sec	T2+	Runbooks, exercise records
GA-IR-05	Incident handling MUST preserve evidence and observe investigation safeguards before disruptive containment where feasible, balancing service criticality (see GA-EX and Section 13.1 proportionality).	Sec	All	Handling records
GA-IR-06	Post-incident review MUST feed corrective actions into the control catalogue and TEVV specifications.	Gov	T2+	Review outputs, resulting changes

Assessment findings resolved: GA-040, GA-053, GA-023 (proportionate response).

8.14 Domain RS: Resilience and continuity

AI dependencies fail in AI-specific ways: silent model degradation, provider withdrawal, tool ecosystem breakage. Fallback is designed, not assumed.

ID	Requirement	Owner	Applicability	Evidence and test
GA-RS-01	Dependencies of each critical AI use case (models, providers, tools, corpora, connectors) MUST be mapped and kept current.	Ops	T3+	Dependency map
GA-RS-02	Degradation modes and safe-state behaviour MUST be defined per critical use case, including what the system does when confidence, availability or integrity drops.	UCO	T3+	Design records
GA-RS-03	Fallback paths (alternate model, manual process, graceful denial) MUST exist and be tested for critical use cases.	Ops	T3+	Fallback test results
GA-RS-04	A provider withdrawal and model deprecation response MUST be defined, linked to the exit plan in GA-SP-06.	UCO	T3+	Response plan
GA-RS-05	Availability objectives and health monitoring MUST be defined per critical use case, with alerting on breach.	Ops	T3+	Objectives, monitoring evidence

Assessment findings resolved: Resilience gaps identified in the coverage matrix; GA-004 (dependency criticality).

8.15 Domain HO: Human oversight and competence

Oversight is required by risk, safety design and provider instructions, not only where a statute says the word. And an overseer who cannot catch failures is decoration.

ID	Requirement	Owner	Applicability	Evidence and test
GA-HO-01	Each use case MUST define its oversight mode (in-the-loop, on-the-loop, or post-hoc review) based on risk tier, safety design and provider instructions for use, not solely on statutory mandate.	UCO	All	Oversight design record
GA-HO-02	Overseers MUST hold defined authority and intervention points, including the power to stop, override or suspend the system's output or action.	UCO	T2+	Authority definition, intervention mechanism
GA-HO-03	Overseer competence requirements MUST be defined and individuals assessed against them before assignment.	P&C	T3+	Competence assessments
GA-HO-04	Oversight effectiveness MUST be tested: seeded failure cases measure whether overseers detect and act within the design assumptions.	Audit	T3+	Effectiveness test results
GA-HO-05	Role-based AI literacy MUST cover system limitations, correct use, escalation routes, oversight duties and affected-person rights, with competence assessment and defined retraining triggers. Annual awareness training alone MUST NOT be treated as satisfying this control.	P&C	All	Curriculum, assessment records
GA-HO-06	High-impact decision workflows MUST implement automation-bias countermeasures, such as rationale prompts, spot-check sampling or blind review, selected per use case.	UCO	T3+	Countermeasure design, operation evidence

Assessment findings resolved: GA-018, GA-019.

8.16 Domain SQ: Safety, quality and human impact

Principles become controls here: harm is modelled before approval, fairness is measured in context, and affected people get notice and a route to contest.

ID	Requirement	Owner	Applicability	Evidence and test
GA-SQ-01	A use-case harm model MUST be documented before approval, identifying who can be harmed, how, at what severity and reversibility, and which controls respond.	UCO	T2+	Harm model record
GA-SQ-02	Fairness evaluation MUST use context-specific metrics with subgroup analysis and confidence reporting, against acceptance criteria approved by the governance body per use case. Generic parity thresholds MUST NOT be applied as pass criteria.	UCO	T3+, decisions affecting persons	Evaluation reports, approved criteria
GA-SQ-03	Accuracy and quality acceptance criteria MUST be defined per use case, with severity-weighted measurement; aggregate pass rates MUST NOT conceal defined zero-tolerance failures.	UCO	T2+	Criteria, evaluation results
GA-SQ-04	In-scope user interfaces MUST meet WCAG 2.2 Level AA where accessibility obligations apply to the deployment's jurisdiction and context.	UCO	Interfaces in scope	Accessibility test results
GA-SQ-05	Affected persons MUST receive transparency appropriate to the case: notice of AI involvement and labelling of AI-generated content, coordinated with GA-LR-06.	UCO	T2+	Notices, labelling evidence
GA-SQ-06	A contestability and redress route MUST exist for persons materially affected by AI-supported outcomes, with defined response times and escalation.	UCO	T3+	Redress procedure, case records
GA-SQ-07	Use cases generating or consuming synthetic media MUST define use-case-specific provenance, authenticity and manipulation-detection measures. No fixed external benchmark is mandated; the discredited composite "NIST IARPA FVRT" requirement from v1.1 is withdrawn.	Sec	AC-05	Measure design, test results

Assessment findings resolved: GA-017, GA-036, GA-054, GA-057.

