

RAIT.

Responsible AI Transfer Framework:

A Framework for Governing Public-Sector AI Acquisition

By: [Samuel Abinsinguza](#)

Founder, [CentPol](#) | Center for Emerging & Next Technologies. - Policy, Strategy & Foresight.

Abstract

As artificial intelligence moves across borders and into public institutions, governments face the challenge of turning technological promise into lasting public benefit. Systems acquired to improve services, support scientific discovery and advance economic development also bring responsibilities: to understand how they operate, protect the people they affect and retain meaningful control over their use. Acquiring the technology does not automatically provide the expertise, resources or authority needed to meet those responsibilities. The central challenge of responsible AI transfer is to build that capacity into the acquisition itself and sustain it throughout the system's use.

The Responsible AI Transfer (RAIT) Framework links the scope of an AI deployment to evidence that the responsible institutions can govern its risks. Six contractual modules connect capacity development, human rights assessment, data governance, auditability, workforce development and exit arrangements. Four assurance stages carry these obligations through assessment, testing, deployment and continued operation, while monitoring examines whether safeguards protect people in practice.

RAIT adapts assurance to the consequences of each use while maintaining essential protections and asking what capability remains with the institution after acquisition. Drawing on documentary evidence and existing governance instruments, this paper presents a proposed framework and an illustrative public-benefits application. Its usefulness, costs and contribution to lasting institutional capability require pilot evaluation.

Central proposition

An institution should acquire the means to govern a system as part of acquiring the system itself. For each consequential use, it should be possible to identify the responsible authority,

the required safeguard, the evidence that the safeguard works, its funding source and the action taken when it fails.

1. The problem and the scope of RAIT

In 2019, it emerged that the Dutch tax authority had spent years running a self-learning algorithm to build risk profiles in childcare-benefit fraud detection. By the time parliamentary inquiry forced the system offline, more than twenty thousand families had been wrongly accused and pushed into debt, with immigrant-background families disproportionately affected. The government resigned. [\[7\]](#)

In Kampala around 2019, a foreign-supplied "smart city" package arrived with facial recognition, licence-plate readers and policing dashboards. No dedicated law governed how the footage would be used, and no independent authority had reviewed the contracts. [\[12\]](#)

Although these cases differ in their causes and institutional settings, they raise a common practical question: who can investigate, challenge and correct a system once it becomes part of ordinary administration? A ministry may have a contract and trained users while lacking the records needed to explain a decision; a regulator may have formal powers but insufficient technical support; and a person seeking correction may encounter several agencies, each responsible for only part of the process.

RAIT organizes this work around the acquisition relationship, where buyers, suppliers, integrators and funders can agree on information access, oversight responsibilities, recurring costs and exit arrangements. These decisions matter wherever an institution relies on externally supplied technology, including domestic suppliers and systems built with open-source components.

1.1 Intended users and uses

The primary users are public-sector implementing teams, procurement authorities, sector regulators, data protection bodies, development partners and independent reviewers. RAIT can inform an acquisition plan, procurement specification, implementation agreement, assurance review or financing arrangement. Its requirements must be adapted to the applicable law and the authority of the organizations involved.

For this paper, responsible transfer includes the acquisition, licensing, integration or substantial adaptation of an AI component for institutional use. A cross-border transaction is

a common setting, not a necessary condition. The relevant unit is the deployment: a defined system, purpose, population, data environment and institutional arrangement. A country does not receive a permanent readiness grade authorizing every ministry or supplier to deploy high-risk systems.

1.2 Connection to digital public infrastructure

Assessing an AI deployment also means examining the infrastructure it relies on. Digital public infrastructure (DPI) provides shared capabilities – such as identity verification, payments and data exchange – that connect public services. An AI application may draw on these capabilities without itself becoming part of that shared infrastructure, yet its use can change how people are identified, assessed or served, and how they challenge a decision. RAIT therefore considers the application together with the infrastructure and institutional responsibilities on which it depends.

The Universal Digital Public Infrastructure (DPI) Safeguards Framework addresses responsible authorities and protections across the DPI lifecycle, including exclusion, accountability and redress. RAIT is intended as a complementary implementation approach for AI acquisition and integration, rather than a replacement framework or a claim of formal alignment certification. [\[1\]](#)

1.3 Boundaries

RAIT applies to a defined institutional acquisition and deployment. General platform governance, military AI and many private-sector uses require other instruments; autonomous weapons also raise distinct questions of international humanitarian law. A public institution's procurement of a platform service can nevertheless fall within RAIT when its purpose, responsibilities and contractual relationship can be defined.

Open-source licensing does not remove these responsibilities. Where a paid developer, host or integrator supplies an open-source system, the relevant obligations belong in that agreement. Where an institution deploys directly, it must assign and fund the documentation, evaluation, training and continuity work internally or through separately procured support. In either case, applicable legal prohibitions remain in force: a use that is unlawful, unnecessary or incapable of adequate safeguards should not proceed.

2. Evidence, regional context and method

The evidence informing RAIT points to several ways public-sector automation can fail, including unlawful policy, poor data, weak review arrangements and institutional incentives. Acquisition agreements can address some of these problems by securing information, resources and cooperation; others require legislation, administrative reform or a decision not to automate. The discussion below uses these distinctions to connect documented failures with the capabilities institutions need.

2.1 What the documented record shows

The cases in this paper illustrate governance problems rather than establish a single procurement-related cause. RAIT's proposed responses are design inferences drawn from that record; the method and its limits are set out in §2.4.

Automated decisions in public benefits. Australia's Royal Commission into the Robodebt Scheme recommended clearer routes to review automated decisions, public explanations of how automation works, access to business rules and algorithms for independent scrutiny, and stronger governance of inter-agency data exchanges. [\[6\]](#) In the Netherlands, a risk-profiling algorithm operated for years inside a benefits process before its logic became publicly known. [\[7\]](#) In Denmark, Amnesty International's 2024 review of AI-supported welfare fraud detection identified discrimination risks for marginalized groups, discriminatory design features in particular models, and an atmosphere of surveillance and fear among claimants. [\[8\]](#)

Together, the benefits cases show why established legal institutions need operational support to make oversight effective. The Robodebt recommendations emphasize review, independent scrutiny and inter-agency data governance, while the Dutch and Danish cases raise concerns about transparency and contestability. RAIT responds by making the relevant responsibilities, resources and tests explicit before a service relies on the system.

