

REQUEST FOR EXPRESSIONS OF INTEREST

PROJECT PREPARATION UNDER THE GRANT FACILITY FOR PROJECT PREPARATION (GRANT NO. E6900-001) AI FOR ECONOMIC TRANSFORMATION IN RWANDA (P517443)

Assignment title: Proof of Concept for AI-Enabled Agentic Workflows for Existing Government Services and User Journeys (Structured into Three Lots)

Reference No. (as per Procurement Plan): RW-RAIA-577862-CS-QCBS

The Government of Rwanda has received financing from the World Bank toward the cost of project preparation for the AI for Economic Transformation in Rwanda operation, through the Grant Facility for Project Preparation (GFPP). The Government has appointed the Rwanda Artificial Intelligence Agency to be the implementing entity for the GFPP and the institutional host of the Project Implementation Unit (PIU).

About the Rwanda Artificial Intelligence Agency

The Rwanda Artificial Intelligence Agency (RAIA) advances the responsible and inclusive use of Artificial Intelligence and other emerging technologies in support of Rwanda's national development priorities. The Agency brings public AI initiatives under one umbrella, covering AI deployment and government transformation, research, innovation and talent, governance, ethics and safety, industry development and investment, and international engagement. It was approved by Cabinet in June 2026 and continues the work previously carried out under the Centre for AI Policy and Innovation Ltd, trading as the Rwanda Centre for the Fourth Industrial Revolution.

About the Consultancy

The GFPP intends to apply part of the funds for consulting services ("the Services"). The Services are divided into three lots, each covering one use case for citizens or businesses on a live national platform:

- **Lot 1: Mbaza citizen engagement platform (MINALOC).** Agentic complaint formulation, categorisation and routing, and trend insights for leaders.
- **Lot 2: Food and cosmetics licensing on IremboAI (Irembo and MINICOM).** AI-guided self-assessment and document checks for manufacturers of food, beverages and non-pharmaceutical cosmetics applying for licences, with officers deciding every application.
- **Lot 3: Building permits on Kubaka (MININFRA and RHA).** Agentic compliance review of permit applications and improved applicant support.

A separate contract will be awarded for each lot. **Firms may express interest in a maximum of two lots** and must state clearly which lot or lots they are applying for.

For its lot, the selected firm will design, build and evaluate a working proof of concept of AI-enabled agentic workflows. Firms are expected to propose their own technical approach to agent architecture, orchestration, model selection, human oversight and integration. The Services cover an agent-readiness assessment (including process maps and baselines, API discovery and a data catalogue), validation and application of RAIA's risk-based autonomy framework, solution design and guardrails, the proof of concept and its evaluation harness, a security, privacy and data protection assessment, a controlled pilot, a costed, phased scale-up blueprint for the follow-on operation with a sustainability pathway for long-term operation, and handover. The three lots run in parallel over about 7.5 months, from an expected start in late March 2027, with all final deliverables due by 8 November 2027.

The detailed Terms of Reference (ToR) for the assignment are in Annex I of this Request for Expressions of Interest, and descriptions of the three use cases are in Annex II.

Shortlisting criteria

RAIA now invites eligible consulting firms ("Consultants") to indicate their interest in providing the Services. Interested Consultants should provide information demonstrating that they have the required qualifications and relevant experience to perform the Services. Shortlisting will be done separately for each lot, against the following criteria:

- **Core business and track record:** core business in delivering AI and machine learning solutions or digital government services, with evidence of projects delivered, and at least two assignments completed in the last five years delivering AI systems for government or large institutions, including at least one involving generative AI models, specialised AI solutions for automating administrative processes, or agentic workflows with evaluation and guardrails. Products that the firm has built and operates may be presented as assignments.
- **Local-language AI:** experience building or adapting AI for low-resource languages, preferably Kinyarwanda, across text, voice and documents.
- **Integration with government systems:** experience integrating AI or digital solutions with government platforms, interoperability layers, or national identity and payment systems, preferably in Rwanda or a comparable public-sector context.
- **In-country delivery and data residency:** demonstrated ability to deploy and operate solutions on national or on-premise infrastructure with data kept in the country, and to field a core team on site in Kigali for the duration of the assignment.
- **Technical and managerial capability:** in-house capacity in AI and machine learning engineering, systems integration and data architecture, public service process design, application security and data protection, and user research, and experience transferring skills, code and operational ownership to government or institutional teams.

