

GenAI Assure™ Framework

Version 1.1

Practical, Achievable, Auditable AI Governance

Document Control

Field	Value
Document Title	GenAI Assure™ Framework
Document Type	Industry Framework Specification
Version	1.1
Status	Production Ready – Enterprise Implementation
Date	31 January 2026
Author	Faisal Ali (CISM, CRISC) Founder, ELSA AI
Contributor	Security Architecture Review Team
Owner	ELSA AI LTD
Scope	AI deployers (organisations using third-party AI tools; no model development)
Applies To	SMB to Enterprise; tiered controls for fast adoption
Classification	Open Framework (CC BY-ND 4.0)

Document Control	2
Version History.....	7
1. Executive Preface	8
1.1 Purpose and Scope	8
In Scope	8
Out of Scope.....	8
1.2 Intended Audience.....	9
1.3 Standards Alignment	9
2. Core Governance Principles	10
2.1 Security & Data Protection.....	10
Threat Examples.....	10
Key Controls	10
A) Identity & Access Management.....	10
B) Secrets & Tokens.....	10
C) Data Controls.....	10
D) Network / Egress Controls	11
E) Logging & Telemetry	11
F) Incident Response & SOAR	11
Implementation Evidence (Examples).....	12
KPIs	12
2.2 Governance & Regulatory Assurance.....	13
Key Controls	13
Implementation Evidence (Examples).....	13
KPIs	13
2.3 Ethical & Human Impact	14
Key Controls	14
Evidence & KPIs	14
2.4 Accountable Operations	14
Key Controls	14
Evidence & KPIs	14
2.5 Trust & Safety Culture	15
Key Controls	15
Implementation Evidence.....	15
KPIs	15
3. Governance Domains & Controls	16
3.1 Pre-Implementation Readiness Assessment	16

3.1.1 Value & Risk Assessment Framework.....	16
3.1.2 Objective Tiering Rubric.....	17
3.2 Implementation Methodology (30-60-90 Day Plan).....	18
Days 1-30: Foundation.....	18
Days 31-60: Core Controls.....	18
Days 61-90: Optimisation & Scale.....	18
3.3 Shadow AI Operations Playbook	19
Playbook Steps.....	19
4. Control Catalogue with Implementation Profiles	20
4.1 GA-PG-001 Policy & Governance.....	20
Minimums	20
Configuration	20
4.2 GA-TP-001 Technical Protection	20
Identity & Access	20
Secrets & Tokens.....	20
DLP (AI-aware)	20
Network/Egress.....	20
Cryptography	21
Configuration Baselines	21
4.3 GA-DM-001 Detect & Monitor	21
Event Schema (Log Fields).....	21
Core Detections	21
Dashboards	21
4.4 GA-DC-001 Documentation & Compliance	22
Evidence Pack (Tiered SLA)	22
Contents	22
Lawful Basis by Common Use	22
Data Subject Rights in AI Context	22
4.5 GA-RR-001 Response & Remediation.....	23
Runbooks.....	23
SOAR Actions.....	23
Redress	23
4.6 GA-RB-001 Resilience & Business Continuity	23
Requirements	23
5. TEVV-Lite Testing Playbooks.....	24
5.1 Avatar/Video Generation	24

5.2 Workflow Automations	24
5.3 Developer Copilots	24
5.4 Customer Service Chatbots	24
5.5 Document Analysis/Generation	24
5.6 HR Screening Tools	25
5.7 Financial Analysis Tools	25
5.8 Marketing Content Creation	25
6. Standards Alignment and Regulatory Implementation	26
6.1 EU AI Act - Deployer Duties (Article 26)	26
Compliance Checklist.....	26
6.2 FRIA (Fundamental Rights Impact Assessment)	26
Assessment Components	26
6.3 GDPR/UK GDPR Practical Implementation	27
Key Requirements	27
6.4 NIST CSF 2.0 Mapping.....	27
7. Evidence Automation and Auditability	28
7.1 Architecture Pattern	28
7.2 Correlation Keys	28
7.3 Evidence YAML Schema Example	28
7.4 Validation Requirements.....	29
8. Continuous Improvement Management System	30
8.1 AIMS (ISO/IEC 42001) Cycle.....	30
8.2 Feedback Loop Integration	30
9. Resource Planning and Implementation Support	31
9.1 Resource Requirements by Tier	31
9.2 Change Management Approach	31
Stakeholder Engagement.....	31
Resistance Management	31
Adoption Tracking	31
10. Next Steps	32
10.1 Framework Adoption Recommendations	32
10.2 Success Metrics and Milestones.....	32
Day 30 Success Indicators.....	32
Day 60 Success Indicators.....	32
Day 90 Success Indicators.....	32
10.3 Long-term Maturity Evolution	33

