



Governance and Responsible AI Framework

Partner and client assurance

Finaffairs Consulting
Version 1.0
11 October 2026

Proposed for adoption
Approval and effective date to be recorded by the adopting legal entity

This framework defines how Finaffairs should govern AI and related delivery risks, protect information, retain human accountability and substantiate assurances to partners and clients. It includes decision rights, risk controls, approval gates, evidence requirements and reusable templates.

Use the main policy for adoption and internal delivery. Use the companion assurance schedule to define the scope and evidence of a specific client or partner review. Neither document alone establishes operating effectiveness or external certification.

Navigation

Sections 1 to 7 Governance risk and lifecycle
Sections 8 to 16 Information security delivery and operations
Sections 17 to 20 Assurance obligations reporting and adoption
Appendices A to G Control register and working templates
References Primary sources and glossary

1 Purpose scope and adoption

This framework establishes the governance and responsible AI requirements proposed for Finaffairs Consulting. Its purpose is to protect clients, partners, employees and affected people while enabling useful, proportionate adoption of AI. Assurance must be based on identifiable evidence of controls operating within a stated scope and period.

It applies to internal AI use, consulting deliverables, client solutions, demonstrations, training, embedded vendor AI, predictive models, generative AI and agents acting on systems. Contractors, associates, subcontractors and delivery partners are included through contractual obligations. Deterministic financial calculators and analytics remain subject to quality and model controls even when they are not AI.

Adoption and document control

Field	Controlled value
Policy identifier	FA GOV RAI 001
Version and issue date	1.0 11 October 2026
Status	Proposed for adoption Approval pending
Executive sponsor	Founder and Managing Principal
Policy owner	Designated Governance and Responsible AI Lead
Approval authority	Governing body or duly authorised entity director
Effective date	To be entered after formal approval
Review	Annual and on material legal or operating change

Each Finaffairs legal entity must adopt the framework through its own authorised governance process and record its trading name, jurisdiction and accountable director. Adoption does not establish certification, regulatory approval or evidence that every control already operates. The initial implementation status of controls is Not assessed until evidence is reviewed.

The signed statement of work, applicable law and sector obligations take precedence where they impose stricter requirements. Only lawful, documented internal variations may be approved; commercial urgency cannot waive a legal obligation.

2 Governance commitments and risk appetite

Finaffairs will require an identified business purpose, an accountable human owner and an approved operating boundary for every AI use case. A non-AI alternative must be considered where it achieves the objective with lower risk or cost. No benefit forecast justifies an unlawful or unacceptable use.

Principles translated into decisions

- **Accountability.** A named owner remains responsible for the outcome, including work assisted by a vendor model or subcontractor. AI does not sign advice, approve reporting or assume professional responsibility.
- **Purpose and proportionality.** Use only the data, capability and autonomy required for an approved purpose. Evaluate the effect on people as well as on the client business case.
- **Privacy and confidentiality.** Establish rights to use information before processing. Restrict training reuse, access, retention and cross-client retrieval.
- **Fairness and human rights.** Test foreseeable disadvantage, accessibility and opportunities to challenge material decisions. Address harmful proxies and distributional effects.
- **Reliability and security.** Demonstrate performance in the intended context, test misuse and foreseeable failure, and retain a workable fallback.
- **Transparency and professional judgement.** Disclose material AI involvement and limitations appropriately. Verify factual claims and financial outputs against authoritative evidence.
- **Sustainability and improvement.** Assess workforce effects, vendor concentration, computational cost and relevant environmental impacts alongside benefits.

Non-negotiable boundaries

No unlawful discrimination, prohibited surveillance, deceptive impersonation, fabricated evidence, exploitative manipulation or unauthorised data extraction. No unapproved upload of client secrets or personal data into consumer AI accounts. No autonomous adverse employment, credit, insurance, public-benefit or similar material decisions without a separately approved lawful decision process and meaningful safeguards.

High-risk deployment requires independent challenge and executive acceptance. Critical residual risk blocks production. A policy exception cannot authorise a prohibited use or bypass a failed mandatory release gate.

3 Governance structure and accountability

The model scales to a consulting firm. Roles may be assigned to existing leaders or competent retained specialists; new departments are not assumed. Names, competence, deputies, time allocation and conflicts must be recorded before the roles are treated as operational.

Role	Primary responsibility
Governing body or entity director	Approve policy and risk appetite; oversee material exposures and resourcing.
Managing Principal	Executive sponsor; accept high-tier residual risk; ensure corrective actions are funded.
Governance and Responsible AI Lead	Maintain inventory, policy, gate decisions, risk register and assurance coordination.
Engagement or AI System Owner	Own purpose, client consent, delivery controls, results and monitoring.
Technical and Data Leads	Implement architecture, data controls, testing, deployment and traceability.
Privacy and Security Leads	Challenge data protection, transfers, security design and incident response.
Legal and Commercial Reviewer	Assess applicable law, terms, IP, liability, sector constraints and representations.
Independent Reviewer	Challenge evidence and model limitations; report findings without delivery pressure.
Client Accountable Owner	Approve client purpose, data access, use boundary, acceptance and local deployment.

Governance forum

Convene the Governance and AI Assurance Forum monthly during implementation and at least quarterly thereafter. Membership must include executive, delivery, AI/data, privacy/security and assurance perspectives. For a high-tier decision, require the executive delegate, Governance Lead and relevant specialist reviewers; record attendance, conflicts, dissent and decisions.

A builder must not independently validate their own work. Where staffing prevents separation, appoint a competent external reviewer and record the arrangement. No risk acceptance substitutes for required privacy, security or legal clearance. Stop-work decisions must be actionable immediately.

4 Decision rights and escalation

The decision owner below is accountable for the decision. Responsible staff assemble the evidence. Required clearances are separate veto gates, not optional consultation. The Governance Lead records the signed decision and its conditions.