8.17 Domain TE: Testing, evaluation, verification and validation

TEVV controls anchor the methodology in Section 11. Labels without protocols produced v1.1's unreliable results; specifications produce repeatable ones.

ID	Requirement	Owner	Applicability	Evidence and test
GA-TE-01	Every applicable control and every Tier 2+ use case MUST have a TEVV specification defining method, environment, dataset or corpus, sample size, expected result, tolerance, reviewer and evidence output.	Sec	T2+	Specification register
GA-TE-02	A pre-deployment TEVV gate MUST pass before go-live, at the depth the tier demands, covering security, privacy, quality and oversight tests applicable to the use case.	Gov	T2+	Gate results linked to approval
GA-TE-03	Adversarial testing MUST be performed for Tier 3 and Tier 4 use cases, scoped to include prompt injection, jailbreak, data extraction and, where agentic, tool misuse and containment escape.	Sec	T3+	Red-team reports
GA-TE-04	Regression TEVV MUST run on model, provider or material configuration change before continued production use (with GA-CM-03).	Sec	T2+	Regression results
GA-TE-05	Ongoing production evaluation MUST monitor for drift against acceptance criteria, with defined triggers for re-test and escalation.	UCO	T3+	Monitoring outputs, trigger records
GA-TE-06	For Tier 4 use cases, TEVV results MUST receive independent review before the approval decision.	Audit	T4	Independent review record

Assessment findings resolved: GA-037, GA-038, GA-017 (methodology aspects).

8.18 Domain EV: Evidence and assurance

The evidence-pack concept survives from v1.1; the privacy and conclusion discipline around it is new.

ID	Requirement	Owner	Applicability	Evidence and test
GA-EV-01	Each use case MUST maintain an evidence pack with a manifest, integrity hashes and correlation keys (use_case_id, control_id, evidence_id).	UCO	T2+	Manifest, hash verification
GA-EV-02	Evidence stores MUST enforce classification, role-restricted access and minimisation; evidence containing personal data follows GA-DP-09.	Sec	All	Store controls evidence
GA-EV-03	Evidence retention and secure deletion schedules MUST be defined and executed; WORM protection is applied where integrity requires it and MUST NOT override lawful deletion obligations.	Ops	All	Schedule, deletion records
GA-EV-04	Assurance conclusions MUST state which level they support: design, implementation, or operating effectiveness over a stated period, using the Section 10 model. Artefact existence alone MUST NOT be reported as effectiveness.	Audit	All	Conclusion records
GA-EV-05	Sampling approach and evidence-sufficiency rules MUST be documented as internal methodology. Sampling rates, retrieval timeframes and CAPA deadlines are internal service objectives; they MUST NOT be attributed to ISO/IEC 42001, ISO/IEC 27001 or SOC 2 as source mandates.	Audit	T2+	Methodology document
GA-EV-06	An internal assurance programme MUST test controls on a risk-based cycle, with findings tracked to closure through corrective actions.	Audit	T2+	Programme plan, findings register

Assessment findings resolved: GA-032, GA-043, GA-045, GA-052, GA-056 (conclusion side), GA-060.

8.19 Domain EX: Exceptions and risk acceptance

Exceptions are how a control system stays honest under pressure. Some requirements, though, are not the organisation's to waive.

ID	Requirement	Owner	Applicability	Evidence and test
GA-EX-01	Every departure from a MUST requirement MUST be recorded in an exception register with a unique ID, scope, rationale, compensating controls and expiry date.	Gov	All	Exception register
GA-EX-02	A non-exceptionable list MUST be maintained, including at minimum: statutory obligations (GDPR, UK GDPR, EU AI Act duties where applicable), Stage 1 prohibited uses, the agent irreversible-action deny-list, and defined zero-tolerance safety events. Exceptions against these MUST NOT be granted at any authority level.	Legal	All	Published list, rejected-request records
GA-EX-03	Approval authority MUST scale with tier: use-case owner for Tier 1 and 2 within limits, governance body for Tier 3, accountable executive for Tier 4. Maximum durations apply per tier.	Gov	All	Approval records
GA-EX-04	Exceptions MUST be revalidated before expiry; breach of exception conditions escalates automatically to the next authority level.	Gov	All	Revalidation records
GA-EX-05	Aggregate exception exposure MUST be reported to the governance body on a defined cadence.	Gov	T2+	Reporting packs

Assessment findings resolved: GA-055.

9. AI security threat model

Threat-led assurance needs a threat model. This section maps the current AI attack classes to the controls that respond, so an assessor can start from either end. References: OWASP Top 10 for LLM and GenAI Applications 2025, OWASP Agentic AI Threats and Mitigations (February 2025 and subsequent project resources), and MITRE ATLAS (current knowledge base). The mapping shows which controls respond to which threat; it does not claim the threat is eliminated (GA-027).