Key Experts will not be evaluated at the shortlisting stage, and curricula vitae are not required. For each lot, a shortlist of not fewer than five and not more than eight eligible firms will be established.

Content of expressions of interest

Expressions of interest should include:

- a cover letter stating the lot or lots applied for (maximum two);
- a firm profile stating the firm's core business and years in business;
- summaries of no more than five relevant assignments completed in the last five years, each stating the client, country, contract value, duration and year of completion, with a brief description of the services provided and the lot to which it is relevant;
- a description of the firm's technical and managerial capability, including staff numbers by relevant discipline and, for firms applying for two lots, its capacity to deliver both in parallel;
- evidence of legal registration; and
- where applicable, details of any proposed association and the role of each member.

Expressions of interest should not exceed 15 pages, excluding the cover letter and registration documents. Pages beyond this limit will not be considered. Firms applying for two lots should submit a single expression of interest.

Procurement terms

The attention of interested Consultants is drawn to Section III, paragraphs 3.14, 3.16, and 3.17 of the World Bank's "Procurement Regulations for IPF Borrowers" September 2025 ("Procurement Regulations"), setting forth the World Bank's policy on conflict of interest. In addition, the following specific information on conflict of

interest applies to this assignment: firms must disclose any affiliation with providers of AI models, cloud services or software platforms that could be proposed for use in the proof of concept. The selected firm for any lot, and any affiliate, will not be eligible to provide goods, works or non-consulting services resulting from or directly related to this assignment, including scale-up of the proof of concept.

Consultants may associate with other firms to enhance their qualifications, but should indicate clearly whether the association is in the form of a joint venture and/or a sub-consultancy. In the case of a joint venture, all the partners shall be jointly and severally liable for the entire contract, if selected.

A Consultant will be selected in accordance with the Quality and Cost-Based Selection (QCBS) method set out in the Procurement Regulations, with a technical to financial weighting of 80:20. The assignment will be carried out in English.

Further information can be obtained at the address Block 5, Norrsken House, 1 KN 78 St, Kigali - Rwanda or via email procurement@raia.rw during office hours 09:00 to 17:00 hours, Kigali time, Monday to Friday.

Expressions of interest must be delivered in a written form to the address; Block 5, Norrsken House, 1 KN 78 St, Kigali – Rwanda, in person by 09 November 2026 at 17:00 hours, Kigali time. The envelope subject line should read: "REOI — Proof of Concept for AI-Enabled Agentic Workflows — Lot(s) [number(s)] — RW-RAIA-577862-CS-QCBS".

Procurement-related complaints. A Consultant that considers it has been harmed by an alleged breach of obligation in this procurement process may submit a procurement-related complaint in writing at complaints@raia.rw within the periods and following the procedures set out in Annex III (Procurement-Related Complaints) of the World Bank Procurement Regulations.

Annex I: Terms of Reference

Package 4 — Proof of Concept for AI-Enabled Agentic Workflows for Existing Government Services and User Journeys

1. Background

The Government of Rwanda has received a Grant Facility for Project Preparation (GFPP) from the World Bank toward the preparation of the follow-on operation "AI for Economic Transformation in Rwanda", which covers AI-ready data infrastructure, AI solutions in priority sectors, AI-ready digital skills, and project preparation and institutional capacity. The Government has appointed the Rwanda Artificial Intelligence Agency as the implementing agency.

Government services in Rwanda are delivered through a number of existing digital platforms. This assignment tests whether AI-enabled agentic workflows can improve the speed, quality and accessibility of selected services on those platforms, and produces the design and evidence needed to decide on scale-up under the follow-on operation. Although delivered through multiple channels, these platforms and their underlying workflows are often staff-operated, form-based transactions. To support AI agents, their data and functions must become discoverable and securely callable through governed APIs.

The assignment is exploratory by design. Firms are expected to propose their own technical approach to agent architecture, orchestration, model selection, human oversight and integration.