Decision	Accountable decision owner	Required clearance
Entity policy adoption	Entity director or governing body	Managing Principal and legal review
Use-case risk classification	Governance Lead	System Owner and relevant specialists
Approved tools and vendors	Governance Lead	Security privacy and commercial
Low or moderate release	Engagement Partner or delegated director	Gate evidence and relevant specialists
High-tier release	Managing Principal	Independent review and all relevant specialists
Client production acceptance	Client Accountable Owner	Finaffairs release and agreed client gates
Policy exception	Managing Principal	Governance Lead and affected control owner
Emergency suspension	Security Lead or System Owner	Notify Governance Lead immediately
External incident notification	Authorised entity director	Privacy legal and client coordination
Assurance statement publication	Managing Principal	Evidence verification and commercial review

Escalation and conflict handling

Anyone may report a suspected breach or request a pause without retaliation. Raise delivery concerns with the Governance Lead; escalate unresolved material concerns directly to the Managing Principal or entity director. Allegations concerning that individual go to an unconflicted director or retained independent adviser.

Record referral fees, reseller incentives, vendor interests, personal relationships and advisory or political conflicts. Declare commissions and material incentives to the client where relevant. A commercial sponsor cannot suppress an unfavourable validation finding. Separate consultancy recommendations from any assurance opinion where independence is impaired.

Maintain an approved escalation contact list and out-of-hours route. No new AI use case starts until its owner and escalation route are identifiable.

4A Professional integrity and client acceptance

Responsible AI sits within professional governance. Before accepting an engagement, Finaffairs must establish that the proposed work is lawful, within its competence, appropriately insured where required and consistent with independence and integrity obligations.

Client and engagement acceptance

- Identify the contracting entity, authorised client representative, intended users and delivery scope. Apply proportionate identity, sanctions and other legally required checks; record competence, capacity, commercial terms and delivery risks.
- Declare financial, vendor, personal and advisory conflicts before recommending technology or accepting review work. Explain material referral or reseller incentives; obtain appropriate disclosure and mitigation rather than assuming client consent cures every conflict.
- Reject requests to fabricate evidence, conceal material limitations, evade controls, misuse protected information or make unsupported compliance claims. Record the reason and escalate material concerns.

Conduct and quality controls

- Apply anti-bribery, gifts/hospitality and procurement rules appropriate to the entity and engagement. Record approvals and prohibit improper inducements, secret commissions and manipulation of procurement or evaluation.
- Use a signed scope, accountable delivery lead, competent reviewers and documented acceptance criteria. Monitor subcontractor conduct and information use. Preserve professional scepticism when commercial incentives favour a particular vendor or result.
- Protect whistleblowers, maintain confidential reporting and investigate allegations through an unconflicted person. Client concerns must be tracked to a reasoned resolution, with remedial action where warranted.

Governance evidence

Maintain client acceptance, conflict declarations, conduct training, approval records, quality review and complaint/action logs. Coordinate with existing entity policies; the framework does not replace employment, anti-corruption, security or professional regulatory obligations.

Separate internal quality control from an independent assurance engagement. Define who may sign a report, the scope of responsibility and the evidence on which the signature relies.

5 AI inventory and classification

Maintain a central inventory covering purchased, embedded, internally developed and client-delivered AI. Link each record to the engagement, legal entity, supplier/model/version, data classes, jurisdictions, affected people, connected tools, lifecycle status and approval history. Record non-AI models separately when material.

Internal risk tiers

Tier	Typical exposure	Minimum treatment
Prohibited	Unlawful or explicitly forbidden use	Reject or suspend; no exception.
High	Material effects on rights livelihood finance safety or essential services; sensitive data at scale; significant autonomy	Full impact assessment; independent validation; executive approval; frequent monitoring.
Moderate	Client advisory research or operational support with bounded data and human approval	Documented assessment; specialist checks; acceptance testing; controlled deployment.
Low	Public or synthetic information; reversible internal assistance with limited impact	Register; approved tool; trained user; proportionate review.

These are Finaffairs internal tiers, not EU AI Act classifications. Record legal classifications and actor roles independently. An internal Low rating cannot override a legal obligation.

Mandatory escalation rules

- Classify at least High when AI materially influences recruitment, workforce ranking, lending, underwriting, claims denial, public-service eligibility or critical infrastructure decisions.
- Escalate sensitive personal data, children or vulnerable groups, cross-client retrieval, broad write access, autonomous external actions, novel model behaviour or unresolved supplier opacity.
- A sandbox using synthetic data is not a production approval. A demonstration becomes a new use case when real client data, an operational decision or system write access is introduced.

Review classification before design, release and material change. Document the basis for any reduction in tier; the Governance Lead must approve it after relevant specialist challenge.

6 Risk assessment and acceptance

Assess each credible harm scenario separately. Identify the cause, affected party, impact, likelihood, existing controls, uncertainty and accountable risk owner. Avoid an average score that hides a serious privacy, fairness or safety exposure.

Internal scoring method

Use impact 1 to 5 and likelihood 1 to 5. Impact runs from minor reversible inconvenience to severe harm to rights, safety, financial position or public trust. Likelihood runs from rare with credible evidence to frequent or already observed. Score each scenario as impact multiplied by likelihood; take the highest material scenario as the use-case residual rating.

Residual score	Treatment	Authority
1 to 4 Low	Within approved controls; routine review	Delegated release owner
5 to 9 Moderate	Documented treatment and monitoring	Engagement Partner
10 to 16 High	Independent challenge; explicit conditions	Managing Principal
17 to 25 Critical	No production; redesign or discontinue	No delegated acceptance

These thresholds are proposed internal policy defaults, not legal standards or statistically calibrated probabilities. Any impact 5 scenario remains at least High even if the product score is lower. Missing evidence cannot be scored as proof of low likelihood. A prohibited use is rejected regardless of score.

Example classifications

- Public research drafting can be Low if it uses approved tools, public sources and a qualified reviewer. Confidential client advisory drafting is generally Moderate or higher depending on data and consequences.
- Recruitment ranking or claims triage that influences material outcomes is High even where staff can override recommendations. An agent that can initiate payments or publish reports requires High-tier controls and transaction approval.
- An Academy calculation prototype requires analytical validation and disclosed assumptions. It must not be represented as an AI system merely because AI helped create it, or as an approved regulatory engine because its examples reconcile.

Acceptance records must state scope, residual scenarios, compensating controls, owner, expiry and review triggers. Reduce risk through evidence of effective controls; signing a form does not reduce a score.