Threat class	Reference	What it looks like in deployment	Primary responding controls
Prompt injection (direct and indirect)	OWASP LLM01	Instructions smuggled through user input, documents, web pages or retrieved content override intended behaviour.	AI-01, AI-02, RG-02, RG-05, LG-04
Sensitive information disclosure	OWASP LLM02	Model reveals personal data, secrets or confidential corpus content in outputs.	AI-05, DP-04, TS-06, LG-01
Supply chain compromise	OWASP LLM03	Compromised models, datasets, plugins or connectors enter the stack.	SP-01, SP-02, SP-07, AA-02
Data and model poisoning	OWASP LLM04	Manipulated training, fine-tuning or retrieval data biases or backdoors behaviour.	RG-01, RG-02, RG-05, SP-07
Improper output handling	OWASP LLM05	Generated content executed or rendered downstream without validation.	AI-04, CM-05
Excessive agency	OWASP LLM06	Agents hold more tools, permissions or autonomy than the use case needs.	AA-01 to AA-05, GV-05
System prompt leakage	OWASP LLM07	Confidential instructions extracted, exposing logic and embedded data.	AI-03, TS-03
Vector and embedding weaknesses	OWASP LLM08	Vector store abused: cross-tenant retrieval, embedding inversion, poisoned entries.	RG-03, RG-04, RG-05
Misinformation and overreliance	OWASP LLM09	Confident wrong output drives decisions unchecked.	AI-08, HO-01 to HO-06, SQ-03
Unbounded consumption	OWASP LLM10	Runaway usage: denial of wallet, resource exhaustion.	AI-06, AA-04, LG-04
Agent goal and tool misuse	OWASP Agentic	Agents pursue goals through unintended tool use, credential abuse or unsafe chains.	AA-02 to AA-06, LG-04
Agent memory manipulation	OWASP Agentic	Persistent memory seeded with malicious or stale instructions.	AA-07, RG-02
Multi-agent cascade	OWASP Agentic	Delegation loops, authority creep or cascading failure across agents.	AA-08, RS-02
Adversarial ML techniques	MITRE ATLAS	Reconnaissance, evasion, extraction and inference attacks against deployed models.	TE-03, LG-04, SP-02
Shadow deployment	Internal	Unsanctioned tools bypass the entire control system.	LG-05, GV-03, Section 13.1

Tier 3 and Tier 4 use cases MUST evidence threat-model coverage: each applicable threat class above maps to at least one tested control or a recorded, accepted gap (enforced through GA-TE-01 and GA-TE-03).

10. Assurance model

10.1 The distinction v1.1 collapsed

A configured dashboard proves someone configured a dashboard. v1.1 repeatedly treated artefact presence as effectiveness, which the assessment flagged as its deepest assurance defect (GA-052). This framework separates the questions an assessor must answer in order:

Assurance level	Question answered	Typical method	What it cannot tell you
AL1: Design	Would this control, as specified, address the risk if it operated?	Specification review, walkthrough, threat-model tracing.	Whether it exists in the environment.
AL2: Implementation	Does the control exist and function at a point in time?	Configuration inspection, single-transaction test, re-performance.	Whether it kept working.
AL3: Operating effectiveness	Did the control operate as designed over a defined period?	Testing a defined sample of occurrences across the period, plus exception analysis.	Nothing further within its period and scope.

10.2 Conclusion rules

1. Every assurance conclusion states the control ID, the level reached (AL1 to AL3), the period covered, the evidence relied on and the exceptions found.
2. A conclusion at AL1 or AL2 **MUST NOT** be reported as effectiveness. Reporting language is controlled: “designed adequately”, “implemented at [date]”, “operated effectively over [period]”.
3. Exceptions are evaluated for severity and pervasiveness before a conclusion is drawn; one high-severity exception can fail a control that passes on aggregate counts.
4. Tier determines the minimum level: Tier 1 AL1, Tier 2 AL2, Tier 3 AL2 with AL3 on security-critical controls, Tier 4 AL3 on applicable **MUST** controls, subject to the legal-determination assurance ceiling below.

Legal-determination assurance ceiling. For controls requiring legal interpretation or a legal conclusion, AL1 to AL3 describes the design, implementation and operation of the decision process and its evidence. It does not attest to the correctness of the legal determination or to legal compliance. The organisation’s authorised legal or privacy decision-maker owns that determination; assessors report missing or unsupported determinations and refer them for specialist review.

5. Sampling design (population definition, method, size rationale) is internal methodology under GA-EV-05 and is documented, not borrowed from a standard that never specified it.