2. Objective

For each lot, the objective is to demonstrate, through a working proof of concept, whether and how agentic AI workflows improve service delivery against an agreed baseline; to assess the readiness of the underlying systems, data and APIs to support AI agents; to validate and apply RAIA's risk-based autonomy framework; and to produce a costed, phased scale-up blueprint for the follow-on operation, with a sustainability pathway setting out viable business models for long-term operation.

3. Lot structure and use cases

The assignment is divided into three lots, one for each use case. Each lot is a separate contract and covers the full scope in Section 4 and the deliverables in Section 5 for its own use case. A firm may be awarded a maximum of two lots.

Each use case is a service on a live national platform that already has AI work or digitised processes in place, so the firm starts from a working base rather than a blank page. All three follow a *supervised by default* principle: agents handle routine steps, citizens and businesses confirm actions taken on their behalf, and authorised officials keep decisions. Full descriptions are in **Annex II**, covering each problem, work to date, what the proof of concept must add, and success measures.

1. **Lot 1: Mbaza citizen engagement platform (MINALOC).** Mbaza is the national channel for inquiries, complaints and tip-offs. All institutions are onboarded and it has processed over one million cases. Cases are still read and routed by hand, which leads to delays and misrouting. The proof of concept will add agents working in Kinyarwanda, English and French for complaint formulation, categorisation and routing, and trend insights for leaders. Officers remain responsible for resolving cases.
2. **Lot 2: Food and cosmetics licensing on IremboAI (Irembo and MINICOM, with RFDA, RICA, RSB and REMA).** Manufacturers of food, beverages and non-pharmaceutical cosmetics need licences and certificates from several institutions: around 4,000 applications across five services in 2025. Applying to

each institution in turn could take up to 18 months; a harmonised service led by MINICOM, now being tested, lets businesses submit one application to all institutions on IremboGov, followed by a joint inspection. Checklists remain hard to read, required documents come from many issuers, and officers return incomplete files for re-upload by hand, one round at a time. The proof of concept will add AI-guided self-assessment on IremboAI: by chat, an agent explains the checklist, flags gaps and checks documents before submission, and a review agent summarises each file for the institution's officer. Officers approve, return or order an inspection for every application.

3. **Lot 3: Building permits on Kubaka (MININFRA and RHA).** Kubaka is the national building permit platform. Officials check applications against the Building Code by hand, which causes late feedback, repeated resubmissions and missed 15-day decision windows. An in-house compliance agent is in pilot at about 75% accuracy, against a 90% target. The proof of concept will improve compliance review using machine-readable rules and an evaluation set built with RHA reviewers, and will improve the support agent. Permit decisions stay with the reviewing officer.

Common requirements. Citizen and business data stays in Rwanda, hosted on the national compute platform. The firm builds on existing systems rather than replacing them. Each use case is built on a reusable agentic architecture, with components designed so that other government services can reuse them, within or beyond the GFPP. Examples include language and voice, identity and delegated authority, API wrappers and tool registry, and an evaluation harness. These components are handed over as Government-owned assets. Each workflow step is assigned its autonomy level under the validated risk-based autonomy framework. Results are measured against a baseline agreed at inception.

4. Scope of services

The scope below applies to each lot, for that lot's use case.

Phase 1 — Confirmation and design (months 0–2.5)

- **Process maps, baselines and feasibility:** develop as-is process maps and baseline performance measures (e.g. processing times, transaction volumes, rework or rejection rates, and cost to serve) for the use case, working with RISA and the service-owning institution and building on any existing RISA documentation. Review and validate the agentic AI readiness assessment framework developed by RAIA, covering rules-based logic, deterministic workflows, API readiness, discoverability and service ownership, and propose refinements where needed. Then apply the validated framework to assess the feasibility of the use case.
- The readiness assessment shall distinguish between: (i) gaps that must be addressed before the proof of concept proceeds; (ii) gaps that can be addressed through documentation, configuration or lightweight adapters during the assignment; and (iii) gaps requiring longer-term technical, institutional or legal reforms, to be included in the scale-up blueprint. The assessment shall be applied as a maturity pathway rather than as a binary national readiness test.
- **Service and user journey:** review and validate the end-to-end as-is service and user journey, covering service ownership, institutional handoffs, information and documentary requirements, decision points, approvals, exception pathways, processing times, user pain points, and dependencies on shared government systems. Based on this assessment, develop a simplified to-be service and user journey before designing the agentic workflow. The to-be design should identify steps, documentary requirements, approvals and data requests that can be eliminated, consolidated, pre-populated or obtained directly from authoritative government sources, including through application of the once-only principle. The proof of concept should not replicate an existing process without first assessing opportunities for simplification.
- **API discovery:** undertake a structured inventory of the services, application programming interfaces (APIs), system functions, events and integration points supporting the use case, including those available