7 Lifecycle gates and controlled delivery

Every stage produces retained evidence. Gates may return work for remediation or reject it. The client acceptance gate and Finaffairs release gate are distinct. A commercial milestone is not a control clearance.

Gate	Evidence required	Decision
G0 Intake	Purpose owner inventory entry non-AI alternative and intended benefit	Register or reject
G1 Assess	Risk tier impact assessment legal roles data rights and misuse scenarios	Permit design or redesign
G2 Design and procure	Architecture supplier review contract controls evaluation plan and fallback	Permit build or pilot
G3 Validate	Versioned tests privacy/security checks fairness findings independent review where required	Pass or remediate
G4 Release	Signed evidence index client approval operating limits training and rollback test	Approve within conditions
G5 Operate and change	Monitoring complaints incidents drift access reviews and change assessments	Continue restrict or suspend
G6 Retire	Client handover revoked access deletion/retention evidence and dependency closure	Close and archive

Release hard stops

No release where purpose/data rights are unresolved; required reviewers have withheld clearance; a material security or fairness finding remains open; outputs fail a critical test; the operator cannot stop the system; or the fallback has not been demonstrated. A pilot with controlled restrictions must be labelled as such and cannot silently become production.

The approval applies to the recorded model, configuration, retrieval corpus, prompt templates, connectors, users and deployment context. Any material change returns to the appropriate gates. Maintain separate development, test and production environments and prevent uncontrolled promotion.

The operator may suspend use at any point. Restoration requires the failed controls to be retested, the owner to accept the remedial evidence and the original release authority to approve resumption.

8 Data governance privacy and confidentiality

The Data Lead must document provenance, permitted purpose, lawful basis where applicable, sensitivity, quality, retention, access and approved geography for each data source. Public availability does not establish permission for reuse or remove personal-data obligations.

Data classes and handling

Class	Examples	AI handling rule
Public	Published reports and authorised marketing	Approved tool and source verification.
Internal	Operating documents and non-public strategy	Approved business account; access restrictions.
Confidential	Client records pricing proposals unpublished financials	Client authority and approved isolated environment; contracted restrictions.
Restricted	Sensitive personal data credentials regulated secrets	Specific privacy/security approval; minimisation; stringent isolation and access.

Mandatory controls

- Separate controller/processor roles, data subject notices and client instructions. Conduct privacy impact assessment when required or when elevated risk warrants it; obtain specialist clearance. Determine whether a statutory DPO or registration is required and appoint/register as applicable.
- Use minimisation, redaction or synthetic data first. Pseudonymised data remains personal data where re-identification is possible. Test anonymisation claims rather than assuming masking is sufficient.
- Enforce client and tenant boundaries in storage, retrieval indexes, embeddings, caches, logs, training files and exports. Prevent one client from retrieving another client material.
- Verify vendor retention, training use, human review and deletion settings by contract and configuration. A business subscription alone does not prove no-training, residency or confidentiality guarantees.
- Document transfers, subprocessor locations and required transfer safeguards before routing data abroad. Honour rights requests and legal holds; reconcile deletion with necessary audit retention.

Retention is an approved schedule by record type, jurisdiction and contract, with maximum periods and deletion owners. Store evaluation and decision evidence without unnecessarily retaining sensitive prompts or complete source records. Access to assurance evidence is controlled and logged.

9 Security architecture and AI agents

Treat prompts, retrieved documents, model responses and tool output as potentially untrusted. A model cannot grant itself permissions or override the client authority boundary. Threat modelling must cover the AI application and the systems it can reach.

Application security

- Apply identity controls, multifactor authentication, least privilege, encryption, approved secrets management and access review. Use service identities with scoped permissions instead of embedded credentials or shared administrator accounts.
- Test prompt injection, data leakage, poisoned sources, supply-chain changes, unsafe output handling, insecure retrieval and cost-exhaustion abuse. Enforce access checks outside the model and before every retrieval or action.
- Validate output schemas and parameters. Escape displayed content. Do not execute model-generated code, SQL or shell instructions directly in a privileged production context. Scan dependencies and protect source repositories.

Additional agent controls

- Define an allowlist of tools, destinations, data classes and permitted actions. Set limits on transaction value, volume, duration, tokens and retries. Require human approval for material external actions.
- Bind approval to the exact proposed action, recipient, account, amount and time window. A broad instruction to help does not authorise payments, external publication or unrestricted deletion.
- Use isolated execution, network egress restrictions, dry-run previews and reversible actions where possible. Prevent agents from changing their own control configuration or approving their own outputs.
- Record actor identity, model/configuration version, tool calls, approval decisions and action results. Test emergency shutdown, credential revocation, rollback and restoration.

Default client agents to read-only until action permissions are explicitly granted and tested. Security review must cover connectors, plug-ins, retrieval stores, memory, browser access and outsourced model endpoints. Use OWASP AI security resources as a threat reference, not as proof that testing is complete. [R8]

10 Responsible generative AI use

Use approved business tools for authorised tasks. Users must understand that plausible text is not evidence of truth. The accountable professional checks material facts, citations, computations, confidentiality and suitability before an output enters a client deliverable.

Permitted use within approved boundaries

- Drafting, summarisation, translation, coding assistance and research may support qualified staff where information rights, review and disclosure controls are satisfied. Use retrieval from approved sources where current or specialised evidence is needed.
- Record material prompts or their controlled templates, retrieval sources, model/version and review decisions where reproducibility or client assurance requires it. Avoid retaining sensitive raw content when a reference or redacted trace is sufficient.

Output verification

- Open and verify cited sources; confirm authorship, dates, jurisdiction and that each source supports the claim. Do not publish invented statistics, quotations, references, credentials or case studies.
- Recalculate financial amounts independently using controlled logic or approved systems. Review generated code, test failure modes and scan for credentials, unsafe dependencies and licensing concerns before release.
- Treat legal, financial, tax and regulatory guidance as requiring competent professional review and current primary sources. Do not describe a draft or prototype as an approved professional opinion.

Disclosure and consent

Disclose material AI assistance where it changes how a client should evaluate the work, where agreed by contract, or where required by law. Inform users of AI interaction and provide an alternative escalation route where appropriate. Obtain client permission before confidential data enters any external model service.