Surveillance and biometric deployment. Research on the international trade in facial recognition and related surveillance systems has examined patterns in which importing states with weaker democratic accountability acquire such systems, and has raised concerns about the relationship between those acquisitions and subsequent restrictions on political activity. [\[13\]](#) Analysis of biometric expansion across several African states has identified rapid

deployment ahead of governing legal frameworks, with consequent exposure to privacy violations and data breaches. [\[14\]](#) These sources identify patterns and institutional concerns rather than establish that a particular acquisition caused a particular violation.

The Myanmar and Ethiopia platform cases offer related lessons about algorithmic amplification and local review capacity, but require platform-governance instruments beyond the acquisition focus of this paper. [\[9\]](#) [\[10\]](#) [\[11\]](#)

2.2 RAIT's contribution alongside existing instruments

RAIT builds on an existing body of procurement guidance, contractual clauses and AI governance frameworks. Its proposed contribution is to connect their safeguards with committed oversight resources, demonstrated institutional capability and decisions about deployment and continued operation.

The UK Guidelines for AI procurement already address impact assessment, oversight, contractual provisions, knowledge transfer, supplier dependence and continuing evaluation. EU AI model contractual clauses provide full and lighter approaches for public buyers. RAIT uses these as close comparators, bringing funded responsibilities, observed capability tests, deployment authorization and capability retained across acquisitions into one process. Appendix A maps the relationships in more detail. [\[25\]](#) [\[26\]](#)

Alongside these procurement tools, research on AI regulation highlights the importance of specifying the outcomes governance should achieve. In its 2023 analysis of more than 1,600 policy initiatives across 69 countries and the EU, Deloitte reported that about one per cent of regulations were outcome-based or risk-weighted, with none in its dataset combining both approaches. [\[16\]](#) Although the study also acknowledged protections in existing laws beyond AI-specific regulation, its findings provide useful context for RAIT's emphasis on demonstrable safeguards: what an institution must be able to do, how that capability is assessed and what happens when it falls short.

The technology-transfer tradition adds a second concern: what expertise and control the acquiring institution gains through the transaction. UNCTAD's Draft International Code of Conduct on the Transfer of Technology addresses intellectual property, licensing and technical training. RAIT adapts these concerns to operational oversight, access to remedy and the ability to maintain a service beyond a particular supplier relationship. [\[17\]](#)

2.3 Uganda and Kenya: build on existing protections

Uganda and Kenya illustrate why an assessment of governance capacity should begin with the protections and institutions already in place, before examining how they operate within a particular service.

Uganda's Personal Data Protection Office was operationalized in 2021. In 2022, the Office and UNCDF launched a portal supporting complaints, breach reporting and registration, including SMS/USSD functionality. The portal offers an institutional foundation for further safeguards, whose effectiveness must be assessed through complaint handling and enforcement in practice. [\[18\]](#)

Uganda's 2019 law also addresses decisions based solely on automated processing. [\[21\]](#) A comparative analysis by CIPIT describes reconsideration protections and differences between the Ugandan and Kenyan regimes. Kenya's Data Protection Act contains provisions concerning impact assessment, automated decisions, privacy by design and complaints. These provisions provide a legal foundation for assessing how safeguards operate in a specific deployment. [\[19\]](#) [\[20\]](#)

RAIT would examine review staffing, audit access, complaint handling, lawful data sharing and record correction with the responsible authorities and affected communities. This makes the assessment specific to an institution and service, rather than a national readiness grade.

2.4 Method and evidentiary limits

The paper combines targeted documentary synthesis with normative framework design. Public safeguards frameworks, risk-management guidance, legislation, institutional publications, civil society investigations and an official inquiry inform specific design questions. The cases are illustrative rather than a systematic sample, and the analysis does not establish that different acquisition terms would have prevented the harms described.

The modules, assurance stages, indicators and worked examples are the author's proposed synthesis. They have not been field-validated or adopted as RAIT by the institutions cited. Numerical examples are illustrative, and the framework's costs, usefulness and effects require

pilot evaluation. This distinction between documented evidence and proposed design underpins the principles that follow.

3. Governing principles and risk assessment

RAIT's six commitments – public accountability, respect for human rights, local agency, proportionate assurance, shared responsibility and sustained capacity, guide decisions about both whether and how to deploy AI. They establish protections that must be present from the start, while allowing the evidence burden to reflect the consequences of the use.

3.1 Essential protections throughout deployment

Before a system affects people, the implementing authority must establish a lawful purpose, a responsible decision-maker and appropriate controls over personal data. It must explain the system's role, provide usable channels for concerns and maintain a means of correcting consequential errors. These obligations apply to pilots and experiments as well as established services, including testing that processes personal data.

For consequential decisions, human review must give reviewers the records, time, competence and authority needed to change an outcome. A person placed in a workflow cannot provide meaningful oversight without these conditions. People affected by the service must therefore have access to remedy from the outset, supported by a review process that can investigate and respond.

3.2 Assess the use in its setting

The assessment should document the consequence of error, the affected population, the sensitivity and reliability of data, the opportunity to contest a decision and the degree of reliance on the system. It should examine cumulative harm, misuse, foreseeable repurposing, infrastructure failure and the ability to recover. A small deployment can still create serious harm; a large deployment can make a modest error rate consequential.

A language tool used to draft internal notices has a different risk profile from the same tool explaining legal entitlements without review. A matching model used to flag possible

duplicate records becomes more consequential if its score automatically suspends a payment. Risk classification must therefore follow actual authority and operational effects, including delays and informal reliance by staff.

Even apparently advisory uses require this contextual assessment. A crop-disease model may be a lower consequence where farmers can verify its suggestions and recover from mistakes, yet pose serious livelihood risks where they depend on it. Using the same output to decline credit or influence a contested land claim changes the consequences again.

NIST's AI Risk Management Framework similarly treats context, organizational responsibilities and third-party dependencies as part of risk management. RAIT carries that approach into procurement evidence and deployment decisions while retaining technical evaluation and any applicable statutory classifications. [\[2\]](#)

3.3 Local agency and shared obligations

Recipient institutions and affected communities must help define priorities, acceptable service arrangements and review procedures, including the option to narrow a deployment, choose a simpler system or reject an acquisition. Suppliers must disclose limitations, support testing and fulfil capacity-building obligations, while funders should resource the oversight their projects require.

Building this capacity should make beneficial adoption more feasible. Requirements imposed by wealthy exporters should therefore be accompanied by support for the safeguards demanded, with environmental resource use, labour conditions and participation costs assessed where material. These commitments become operational through decisions about whether a proposed use is justified, ready for testing and suitable for public service.