through the government interoperability layer. For each relevant interface, document its owner, purpose, users, authentication and authorisation requirements, inputs and outputs, data definitions, dependencies, error responses, usage limits, version, security controls, and fitness for use as an agent tool.

- **Data catalogue:** prepare a use-case-specific, machine-readable catalogue of the data required for each selected workflow. For each dataset or data element, document its authoritative source, institutional owner, purpose, classification (including personal data), legal basis for use, provenance, meaning, format, quality, update frequency, access conditions, retention requirements, version history, and applicable security and privacy controls. The catalogue shall align with relevant national metadata standards and data-governance arrangements and shall be designed for reuse beyond the proof of concept. Where data-quality or metadata gaps could materially affect the workflow, the firm shall identify the responsible data owner, assess the implications, and propose proportionate remedial actions and monitoring arrangements. Direct access by an agent to source databases shall not be permitted.
- **Rules as code:** identify and document the legal, policy, eligibility, procedural and business rules governing each selected service and workflow. Where appropriate, translate these rules into explicit, versioned, machine-readable decision logic that can be tested and audited.
- Confirm the workflows within the use case to be implemented in the proof of concept, taking account of the readiness of their data and APIs.
- **Solution architecture and guardrails:** propose and document the agent architecture, orchestration approach, model selection rationale, guardrails and human-in-the-loop design, including how agents access existing systems through a governed API and tool layer (agent identity, least-privilege permissions, audit trail) and the minimum changes required to existing modules. Agent-to-system and agent-to-agent interfaces shall use open standards, such as OpenAPI for API descriptions, the Model Context Protocol for tool access, and agent-to-agent protocols where agents interact, to support interoperability and avoid vendor lock-in. The design shall distinguish between: deterministic rules executed by government systems; model-based functions used to retrieve, classify, extract, summarise, explain or recommend; actions that an agent may undertake within explicitly delegated authority; and decisions or actions requiring confirmation, review or approval by an authorised person.
- **Risk-based autonomy pathway:** validate the risk-based autonomy framework developed by RAIA and use it to classify each workflow step by risk. The risk class sets the maximum autonomy permitted for that step, from assistive to supervised, proactive, and autonomous with oversight. Steps move up only when defined entry criteria are met, and decisions with legal or financial effect remain with an authorised official. The proof of concept should reach at least the supervised level, with higher levels demonstrated on steps where the risk classification permits.

Phase 2 — Build and assurance (months 2.5–5)

- Build the proof of concept on the national compute platform, with the agreed safeguards in place, integrated with the existing platform through its APIs or lightweight adapters in non-production environments.
- Develop the reusable agent-enablement components the proof of concept needs (e.g. API wrappers and tool definitions, connectors, a tool registry and data catalogue entries), documented so that other services can reuse them.
- Where feasible and where the risk classification permits, demonstrate at least one proactive workflow in which the agent initiates an action or notification from an event or data signal rather than a user request.
- For each agent and workflow, specify: which person or institution the agent represents; the purpose and scope of the delegated authority; the data, systems and tools the agent may access; the actions the agent may perform independently; the actions requiring user confirmation or official approval; the duration and conditions of the authority; and how the authority can be reviewed, limited, suspended or revoked.

- Build an evaluation harness measuring accuracy, latency, end-to-end processing time, staff effort per case, cost per transaction, escalation rates and user satisfaction.
- Complete a security, privacy and data protection assessment, including audit logging of agent actions, agent access permissions, prompt-injection and tool-misuse risks, and treatment of personal data in line with Law No 058/2021 of 13/10/2021 relating to the protection of personal data and privacy.

