

RAIT.

Responsible AI Transfer Framework — a framework for governing public-sector AI acquisition.

Samuel Abinsinguza · Founder, CentPol

Acquiring an AI system does not come with the capacity to govern it.

Systems acquired to improve services also bring responsibilities: to understand how they operate, protect the people they affect and keep meaningful control over their use. The technology alone provides none of the expertise, resources or authority those responsibilities demand.

In the Netherlands, a self-learning risk-profiling algorithm ran for years inside childcare-benefit fraud detection. More than twenty thousand families were wrongly accused and pushed into debt, and the government resigned. In Kampala, a foreign-supplied "smart city" package arrived with facial recognition and policing dashboards — with no dedicated law governing the footage and no independent review of the contracts.

The cases differ, but they raise one 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 but not the records to explain a decision; a regulator may have powers but no technical support; a person seeking correction may meet several agencies, each responsible for only part of the process.

The 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.

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. It applies wherever an institution relies on externally supplied technology, including domestic suppliers and open-source systems. The unit of assessment is the **deployment** — a defined system, purpose, population and institutional arrangement — never a national readiness grade.

Who it is for

Public-sector implementing teams, procurement authorities, sector regulators, data protection bodies, development partners and independent reviewers. RAIT can inform an acquisition plan, a procurement specification, an implementation agreement, an assurance review or a financing arrangement — adapted to the applicable law.

SIX GOVERNING COMMITMENTS

Public accountability

Respect for human rights

Local agency

Proportionate assurance

Shared responsibility

Sustained capacity

Some protections apply from the start — including to pilots: a lawful purpose, a responsible decision-maker, controls over personal data, an explanation of the system's role, usable channels for concerns and a way to correct consequential errors. Human review counts only when reviewers have the **records, time, competence and authority** to change an outcome.

Four stages of deployment assurance

Each stage builds on the last, and approval stays specific to the purpose, system version, data dependencies, users and review date. As conditions change, a deployment can be reassessed, reduced or paused.

1 • Establish	Purpose, legal basis, affected people and responsible institutions. Compare AI with improving the existing service, and record a capability baseline : what the institution can already do without the supplier. Cannot authorize consequential decisions about people.
2 • Test	Test the system <i>and</i> the safeguards together in the local setting — data, errors, staff behaviour, complaints, correction, logging and fallback. Supplier benchmarks and training completion cannot establish readiness on their own.
3 • Authorize	A signed, bounded deployment decision: defined use, operating limits, stop conditions, recurring funding and accessible remedy in place. Accepting a supplier's deliverables is not the same as authorizing public use.
4 • Sustain	Monitoring, incident response, change control and remedy. Material changes trigger reassessment, and reassessment returns to the baseline — has the institution become <i>less</i> able to govern the system? When exit is necessary, protect people and preserve records.

Proportionate track. Eligible low-consequence uses follow a lighter version of the four stages — but keep every essential protection. The track is unavailable wherever an output influences access to a service, entitlement, liberty, treatment or legal status, or runs on shared identity or population-scale infrastructure. Delivery pressure or a thin budget cannot justify lowering essential protections.

Capability across acquisitions. An institution's *governability frontier* — the uses it can oversee with adequate protection — can move as it retains tested contract terms, experienced reviewers and evidence of oversight costs. An acquisition that replaces scarce expertise without transferring knowledge may improve performance today while deepening dependence.

Six contract modules

The modules turn assurance requirements into obligations buyers and suppliers can deliver and verify. Each names the responsible party, deliverable, deadline or trigger, acceptance evidence, funding and consequence of non-performance.

Capacity plan	The people and responsibilities governance needs — legal, evaluation, security, supervision, correction, complaints — with competence shown through real tasks, not attendance.
----------------------	---

Rights & impact	Necessity, proportionality, discrimination, exclusion, accessibility, privacy and redress across the whole service pathway, with independent review for high-consequence uses.
Data governance	How information moves between agencies, who can change it, and proof that corrections and access withdrawals take effect everywhere downstream.
Auditability	Documentation, logs, version histories, audit access and change control — enough for the buyer and independent reviewers to investigate a failure.
Workforce	Training tied to the responsibilities people actually hold, with completion, competence, retention and progression tracked separately.
Exit & continuity	A tested way to stop using the AI component, change supplier or run a simpler process — preserving pending cases, complaints and records.

Responsibility, verification and funding

The implementing authority keeps the public-service decision and resources the controls; data-holding agencies stay responsible for their records; suppliers fulfil their agreements. An assurance panel with a clear mandate can bring legal, technical, operational and community expertise — complementing regulators and courts, never displacing them. A supplier must not select the sole reviewer of its own work.

Verification means **observed exercises**, not just documents: retrieving the record behind a disputed output, investigating a mismatch, processing a complaint, correcting data and operating without the AI component.

Oversight is treated as an acquisition cost across the whole operating period. Cost the work, not a budget percentage: if a model produces 200 flags a week and first-line review takes 30 minutes each, that is **100 staff-hours weekly** before supervision, appeals or incident peaks. Payment milestones can follow accepted documentation, completed tests and demonstrated capability.

WORKED EXAMPLE — A BENEFITS SERVICE

An agency acquires a model to flag possible duplicate benefit applications. A flag routes a case for review but can never suspend or terminate a benefit on its own. Before authorization, a reviewer must show they can reconstruct a flagged case, identify the system version and inputs, record a reasoned decision and start a correction. If false matches later rise, the agency restricts

that function, reviews affected cases and reinstates it only after re-evaluation and independent scrutiny.

What remains open

RAIT is a proposed framework, drawn from documentary evidence and existing instruments such as the UK Guidelines for AI procurement, the EU AI model contractual clauses, the NIST AI RMF and the Universal DPI Safeguards Framework. It has not been field-validated: its usefulness, costs and contribution to lasting capability require **pilot evaluation** — including open tests of supplier leverage, affordability, reviewer independence and whether affected communities gain real influence.

About CentPol

CentPol is a **global community of people — from wildly different professional backgrounds — curious enough to understand, shape, and build the future of AI** rather than watch it happen to them.

We produce **policy intelligence for the next technology era** through community, stakeholder dialogue, and research. RAIT is part of CentPol's research on how institutions can adopt AI while keeping the capacity to govern it.

This brief summarizes *RAIT: Responsible AI Transfer Framework* by Samuel Abinsinguza. AI tools supported drafting under the author's direction and review; the author takes responsibility for the framework, analysis and final text.

Read the full framework.

The complete paper — evidence, the worked benefits example, monitoring measures and the mapping to existing frameworks, with 26 references.

centpol.com/research/white-papers/rait-framework