4. Four stages of deployment assurance

RAIT uses four stages to connect acquisition decisions with evidence of readiness and continuing oversight. Each stage builds on the previous one, but approval remains specific to the purpose, system version, data dependencies, permitted users and review date. As conditions change, the deployment may be reassessed, reduced or paused.

4.1 Proportionate assurance for low-consequence uses

The evidence needed to govern a system should reflect the consequences of its use. A proportionate track therefore simplifies the four stages for eligible low-consequence applications while preserving all the protections in §3.1. This approach draws on the idea of minimum viable governance in the enterprise literature, adapted to the responsibilities of public institutions. [\[24\]](#)

Eligibility depends on actual use. The track is unavailable where an output influences a person's access to a service, entitlement, liberty, treatment or legal status, including through advice to a decision-maker. It is also unavailable for systems operating on shared identity or population-scale infrastructure, or wherever law requires stronger assurance. Processing personal data requires review by the responsible data protection function rather than self-assessment alone.

For an eligible use, Stage 1 records the purpose, legal basis, data assessment and eligibility reasoning, with classification approved by the designated authorizing authority. Stage 2 provides a short record demonstrating the relevant minimum controls. Stage 3 produces a signed approval specifying the use, responsible officer, review date and stop conditions. Stage 4 maintains the register of changes, complaints and reassessment decisions. A full impact assessment, independent assurance review, shadow evaluation and the complete module set are not automatically required.

The classification remains open to challenge and must be reviewed when actual use changes. A drafting assistant can acquire an advisory role, and staff may begin relying on an output in ways the original approval did not anticipate. Periodic sampling of the low-consequence register should test for such changes. Proportionate assurance is justified by the consequences assessed in advance; delivery pressure or an inadequate budget cannot justify lowering essential protections.

4.2 From assessment to continuing operation

Stage 1. Establish purpose, authority and feasibility

The first stage establishes the service problem, legal basis and people likely to be affected, then maps the institutions responsible for decisions, data correction, security and remedy. The assessment identifies essential safeguards, material uncertainties and the resources needed to resolve them.

The acquisition brief compares AI adoption with improvements to the existing service and continuation of current practice. Each option is assessed for service quality, staff and user burdens, oversight costs, dependence and opportunities to build capability. Better records, simpler rules, staff development and shared expertise may offer useful alternatives, while controlled shadow evaluation can support learning without operational deployment.

Alongside this comparison, the institution records what it can currently do without supplier assistance: reconstruct a decision, reproduce an evaluation, investigate an inter-agency data mismatch, resolve a complaint and operate the service without the proposed system. This capability baseline provides the factual starting point for reassessment at Stage 4.

The resulting acquisition and assessment brief may authorize research, procurement preparation or controlled testing, but cannot authorize consequential decisions about people. An unlawful use, unjustifiable data basis or inability to support essential safeguards requires redesign or a decision to stop.

Stage 2. Test the system and the safeguards

Testing establishes whether the proposed system and the institution's safeguards work together in the local setting. It examines representative data and user needs, security, error patterns and staff behaviour, alongside complaint channels, correction, logging and fallback arrangements. Where feasible, shadow evaluation compares outputs with an independently reviewed baseline while keeping them separate from service decisions; lawful processing and data protection remain necessary.

The evaluation report, updated impact assessment and remediation record support the next decision. Unlawful uses, failures of minimum protections and risks above documented acceptance tolerances block progression. Other residual risks may proceed to Stage 3 only with evidence, mitigations, monitoring and an accountable acceptance decision. Supplier benchmarks and training completion can contribute evidence, but cannot establish operational readiness on their own.

Stage 3. Authorize a bounded deployment

Once the required legal and regulatory reviews are complete, the responsible authority may approve a defined use with operating limits, performance tolerances, reporting duties and stop conditions. Staff must demonstrate that they can use the controls, with recurring funding, required independent scrutiny and accessible remedy in place before the service relies on the system.

A signed deployment decision brings this evidence together and identifies who accepts residual risks, what remains under observation and when approval expires or must be reconsidered. Contractual acceptance of a supplier's deliverables remains separate from the public authority's decision to put the system into service.

Stage 4. Sustain, reassess and exit

Approval is sustained through monitoring, incident response, supplier change control and effective remedy. Material changes to the purpose, model, data source, population or integration trigger reassessment proportionate to their effects; expansion requires evidence that the institution can manage the additional risks.

Reassessment also returns to the Stage 1 capability baseline. Improving service indicators can conceal a growing dependence on the supplier if the institution becomes less able to investigate decisions, evaluate performance or operate an alternative process. Any decline belongs in the assurance record, together with a response such as renegotiation, additional internal capacity or a change to the arrangement.

The continuing assurance record combines review decisions, incidents, complaints, corrections, supplier performance, continuity tests and the revised baseline. It supports decisions to continue, restrict, return to an earlier stage or retire the use. When exit is necessary, the transition must protect affected people and preserve records needed for review.

4.3 Building capability across acquisitions

Beyond governing an individual deployment, RAIT aims to help institutions retain the knowledge and practical capabilities needed for future acquisitions. Its governability frontier is the boundary between uses it can oversee with adequate protection and uses whose consequences exceed its current capacity. That boundary belongs to an institution, service and moment, rather than to a country as a whole.

The frontier can move as institutions retain tested contract terms, experienced reviewers, correction pathways and evidence of oversight costs and failures. This requires stable staffing, usable documentation, repeated exercises and ownership of the resulting knowledge. Controlled evaluation, shared expertise and procurement preparation can also build these assets; pilots must establish how much lasting capability deployment itself contributes.

On this view, progress includes the range of beneficial uses an institution can responsibly govern, adapt and eventually specify for itself. An acquisition that replaces scarce expertise without transferring knowledge may improve immediate performance while deepening dependence. This is why the baseline is reassessed and why a contractual exit right must be supported by the practical ability to maintain service.

Where several acquisitions are planned, their sequence should reflect what each can help the institution learn. A recoverable use within existing oversight capacity may provide experience for a more consequential one, while beginning with the harder use requires a specific justification. The contract modules that follow secure the information, assistance and resources on which these assurance decisions depend.

5. Contract modules: making governance deliverable

The six modules translate RAIT's assurance requirements into obligations that buyers and suppliers can deliver and verify. Each must be adapted to the governing law and identify the responsible party, deliverable, deadline or trigger, acceptance evidence, funding and consequence of non-performance. Together they connect the public institution's continuing duties with the support it acquires.

Module 1. Governance capacity-building plan

An institution needs more than trained users to govern a consequential system. Legal review, procurement oversight, evaluation, security operations, frontline supervision, correction and complaints all require people with defined responsibilities. The capacity plan sets out this work and separates public duties from assistance supplied by the vendor or an independent provider; accepting a contractual task does not confer public decision-making authority.