Phase 3 — Pilot, blueprint and handover (months 5–7.5)

- Run a controlled pilot with a reasonable group of users and evaluate it against the baseline. In this sandbox environment, integrate the proof of concept with the existing platform through governed APIs, managed interfaces or lightweight adapters in non-production environments. Models and agents shall not connect directly to production systems, source databases or unrestricted tool collections.
- Produce a scale-up blueprint covering the target agentic architecture and shared components (API and tool registry, data catalogue, agent identity and monitoring), changes required to existing systems and modules, operating model, skills, governance, risks, and a costed, phased roadmap for extending agentic and proactive workflows to other services.
- The blueprint shall assess total cost of ownership, including model use, compute and hosting, software licensing, integration, monitoring, security operations, human review, evaluation and testing, data and API maintenance, specialist staffing, user support, and onboarding of additional institutions and services. It shall explain how costs are expected to change as reusable components are applied to additional services.
- **Sustainability pathway:** complement the costed scale-up roadmap with viable business models for the long-term operation of the solution beyond the proof of concept, for example funding through institutional budgets, shared-service arrangements across institutions, cost recovery where appropriate, or partnerships. For each option, set out the institutional owner, funding source, operating responsibilities and the conditions under which it is financially viable, drawing on the total cost of ownership analysis.
- Transfer knowledge to RAIA technical staff and hand over all code, configurations, prompts, tool definitions, API inventory, data catalogue entries and documentation. All of these are Government-owned intellectual property and must be delivered to a Government-designated repository, so they can be reused in scale-up.

Out of scope

- Production rollout.
- Changes to the owning institutions' production systems; required changes are specified in the scale-up blueprint.

5. Deliverables and indicative timing

Each lot delivers D1 to D8 for its use case. Timing is the same for all lots.

Deliverable	Due (indicative)
D1 Inception report and confirmation of the workflows in scope	Month 0.5 — 8 Apr 2027
D2 Process maps and baselines; validated readiness assessment framework; feasibility and readiness assessment (incl. API discovery and data catalogue); Security, privacy, and data protection assessment	Month 2 — 21 May 2027
D3 Validated risk-based autonomy framework and autonomy level per workflow step; solution architecture and guardrail design	Month 2.5 — 8 Jun 2027

Deliverable	Due (indicative)
D4 Working proof of concept, evaluation harness and reusable agent-enablement components	Month 5 — 24 Aug 2027
D5 Pilot evaluation report against the baseline	Month 6.5 — 11 Oct 2027
D6 Scale-up blueprint, costed, phased roadmap and sustainability pathway, incl. required changes to owning institutions' systems	Month 7 — 25 Oct 2027
D7 Handover and knowledge transfer report, confirming delivery of all assets to the Government repository	Month 7.5 — 8 Nov 2027

All deliverables are subject to written acceptance within ten business days of submission. All final deliverables must be submitted no later than 8 November 2027, so that they can be reviewed, accepted and paid before the grant closes; services performed after 31 December 2027 cannot be financed by the grant. The schedule will be re-based from the actual date of contract signature, but the final deadline will not move.

6. Key Experts

The following Key Experts will be evaluated for each lot at the Request for Proposals stage, not at shortlisting. Firms may propose additional non-key experts.

Position	Minimum qualifications and experience
Team Leader / AI Solutions Lead	Master's in computer science or related field; 10+ years, including 5+ delivering production AI/ML systems; public sector delivery experience
Agentic AI / ML Engineer	5+ years; demonstrable delivery of LLM-based or agentic systems, including evaluation and guardrails
Integration and Data Architect	8+ years in enterprise or government systems integration (APIs, interoperability platforms) and data management, including metadata and data cataloguing
Government Service Design / Process Specialist	8+ years redesigning public service delivery processes
Security and Data Protection Specialist	8+ years in application and data security; privacy-by-design experience
User Research / UX Specialist	6+ years, including digital public services

7. Client inputs and facilities

- RAIAs agentic AI readiness assessment framework and risk-based autonomy framework, for validation.
- Any existing process documentation and baseline data held by RISA and the owning institution.
- Access to the owning institution's platform, data and staff under agreed arrangements, including existing API documentation, interface specifications, data dictionaries and test environments.
- Designated technical focal points in RAIAs and the owning institution.
- A development environment and access to the national compute platform.
- Technical review by the PIU AI Solution Architect.