10.4 Community and Support	33
11. Licence and Attribution.....	34

Version History

Version	Date	Author	Summary of Changes
1.0	27 August 2025	Faisal Ali	Initial framework release
1.1	31 January 2026	Faisal Ali, Security Architecture Review Team	Critical updates based on security architecture review: • Clarified accountability boundaries (organisations MUST implement controls) • TLS 1.3 mandatory, TLS 1.2 deprecated with 12-month migration (NIST SP 800-52 Rev. 2) • Added training data exposure assessment (NIST AI RMF MEASURE 2.7, 2.11) • Absolute evidence pack SLA with failure escalation (SOC 2 CC7.2) • Added EU AI Act Article 73 serious incident reporting (15-day) • Enhanced DLP, logging, TEVV-Lite, audit scope, and resource specifications • Converted to British English throughout

1. Executive Preface

1.1 Purpose and Scope

Purpose. GenAI Assure™ is a concise, security-led governance framework for organisations that deploy third-party AI tools and services. It links core principles → actionable controls → measurable evidence, enabling auditable adoption and a pragmatic 30–60–90-day implementation path for SMBs through to enterprises.

This Framework provides guidance only. Organisations **MUST** implement and operate all controls. ELSA AI delivers governance oversight, reviews, and approvals through documented change management processes. ELSA AI does not directly deploy, configure, or administer systems.

Scope. The framework focuses on the operational realities of AI deployer organisations governing how external AI tools are selected, configured, monitored, and evidenced within business workflows. It is technology-agnostic, risk-based, and aligned to leading standards (EU AI Act deployer duties, NIST AI RMF, ISO/IEC 42001, GDPR/UK GDPR, SOC 2).

In Scope

- Use of third-party AI tools/services (e.g., workflow automations/orchestrators, developer copilots, content/voice/video generators, chatbots/assistants, document intelligence).
- Data protection and privacy controls, transparency and labelling of AI use, and vendor risk management.
- Logging/monitoring, evidence management (including WORM/attestations), and Trust & Safety culture development.

Out of Scope

- Building, training, or fine-tuning AI models (provider or internal model development).
- MLOps and provider-side safety techniques (e.g., model evaluations and safety tuning performed by vendors).
- General cybersecurity baselines beyond AI-specific controls (addressed indirectly via SOC 2/ISO 27001 alignment).

1.2 Intended Audience

This framework is designed for AI deployer organisations ranging from SMBs to enterprises, with specific focus on:

- Business unit leaders implementing AI solutions
- Chief AI Officers (CAIOs) and Product Leaders
- Chief Information Security Officers (CISOs) and Security Engineering
- Compliance and Legal teams
- Risk and Governance functions
- Where applicable: Data Protection Officer (DPO/Privacy), Procurement/Vendor Management, IT Operations, and Internal Audit.

1.3 Standards Alignment

The framework supports compliance with and maps to:

- GDPR/UK GDPR (data protection requirements)
- EU AI Act (deployer duties under Article 26)
- NIST AI RMF (Govern/Map/Measure/Manage functions)
- ISO/IEC 42001 (AI Management System)
- NIST CSF 2.0 (cybersecurity framework)
- SOC 2 Trust Services Criteria
- NIST SP 800-53 Rev. 5 (Security and Privacy Controls)
- PCI DSS 3.2.1 (Payment Card Industry Data Security Standard)
- ISO/IEC 27559 (Privacy enhancing data de-identification framework)

2. Core Governance Principles

Each principle follows the structure: Objective → Key Controls → Implementation Evidence → KPIs → Owner.

2.1 Security & Data Protection

Objective (Protect). Reduce data loss, token abuse, and unauthorised egress across AI tools and automations.

Threat Examples

- PII leakage through AI prompts/outputs
- Compromised connectors/API keys
- Unsanctioned webhooks and data exfiltration
- Automation loops and bulk data transfers

Key Controls

A) Identity & Access Management

- All AI tools behind IdP SSO + MFA
- Least-privilege role assignments
- Privileged Access Management (PAM) for elevated roles/sessions
- SCIM provisioning and de-provisioning where available

B) Secrets & Tokens

- Secrets management in an enterprise vault (no secrets in code/repos)
- Enforced token rotation SLA (≤ 90 days); revoke on role change
- Alerts for unused tokens (> 30 days) and stale integrations

C) Data Controls

- Data classification for AI inputs and outputs
- DLP controls on endpoints, email, and web/CASB
- Prompt/response redaction for sensitive fields (PII, credentials, regulated data)
- Pattern matching for PII/credentials/regulatory keywords with block or quarantine actions

UPDATED REQUIREMENT: Organisations MUST implement DLP rules detecting: PCI DSS compliant PANs (13-19 digits passing Luhn algorithm check per PCI DSS 3.2.1 Req 3.4), national identifiers per ISO/IEC 27559, HIPAA identifiers per 45 CFR 164.514(b). Reference: PCI DSS 3.2.1 Req 3.4, ISO/IEC 27559, HIPAA 45 CFR 164.514(b).

D) Network / Egress Controls

- Allow-list for outbound AI API FQDNs and sanctioned webhook destinations
- Block unknown webhook domains and unsanctioned API endpoints
- TLS enforcement with strong cipher suites

UPDATED REQUIREMENT: Organisations MUST enforce TLS 1.3 (RFC 8446). TLS 1.2 MAY be supported only for legacy systems with documented risk acceptance expiring within 12 months. Reference: NIST SP 800-52 Rev. 2.