Competence should be demonstrated through the tasks the service requires, such as reproducing an evaluation, investigating an incident, retrieving a decision record and operating a fallback without supplier intervention. Acceptance depends on these exercises and the resources available for continued operation.

The plan should state how it will affect the Stage 1 capability baseline and what will remain after the supplier relationship ends. Supplier-employed staff, temporary evaluation support and documentation the institution cannot maintain may be useful, but should not be counted as retained internal capability. Any decision not to improve the baseline needs an explicit justification. Relevant payments should follow verified deliverables, with continuity protected during disputes.

Module 2. Human rights and service impact assessment

The effects of an AI system extend through the service pathway, from the data collected to the remedy available after a decision. The impact assessment therefore examines necessity, proportionality, discrimination, exclusion, accessibility and privacy alongside the ability to obtain redress. It should explain whose experience informed the assessment, how concerns changed the design and which risks remain beyond adequate mitigation.

The public authority is responsible for commissioning a credible assessment and responding to its findings. High-consequence uses require independent review of assumptions and mitigations, while an accessible public account should explain the purpose, risks, decisions and remedies without exposing personal data or sensitive security details. Material changes or evidence of harm trigger reassessment. UNESCO's recommendation provides a foundation for this continuing human responsibility. [\[4\]](#)

Module 3. Data governance and institutional control

When several agencies rely on the same record, an error can continue to affect a person after the originating agency has corrected it. The agreement therefore needs to establish how information moves between institutions, who can change it and how receiving systems confirm that a correction has taken effect. It must also identify controllers, processors, permitted purposes, lawful access, storage, retention and deletion, supported by source provenance and known quality limitations.

These responsibilities extend to cross-border processing, onward transfers and subcontractors, as well as unauthorized reuse, model training and access revocation. Localization decisions must follow applicable law and a documented risk assessment: keeping data in-country does not by itself prevent misuse, and institutional control cannot become an unlimited ownership claim over personal information.

A data-flow map, access-control configuration, retention schedule, correction procedure and inter-agency agreements make these arrangements reviewable. Acceptance tests should demonstrate that users cannot access data beyond their roles and that corrections and access withdrawals are respected throughout the integration.

Module 4. Transparency, auditability and change control

To investigate a failure, the buyer and authorized independent reviewers need evidence of how the system operated in the setting under review. The agreement should secure system documentation, evaluation results, limitations, version histories, relevant decision logs and subcontractor information, together with audit access, testing permissions, cooperation and record retention.

Source code may be necessary in some settings, but reviewers may also need reproducible test environments, configuration records, input and output samples, data documentation and access to relevant personnel. Intellectual property arrangements should permit essential scrutiny while protecting sensitive material from indiscriminate disclosure. If supplier restrictions make that scrutiny impossible, the use must be redesigned or rejected.

The same access must continue as the system changes. Material changes require notification and assessment before production use, with a defined exception process for emergency security updates, prompt documentation and subsequent review. Incident provisions should establish severity categories, reporting triggers, evidence preservation and lawful notification duties. Public summaries can disclose material findings while protecting personal records and security.

Module 5. Technical training and workforce development

Training becomes useful when people can apply it to the responsibilities they hold. Operators, engineers, procurement teams, regulators and complaint handlers need different levels of technical understanding and practical authority, so training should be linked to demonstrated tasks, refresher needs, staff turnover and accessible local-language materials.

University partnerships, supervised placements, fellowships and local maintenance arrangements can support this development where demand and funding justify them. Commitments should identify the curriculum, participants or funded places, delivery dates and acceptance tests, with completion, competence, retention and career progression monitored separately. Employment outcomes must be evaluated rather than assumed.

Where labour-market development is material to an acquisition, the agreement may also specify funded placements, maintenance capability or preferential consideration for locally developed components. Participation should address actual barriers such as inaccessible materials, transport costs and unpaid work. Fair conditions and support for workers handling harmful content or sensitive data belong in the agreement where relevant; demographic targets alone cannot establish accessibility.

Module 6. Exit, transition and service continuity

A public institution needs a workable alternative when a system no longer serves its purpose. Before deployment, the agreement should therefore specify how it can stop using the AI component, change suppliers or maintain the service through a simpler process. Usable data and configuration exports, interface documentation, knowledge transfer and transition assistance must be supported by cost and time estimates, lawful retention and deletion of supplier-held copies.

Portability and fallback should be tested while there is time to repair failures. The transition must preserve pending applications, unresolved complaints and evidence of past decisions, with termination fees and licensing restrictions reviewed for their effect on practical choice. Open-source licensing can reduce some dependencies while leaving hosting, specialist skills or integration support essential. Making these obligations effective requires the clear allocation of authority and scrutiny described next.

6. Institutional responsibilities and verification

Contracts can assign deliverables, but public accountability depends on who operates the service, who authorizes it and who can scrutinize the evidence independently. The implementing authority retains the public-service decision and ensures controls are resourced; technical operators manage integration, data-holding agencies remain responsible for their records, and suppliers fulfil their agreements. These relationships must be mapped before the contract is finalized, alongside the existing powers of regulators and courts.

6.1 A review arrangement with a clear mandate

A project or sector assurance panel can bring legal, technical, operational and community expertise into the review. Its membership should reflect the service and affected population, with a mandate to examine evidence, record dissent, identify unresolved risks and recommend authorization, conditions, reassessment or suspension.

Where the panel holds delegated decision authority, the instrument establishing it must define its scope, accountability and appeal arrangements. Otherwise its role is advisory. In either case, the review must complement statutory regulation and judicial scrutiny rather than displace them.

6.2 Independence and participation

Credible scrutiny depends on reviewers being able to reach and publish adverse findings. Transparent appointments, disclosure of interests and recusal rules should therefore be

supported by funding arrangements that prevent suppliers or sponsors from suppressing results. A supplier must not select the sole reviewer of its own work; pooled or independently administered funding can reduce direct dependence, provided its governance is also open to scrutiny.

Affected communities and civil society need accessible information, reasonable notice and opportunities to submit evidence and challenge assumptions, with participation costs supported where appropriate. Reasoned summaries should explain how concerns were addressed while protecting confidential complainants and personal data. Challenges to the assurance process require a channel independent of the original assessment team.

6.3 Test actual capability

Verification brings documents together with observed exercises: retrieving the record behind a disputed output, investigating a mismatch, processing a complaint, correcting data and operating without the AI component. Frontline workers and service users can identify gaps that records alone miss, while dated evidence tied to the assessed configuration allows reviewers to judge what has actually been demonstrated.