11. TEVV methodology

11.1 Test specifications, not test labels

v1.1's TEVV-Lite playbooks named tests without defining them, so two assessors could run “the same test” and reach opposite conclusions (GA-037). Every test in this framework is a specification with the following mandatory fields, registered under GA-TE-01:

Field	Requirement
Method	The procedure, stated so a second tester can repeat it.
Environment	Production, pre-production or sandbox, with configuration parity noted.
Dataset / corpus	The inputs used, including adversarial corpora, with version and provenance.
Sample size and basis	How many cases and why that number supports the conclusion sought.
Expected result and tolerance	The pass condition, the tolerance around it, and any zero-tolerance events.
Reviewer	Who evaluates results, and independence requirements by tier.
Evidence output	The records produced, keyed for the evidence pack.

11.2 Acceptance criteria rules

The generic thresholds in v1.1 (a 2 per cent harm rate, a 5 to 10 per cent parity gap, 95 per cent escalation) are withdrawn (GA-017, GA-038). Acceptance criteria are now governed by rule rather than by number:

- Criteria are set per use case, derived from the harm model (GA-SQ-01) and the legal position, and approved by the authority the tier demands.
- Where sampling supports a criterion, results carry confidence intervals; a point estimate without one is not a pass.
- Subgroup analysis applies wherever outcomes affect people, so aggregate rates cannot conceal concentrated failure.
- Zero-tolerance events are defined in advance per use case (for example: disclosure of special category data, execution of a denied irreversible action). One occurrence fails the test regardless of the aggregate.
- Where a legal precondition applies (DPIA, FRIA, registration), the criterion is completion, at 100 per cent, as a blocking gate. Percentage targets below 100 exist only for genuinely discretionary practices.

12. Legal and regulatory implementation

Status of this section. These legal summaries are stated as at July 2026 and support workflow and evidence design. They are not legal advice. Organisations own legal determinations using current authoritative sources and their authorised advisers. Assurance covers process and evidence within Section 10.2 and does not certify statutory compliance. GA-GV-07 maintains the legal position as legislation and guidance change.

12.1 EU AI Act deployer obligations (high-risk systems)

For each high-risk system deployed, GA-LR-04 requires an obligation map completed at the level below. This replaces the v1.1 generic checklist that blurred or omitted obligations (GA-039).

Article 26 obligation area	What the deployer evidences
Use per instructions	Technical and organisational measures ensuring use in accordance with the provider's instructions for use; instruction documents retained and mapped to configuration.
Human oversight	Oversight assigned to persons with the necessary competence, training, authority and support (controls HO-01 to HO-04).
Input data control	Where the deployer exercises control over input data, evidence that inputs are relevant and sufficiently representative for the intended purpose.
Monitoring and suspension	Operation monitored per instructions; use suspended and the provider or distributor and relevant authorities informed where the system presents a risk within the meaning of the Act; incidents fed into the Section 12.3 matrix.
Log retention	Logs automatically generated by the system, to the extent under the deployer's control, retained at least six months, subject to applicable law (GA-LG-06).
Worker notification	Where the system is used in the workplace, workers and their representatives informed before putting into service.
Public-authority duties	Registration and additional duties identified and completed where the organisation deploys as a public authority or in listed contexts (GA-LR-07).
Information use for DPIA	Provider-supplied information used to meet DPIA obligations where applicable (GA-DP-02).
Cooperation	Cooperation with competent authorities evidenced through defined contact routes and record availability.

12.2 FRIA and the UK overlay

FRIA. Article 27 applies its fundamental rights impact assessment duty to specified deployers of specified high-risk systems. That statutory trigger is narrower than v1.1 claimed, and this framework applies it as written: where the trigger is met, the FRIA and its notification duties are a blocking gate (GA-LR-05). Where the trigger is not met but the harm model shows material rights impact, a voluntary HRIA SHOULD be performed and is labelled voluntary throughout (GA-016).

UK position. The Data (Use and Access) Act 2025 restructured the UK GDPR treatment of automated decision-making into Articles 22A to 22D. In broad terms, significant decisions based solely on automated processing are permitted subject to prescribed safeguards (information to the individual, human intervention on request, the ability to make representations and to contest the decision), with stricter conditions where special category data is involved. ICO guidance in this area was under consultation, last updated 31 March 2026, at the time of writing. GA-DP-03 implements the safeguards; GA-GV-07 tracks commencement and guidance movement (GA-041).

12.3 Incident notification decision matrix

v1.1 converted Article 73 into a universal 15-day deployer deadline. The actual position is route- and circumstance-specific, and the matrix reflects the outer limits as enacted (GA-040). For every qualifying incident, GA-IR-02 requires each row to be answered with a documented decision.