E) Logging & Telemetry

- Prompts, outputs, and key actions logged to SIEM (with UEBA where available)
- Integrity protection via append-only / WORM storage
- Hash/tokenise high-value data elements in logs
- Standard correlation keys (e.g., use_case_id, control_id, vendor_id, token_id, connector_id, user_id, timestamp)

UPDATED REQUIREMENT: Organisations MUST hash sensitive log elements using SHA-256 or stronger (NIST FIPS 180-4). Tokenisation MUST use format-preserving encryption per NIST SP 800-38G where log structure must be preserved. Reference: NIST 800-53 r5 AU-9, AU-11.

UPDATED REQUIREMENT: Add mandatory log fields: incident_category (per NIST CSF 2.0 DE.AE-2), regulatory_trigger (boolean), breach_notification_required (boolean), affected_data_subjects_count (integer). Reference: GDPR Art. 33, NIST CSF 2.0 DE.AE-2.

F) Incident Response & SOAR

- AI-specific runbooks (PII exfil, token compromise, webhook misuse, bulk sends)
- Automated actions: token revocation, connector disable, case creation, stakeholder notifications
- Post-incident CAPA with time-bound remediation

Implementation Evidence (Examples)

- SSO/MFA coverage report for sanctioned AI tools; SCIM provisioning logs
- Vault policy and token inventory with last-used/last-rotated timestamps
- DLP policy set (patterns/rules), recent violation reports, and tuning records
- CASB/Proxy allow-list exports; change-control tickets for egress rules
- SIEM event schema and sample log extracts (prompts/outputs/actions) stored in WORM
- SOAR runbooks/playbooks and execution transcripts
- Exception register (if any) with expiry and approved mitigations

KPIs

- Shadow-AI coverage (%): $(\text{sanctioned} + \text{discovered}) \div (\text{sanctioned} + \text{discovered} + \text{ungoverned identified}) \times 100$. Trajectory: Q1 $\geq 80\%$, Q2 $\geq 85\%$, Q3 $\geq 90\%$
- DLP effectiveness (%): $\text{blocked violations} \div \text{total violations} \times 100$. Target: Q1 $\geq 90\%$, tune +2–3 pts q/q towards $\geq 95\%$
- MTTD / MTTR (hours): mean time to detect / respond to AI incidents. Targets: $\leq 24\text{h}$ / $\leq 72\text{h}$
- Token hygiene (%): $\text{connectors within rotation SLA} \div \text{total connectors} \times 100$. Target: $\geq 95\%$

Owner: CISO / Security Engineering

2.2 Governance & Regulatory Assurance

Objective. Maintain continuous conformity across jurisdictions and retain audit-ready evidence for all AI-enabled use cases.

Key Controls

- Regulatory mapping & change tracking (jurisdictions, obligations, owners, effective dates)
- Lawful basis documentation and periodic validation per use case
- Transparency & AI labelling in notices, interfaces, and user communications
- Data minimisation & retention controls aligned to policy
- Data subject rights handling (SAR, erasure, restriction, objection) covering prompts, outputs, and logic summaries
- Cross-border management: SCC/IDTA and a maintained Transfer Register (locations, processors, sub-processors)
- DPIA/FRIA triggers & execution with recorded mitigations and approvals prior to go-live
- Vendor due diligence & ongoing monitoring (attestations, locations, sub-processors, reassessment cadence)
- Lifecycle gates & exceptions enforced via GRC/workflow tools (create/change/decommission), with exception register (expiry, compensating controls)

DEFINITIONS: DPIA (Data Protection Impact Assessment) addresses GDPR Article 35 requirements. FRIA (Fundamental Rights Impact Assessment) addresses EU AI Act Article 27 requirements for high-risk AI systems. Reference: GDPR Art. 35, EU AI Act Art. 27.

Implementation Evidence (Examples)

- GRC/workflow lifecycle gates with approvals and audit trails
- RoPA entries with chain-of-custody for AI processing
- Transfer Register plus approved SCC/IDTA repository (versions, effective dates)
- Completed DPIA/FRIA (risk ratings, mitigations, approvals, residual risk decision)
- Vendor due-diligence artefacts (SOC 2/ISO/42001, sub-processors, data residency) and monitoring cadence records
- Transparency materials (privacy notice AI section, UI labels/screenshots)
- Rights-request logs mapping requests to AI prompts/outputs and actions taken

KPIs

- Regulatory control coverage: 100% of applicable controls mapped and implemented
- Evidence retrieval: within tiered SLA (see Evidence & Auditability section)
- Vendor reassessment on schedule: ≥95%
- DPIA/FRIA completion before go-live (high-risk): ≥95%
- Transparency coverage (labelled AI interactions): ≥98%
- Exception closure time: ≤10 business days (or documented extension with approval)

Owner: Compliance/Legal with GRC

2.3 Ethical & Human Impact

Objective: Reduce harm, promote dignity, inclusion, and accessibility whilst ensuring appropriate explainability.