The approval record must state the decision, authority, scope, conditions, uncertainties and review triggers. Material failures require a documented response, which may include correction, enhanced monitoring, restriction or suspension. Because abrupt withdrawal can also harm people, essential services need continuity arrangements that remain usable during investigation.

Regional cooperation can make specialist evaluation and peer review more accessible by sharing fixed costs. Recognition of another institution's assessment should still depend on relevant evidence and comparable safeguards for the deployment concerned. Whatever review arrangement is chosen, its responsibilities need funding that lasts as long as the service relies on them.

7. Financing safeguards and creating incentives

Assigning responsibility makes oversight accountable; sustaining it requires a budget that reflects the work involved. RAIT therefore treats evaluation, complaints, maintenance,

independent scrutiny and exit as acquisition costs extending across the expected operating period, including duties that continue after supplier support or external funding ends.

7.1 Cost the work that must be performed

Costing begins with service workload and the functions needed to handle it. Initial assessment and integration testing create one-time costs, while access reviews, security monitoring, complaints, correction, reassessment and staff replacement recur. Accessible communication, infrastructure resilience, inter-agency correction and environmental resource use should be included where material.

Expected workload multiplied by the time and cost of each activity provides a starting estimate, to which fixed technical and institutional costs, contingency and transition provision must be added. Planning should allow for complex investigations, supervision, staff absence and peaks in demand, including an increase in complaints after an incident. A budget percentage cannot substitute for understanding this workload.

Each budget line should name an accountable institution and a committed funding source. Suppliers may finance documentation and training, the state may fund statutory decisions and remedy, and development partners may support evaluation or inclusion. Temporary contributions require a replacement plan, and supplier funding must preserve independent oversight.

7.2 Link financing to verifiable delivery

Payment milestones can reinforce these arrangements when they depend on accepted documentation, completed tests and demonstrated capability. Contractual remedies must also protect essential service continuity, so that a payment dispute or failed deliverable does not leave affected people without support.

The World Bank's Environmental and Social Framework provides an institutional precedent for risk-based oversight, stakeholder engagement, grievance mechanisms and borrower capacity building in applicable projects. RAIT proposes adapting this implementation discipline to recurring AI safeguards, subject to each project's governing requirements. [\[22\]](#)

7.3 Incentives and practical limits

Predictable evidence requirements may help suppliers reuse documentation and help public buyers compare offers and plan transitions. These gains must be weighed against longer procurements and the advantage that extensive documentation requirements can give larger firms. Common templates, proportionate requirements and independent technical assistance should be tested for their ability to reduce those burdens.

The immediate route is a lawful acquisition agreement backed by competent institutions and credible enforcement. Broader trade or export measures would require separate mandates and analysis, as discussed in Appendix B. For the acquisition itself, monitoring must establish whether the resources committed are producing effective protection.

8. Monitoring, evaluation and learning

Funding a safeguard and recording its existence are only the beginning of evaluation. Monitoring must establish whether it operates and what difference it makes to people: attendance at training may not produce competence, while a low complaint count may reflect barriers to reporting. An increase in incidents can even indicate better detection. These patterns need interpretation alongside service-user experience and independent evidence.

The following proposed measures connect operational records with the questions an implementing team needs to answer. Before deployment, each measure should have a baseline, denominator, data source, reporting frequency, responsible owner and escalation threshold. Disaggregated analysis requires lawful collection and careful handling of demographic information; the measures are starting points for adaptation rather than validated benchmarks.

Measure	Operational definition	Evidence and interpretation
Safeguard coverage	Active in-scope deployments with current required approval / all active in-scope deployments.	Register checked against procurement and service inventories. Missing systems can inflate apparent coverage.
Review timeliness	For completed complaints, median and upper-percentile time from receipt to reasoned decision; separately, number and age distribution of unresolved complaints.	Report completion times and pending backlog separately, with overdue cases and channels used. Excluding unfinished cases can conceal delays.
Correction completion	Validated corrections applied to all relevant downstream systems / all validated corrections due.	Source and receiving-system acknowledgements. Closing a ticket alone is insufficient.
Error and exclusion	Incorrect adverse outcomes in a reviewed sample / all adverse outcomes reviewed.	Independent reference review; assess errors among unflagged cases separately. Report uncertainty and subgroup differences.
Incident response	Time to detect, contain and remedy incidents, separated by severity.	Incident records checked against service failures and external reports. Fewer reports do not necessarily mean less harm.
Staff capability	Staff demonstrating required operational tasks / staff assigned those tasks.	Observed exercises and refresher assessments, including contractors and staff turnover.

Measure	Operational definition	Evidence and interpretation
Independence and exit	Completion and findings of independent reviews, portability tests and fallback exercises.	Review reports and observed drills. Report unresolved failures, not only completion rates.
Capability retention	Change in the Stage 1 capability baseline, by item, at each reassessment.	Observed exercises performed without supplier assistance. Distinguish capability held by the institution from capability available only under contract.

8.1 Turn findings into decisions

Operational measures should be read alongside service access, delay, administrative burden and people's ability to obtain help, including the experience of those who abandoned applications or never complained. Where possible, evaluation should compare the service with its previous operation and a credible non-AI alternative, accounting for changes in workload and eligibility rules. Attributing improvement to RAIT requires a design capable of separating its effects from those changes.

Periodic public reporting should explain findings and corrective action in accessible terms, using aggregation or suppression where necessary to protect small groups. Independent reviewers should check that unresolved problems are represented fairly. Each review must lead to a decision to continue, change, restrict or stop the use, with someone responsible for carrying it out. The benefits-service example below shows how this cycle connects acquisition to everyday administration.

9. Worked example: AI connected to a benefits service

Consider a hypothetical public benefits agency acquiring a model to flag possible duplicate applications. It verifies identity through shared infrastructure, exchanges limited records with another agency and uses an existing payment channel. The model sends cases for review while entitlement remains a public decision, allowing the example to show how even an advisory output can create consequential demands on staff and applicants.

9.1 Establish the service boundary

The acquisition brief defines the problem as duplicate applications requiring investigation. It compares AI-assisted review with simpler matching rules and improvements to record quality. It identifies potential harms: false matches, delayed payments, disproportionate scrutiny of particular communities, unauthorized data linkage and an inability to contest the flag. The assessment treats review queues and administrative delay as part of the model's consequences.

Only data necessary and lawfully available for the stated purpose may be accessed. The design should use limited verification responses where feasible rather than copy an entire identity database. A discrepancy in identity records must lead to an accessible correction route, not an automatic inference of fraud. The supplier cannot reuse service data to train other models without a separately justified and lawful arrangement.