AI-generated media must not impersonate individuals without authority or misrepresent real events. Assess IP, training-data and output rights in context; a vendor permission clause does not eliminate infringement risk. Respect source restrictions and third-party licence terms.

Do not permit unregistered personal accounts, covert meeting transcription, automated final advice, undisclosed subcontractor AI processing or unapproved fine-tuning on client data.

11 Validation quality and reproducibility

Validation must establish whether the whole system is fit for the approved use, not whether the base model has a strong public benchmark. Define acceptance thresholds before final testing and retain the rationale, data, versions, results and unresolved limitations.

Evaluation design

- Use representative tasks, populations, languages, product types, rare events and boundary cases. Separate development and final evaluation material. Document label quality, leakage controls and independence of the test set.
- Compare against an appropriate human or non-AI baseline. Include uncertainty, missing inputs, unfamiliar cases, retrieval failures, malicious content and operational overload. Test the fallback and human escalation route.
- For generative AI, assess unsupported material claims, citation accuracy, source fidelity, refusal/abstention, information leakage and task success. A second model may assist evaluation but cannot replace qualified review of critical cases.
- For predictive systems, examine errors, calibration, false positives/negatives, subgroup behaviour and stability. Explain how thresholds change affected outcomes and who bears the cost of mistakes.

Release evidence

Record expected result, actual result, pass/fail, severity, reviewer and evidence link for each test. No observed leakage in a finite test set is not a guarantee of zero leakage. Report sample sizes, coverage gaps and relevant uncertainty; avoid unsupported headline accuracy claims.

Material security, rights or financial defects must be resolved before release. High-tier systems require a reviewer independent of the builder. For vendor black-box systems, document inaccessible internals and evaluate the application externally; restrict or reject use where evidence cannot support the risk.

Model and analytical controls

Keep code/configuration versions, dependency records, model identifiers, approved datasets, parameter sets and test evidence. For financial models, reconcile independent implementations, balances, sensitivity and assumptions; distinguish arithmetic correctness from accounting or regulatory suitability. Academy demonstrations require visible educational limits and synthetic cases.

12 Fairness human oversight and societal effects

Assess who benefits, who may be excluded and who bears error costs. Do not assume that a technically accurate model is fair, accessible or legitimate for its intended decision. Include domain experts and relevant user or affected-person perspectives in high-tier impact assessment.

Fairness and accessibility

- Identify legally relevant protected characteristics, plausible proxy variables and vulnerable groups. Establish a lawful basis and safeguards for any sensitive data used in fairness testing.
- Select metrics appropriate to the decision, including differences in error, service access and adverse outcomes. Examine intersectional and locally relevant groups where evidence permits; state small-sample limitations.
- Document trade-offs, mitigations and remaining disparity. No single ratio or numerical threshold proves the absence of discrimination. Provide accessible explanations and alternatives for people with language, disability or connectivity barriers.

Meaningful human control

A reviewer must have relevant competence, adequate time, contextual evidence and authority to challenge, override and stop the system. Avoid interfaces that make acceptance the default or hide uncertainty. Record reasons for material overrides and assess automation bias, reviewer workload and disagreement.

People affected by material decisions must have an appropriate route to explanation, correction and reconsideration by an authorised person. A complaint cannot be answered solely by the same AI whose output is being challenged.

Workforce and public interest

Consult affected workers when deployment changes roles or monitoring. Assess training, workload, job quality and distributional effects, not only cost reduction. Public-sector work must consider due process, access to services, procurement integrity and the consequences of exclusion.

Do not use client or public-service data for partisan targeting, covert influence or discriminatory segmentation. Assess energy, compute and vendor concentration proportionately; use verified operational data for environmental claims.

13 Partner and supplier governance

Partner status, vendor training or a signed programme agreement does not establish that Finaffairs services are certified, that every product is secure or that client deployment is compliant. Validate the precise status and permissions before using logos, badges or joint marketing.

Due diligence before commitment

- Identify legal supplier, product, AI subprocessors, model families and accountable support route. Verify deployment locations, data use/retention, training policy, access controls, incident history and service dependencies.
- Review available audit reports and certifications for the exact entity, scope, period and exclusions. Record whether evidence is independently verified, self-attested, unavailable or subject to confidentiality.
- Test practical integration, export, deletion, configuration control and fallback. Evaluate lock-in, financial viability, pricing exposure, capacity and the impact of supplier/model withdrawal.

Contract schedule

Area	Required negotiation outcome
Purpose and data	Permitted processing; no unauthorised training or reuse; tenant separation.
People and geography	Approved subprocessors; transfer safeguards; change notification.
Change and security	Material model changes; minimum controls; notice and revalidation rights.
Incident and assurance	Prompt notice; cooperation; proportionate evidence and audit access.
Rights and liability	IP allocation; confidentiality; warranties; liability and insurance appropriate to risk.
Exit and continuity	Data portability; deletion evidence; transition assistance; service fallback.

Contract language requires legal negotiation for the applicable entity and client. Record unresolved gaps and their effect on use restrictions. Marketing promises and supplier features must not be substituted for enforceable terms or tested configuration.

Review high-risk suppliers at least annually and on material change; verify relevant controls before each new high-tier engagement. Flow equivalent obligations to associates and subcontractors and track their AI use.

14 Sector and engagement safeguards

Sector schedules supplement the common framework. The Engagement Owner must identify the client regulator, professional obligations, operating jurisdiction and approval process before solution design.

Banking insurance and financial advisory

Separate AI assistance from approved risk models, accounting policy and final advice. Apply client model-risk, outsourcing and reporting controls. Credit, underwriting and claims decisions require heightened fairness, explainability and challenge. Never permit unapproved autonomous journal posting, regulatory submission, premium changes, claims rejection or payment release.

For IFRS 9, IFRS 17 and Basel tools, test inputs, assumptions, model scope, scenario weighting, calculations and reconciliations. Independent arithmetic benchmarks do not prove compliance with a standard, a regulator rulebook or an entity policy. Production use requires separate calibration, integration, governance and acceptance.