Key Controls

- Bias and harm testing across protected characteristics
- Explainability profiles by use case and risk tier
- Safe-use pattern documentation and enforcement
- Human-in-the-loop requirements where mandated
- Contestability and redress channels
- Accessibility compliance (WCAG 2.1 AA where relevant)

Evidence & KPIs

- Explainability profiles by use case/tier
- Impact assessments (pre/post/periodic) with documented mitigations
- Bias parity gap $\leq 5-10\%$ (domain-specific thresholds)
- Redress SLA compliance $\geq 95\%$

Owner: CAIO/Product with Ethics/UX

2.4 Accountable Operations

Objective: Ensure traceability, clear responsibility, and effective oversight.

Key Controls

- Decision and data-flow logging with immutable audit trails
- Clear approval workflows and RACI matrix implementation
- Independent oversight and governance committee structure
- Grievance and escalation procedures
- Performance and SLO monitoring

Evidence & KPIs

- Decision records (model/provider/version, prompts/data, human reviewer, outcome)
- RACI documentation with named owners and committee minutes
- Governed use cases with complete records: $\geq 98\%$
- Audit findings remediated: ≤ 30 days
- Exception closure: ≤ 10 business days

Owner: Governance Committee / Business Units

2.5 Trust & Safety Culture

Objective (Empower): Foster a proactive Trust & Safety culture where employees are the first line of defence, empowered to use AI responsibly and report concerns without fear of reprisal.

Key Controls

- Mandatory, role-based annual training on AI Use Policy, data handling, and threat identification
- Quarterly awareness campaigns and AI-themed phishing simulations
- Clear, non-punitive 'AI Help & Reporting' channel for questions and incident reporting
- Regular culture assessment and feedback collection

Implementation Evidence

- Training completion records and attestations
- Phishing simulation results and analysis
- Help channel usage logs with query resolution tracking
- Culture assessment survey results

KPIs

- Training completion rate: $\geq 95\%$
- AI-themed phishing simulation click-through rate: $< 5\%$
- Employee-reported AI incidents vs. technically-detected incidents ratio (positive trend indicates cultural maturity)

Owner: Governance Committee with HR/Communications

3. Governance Domains & Controls

3.1 Pre-Implementation Readiness Assessment

3.1.1 Value & Risk Assessment Framework

Before implementation, organisations must perform a high-level assessment using this rubric, which also serves as the first gate for all new AI use case requests.

Assessment Question	Low Risk (1)	Medium Risk (2)	High Risk (3)
Data Sensitivity	Public/internal non-confidential data only	PII, customer, or confidential IP data	Special category data (GDPR), regulated data
Decision Impact	No direct impact on individuals (e.g., content summarisation)	Indirect impact (e.g., marketing personalisation)	Legal or similarly significant effects (e.g., hiring, credit)
Output Scope	Internal use only	B2B/partner-facing	Public-facing or directly customer-facing
Blast Radius	Low (single user/team impact)	Medium (departmental impact)	High (enterprise-wide, brand, or systemic impact)
Training Data Exposure (NEW)	No training on organisational data	Aggregate/anonymised training	Individual-level data used for fine-tuning

UPDATED: Added Training Data Exposure dimension. Reference: NIST AI RMF 1.0 MEASURE 2.7, 2.11.

Scoring: Sum the scores. 5-8 ⇒ Tier-1 (Fast Lane); 9-12 ⇒ Tier-2 (Standard Review); 13-15 ⇒ Tier-3 (Full Governance).

3.1.2 Objective Tiering Rubric

Factor	Tier-1	Tier-2	Tier-3
Compliance Requirement	Internal only	Buyer questionnaires/some regulator	Formal audit (e.g., SOC 2 Type II)
Data Sensitivity	No PII/pseudonymised test	PII/Customer data; no special categories	Special category/regulated
Risk Appetite	Moderate acceptance	Risk-based gating	Low tolerance; zero-trust gating
Trust Signal	None	B2B desired	Mandatory for sales/regulatory
Resources/Budget	≈1.0-1.2 FTE; minimal tools	≈1.5-2.5 FTE; modest tools	≥3.0 FTE; enterprise tools & audits
Timeline	<30 days	Quarter	Multi-quarter + attestations



ELSA AI

Ethical • Legal • Societal • Accountable AI Operations

3.2 Implementation Methodology (30-60-90 Day Plan)

Days 1-30: Foundation

- Sponsor identification and charter establishment
- AI Use Policy development and approval
- Value & Risk Assessment rubric implementation
- Exception workflow definition
- Shadow AI discovery initiation
- Sanctioned tool catalogue creation
- SSO/MFA deployment for approved tools
- AI log routing to SIEM (WORM configuration)
- Baseline DLP policy implementation
- DPIA/FRIA trigger list establishment
- RoPA initiation
- Trust & Safety awareness campaign launch