Route	Trigger	Clock and outer limit
Provider notification	Deployer identifies a serious incident involving a high-risk system it deployed.	Inform the provider without delay; provider then carries Article 73 reporting. Where the provider cannot be reached, reporting obligations can fall to be handled directly; take legal advice immediately.
EU AI Act serious incident (general)	Serious incident as defined by the Act.	Report immediately after establishing the causal link or its reasonable likelihood, and no later than 15 days after awareness.
EU AI Act, death of a person	Serious incident involving the death of a person.	No later than 10 days after awareness of the suspected causal relationship.
EU AI Act, widespread infringement or critical-infrastructure disruption	Widespread infringement, or serious incident involving serious and irreversible disruption to critical infrastructure.	Immediately, and no later than 2 days after awareness.
GDPR / UK GDPR personal data breach	Breach of personal data likely to result in risk to individuals.	Supervisory authority without undue delay and where feasible within 72 hours; affected individuals without undue delay where high risk.
Sectoral and contractual	Sector regulator rules, or contract notification clauses (GA-SP-03).	Per instrument; recorded per incident.

13. Operational playbooks and implementation

13.1 Shadow AI operations playbook

Retained from v1.1 as its strongest section, with the proportionality the assessment required (GA-023). The workflow remains discovery, triage, containment, sanction. What changes is the triage step: automatic blocking within a fixed hour is withdrawn.

1. **Discover.** Continuous signals under GA-LG-05 feed a shadow AI queue.
2. **Triage.** Each discovery is classified: hostile or exfiltrating (contain now), business-critical dependency (contain with continuity plan), policy breach without active harm (engage and migrate), duplicate of a sanctioned tool (redirect). Business-continuity impact is checked before any block, and investigation and evidence-preservation safeguards apply before disruptive action where feasible.
3. **Contain.** Containment is proportionate to the triage class. Not every discovery is an incident; incident status attaches only where the taxonomy in GA-IR-01 says it does, which keeps reporting honest and staff treatment consistent.
4. **Sanction or retire.** Surviving tools enter the inventory through the full risk method; the rest are retired with data recovered and access closed.

13.2 Evidence architecture

The evidence pack per use case carries a YAML manifest, integrity hashes and correlation keys (use_case_id, control_id, evidence_id), unchanged in concept from v1.1. Around it now sit the privacy controls the assessment found missing: store-level classification and access restriction (GA-EV-02), retention and secure deletion that WORM protection cannot override where deletion is legally owed (GA-EV-03), and the metadata-default logging posture of domain LG, which stops the evidence system becoming the organisation's largest unmanaged personal-data store (GA-012, GA-043).

13.3 Measurement and reporting

Indicators are separated so that activity cannot impersonate effectiveness (GA-056, GA-060):

- **Implementation indicators:** coverage of inventory, controls deployed, training completed. Useful for programme tracking, silent on risk.
- **Performance indicators:** detection volumes, response times, exception counts and ages.
- **Effectiveness indicators:** TEVV pass rates against approved criteria at stated confidence, oversight catch rates from seeded tests, assurance conclusions at level AL3.
- **Risk indicators:** residual-risk positions, non-exceptionable near-misses, threat-model coverage gaps.

Success declarations to the governance body draw on the last two categories. A deployed dashboard and a completed DPIA are entry tickets, not achievements.

13.4 Implementation profiles

The fixed 30-60-90 day methodology and the FTE table attributed to SABSA are withdrawn (GA-058, GA-059). Implementation planning now uses illustrative profiles with entry and exit criteria: a profile is complete when its exit criteria pass, not when a calendar date arrives. Duration and staffing depend on inventory size, tier mix, jurisdiction and existing capability; any published figures are labelled illustrative with their assumptions stated. Tier 4 use cases take the time their evidence takes.

Profile	Entry criteria	Exit criteria (illustrative)
Foundation	Owner mandate, accountable executive named.	Policy approved (GV-01), inventory live (GV-03), risk method operating Stages 1 to 4, shadow AI discovery running (LG-05).

Profile	Entry criteria	Exit criteria (illustrative)
Control build	Foundation exit met.	Applicable Tier 1 and 2 controls at AL2; DPIA and lawful-basis method operating (DP-01, DP-02); logging posture converted to metadata-default (LG-01).
Assurance	Control build exit met.	TEVV specifications registered for T2+ cases (TE-01), pre-deployment gate operating (TE-02), first internal assurance cycle complete (EV-06).
Advanced	Assurance exit met; agentic or high-risk use cases proposed.	Domain AA controls at AL2+ for agent cases, threat-model coverage evidenced for T3+ (Section 9), Tier 4 controls at AL3.

14. Standards and regulatory crosswalk

The crosswalk records where this framework supports alignment with external instruments. It establishes neither certification nor legal compliance, and no listed instrument endorses this framework (GA-006). Voluntary guidance is marked V; binding law where applicable is marked L.