Public sector and workforce services

Assess public-benefit eligibility, tax enforcement, hiring, promotion and succession decisions as High when materially influenced by AI. Preserve reasoned decisions, merit-based criteria, auditability, appeal and lawful data use. Check for political, geographic, disability, gender and other inappropriate proxies.

Energy and infrastructure

Use AI initially for bounded advisory or analytics work. Safety-critical operational control requires specialist safety assurance, approved engineering constraints and client authorisation. An advisory chatbot must not issue plant-control commands outside an approved, tested control architecture.

Data analytics training and demonstrations

Protect client segregation, semantic definitions, lineage, row-level access and reconciliation in analytics platforms. Use synthetic or explicitly authorised information in demonstrations. Distinguish learning outputs from operational products, and ensure trainers do not reuse participant/client records in unrelated models.

Where Finaffairs advises on or builds a system and also assesses it, disclose the relationship. Do not represent delivery-team checking as independent external assurance. Appoint an independent reviewer or define a clearly limited quality-review engagement.

15 Monitoring change and retirement

Before release, the System Owner must set monitored measures, collection method, frequency, escalation thresholds, responsible staff and stop conditions. Thresholds depend on use-case harms and accepted performance; a generic accuracy percentage is insufficient.

Operating review defaults

Tier	Routine monitoring	Formal review
High	Continuous alerts where feasible; weekly outcome review	Monthly owner review; quarterly governance challenge.
Moderate	Monthly performance and incident review	Quarterly owner review.
Low	User reporting and proportionate usage sampling	At least six-monthly review.

Increase frequency where law, contract or exposure demands it. Monitor data/behaviour drift, unsupported claims, retrieval integrity, subgroup outcomes, overrides, complaints, privacy/security events, usage boundaries, cost and dependency health. Maintain a reviewed baseline and investigate changes rather than simply resetting it.

Material change triggers

- A new model/version, tuning, prompt template, retrieval corpus, data class, country, user group, decision purpose, connector or write permission requires impact review and proportionate revalidation.
- Supplier upgrades must not silently change the approved operating boundary. Pin versions where possible; otherwise use change detection, regression tests and a viable suspension/fallback route.
- A material incident, unexpected disparate outcome, poor review quality or new legal requirement can require an immediate pause and gate reassessment.

Exit and closure

Confirm client handover and outstanding liabilities; revoke accounts, tokens and connectors; retire retrieval indexes and caches; apply retention/deletion schedules and legal holds; record supplier deletion limitations; and verify that dependent workflows have a functioning replacement. Archive only the evidence justified by the approved schedule.

Review value as well as harm. Continue a system only while its documented benefits and controls justify its cost and exposure.

16 Incidents complaints and remediation

An AI incident includes harmful or misleading output, inappropriate adverse decision, data leak, cross-client retrieval, unauthorised agent action, material bias, model failure or a breach of approved boundaries. Near misses and complaints enter the same register, with severity appropriate to actual or potential harm.

Severity	Examples	Internal response target
S1 Critical	Serious harm ongoing data exposure or unauthorised material transaction	Suspend/contain immediately; alert incident lead and executive within one hour of awareness.
S2 Major	Material error or control failure with plausible client impact	Restrict affected use; notify incident lead within four hours.
S3 Minor or near miss	Bounded defect without material observed harm	Record and triage within one business day.

These are internal proposed targets, not statutory reporting periods. The Privacy/Legal Lead must determine the applicable notice threshold, deadline, authority and client contract immediately; the earliest applicable requirement controls. Do not wait for root cause before a required notice.

Response sequence

- **Contain.** Stop affected actions, isolate access, preserve necessary logs and prevent further harm. Avoid destroying evidence or unnecessarily replicating sensitive information.
- **Assess.** Establish scope, affected people/data, model/configuration versions, actual and plausible harm, notification duties and client dependencies.
- **Communicate.** Use an authorised spokesperson and factual notice stating what is known, unknown, being done and how affected parties can obtain help.
- **Remediate.** Correct outputs or decisions where possible, assist affected clients/people and address configuration, process or supplier failures.
- **Validate and learn.** Independently retest material fixes before restoration. Set action owners and deadlines, review recurrence and update controls/training.

Target an initial management review within two business days and a root-cause/action report within ten business days for S1/S2 events, unless the complexity requires an approved revised plan. Complaints receive acknowledgement, tracking and human resolution within applicable legal/contractual periods.

17 Assurance for partners and clients

Provide an engagement assurance pack that states the legal entity, services, system boundaries, assessment period, evidence examined, limitations and accountable contacts. A policy describes required practice; an assurance conclusion requires evidence that relevant controls operated.

Four distinct types of assurance evidence

Evidence type	What it can establish	What it cannot establish
Policy adoption	Approved expectations and accountability	Operating effectiveness.
Management attestation	Responsible owner statement within scope	Independent assurance.
Independent review	Reviewer findings under stated procedures	Certification unless explicitly awarded.
External certification or audit	Verified entity scope period and opinion	Compliance of all products and engagements.

Minimum engagement pack

- Approved policy extract, named roles and scope; completed intake/risk assessment; system/model card; data-flow and processing schedule; supplier/model list and permitted locations.
- Relevant validation and security summary; human oversight/appeal route; release approvals; monitoring measures; incident escalation; retention, exit and change commitments.
- Control evidence index with owner, artifact identifier, version, date, test outcome, exceptions and the approved disclosure class.

Respond to questionnaires using Verified, Partially implemented, Planned, Not applicable or Not assessed. Give the evidence reference, scope and period for Verified; explain gaps and the action plan for partial/planned controls. Do not mark a requirement Yes merely because this document contains it.

Share only approved redacted evidence through secure channels. Protect other clients, employee data, trade secrets and security details. A refusal to disclose unrestricted raw records should be paired with an appropriate alternative evidence method, subject to legal/contractual obligations.

Have the Managing Principal and relevant evidence owners approve external claims. Do not claim ISO certification, regulator endorsement, guaranteed unbiased output or complete compliance without verifiable support.

18 Legal applicability and standards

The Legal/Privacy Reviewer maintains an engagement-specific obligations register. Record applicable entity, service, affected people, data flow, jurisdiction, regulator, actor role, requirement, source/version, effective date, control, evidence and accountable owner. Refresh before release and on relevant change.