Days 31-60: Core Controls

- DPIAs for top 10 use cases
- Vendor risk assessments and reviews
- Shadow AI Triage Playbook development and deployment
- Control deployment (GA-PG/TP/DM/DC/RR/RB)
- Output and bias monitoring implementation
- Role-based training and attestations
- Explainability profile development
- Transparency labels deployment to production

Days 61-90: Optimisation & Scale

- Discovery process automation
- Dashboard deployment and configuration
- Evidence pack creation per tiered SLA
- Internal audit dry-run execution
- Tier finalisation and documentation
- Scale roadmap development
- Continuous improvement feedback loop implementation

3.3 Shadow AI Operations Playbook

Trigger: Automated discovery (CASB/DNS) or manual report (e.g., expense claim) of an unsanctioned AI tool.

Playbook Steps

1. Contain (Automated ≤1h):

- Add tool domain to 'monitor-only' or 'block' list in proxy/DNS filter
- Create incident ticket and assign to SecOps
- Notify relevant stakeholders

2. Triage (≤24h):

- SecOps, in consultation with user's line manager, runs the tool through Value & Risk Assessment
- Document decision rationale and supporting evidence

3. Remediate (≤5 business days):

- If Low Risk: Move to fast-track sanctioning queue with pre-approved change to add to SSO/DLP/Allow-list policies
- If Medium/High Risk: Maintain block and assign to business unit for decision: A) cease use and find approved alternative, or B) submit full governance request via DPIA/FRIA process

Evidence Requirements: All triage decisions, communications, and resolutions documented in ticketing system with immutable audit trail.

4. Control Catalogue with Implementation Profiles

4.1 GA-PG-001 Policy & Governance

Minimums

- AI Use Policy covering scope, roles, approved tools, disallowed data, labelling requirements, approval processes, and exception handling
- Lifecycle gates for Create/Change, Publish, and Decommission phases
- Exception workflow with risk acceptance, expiry dates, and mitigation requirements

Configuration

- Gate forms integrated in GRC/ticketing systems
- Unique Use-Case ID and Control ID assignment
- Evidence storage under Case-YYYY-NNN format with hash verification and timestamps

4.2 GA-TP-001 Technical Protection

Identity & Access

- SSO + MFA enforcement with local account disablement
- SCIM provisioning and least-privilege role assignment
- PAM for administrative tokens and elevated access

Secrets & Tokens

- Centralised secrets management in vault
- Token rotation ≤90 days with automated alerts
- Token revocation on role changes
- Unused token alerts >30 days

DLP (AI-aware)

- Pattern detection: payment cards (PAN/Luhn), national IDs, health terms, credentials (API keys/JWT), personal identifiers
- Multi-channel coverage: endpoint (clipboard/upload), web/CASB (AI domains), email, SaaS DLP
- Prompt/output redaction with pattern matching, masking, and logging

UPDATED: Organisations MUST implement DLP rules detecting: PCI DSS compliant PANs (13-19 digits passing Luhn algorithm check per PCI DSS 3.2.1 Req 3.4), national identifiers per ISO/IEC 27559, HIPAA identifiers per 45 CFR 164.514(b).

Network/Egress

- Proxy/CASB allow-list for approved AI API FQDNs
- Unknown webhook destination blocking
- TLS enforcement with HSTS where applicable

© 2025 ELSA AI LTD. All rights reserved.

GenAI Assure™ v1.1 Licensed under CC BY-ND 4.0 International.
Practical, Achievable, Auditable AI Governance.

Page 20 of 34

UPDATED: Organisations MUST enforce TLS 1.3 (RFC 8446). TLS 1.2 MAY be supported only for legacy systems with documented risk acceptance expiring within 12 months.

Cryptography

- At rest: AES-256-GCM or provider equivalent
- In transit: TLS 1.3 with modern ciphers and PFS preferred
- Customer-managed keys where offered, stored in KMS/HSM

Configuration Baselines

- AI SaaS tools: disable training on enterprise data, restrict retention, limit public sharing, disable personal spaces for work use

4.3 GA-DM-001 Detect & Monitor

Event Schema (Log Fields)

- User, role, IP/device identification
- Tool/application, use_case_id, action (prompt/output/upload/webhook)
- Data classification tags, decision (allow/block/redact), connector/token ID

UPDATED: Add mandatory fields: incident_category (per NIST CSF 2.0 DE.AE-2), regulatory_trigger (boolean), breach_notification_required (boolean), affected_data_subjects_count (integer). Reference: GDPR Art. 33, NIST CSF 2.0 DE.AE-2.

Core Detections

- Out-of-hours and anomalous request spikes
- PII patterns in responses and prompts
- New/changed webhooks to external domains
- Bulk data transfers and excessive error/loop counts

Dashboards

- Shadow-AI coverage metrics
- Policy violation trends
- Detection alerts over time
- Token rotation status
- Exception backlog monitoring

4.4 GA-DC-001 Documentation & Compliance

Evidence Pack (Tiered SLA)

UPDATED REQUIREMENT: Evidence pack MUST be retrievable within absolute timeframes. SLA clock starts upon documented request receipt. Failures require incident report within 2 business hours. Reference: SOC 2 CC7.2, ISO/IEC 27001:2022 8.16.