9.2 Test before operational reliance

During shadow testing, outputs are isolated from staff making benefit decisions. Independent reviewers compare flags with assessed cases and also examine unflagged records to detect missed duplicates. Evaluation includes incomplete records, name variations, shared contact details, limited connectivity and relevant language needs. It measures subgroup effects where lawful and feasible, with attention to sample size and uncertainty.

Alongside these evaluations, reviewers must demonstrate that they can retrieve a record, explain a flag, correct an error and handle an appeal. They must be able to disregard the score and record their reasons, while technical checks establish that the model cannot issue payment instructions or alter authoritative identity records.

9.3 Authorize a bounded use

Any live pilot has a defined purpose, population, duration, model version and review capacity. A flag alone cannot terminate entitlement or suspend payment. The agency establishes the lawful procedure for investigating suspected duplicates and the protections applicable while a decision is disputed. Users receive a plain-language explanation of any consequential decision and a route to human review through accessible digital and non-digital channels.

Authorization depends on the agency's ability to investigate the expected queue fairly and promptly, supported by tested correction processes and committed funding. If demand exceeds that capacity, the agency must narrow the use, add resources or reject deployment. The review burden is part of the service's cost and cannot simply be transferred to applicants or frontline workers.

9.4 Record responsibilities and acceptance evidence

The agency owns eligibility decisions, notices and appeals; the identity authority owns corrections to its authoritative records; and the data-exchange operator controls authorized interfaces and access logs. The supplier provides model documentation, evaluation support and agreed updates, while regulators retain their statutory roles and an assurance reviewer examines the evidence independently.

Before authorization, six linked records bring these duties together: a lawful-purpose and alternatives assessment, a data-flow and correction map, an impact assessment, a version-specific evaluation report, a staffing and operating budget, and a signed authorization with review and stop conditions. The table illustrates how the same record connects each obligation to its owner, funding, test and response to failure.

Assurance field	Case review	Record correction	Exit and continuity
Obligation	Reconstruct and review a flag	Apply validated corrections downstream	Preserve service and review records

Assurance field	Case review	Record correction	Exit and continuity
Responsible party	Agency review lead; supplier supplies logs	Identity authority and receiving agencies	Agency service owner; supplier supports export
Funding	Agency operating budget; contracted logging	Agency data budgets; funded interface support	Costed transition reserve and supplier assistance
Acceptance test	Reviewer explains a sampled case; flag cannot stop payment	Receiving agencies acknowledge a test correction	Export and fallback drill preserve pending cases
Decision authority	Agency deployment authority; separate contract acceptance	Agency data owners confirm correction; deployment authority clears use	Service owner verifies continuity; authority approves transition
Failure response	Block authorization or restrict flagging; remedy affected cases	Contain affected data use; correct and retest	Delay cutover; maintain safe service and repair export

9.5 Illustrative acceptance provision

"The supplier shall demonstrate the contracted logging, documentation and correction interfaces against the agreed test protocol. Contractual acceptance of those deliverables does not authorize operational deployment. Before authorization, the implementing authority shall demonstrate that a reviewer can reconstruct a flagged case, identify the system version and relevant inputs, record a reasoned decision and initiate correction. It shall verify that a flag alone cannot suspend or terminate a benefit. Failure of these minimum controls shall prevent operational authorization."

The final agreement must adapt this provision to the governing law and specify test protocols, delivery dates, payment consequences and remedies. It must preserve the distinction between accepting supplier deliverables and authorizing public use, together with any stronger legal protections.

9.6 Respond to a failure

Suppose monitoring identifies an increase in false matches among applicants with shared household contact details. The agency restricts the affected flagging function, preserves the relevant logs and routes cases through the established alternative process. It identifies people whose cases may have been delayed or wrongly decided, contacts them through usable channels and reviews the affected outcomes. A data or model correction is followed by re-evaluation before the function resumes.

The assurance record distinguishes confirmed findings from suspected failures, records uncertainty and explains the containment decision. The designated authority may permit reinstatement only after re-evaluation and independent scrutiny appropriate to the risk, rather than relying solely on the supplier's assurance that the issue has been fixed.

9.7 Cost and test the transition

Suppose the model produces 200 flags each week and first-line review takes 30 minutes per flag. That would require 100 staff-hours weekly before supervision, appeals, absence or incident peaks are included. These illustrative assumptions make the review burden visible in the business case; actual staffing must follow evidence of local demand.

Before acceptance, the agency tests the export of required records and a temporary switch to its fallback process, confirming that pending applications remain accessible, complaints retain their histories and statutory records survive supplier exit. Continued monitoring then tests whether demand and costs remain within the approved plan. This example makes RAIT operational, but it also exposes the questions of bargaining power, resources and independence that implementation must address.

10. Implementation tensions

The arrangements in the example depend on choices that a framework cannot settle in advance. Buyers need terms suppliers will accept, institutions need recurring resources, and

reviewers need sufficient independence to challenge a project. These conditions shape both the feasibility of RAIT and the questions its pilots should examine.

10.1 Supplier leverage and deliverable terms

A public buyer with credible alternatives may negotiate stronger terms, but a supplier can still decline them. Early market engagement should establish whether proposed obligations are feasible, while predictable requirements and reusable evidence may reduce the cost of meeting them. Where development finance supports the project, the funder may also influence the contractual floor.

Documented due diligence and common requirements across markets can give suppliers reasons to participate, but neither removes an imbalance in bargaining power. If essential scrutiny cannot be secured, the institution must consider a narrower use, a different arrangement or a decision not to acquire the system.

10.2 Affordability and useful alternatives

Oversight work has a cost whether or not it appears in the acquisition budget. Leaving review, correction and complaints unfunded shifts that burden to frontline staff and service users through delay and unresolved errors. Costing the work explicitly allows an institution to compare the proposed system with other ways of improving the service.

At the same time, demanding a full evidence package for every use can discourage useful adoption. The proportionate track in §4.1 reduces unnecessary requirements while preserving essential protections. Where a consequential use remains unaffordable, an institution can improve records, develop staff, share expertise or run controlled evaluations before committing to deployment. The relevant decision concerns the particular service and its alternatives, rather than AI adoption in the abstract.

10.3 Independence under institutional pressure

Transparent appointments, declared interests, recusal rules and independently administered funding can help reviewers withstand pressure from suppliers and project sponsors. Fixed

and staggered terms, published reasoning with recorded dissent, and an independent challenge route provide further support. These arrangements reduce opportunities for capture without guaranteeing that the review will remain effective.