8. Reporting and management

The firm reports to the PIU Project Director, with technical direction from the PIU AI Solution Architect. Progress meetings will be held every two weeks. RAlA will convene the firms holding the three lots at key points, including inception and blueprint preparation, so that reusable components, shared interfaces (such as the Mbaza–Irembo link) and scale-up recommendations are consistent across lots.

Annex II: Use case descriptions

Purpose and how to read this annex

Each of the three use cases below is a separate lot. Each has existing AI work or digitised processes that the GFPP will take further, so firms start from a working base rather than a blank page. The firm selected for each lot will, with RAI and the service-owning institution, confirm a bounded scope that can be delivered and evaluated in about 7.5 months.

For each use case, the sections below state the problem, the agentic solution the owner is seeking, the work done so far, and the outcomes the proof of concept should evidence. Firms should treat scope as indicative and expect to validate it during inception (D1) and the feasibility and readiness assessment (D2).

All three share one design principle: supervised by default. Agents handle routine steps, citizens and businesses confirm actions taken on their behalf, and authorised officials keep decisions.

Lot 1: Mbaza citizen engagement platform (MINALOC)

Platform and owner. Mbaza is Rwanda's national channel for inquiries, complaints and anonymous tip-offs. It routes each case to the responsible institution and gives citizens a tracking number and SMS or email updates. It is live nationally, with all government institutions onboarded and more than one million cases processed.

Problem. Cases arrive through many channels and are sorted by hand. Staff spend hours reading, categorising and assigning each one; some reach the wrong office or level of government, deadlines are hard to enforce, and leaders get little insight from the volume of submissions. Citizens also submit incomplete or unclear complaints that then need follow-up.

Agentic solution sought. A tested proof of concept of one shared core model working in Kinyarwanda, English and French, supporting agents for intake, knowledge, triage, follow-up and insight. Three capabilities are the priority:

- **Complaint formulation:** the agent helps a citizen write a complete, clear complaint before submitting it, and gives feedback they can act on. Voice input, where the citizen records and reviews a transcript, is a related option.
- **Sorting and routing:** the agent detects language, translates and corrects text, then categorises and routes each case to the right institution, with urgent cases and tip-offs flagged on arrival.
- **Insights and reporting:** conversational summaries of trends, sentiment and bottlenecks for ministers and mayors, for example for weekly review meetings.

Officers remain responsible for resolving cases.

Work done so far. The production platform is live nationally. Citizens submit inquiries, complaints and anonymous tip-offs, receive a tracking number and SMS or email updates, and cases are routed to the responsible institution. Mbaza has run for about a year, with growing institutional adoption. AI capabilities are not yet part of the production service, so the proof of concept starts from the live platform and its case data.

What the proof of concept must add. The proof of concept will add five agentic workflows to the live platform. A citizen agent works with citizens on chat, WhatsApp, SMS, USSD and voice as they write a complaint, query or tip-off, so that it is complete and error-free before submission. A knowledge agent answers routine inquiries from a knowledge base built from each institution's own guidance, so officers can spend their time on real complaints. A triage agent sorts, prioritises and assigns each case to the right officer, flagging urgent cases and tip-offs on arrival. A follow-up agent tracks deadlines, reminds officers and escalates late cases. An insight agent reads sentiment, trends and patterns across all cases, so leaders can spot failing services and rising

issues early. These agents depend on labelled Kinyarwanda text and voice data; hosting on national compute, with citizen data kept in Rwanda in line with national data-protection requirements; validation of accuracy and error rates by language and channel; and integration with Irembo so agents can follow a service case end to end.

Expected outcomes. Instant acknowledgement and faster, more accurate routing; more cases resolved within deadlines; less staff time per case; and earlier warning of emerging issues for leaders. Measures include categorisation and routing accuracy, time to produce a categorisation, share of outputs approved by staff without revision, and reduction in staff processing time.