Instrument	Type	Nature of support	Known limitation
Regulation (EU) 2024/1689 (EU AI Act)	L	Actor classification (S.5), prohibited-use screening (LR-01), deployer obligation map (S.12.1), FRIA gate (LR-05), notification matrix (S.12.3).	Application timetable phases in; provisions and amendments tracked under GV-07. Provider obligations out of scope.
GDPR / UK GDPR with DUAA 2025	L	Per-case lawful basis (DP-01), DPIA gate (DP-02), rights procedures (DP-05, DP-06), UK ADM safeguards (DP-03), transfer register (DP-07).	ICO guidance under consultation as at 31 March 2026 update; UK commencement tracked.
ISO/IEC 42001:2023	V	Management-system elements: policy, roles, operation, evaluation, improvement (GV, EV domains).	Not a complete AIMS; no clause-level conformity claim. Internal sampling and CAPA figures are not ISO requirements.
ISO/IEC 23894:2023	V	Risk-method staging, override discipline, residual-risk acceptance (S.7).	Anchored descriptors are the organisation's, not the standard's.
ISO/IEC 5338:2023	V	Lifecycle gates, change and decommissioning concepts (CM, RS, SP-06).	Lifecycle process outputs not comprehensively mapped.
NIST AI RMF 1.0 and GenAI Profile (AI 600-1)	V	Govern, Map, Measure, Manage support across S.6, 7, 11, 13.3.	No subcategory-level traceability table in this version.
NIST CSF 2.0	V	All six functions including Govern, mapped at function level to GV, TS, LG, IR, RS.	Category and subcategory crosswalk deferred to the control register.
NIST SP 800-218A (July 2024)	V	Acquisition and supplier expectations (SP domain).	Secure development practices apply to providers; deployer view only.
OWASP LLM Top 10 2025; OWASP Agentic AI	V	Threat-to-control mapping (S.9); domains AI, RG, AA.	Threat landscape moves faster than framework versions; GV-07 cadence applies.
MITRE ATLAS	V	Adversarial technique awareness feeding TE-03 scoping and LG-04 detections.	No technique-level coverage claim.
CSA AI Controls Matrix v1 (July 2025, updated October 2025)	V	Domain-level correspondence recorded in the control register.	The v1.1 reference to unpublished CSA benchmarks and an industry-median target is withdrawn (GA-046).
DSIT AI Cyber Security Code of Practice (31 January 2025); NCSC secure AI guidance	V	IAM, inventory, monitoring, incident and supplier controls (TS, LG, IR, SP).	Secure-by-design and end-of-life mapping partial.
PCI DSS v4.0.1	L (in scope)	TS-06 applies PAN detection where cardholder data scope exists.	Applies only within PCI scope; no framework-wide claim (GA-007, GA-008).
WCAG 2.2 (December 2024)	V/L by jurisdiction	SQ-04 accessibility baseline for in-scope interfaces.	Jurisdictional applicability assessed per deployment (GA-057).
ISO/IEC 27001:2022, 27002:2022, 27559:2022	V	Security management and logging support (TS, LG, EV); 27559 informs de-identification risk method in DP-08.	27559 is not a national-identifier catalogue (GA-009); no clause-level conformity claim.

14.1 Source register

Every external claim in this framework traces to the controlled source register maintained with the control register: instrument name, version, date, status (in force, consultation, under review), access date and the sections relying on it. Two register rules respond directly to assessment findings: NIST SP 800-52 Revision 2 is recorded as under review as at May 2026 and never cited as a universal mandate (GA-011), and no source may be cited for a requirement it does not establish, verified at each Publication Gate (GA-051). The deepfake benchmark citation from v1.1 had no verifiable source and is removed rather than replaced (GA-036).

15. Continuous improvement

The improvement cycle keeps its Plan, Do, Check, Act shape from v1.1, now fed by defined inputs and honest about what ISO/IEC 42001 does and does not require (GA-044, GA-045): findings from the internal assurance programme (EV-06), post-incident reviews (IR-06), TEVV drift triggers (TE-05), regulatory watch outputs (GV-07), exception-register trends (EX-05) and provider change notifications (SP-05). Corrective actions carry owners and target dates set by the governance body as internal policy. Management review runs on a defined cadence and covers effectiveness and risk indicators, not implementation counts.

16. Publication Gate, licence and attribution

16.1 Publication Gate

Publication approval for Version 2.1 is recorded on 30 September 2026 under the owner's instruction to prepare this release. The publication conditions and their disposition are recorded below (GA-001, GA-049).

1. Assurance boundary approved by the owner. The framework is an operational and technical assurance scaffold. Framework-level qualified-counsel sign-off is removed as a release condition. Organisations using it retain responsibility for statutory interpretation, actor status, applicability, lawful basis and notification decisions. Use-case legal gates and specialist referrals remain mandatory where applicable; publication does not certify legal compliance.
2. Citation verification completed, as confirmed by the owner: external references checked against the source register, including version and status checks.
3. Independent technical review completed, as confirmed by the owner: the control catalogue, threat model and TEVV methodology reviewed by a reviewer independent of the author.
4. Branding and document QA completed, as confirmed by the owner. Publication edits set the author credit, copyright year, status and licensing terminology consistently.
5. All three owner decisions closed in Section 1.3: contributor attribution removed, release year fixed at 2026 and licence wording adopted.
6. Final Word release QA: document rendered and inspected for layout; headings, document metadata and publication wording checked. Any separately exported PDF requires its own layout and accessibility checks.