For that reason, assurance must combine documentary review with observed capability and evidence from staff and service users. A clean compliance record deserves investigation when it may reflect weak reporting or suppressed concerns. Pilots should examine whether reviewers can obtain information, report adverse findings and secure a response in practice.

10.4 Local agency and unequal bargaining power

A framework imposed by wealthy exporters without recipient participation could reproduce the dependence it seeks to address. RAIT therefore makes assessment specific to the deployment, supports locally defined priorities and treats capacity building as a funded acquisition obligation. The data governance and exit modules are intended to increase the institution's ability to understand and change its arrangements.

These mechanisms still operate within unequal relationships, and human-rights requirements can generate political disagreement. Evaluation should examine who sets the requirements, whose evidence influences decisions and whether the institution and affected communities gain practical influence. Formal participation alone cannot establish that power has shifted.

11. Piloting RAIT and evaluating its contribution

A pilot must test both the proposed AI service and the value of using RAIT to govern its acquisition. A functioning deployment can show that particular safeguards are feasible; it cannot alone establish that RAIT improved outcomes or that its full process was worth the cost. The evaluation should therefore examine the framework's effects on procurement, oversight, public protection and retained institutional capability.

11.1 Design a bounded evaluation

The first pilot should have an identified public owner, a lawful purpose, agreed responsibilities and sufficient resources for independent evaluation. Before operational use, participating institutions should document the existing service, data dependencies, harms and capability baseline, agree the relevant modules and assurance decisions, and establish how benefits and burdens will be assessed.

Evaluation should track the cost of assessment, documentation, training, review, complaints and independent scrutiny alongside service quality, exclusion, correction time, staff workload, supplier participation and dependence. It should examine which requirements changed decisions or revealed problems, which duplicated existing work and which proved difficult to sustain. Comparison with prior practice or a credible alternative should account for differences in workload, policy and resources.

The pilot must test institutional processes together with the system and proceed to bounded use only after required approvals and safeguards are in place. Accessible participation and remedy should remain available throughout. Findings should include costs, negative results and unresolved uncertainty, with a public account of the deployment decision and what the evaluation can and cannot establish.

11.2 Test capability retention and proportionality

The capability baseline should be reassessed after deployment and, where possible, across successive acquisitions. Observed exercises should establish whether the institution has retained the ability to reconstruct decisions, evaluate systems, correct records and operate alternatives. Analysis should distinguish its own knowledge from temporary supplier support and consider learning gained through training, shared expertise or other projects.

Evaluation of the proportionate track should examine whether low-consequence classifications still match operational use. Reclassifications need explanations: legitimate scope change, weak initial assessment, informal reliance or deliberate avoidance call for different responses. The rate alone cannot show whether the boundary is weak or monitoring is effective.

These evaluations should also ask whether audit access improves error detection, funded review reduces time to remedy, tested exit lowers dependence and participation changes design choices. The possibility that costs or restrictions prevent useful services from reaching underserved communities belongs in the same assessment.

11.3 Learn across institutions

If initial pilots provide useful evidence, institutions in different regions and administrative traditions could run comparable evaluations and publish their findings together. Comparison would help distinguish features that transfer between settings from those that depend on one service or legal system.

Regional cooperation could then support shared expertise and more consistent evidence requirements, reducing the need for suppliers and buyers to repeat equivalent work. Any harmonization should follow the evidence and the choices of participating institutions. Broader formalization would come later, if justified; individual pilots and shared learning do not depend on a new international treaty.

11.4 Limits to test openly

RAIT cannot create independent oversight through contract language alone or eliminate political interference, reviewer capture and funding failure. Its administrative burden may also disadvantage smaller institutions and suppliers. Pilots should test whether simpler evidence requirements preserve protection while making participation feasible.

Greater transparency can expose personal information or security weaknesses, while local control of data can increase state access or costs. These tensions require context-specific decisions and scrutiny of their effects. A successful pilot would justify further testing and adaptation, rather than establish universal effectiveness.

Conclusion

Responsible AI transfer should leave public institutions better equipped to serve people and remain accountable for the systems they use. RAIT proposes making that capacity part of the acquisition itself, with responsibilities that are funded, safeguards that can be demonstrated and knowledge that remains available when a supplier relationship ends.

The framework connects those commitments to decisions throughout the service's life: whether a use is justified, whether its controls work, whether continued operation remains acceptable and how to maintain service when the arrangement must change. Its value rests on the institution's ability to explain decisions, investigate failures, correct harm and exercise practical choice.

Piloting RAIT would establish how well these arrangements work and what they cost to sustain. That evidence can guide a sequence of useful acquisitions in which public benefit and the capacity to govern develop together.

Appendix A. Relationship to existing frameworks

The table maps the foundations on which RAIT draws and the operational uses proposed in this paper. It supports the discussion in §2.2; the relationships indicate areas for adaptation and comparison rather than legal equivalence, certification or institutional endorsement.

Framework or instrument	Relevant foundation	RAIT's proposed use
Universal DPI Safeguards Framework [1]	Lifecycle responsibilities, inclusion, agency and remedy.	Map an AI acquisition to responsible authorities, required evidence and recurring safeguard work.
NIST AI RMF 1.0 [2]	Contextual risk management, organizational roles and third-party dependencies.	Connect risk assessment and testing to contract deliverables and deployment decisions.
OECD AI Principles [3]	Accountability, traceability and human-centred governance.	Specify who supplies records and who can investigate or challenge an outcome.
UNESCO Recommendation on AI Ethics [4]	Human rights, impact assessment, oversight and concrete policy action.	Structure rights assessment, community participation and capacity development.

Framework or instrument	Relevant foundation	RAIT's proposed use
AU Continental AI Strategy [5]	An Africa-centred approach to responsible AI and institutional development.	Support locally defined acquisition priorities and shared regional expertise.
Applicable AI and data protection law [19] [20] [23]	Binding duties and restrictions where legally applicable.	Identify requirements an agreement must satisfy; no private certification overrides the law.
World Bank Environmental and Social Framework [22]	Risk-proportionate project safeguards, engagement and capacity building.	Inform financing of recurring AI safeguards where relevant to a project's applicable requirements.
UNCTAD transfer-of-technology tradition [17]	Licensing terms, technical training and recipient bargaining position.	Extend the same logic to information access, oversight capacity and exit for adaptive systems.
UK Guidelines for AI procurement [25]	Impact assessment, oversight, whole-life costs, knowledge transfer and continuing evaluation.	Connect these practices to funded owners, observed capability tests and a separate deployment decision.
EU AI model contractual clauses [26]	Full and light supplier clauses, with commentary for adaptation to public procurement.	Use relevant clauses alongside institutional funding, authorization and retained-capability records.