- Tier-1: MUST be retrievable within 4 business hours (simple use cases, ≤2 systems)
- Tier-2: MUST be retrievable within 8 business hours (moderate use cases, 3-5 systems)
- Tier-3: MUST be retrievable within 24 business hours (complex cases, 6+ systems); documented exception up to 48h with reviewer approval

Contents

Policies, approvals/exceptions, RoPA, DPIA/FRIA, transfer register, privacy notices, retention schedules, SIEM/DLP exports, WORM proof, sanctioned catalogue, discovery results, vendor files (DPA, SCC/IDTA, SOC 2/ISO attestations, sub-processors, locations), labels/screenshots, explainability profiles.

Lawful Basis by Common Use

- Marketing content: legitimate interests (with opt-out) or consent
- Customer support chatbots: contract/legitimate interests with human option
- HR screening: legitimate interests with DPIA, human review for legal effects
- Employee analytics: legitimate interests with transparency and minimisation

Data Subject Rights in AI Context

- Access: Include representative prompts/outputs and logic summary per explainability profile
- Erasure/Restriction: Delete/flag records in connected systems, purge conversation logs where feasible
- Objection: Suppress AI-based marketing profiling upon request

UPDATED: For AI systems where erasure is technically impossible due to model training, organisations MUST document technical infeasibility per GDPR Recital 65, implement data minimisation at input, and provide affected individuals with restriction of processing as alternative remedy. Reference: GDPR Art. 17(3), Recital 65.

4.5 GA-RR-001 Response & Remediation

Runbooks

- PII Exfiltration: Contain (disable connector/flow), notify stakeholders, investigate prompts/outputs, remediate vulnerabilities, issue user notices
- Token Compromise: Revoke/rotate tokens, search historical usage, notify impacted parties
- Misleading/Deepfake Content: Pull content, apply labels, establish redress contact, issue corrective communications

SOAR Actions

- On DLP 'block': auto-disable connector, revoke token, open incident, notify owner + Security + Compliance, attach evidence export

Redress

- User intake with SLA (acknowledgement ≤ 1 business day; resolution target ≤ 10 days)
- Feedback loop integration into policy and control updates

4.6 GA-RB-001 Resilience & Business Continuity

Requirements

- Dependency mapping: AI workflows, vendors, tokens, data stores
- Fallback modes: manual or non-AI paths for critical services
- RTO/RPO documentation and annual testing
- Vendor continuity evidence collection and validation

5. TEVV-Lite Testing Playbooks

Each playbook follows a Pass/Fail approach with specific evidence requirements.

5.1 Avatar/Video Generation

Tests: Licence/likeness validation, misuse prompt testing, bias sampling, approval workflow

Accept Criteria: Permissions on file, misuse blocked/refused, parity within threshold, human approval recorded

UPDATED: Organisations MUST validate deepfake detection capabilities using NIST IARPA FVRT deepfake detection benchmarks. Detection rate MUST exceed 95% for known manipulation techniques. Retest quarterly. Reference: NIST FVRT, OWASP LLM10.

5.2 Workflow Automations

Tests: Data minimisation validation, exfiltration testing to non-approved domains, loop/volume guards, rollback procedures

Accept Criteria: Block+alert on exfiltration test, rate-limit/kill-switch proven, rollback steps documented

5.3 Developer Copilots

Tests: Canary secret detection, licence/IP guardrails, human review for critical changes, training attestations

Accept Criteria: Canary caught, unsafe suggestions flagged, PR review recorded, staff trained

5.4 Customer Service Chatbots

Tests: AI notice on first interaction, PII redaction, escalation to human, response sampling, refusal handling

Accept Criteria: Harm/off-topic rate $\leq 2\%$, escalation success $\geq 95\%$, PII redaction hit rate $\geq 95\%$, label/notice visible

5.5 Document Analysis/Generation

Tests: Sensitive document handling, local redaction, watermark/label presence, citation/attribution verification

Accept Criteria: Sensitive terms masked, watermark present, citations complete

5.6 HR Screening Tools

Tests: Bias checks across protected groups, human-in-the-loop validation, rationale capture, candidate notice

Accept Criteria: Parity gap $\leq 10\%$, human reviewer recorded, notices stored

5.7 Financial Analysis Tools

Tests: Reproducibility validation, data lineage tracking, anomaly thresholds, segregation of duties

Accept Criteria: Deterministic steps documented, threshold alerts triggered, SoD enforced

5.8 Marketing Content Creation

Tests: Brand-safety filtering, claims verification, copyright/asset licencing, labelling

Accept Criteria: No disallowed topics, claims sourced, licences attached, labels present

6. Standards Alignment and Regulatory Implementation

6.1 EU AI Act - Deployer Duties (Article 26)

Compliance Checklist

- Use AI systems according to provider instructions
- Maintain comprehensive usage logs
- Ensure human oversight where required
- Monitor system operation and performance
- Implement robust data governance measures
- Correct misuse and address identified risks
- Cooperate with providers and authorities
- Place appropriate transparency notices

UPDATED REQUIREMENT: Organisations **MUST** report serious incidents to AI provider and market surveillance authorities within 15 days per Article 73. Incident classification **MUST** follow Commission Implementing Regulation criteria (when published). Reference: EU AI Act Art. 26(3), Art. 73.