16.2 Attribution

Author: Faisal Ali, AAISM, CISM, CRISC. Owner: ELSA AI LTD. The author credential line appears only in credential-bearing locations: the cover, this section and document control. Copyright notice throughout: © 2026 ELSA AI LTD.

16.3 Licence

This framework is licensed under the Creative Commons Attribution-NoDerivatives 4.0 International (CC BY-ND 4.0) License: <https://creativecommons.org/licenses/by-nd/4.0/>. The unmodified legal code governs: <https://creativecommons.org/licenses/by-nd/4.0/legalcode.en>.

You may copy and share the licensed content, in whole or part, including commercially, with appropriate attribution and the required notices. Adapted material may be produced privately but may not be shared under this licence. Necessary format changes are permitted. Statutory exceptions and limitations remain unaffected; no additional restrictions are imposed on licensed rights.

The ELSA AI name, logo and related branding are trademarks of ELSA AI LTD. Trademark rights are excluded from this licence. Use suggesting endorsement, partnership or sponsorship requires prior written permission. Retaining branding in a faithful copy of this document is permitted; it does not authorise separate promotional use of the logo.

Attribution example: “GenAI Assure Framework v2.1 (2026), by Faisal Ali, AAISM, CISM, CRISC. © 2026 ELSA AI LTD. Licensed under CC BY-ND 4.0: <https://creativecommons.org/licenses/by-nd/4.0/>.” Retain supplied copyright, licence and disclaimer notices and link to the source where reasonably practicable. A factual reference such as “Uses the framework developed by ELSA AI” does not imply endorsement. Do not present an adapted framework as an ELSA AI publication or place the ELSA AI logo on a separate product or promotional material without permission.

ELSA AI LTD may issue revised versions of its own work. Each release identifies its version and terms; later releases do not revoke compliant users’ rights in earlier licensed releases. The framework is supplied without warranties to the extent permitted by applicable law; Section 5 of the licence sets out the disclaimer and liability terms. Licensing and trademark enquiries: contact@elsaai.co.uk.

Annex A: Assessment finding traceability

Disposition of all 66 Phase 1 assessment findings in this version. Resolved: the defect is corrected in the cited location. Withdrawn: the defective claim or mechanism is removed without replacement. All three owner decisions are closed; their outcomes are recorded below.

Finding	Disposition	Where and how v2.1 responds
GA-001	Resolved	Publication edition and owner approval recorded; Publication Gate dispositions closed (S.1, S.16.1).
GA-002	Resolved	Normative language and per-control applicability replace “implement all controls” (S.3.1, S.8.1).
GA-003	Resolved	Full actor model with provider-conversion routes (S.5).
GA-004	Resolved	Architecture taxonomy AC-01 to AC-13 in scope (S.4.2).
GA-005	Resolved	Modification decision tree; AC-04 always triggers assessment (S.5.3).
GA-006	Resolved	Cautious crosswalk with limitations; no compliance claims (S.14).
GA-007	Resolved	PCI DSS v4.0.1, applied only within cardholder scope (GA-TS-06).
GA-008	Resolved	Luhn detection reframed as internal guidance, PCI citation removed (GA-TS-06).
GA-009	Resolved	Jurisdiction-specific patterns; ISO/IEC 27559 used for de-identification risk only (GA-TS-06, GA-DP-08).
GA-010	Resolved	Jurisdictional overlays replace universal HIPAA mandate (GA-LR-08, GA-TS-06).
GA-011	Resolved	Risk-based internal cryptographic baseline; NIST attribution removed; SP 800-52r2 status noted (GA-TS-04, S.14.1).
GA-012	Resolved	Metadata-default logging; content by justified exception (GA-LG-01, GA-LG-02).
GA-013	Resolved	Threat-based technique selection; hashing and FPE errors corrected (GA-DP-08).
GA-014	Resolved	Event taxonomy separates events, alerts, incidents, regulatory records (GA-LG-03).
GA-015	Resolved	DPIA/FRIA as 100 per cent blocking gates (GA-DP-02, S.11.2).
GA-016	Resolved	Article 27 statutory trigger applied as written; voluntary HRIA separated (GA-LR-05, S.12.2).
GA-017	Resolved	Generic parity threshold withdrawn; context-specific fairness criteria (GA-SQ-02).
GA-018	Resolved	Oversight required by risk, safety design and instructions, not statute alone (GA-HO-01).
GA-019	Resolved	Role-based competency with assessment and retraining triggers (GA-HO-05).
GA-020	Resolved	Sequential method with mandatory overrides replaces additive scoring (S.7).
GA-021	Resolved	Nine inherent-risk dimensions including autonomy, reversibility, external action (S.7.4).
GA-022	Resolved	Budget and timeline removed from classification; profiles follow tiering (S.7.5, S.13.4).
GA-023	Resolved	Proportionate triage replaces one-hour auto-block (S.13.1).
GA-024	Resolved	115 atomic controls across 17 domains (S.8).
GA-025	Resolved	Mandatory control schema, all fields (S.8.1).
GA-026	Resolved	Agentic AI domain: autonomy tiers, gates, limits, containment, interruption (S.8, domain AA).
GA-027	Resolved	Formal threat model with threat-to-control mapping (S.9).
GA-028	Resolved	Injection, jailbreak, prompt leakage, output handling controls (GA-AI-01 to 04, 07).
GA-029	Resolved	RAG domain: ingestion, provenance, isolation, retrieval authorisation, poisoning tests (domain RG).
GA-030	Resolved	Foundation-model, service, dataset, hosting, change and exit assurance (domain SP).
GA-031	Resolved	Change classification, pinning, regression, reapproval (domain CM).
GA-032	Resolved	Evidence timeframes as internal objectives; SOC 2 / ISO attribution removed (GA-EV-05).
GA-033	Resolved	Per-use-case lawful basis method; class-level defaults prohibited (GA-DP-01).
GA-034	Resolved	Article 17 technical-impossibility claim removed; per-request legal analysis (GA-DP-06).
GA-035	Resolved	Access procedure: discovery, identity, redaction, rights conflicts (GA-DP-05).
GA-036	Withdrawn	Unverifiable deepfake benchmark removed; use-case-specific measures instead (GA-SQ-07).
GA-037	Resolved	Full test specifications replace test labels (S.11.1, GA-TE-01).
GA-038	Resolved	Rule-based acceptance criteria; generic thresholds withdrawn (S.11.2).
GA-039	Resolved	Article 26 obligation map at obligation level (S.12.1, GA-LR-04).
GA-040	Resolved	Route-specific notification matrix with 2, 10 and 15-day outer limits (S.12.3).