Implementation teams should identify the requirements that apply to their service and reuse existing evidence where it is adequate and current. The EU AI Act's prohibitions and risk-based duties remain applicable wherever the law governs the use; RAIT's assurance stages do not alter statutory classifications or implementation requirements. [\[23\]](#) The African Union's Continental AI Strategy supplies regional context for locally defined priorities and capacity development. [\[5\]](#)

Appendix B. Trade, export control and development finance

Acquisition agreements provide the immediate route for testing RAIT. Broader conditions on technology transfer raise additional questions about mandates, trade rules and the interests of states that both import and export systems. These pathways require specialist legal and policy analysis beyond the proposed pilot.

Existing dual-use export control arrangements already condition the transfer of sensitive technology on assessments made by the exporting state. Arms control instruments that require exporters to assess the risk that a transfer will be used to commit serious violations establish a comparable structure: a forward-looking risk assessment, made by the exporter, with denial as an available outcome. Whether AI systems used in policing, migration or benefits administration should attract analogous treatment is a live policy question that this paper does not attempt to resolve.

Conditioning market access on governance capacity raises questions under international trade rules on non-discrimination. General exceptions permitting measures necessary to protect public morals or to secure compliance with domestic law are the provisions that would ordinarily be examined. Whether a given measure qualifies is fact-specific and disputed; nothing in this framework should be read as an opinion that any particular measure would survive challenge.

Development finance offers a nearer implementation pathway where a project's governing arrangements allow it. Multilateral and bilateral funders already use impact assessment, stakeholder engagement, grievance mechanisms and borrower capacity building in relevant investment categories. Recurring AI oversight could be included as a project cost through those mechanisms, subject to the applicable requirements and funding decisions. [\[22\]](#)

States that both acquire and supply AI may weigh the commercial costs of stronger export requirements against their interests as buyers and the reputational consequences of harm abroad. Common requirements could favour suppliers already able to provide credible evidence, while also giving importing institutions a consistent basis for assessing offers. Whether these incentives support wider adoption depends on particular markets and institutions and remains an empirical question.

References

References below support the factual and institutional foundations cited in the text. The stages, contractual modules, indicators and illustrative provisions are proposed design contributions.

- [1] [UNDP and United Nations Office for Digital and Emerging Technologies. Universal DPI Safeguards Framework. Online framework, lifecycle and responsible-authority guidance.](#)
- [2] [National Institute of Standards and Technology. 2023. Artificial Intelligence Risk Management Framework \(AI RMF 1.0\), NIST AI 100-1. Especially the GOVERN and MAP functions.](#)
- [3] [Organisation for Economic Co-operation and Development. AI Principles. Adopted 2019; updated 2024. Accountability, traceability and human-centred values.](#)
- [4] [UNESCO. Recommendation on the Ethics of Artificial Intelligence. Adopted 2021. Official overview of principles and policy actions.](#)
- [5] [African Union. 2024. Continental Artificial Intelligence Strategy. Endorsed by the Executive Council in July 2024; published Aug 9, 2024.](#)
- [6] [Royal Commission into the Robodebt Scheme. 2023. Report. Recommendations 16.1–16.2 and 17.1–17.2 concerning data exchanges and automated decision-making.](#)
- [7] [Politico. Dutch scandal serves as a warning for Europe over risks of using algorithms. Coverage of the childcare benefits risk-profiling system and its consequences.](#)
- [8] [Amnesty International. 2024. Denmark: AI-powered welfare system fuels mass surveillance and risks discriminating against marginalized groups.](#)
- [9] [Amnesty International. 2022. The Social Atrocity: Meta and the Right to Remedy for the Rohingya.](#)
- [10] [Time. 2022. Reporting on the Amnesty International findings regarding algorithmic amplification and violence against the Rohingya.](#)
- [11] [Council on Foreign Relations. 2021. Facebook's Content Moderation Failures in Ethiopia.](#)
- [12] [The Irish Times. 2019. Chinese surveillance systems appear across Ugandan capital Kampala.](#)
- [13] [Polyakova, A. and Meserole, C. 2019. Exporting Digital Authoritarianism: The Russian and Chinese Models. Brookings Institution.](#)
- [14] [Institute for Security Studies. 2024. The Future of Facial Recognition Technology in Africa.](#)
- [15] [World Economic Forum. 2024. Deepfakes in India's elections and the response infrastructure developed to address them.](#)
- [16] [Mariani, J., Eggers, W. D., and Kishnani, P. 2023. The AI Regulations That Aren't Being Talked About. Deloitte Insights. Review of over 1,600 AI policies across 69 countries.](#)
- [17] [UNCTAD. 1985. Draft International Code of Conduct on the Transfer of Technology.](#)
- [18] [United Nations Capital Development Fund. 2022. Uganda Launches Data Protection & Privacy Portal to Streamline Management of Personal Data. Jun 7, 2022.](#)
- [19] [Mulindwa, Calvin. 2026. A Comparative Perspective on Privacy Safeguards in Kenya and Uganda Data Protection Laws. CIPIT, Strathmore University, Feb 5, 2026. See the discussion of data-subject rights.](#)

- [20] [Republic of Kenya. 2019. Data Protection Act, No. 24 of 2019. Sections 31, 35, 41 and 56. Official Gazette text hosted by WHO.](#)
- [21] [Republic of Uganda. 2019. Data Protection and Privacy Act, No. 9 of 2019.](#)
- [22] [World Bank. Environmental and Social Framework. Approved 2016. Official overview of applicable project safeguards and institutional capacity building.](#)
- [23] [European Commission. AI Act. Official overview of Regulation \(EU\) 2024/1689, prohibited practices and risk-based requirements.](#)
- [24] [MIT Sloan. Balance AI Innovation and Risk with "Minimum Viable Governance." Ideas Made to Matter, drawing on MIT Center for Information Systems Research. Enterprise-level treatment of proportionate, adaptive AI governance.](#)
- [25] [UK Government, Office for Artificial Intelligence. Guidelines for AI procurement. Guidance on impact assessment, governance, supplier dependence, knowledge transfer and lifecycle management.](#)
- [26] [European Commission, Public Buyers Community. EU AI model contractual clauses: full and light versions and commentary for responsible public procurement.](#)

1

¹ Research and editorial assistance:

AI tools supported research, consolidation and editorial revision under the author's direction and review. **The author takes responsibility for the framework, analysis and final text.**