High-Risk Use Cases: Ensure trained staff, oversight procedures, and log retention meet both provider and legal requirements.

6.2 FRIA (Fundamental Rights Impact Assessment)

When Required: When deployer context creates risks to fundamental rights (e.g., HR screening, eligibility determinations)

Assessment Components

- Purpose and necessity justification
- Legal basis identification
- Rights at stake analysis
- Affected groups identification
- Risk scenarios (likelihood and impact assessment)
- Mitigation measures (human review, appeals/redress)
- Residual risk evaluation
- Proceed/redesign decision documentation

6.3 GDPR/UK GDPR Practical Implementation

Key Requirements

- Lawful basis documentation per use case
- DPIA execution for likely high-risk processing
- Subject Access Request flow including prompt/response and logic summaries
- Erasure/restriction capabilities where feasible
- Opt-out mechanisms for AI-based marketing profiling
- Transfer controls: maintain current SCC/IDTA, log vendor locations/sub-processors, update notices when vendors change

CLARIFICATION: For AI systems where erasure is technically impossible due to model training, organisations **MUST** document technical infeasibility per GDPR Recital 65, implement data minimisation at input, and provide affected individuals with restriction of processing as alternative remedy. Reference: GDPR Art. 17(3), Recital 65.

6.4 NIST CSF 2.0 Mapping

Function	GenAI Assure™ Controls
Identify	Shadow-AI discovery, sanctioned catalogue, RoPA (GA-DM-001/GA-DC-001)
Protect	SSO/MFA/least privilege, DLP, vault, egress allow-list (GA-TP-001)
Detect	SIEM/UEBA, detections, dashboards (GA-DM-001)
Respond	IR runbooks, SOAR, redress (GA-RR-001)
Recover	Fallback modes, continuity tests (GA-RB-001)

7. Evidence Automation and Auditability

7.1 Architecture Pattern

Sources: SIEM, DLP, CASB/proxy, IAM/IdP, ticketing/GRC, SaaS AI admin APIs, WORM storage, vendor portals

ETL/Collector: Scheduled API pulls with use_case_id and control_ids, normalised to common schema

Storage: Evidence YAML + binaries in WORM/object storage with SHA-256 hashes and NTP time synchronisation

Bundle Generation: Per use-case evidence packs with manifest and hash verification

7.2 Correlation Keys

Essential for audit trail maintenance: use_case_id, control_id, vendor_id, token_id, connector_id, environment, timestamp, decision (allow/block/redact)

7.3 Evidence YAML Schema Example

use_case_id: UC-023

controls:

- id: GA-TP-001

- artifacts:

- type: dlp_policy_export

- system: casb

- uri: s3://evidence/UC-023/dlp-export.json

- sha256: "<hash>"

- id: GA-DM-001

- artifacts:

- type: siem_alert

- system: splunk

- query: "ai_webhook_block"

- uri: s3://evidence/UC-023/alerts-2025-08.json

- sha256: "<hash>"

pack_built_at: "2026-01-31T12:10:03Z"

builder: "evidence-bot@company"

manifest_sha256: "<hash>"

© 2025 ELSA AI LTD. All rights reserved.

GenAI Assure™ v1.1 Licensed under CC BY-ND 4.0 International.
Practical, Achievable, Auditable AI Governance.

Page 28 of 34

7.4 Validation Requirements

- Hash verification on retrieval
- Tamper-evident WORM settings
- Cross-system timestamp reconciliation
- Automated missing-artefact alerts

UPDATED: sha256 field **MUST** contain 64-character hexadecimal SHA-256 digest.
Organisations **MUST** validate hash on evidence retrieval and reject packs with mismatched hashes. Failed validations **MUST** trigger security incident per GA-RR-001. Reference: NIST FIPS 180-4, SOC 2 CC6.1.

8. Continuous Improvement Management System

8.1 AIMS (ISO/IEC 42001) Cycle

Phase	Description
Plan	Policy development, objective setting, risk criteria establishment, roles (RACI) definition
Do	Control implementation, training delivery, TEVV-Lite execution, vendor oversight
Check	Quarterly internal audits and KPI reviews, management review inputs, AI incident analysis, DLP block review, user redress ticket analysis for systemic issues
Act	Corrective and Preventive Action (CAPA), control updates, risk register refresh, policy and training material updates based on Check phase findings

Cadence: Quarterly checks with annual management review and comprehensive internal audit

UPDATED: Internal audits MUST sample minimum 10% of active use cases per quarter, rotating to achieve 100% coverage annually. Audit scope MUST include controls GA-PG through GA-RB. Non-conformities MUST be remediated per documented CAPA within 30 days. Reference: ISO/IEC 42001:2023 9.2, 10.1.