Finding	Disposition	Where and how v2.1 responds
GA-041	Resolved	UK overlay reflects DUAA 2025, Articles 22A to 22D, ICO consultation status (S.12.2, GA-DP-03).
GA-042	Resolved	CSF 2.0 mapping includes Govern (S.14).
GA-043	Resolved	Evidence-store privacy controls: classification, access, retention, deletion (GA-EV-02, 03, GA-DP-09).
GA-044	Resolved	ISO/IEC 42001 represented as management-system elements, limits stated (S.14, S.15).
GA-045	Resolved	Sampling and CAPA figures reframed as internal policy (GA-EV-05, S.15).
GA-046	Resolved	CSA AICM v1 referenced; unpublished benchmark and median target withdrawn (S.14).
GA-047	Resolved	Mandated slogan applied throughout; former tagline removed (cover, footers).
GA-048	Resolved	Owner fixed release year at 2026; © 2026 ELSA AI LTD applied throughout (S.1.3).
GA-049	Resolved	Author credit “Faisal Ali, AAISM, CISM, CRISC” in credential-bearing locations; Publication Gate inserted (S.1, S.16).
GA-050	Resolved	Definitions, abbreviations, normative conventions (S.3).
GA-051	Resolved	Controlled source register with citation-verification gate (S.14.1).
GA-052	Resolved	Three-level assurance model with conclusion rules (S.10).
GA-053	Resolved	Incident taxonomy and notification decision matrix (GA-IR-01, S.12.3).
GA-054	Resolved	Safety, quality and human-impact controls incl. contestability (domain SQ).
GA-055	Resolved	Non-exceptionable list and tiered exception authority (domain EX).
GA-056	Resolved	Indicator categories separate activity from effectiveness (S.13.3).
GA-057	Resolved	WCAG 2.2 AA, jurisdiction-dependent (GA-SQ-04).
GA-058	Resolved	Illustrative profiles with entry and exit criteria replace fixed 30-60-90 (S.13.4).
GA-059	Resolved	SABSA-attributed FTE table withdrawn; figures labelled illustrative with assumptions (S.13.4).
GA-060	Resolved	Success measured on effectiveness and risk indicators (S.13.3).
GA-061	Resolved	Owner removed review-team contributor attribution; author and owner credits retained (S.1.3).
GA-062	Resolved	Owner adopted CC BY-ND 4.0 wording with separate trademark, attribution and version provisions; licence legal-review release condition withdrawn (S.1.3, S.16.3).
GA-063	Resolved	Revision markers removed; changes live in the change log only (S.6.3).
GA-064	Resolved	Single-statement rule: requirements live once in the catalogue (S.6.3).
GA-065	Resolved	One authoritative rights section, cross-referenced (GA-DP-06).
GA-066	Resolved	“Licensing” spelling standardised (S.16.3).

End of document. GenAI Assure Framework v2.1. Publication edition © 2026 ELSA AI LTD.
Ethical • Legal • Safe • Accountable AI Governance.