Jurisdiction	Required applicability review
United Kingdom	UK GDPR and Data Protection Act as amended; current automated-decision safeguards; equality sector and professional obligations. ICO AI guidance is under review following the Data Use and Access Act. [R5]
Nigeria	Nigeria Data Protection Act 2023 and current NDPC directives including GAID; registration and filing applicability; impact assessments transfers rights and sector rules. [R6]
Ghana	Data Protection Act 2012 Act 843; registration security rights and automated decision-taking requirements; sector rules. [R7]
European Union	AI Act actor roles legal risk classification territorial scope transparency and applicable duties; GDPR and sector rules. Verify amended rules and commencement dates for each use case. [R4]
Canada	Applicable federal/provincial privacy rules; accountable privacy lead vendor protection and cross-border handling; sector and public-body requirements. [R9]
United States	Applicable federal state and city privacy consumer employment and sector rules; substantiated AI and privacy claims. No single nationwide blanket assurance claim. [R10]

An internal consent checkbox does not establish a lawful basis or transfer mechanism. Client approval does not relieve Finaffairs of its own obligations. Contractual allocation of responsibility must match the actual processing and decision roles.

Standards positioning

Use NIST AI RMF 1.0 as a voluntary organising reference and its generative AI profile where relevant; NIST reports a revision is in progress. ISO/IEC 42001 provides an AI management-system reference, and ISO/IEC 42005 addresses impact assessment. OECD principles and OWASP security resources inform ethical and technical challenge. This is a design mapping, not a clause-complete conformity assessment or certification. [R1 R2 R3 R8 R11 R12]

19 Assurance testing and management reporting

The independent reviewer tests the design and operation of selected controls using a defined population, period and sampling rationale. High-tier engagements require review before release and periodic follow-up. Management must not select only successful examples.

Status definitions

Not assessed means evidence has not been reviewed. Planned means an owner and action exist. Partially implemented means design or operation is incomplete. Verified means the stated control passed the defined test within scope and period. Not applicable requires a documented rationale and reviewer agreement.

Measure	Definition	Proposed target
Inventory coverage	Registered in-scope systems / discovered systems	100 percent
Approval coverage	Active systems with valid approvals / active systems	100 percent
Critical open defects	Unresolved release-blocking findings in production	Zero
Training coverage	Authorised users with current role training / authorised users	100 percent
Overdue exceptions	Exceptions past expiry or review date	Zero
Monitoring completion	Required reviews completed on time / required reviews	100 percent
Incident response	Events within internal target / events by severity	100 percent; report breaches
Outcome quality	Use-case error fairness/override complaint measures	Approved case-specific thresholds

Do not present these proposed targets as achieved results. State denominator, period, data quality and confidence. Investigate missing inventory discovery and unreported incidents; zero recorded incidents alone is not evidence of safe operation.

Quarterly governance review

Review material exposures, tier changes, incidents, findings, supplier dependencies, client complaints, legal changes, benefits and overdue actions. Record decisions and resources. An annual management review evaluates policy adequacy, competence, audit results and improvement priorities. Use independent external challenge where internal independence or expertise is insufficient.

20 Implementation and adoption roadmap

Implement in phases with an owner, evidence and acceptance criterion. Start with discovery rather than assuming the firm has no AI in use. Existing high-risk use without adequate controls must be paused or restricted while assessed.

Period	Priority deliverables	Acceptance evidence
Days 1 to 15	Approve policy; appoint accountable roles/deputies; discover tools and engagements; restrict unapproved confidential use	Signed adoption; contact list; initial inventory and urgent actions.
Days 16 to 30	Classify uses; approve tool/data rules; review suppliers and contract gaps; train initial users	Risk decisions; approved-tool list; vendor files; training records.
Days 31 to 60	Pilot gates and impact assessments; test security/quality/oversight; assemble client evidence packs	Completed pilot reviews; remediated blockers; evidence index.
Days 61 to 90	Conduct independent control testing; test incident/fallback processes; report status; resolve priority gaps	Review report; exercise log; governance minutes; dated action plan.
After day 90	Embed periodic monitoring legal refresh supplier review and annual improvement	Operating evidence; management review; renewed approvals.

Resources and dependencies

Allocate named staff time for governance, privacy/security, delivery and independent review. Budget for approved tools, secure evidence storage, evaluation datasets, legal support, training and specialist testing. Formalise client responsibilities and ensure suppliers can provide sufficient evidence before committing to delivery.

Adoption record

Adopting legal entity and jurisdiction: _____

Approving director or governing body and authority: _____

Approval date and effective date: _____

Policy owner and deputy: _____

Recorded exceptions and implementation plan reference: _____

Authorised signature: _____

Initial adoption should approve the policy and the dated implementation plan. External statements must accurately identify which controls are verified and which remain planned.

Appendix A Control and evidence register

Use the control identifiers to link policy, use cases, tests and external evidence. The owner names below are role assignments to be confirmed during adoption. Controls begin Not assessed; the table is a requirements register, not evidence of operation.

ID	Control requirement	Role owner	Evidence
G01	Entity adoption and delegated authority	Director	Signed adoption; delegation
G02	Named competent roles and independent challenge	Governance Lead	Role register; conflict record
G03	AI and material model inventory	Governance Lead	Inventory; discovery evidence
G04	Risk scenarios tier and acceptance	System Owner	Risk register; decision
G05	Legal obligations and processing roles	Legal/Privacy	Applicability register
G06	Conflicts incentives and honest claims	Managing Principal	Declarations; approved claims
G07	Approved tools and user training	Governance Lead	Tool list; training record
D01	Data rights purpose and provenance	Data Lead	Data register; client authority
D02	Minimisation sensitive data and impact review	Privacy Lead	DPIA/impact record
D03	Client segregation and retrieval access	Security Lead	Isolation test; access logs
D04	Vendor training retention and transfers	Privacy Lead	Terms; configuration; transfer record
D05	Retention rights requests and deletion	Data Lead	Schedule; deletion/rights log
S01	Identity secrets and least privilege	Security Lead	Access review; configuration
S02	Threat testing and secure outputs	Security Lead	Threat model; test findings
S03	Agent boundaries transaction approval and stop	Technical Lead	Allowlist; approval/stop test

For every control record add: applicability; assigned person; implementation status; evidence identifier/link; scope and period; test method/sample; reviewer; result/findings; action owner/date; next review; disclosure class. Retain version history.