8.2 Feedback Loop Integration

All AI incidents, DLP blocks, and user redress tickets are systematically reviewed to identify patterns and control gaps. Findings are formally integrated into actionable updates to policies (GA-PG), controls (GA-TP), or training materials.

9. Resource Planning and Implementation Support

9.1 Resource Requirements by Tier

Item	Tier-1	Tier-2	Tier-3
Core Staff	≈1.0-1.2 FTE total	≈1.5-2.5 FTE	≥3.0 FTE
Tooling (annual)	Basic CASB/Proxy, SIEM seats, Vault, DLP lite	CASB + SaaS DLP, SIEM tier, SOAR starter	Enterprise CASB/DLP, SIEM/SOAR enterprise, evidence automation
External Support	Ad-hoc DPIA/FRIA support	Readiness + audit preparation	Formal attestation programme
Timeline	30-60 days	60-120 days	120-180 days

UPDATED - Role Breakdown by Tier (Reference: SABSA Business Attributes):

- Tier-1 (1.0-1.2 FTE): 0.5 Security Engineer + 0.3 Compliance Analyst + 0.2-0.4 Business Process Owner
- Tier-2 (1.5-2.5 FTE): Add 0.5 GRC Specialist + 0.5 Data Protection
- Tier-3 (≥3.0 FTE): Add dedicated AI Governance Lead + Legal/Privacy counsel allocation

9.2 Change Management Approach

Stakeholder Engagement

- Executive sponsor identification and communication
- Multi-tiered briefings for different role groups
- Regular milestone celebrations and quick-win demonstrations

Resistance Management

- Transparent communication about benefits and necessary changes
- Clear exception processes with defined SLAs
- Practical solutions for common integration challenges

Adoption Tracking

- Tool adoption behind SSO, sanctioned catalogue completion
- Training completion and assessment pass rates
- DLP false-positive trends and exception processing times

10. Next Steps

10.1 Framework Adoption Recommendations

1. Begin with Readiness Assessment: Use the Value & Risk Assessment framework to determine appropriate tier and scope
2. Secure Executive Sponsorship: Ensure clear commitment and resource allocation aligned with chosen tier
3. Establish Cross-Functional Team: Include Security, Compliance, Legal, Privacy, and relevant Business Units
4. Follow Phased Implementation: Adhere to 30-60-90 day plan with regular checkpoint reviews
5. Maintain Continuous Improvement: Implement quarterly review cycles and annual comprehensive assessments

10.2 Success Metrics and Milestones

Day 30 Success Indicators

- AI Use Policy approved and published
- Sanctioned tool catalogue established with SSO/MFA implementation
- Shadow AI discovery process operational
- Initial SIEM logging and DLP policies deployed

Day 60 Success Indicators

- Top 10 use cases with completed DPIAs
- Vendor risk assessments completed for critical tools
- Shadow AI Triage Playbook operational
- Role-based training programme launched

Day 90 Success Indicators

- Comprehensive dashboards operational
- Evidence automation functional per tiered SLA
- Internal audit readiness demonstrated
- Continuous improvement feedback loop established

10.3 Long-term Maturity Evolution

The framework supports organisational maturity growth through:

- Tier progression: Natural advancement from Tier-1 through Tier-3 as capabilities mature
- Standards evolution: Adaptability to emerging regulatory requirements and industry standards
- Technology integration: Scalable architecture supporting advanced tooling adoption
- Culture development: Progressive Trust & Safety culture maturation with measurable indicators

UPDATED: Organisations SHOULD compare achieved KPIs against Cloud Security Alliance AI/ML Security Benchmarks (when published) and document variance rationale in management review. Target: within 10% of industry median by Year 2. Reference: CSA CCM AI/ML Addendum.

10.4 Community and Support

Organisations implementing this framework are encouraged to:

- Participate in industry working groups and standards development
- Share anonymised lessons learnt and best practices
- Contribute to framework evolution through structured feedback mechanisms
- Engage with regulatory bodies and industry associations for guidance updates

11. Licence and Attribution

This framework is made available under Creative Commons Attribution-NoDerivatives 4.0 International (CC BY-ND 4.0).

You are free to:

- Share — copy and redistribute the material in any medium or format for any purpose

Under the following terms:

- Attribution — You must give appropriate credit, provide a link to the licence, and indicate if changes were made
- NoDerivatives — You may not modify, transform, or build upon this material

Required citation:

Ali, F. (2026). GenAI Assure™ Framework v1.1. ELSA AI LTD. Available at: www.elsaai.co.uk

Full licence text: <https://creativecommons.org/licenses/by-nd/4.0/>

For implementation support or commercial licencing enquiries: contact@elsaai.co.uk

© 2026 ELSA AI LTD • GenAI Assure™ is a trademark of ELSA AI LTD • All rights reserved