Appendix A Control and evidence register continued

ID	Control requirement	Role owner	Evidence
V01	Supplier and subprocessor diligence	Commercial Lead	Due-diligence file
V02	Contract rights incidents and exit terms	Legal Reviewer	Executed schedules; gap record
Q01	Representative independent evaluation	Validation Lead	Test plan; data; results
Q02	Bias accessibility and human oversight	System Owner	Impact and subgroup review
Q03	Source code and numeric verification	Technical Lead	Review record; benchmarks
Q04	Versioned model/system documentation	Technical Lead	Model card; configuration trace
L01	Gate and client release approval	Release Owner	Signed release; client acceptance
L02	Disclosure complaints and appeal	Engagement Owner	Notice; human escalation route
O01	Performance drift and outcome review	System Owner	Monitoring report; action log
O02	Material change and revalidation	Technical Lead	Change record; regression results
O03	Incident containment notice and remediation	Incident Lead	Incident file; notices; retest
O04	Fallback continuity and retirement	System Owner	Recovery test; closure record
A01	Scoped evidence and assurance statements	Governance Lead	Evidence index; verified statement
A02	Independent control testing and findings	Independent Reviewer	Review report; action follow-up
A03	Management review and improvements	Managing Principal	Minutes; KPI/action record

Evidence handling

Use a secure evidence index rather than copying all records into questionnaires. A recommended artifact identifier is entity / engagement / system / control / version / date. Mark each artifact Public, Share under NDA or Restricted. Redact client data and record authorised recipients. Test access rights and record deletion or legal-hold decisions.

Operating evidence must show who performed a control, when, on which population/version and with what result. A screenshot of a setting alone cannot establish that the setting applied throughout the claimed period.

Appendix B AI use case intake form

Complete before piloting or procuring a new use. The Governance Lead assigns the system identifier and initial tier; missing mandatory information prevents progression to G1.

Identity and purpose

Entity / engagement / system ID / owner / date: _____

Business problem and measurable intended benefit: _____

Non-AI alternative considered and reason for AI: _____

Intended users and affected people: _____

Client sponsor and permitted use boundary: _____

Technology and information

Vendor / model / version / hosting / integrations: _____

Data sources / classes / rights / retention / countries: _____

Controller/processor roles and client permission evidence: _____

Training reuse / human vendor review / subprocessors: _____

Tools and actions including write permissions or financial limits: _____

Risk screening

Does use affect employment, credit, insurance, public benefits, safety or rights? Explain:

Sensitive data, vulnerable groups, overseas transfers or broad autonomy? Explain:

Foreseeable failure or misuse and who may be harmed: _____

Internal tier / legal classification / specialist reviewers: _____

Intake outcome and restrictions: Register / Reject / Further evidence

Decision owner / signature / date / next gate: _____

Appendix C Impact and risk assessment template

Use for High-tier cases and other elevated risks. Link, but do not conflate, the AI impact assessment, privacy impact assessment, security threat assessment and any legally required fundamental-rights assessment.

Context and affected interests

Purpose, decision process and deployment boundary: _____

Affected groups, benefits, adverse impacts and consultation: _____

Legal roles, obligations, data basis and transfer analysis references: _____

Alternatives and necessity/proportionality reasoning: _____

Repeat this record for each material harm scenario

Scenario ID / cause / event / affected party / consequence: _____

Inherent impact 1 to 5 / likelihood 1 to 5 / score / basis: _____

Controls and test evidence that reduce exposure: _____

Residual impact / likelihood / score / uncertainty: _____

Mitigation action / owner / due date / verification method: _____

Human and societal safeguards

Fairness metrics, groups, lawful testing basis and limitations: _____

Human competence, review capacity, override and stop authority: _____

Notices, explanations, correction/appeal and accessible alternative: _____

Workforce, public-interest, concentration and sustainability effects: _____

Decision

Highest material residual scenario and tier floor: _____

Reviewer findings / unresolved blockers / conditions: _____

Accept within boundary / Remediate / Reject / Suspend: _____

System Owner / specialists / independent reviewer / authority / date: _____

Expiry and triggers for reassessment: _____

Appendix D System and model card template

Prepare a client-readable card and a restricted technical annex. The card describes the deployed application, not merely the base model. Keep it current and link it to the approved release.

Identity and intended use

System ID / name / entity / owner / release date: _____

Purpose, intended users and expressly excluded uses: _____

AI component, provider, model/version and relevant configuration: _____

Internal tier and legal actor/classification assessment: _____

Data and operating boundary

Data sources, provenance, update cycle and permitted geography: _____

Personal/confidential data handling, retrieval and training restrictions: _____

Connectors, actions, approvals, permissions and transaction limits: _____

Known context limits, unsupported languages/products and assumptions: _____

Validation and transparency

Test period / dataset scope / sample sizes / baseline / thresholds: _____

Quality, fairness and security results with limitations: _____

Independent review reference and remaining findings: _____

User explanation, uncertainty/abstention and AI disclosure: _____

Human review, correction, appeal and accountable contact: _____

Operation and lifecycle

Monitoring, review cadence, alerts and suspension criteria: _____

Fallback, service dependency, incident route and recovery test: _____

Retention/deletion schedule and client exit/handover: _____

Change log and approval reference: _____

Card owner / reviewer / next review date: _____

Appendix E Release and exception records

Release decision form

System / version / entity / client / release boundary: _____

Gate evidence index and legal/data/supplier clearances: _____

Validation findings, critical tests and independent review: _____

Human oversight, notices, training and monitoring readiness: _____

Fallback/stop test and incident contacts: _____

Highest residual risk / restrictions / acceptance expiry: _____

Client acceptance authority and reference: _____

Decision: Approve / Approve bounded pilot / Remediate / Reject

System Owner and release authority signatures / date: _____

Time-limited exception form

Exception ID / affected control / use case / request owner: _____

Reason, alternatives and scope of deviation: _____

Legal assessment and non-waivable gates checked: _____

Compensating controls, test evidence and residual harm scenarios: _____

Remediation owner / action / due date / expiry / review frequency: _____

Specialist clearances and Managing Principal decision: _____

Client notice/acceptance if required and evidence reference: _____

Conditions, suspension trigger and closure verification: _____

No exception is valid beyond its expiry. Prohibited uses, legal duties and failed mandatory release gates are not eligible for exception. Repeated renewals require governance review of the underlying operating model.

Appendix F Supplier and incident templates

Supplier due diligence record

Supplier legal name / service / AI functions / owner / date: _____

Models, subprocessors, countries and change notice: _____

Data use, training, retention, human access and deletion terms: _____

Security/audit evidence with scope, dates and exclusions: _____

IP, liability, insurance, incident notice and audit/evidence rights: _____

Exit, export, business continuity and concentration assessment: _____

Open gaps, use restrictions, decision and next review: _____

Incident and complaint record

Incident ID / awareness date and time / reporter / system/version: _____

Severity / actual and potential harm / people/data/client scope: _____

Containment actions / owner / times / evidence preservation: _____

Applicable legal/client notices / authority / deadlines / decision: _____

Notifications made / factual content / recipient / time: _____

Root cause, supplier involvement and corrective actions: _____

Affected-output/decision correction and client/individual remedy: _____

Retest, independent challenge, restoration authority and date: _____

Lessons, action closure and recurrence monitoring: _____

Treat these forms as controlled records. Store sensitive details in a restricted case file and share only approved summaries with partners or clients.

Appendix G Assurance questionnaire response guide

For each answer state Verified / Partial / Planned / Not applicable / Not assessed; evidence reference; legal entity; system/service scope; period; gaps and action date. The responses below describe what to substantiate, not predetermined Yes answers.

Client or partner question	Evidence needed to support the response
Who owns AI governance?	Signed role/delegation record and escalation contacts.
Is client data used for model training?	Service-specific contract and settings; any authorised reuse scope.
Where is data processed?	Data flow supplier/subprocessor locations and transfer review.
How are client environments separated?	Architecture access policy retrieval/tenant isolation tests.
How is accuracy validated?	Use-case tests baseline sample sizes thresholds and reviewer.
How are bias and adverse decisions addressed?	Impact/subgroup review human oversight and appeal evidence.
Can agents act autonomously?	Action allowlist approval boundaries transaction limits and logs.
How are incidents handled?	Response procedure escalation exercise and scoped event history.
Do you hold ISO or other certifications?	Verifiable certificate entity scope dates and issuer; otherwise state no verified claim.
Can you provide independent assurance?	Unconflicted reviewer scope method period findings and limitations.

Proposed assurance wording after verification

For the stated engagement and assessment period, Finaffairs has reviewed the controls listed in the attached evidence schedule. Each finding identifies the control scope, evidence, reviewer, result and any limitation or remediation. The statement covers the specified service and approved use boundary; it does not certify all Finaffairs activities or guarantee that AI output is error-free.

Use this wording only after the described review has occurred and the statement has been approved. Before then, provide the policy as a proposed framework with a dated implementation plan.

References and interpretation

Primary sources reviewed on 11 October 2026. The operational requirements, approval thresholds, forms and proposed service targets in this framework are Finaffairs policy design choices. They do not reproduce a complete standard or replace an engagement-specific legal assessment.

R1 NIST AI Risk Management Framework and Generative AI Profile

<https://www.nist.gov/itl/ai-risk-management-framework>

Voluntary framework reference; revision in progress; profile link.

R2 NIST AI RMF Core

<https://airc.nist.gov/airmf-resources/airmf/5-sec-core/>

Govern Map Measure Manage; governance across the lifecycle.

R3 ISO IEC 42001 2023 AI management systems

<https://www.iso.org/standard/42001>

Public standard overview; no claim of complete clause mapping.

R4 European Commission AI Act information

<https://digital-strategy.ec.europa.eu/en/policies/regulatory-framework-ai>

Scope roles and changing commencement schedule; verify enacted text.

R5 ICO guidance on AI and data protection

<https://ico.org.uk/for-organisations/uk-gdpr-guidance-and-resources/artificial-intelligence/guidance-on-ai-and-data-protection/>

Guidance under review after UK legislative changes.

R6 Nigeria Data Protection Act and NDPC GAID

<https://ndpc.gov.ng/>

Official regulator and directives index; obtain current full legal texts.

Use the current Nigeria Act and NDPC directives when completing the applicability register. External guidance and proposed amendments must be checked against enacted law and current regulator instructions.

References continued and glossary

R7 Ghana Data Protection Act 2012 Act 843

<https://dataprotection.org.gh/wp-content/uploads/2025/05/Data-Protection-Act-2012-Act-843.pdf>

Official Act; privacy registration security and decision rights.

R8 OWASP Top 10 for LLM and GenAI

<https://genai.owasp.org/initiatives/top-10-for-llm-and-genai/>

Threat reference; record the edition selected for each assessment.

R9 Canadian OPC PIPEDA accountability guidance

https://www.priv.gc.ca/en/privacy-topics/privacy-laws-in-canada/the-personal-information-protection-and-electronic-documents-act-pipeda/p_principle/principles/p_accountability/

Accountability vendor protection and cross-border handling.

R10 FTC AI privacy and confidentiality commitments

<https://www.ftc.gov/policy/advocacy-research/tech-at-ftc/2024/01/ai-companies-uphold-your-privacy-confidentiality-commitments>

Substantiated privacy and confidentiality representations.

R11 OECD AI principles

<https://www.oecd.org/en/topics/ai-principles.html>

Values reference; updated principles.

R12 ISO IEC 42005 2025 AI system impact assessment

<https://www.iso.org/standard/42005>

Impact-assessment reference; use current licensed text where needed.

Glossary. AI System Owner is the accountable person for the deployed use case. A model is one component of that system. DPIA means data protection impact assessment. Residual risk is the remaining exposure after evidenced controls. Independent review requires separation from the activity reviewed. Release is permission to operate within a recorded boundary. Assurance is a scoped conclusion supported by stated evidence.