Lot 2: Food and cosmetics licensing on IremboAI (Irembo and MINICOM, with RFDA, RICA, RSB and REMA)

Platform and owner. The Ministry of Trade and Industry (MINICOM) is the owner of the licensing process and the focal point for this use case. IremboGov, the national portal run by Irembo, is the interface for applicants, and its back end is integrated with the platforms of the institutions involved: the Rwanda Food and Drugs Authority (RFDA), which licenses food and beverages; the Rwanda Inspectorate, Competition and Consumer Protection Authority (RICA), which inspects premises and issues certificates; the Rwanda Standards Board (RSB), for quality certification; and the Rwanda Environment Management Authority (REMA), where plastics are used in production or packaging. IremboAI is the AI assistant embedded in IremboGov. The scope covers food, beverages and non-pharmaceutical cosmetics; drugs and pharmaceutical cosmetics are out of scope.

Problem. Until recently, manufacturers applied to each institution in turn, with no guidance on the sequence, and licensing could take up to 18 months. Volumes are steady: around 4,000 applications across five services in 2025, or about 330 a month, led by single-use plastics (110) and RFDA products (75). Even with the harmonised service, self-assessment checklists are hard to read, so businesses tick boxes without always knowing what each standard requires. Required documents, such as a lease, trademark registration, occupancy permit, RRA classification and environmental and social impact assessment (ESIA), each come from a different body. Officers return incomplete files by hand for re-upload, one round at a time.

Agentic solution sought. AI-guided self-assessment on IremboAI, with two supervised agents:

- **Self-assessment agent:** by chat or voice, explains each checklist question in plain language; flags answers that point to a gap, with what to fix and which document proves it; and checks, before submission, that each required document is present, legible, valid and consistent.
- **Review agent:** after submission, checks the documents and summarises the file for each institution's officer, with issues flagged against the checklist. When a file is returned, the agent tells the business exactly what to re-upload.

Documents checked include trademark registration (or a notarised agreement where the brand is not owned), a notarised lease or ownership contract, an occupancy permit or no-objection letter for large companies, RRA classification as MSME, SME or large enterprise, proof of payment for large enterprises, the completed self-assessment checklist and, where applicable, the REMA ESIA. Every decision stays with the institution's officer, who approves the application, returns it or proceeds to inspection.

Work done so far. IremboAI is live on IremboGov: it has handled more than 169,000 interactions in four months, over 70% in Kinyarwanda, with continuous evaluation, guardrails and jailbreak testing in place. Under the harmonised licensing service, businesses on IremboGov see the licence types and checklists, complete a 17-question self-assessment of eligibility, and submit one application that goes to all institutions, followed by a joint inspection visit; each certificate is issued as it becomes available. The service is being tested on single-product applications. No agentic AI has been built for these workflows, so the proof of concept builds on the harmonised service, its digitised checklists and the IremboAI platform.

What the proof of concept must add. The digitised checklists as the agents' knowledge base, expressed as versioned, machine-readable rules; a link to RRA for business classification; a set of past applications labelled by officers, to train and evaluate document checks; chat and voice support, including in Kinyarwanda; and secure integration with IremboGov and the institutions' review workflows through governed APIs.

Expected outcomes. Businesses understand the checklist and fix gaps before they submit; fewer re-upload rounds; more consistent document checks across institutions; and less officer time per application. Success measures, such as the share of applications complete at first submission, the number of re-upload rounds, officer review time and time to licence, will be set against a baseline agreed at inception.

Lot 3: Building permits on Kubaka (MININFRA and RHA)

Platform and owner. Kubaka (also called KubakaAI) is the national building permit platform, which replaced the earlier BPMIS system. It offers 11 citizen-facing services and is integrated with Irembo and the engineer registration systems. It is owned by the Ministry of Infrastructure (MININFRA), with the Rwanda Housing Authority (RHA) as technical process owner.

Problem. Permit applications grow as Kubaka coverage expands, while the number of reviewers stays flat. Officials still check site plans, floor plans and supporting documents against the Building Code and zoning rules by hand, so applicants learn of missing or non-compliant material late, resubmit repeatedly, and see approval times that are hard to predict. Applicants submit documents in both English and Kinyarwanda, so review must work in both languages. Reviewers can also interpret the same rules differently, and the 15-day decision window has been difficult to meet at current volumes. Citizen and engineer support queries are repetitive and add to the load.

Agentic solution sought. Two supervised agentic workflows:

- **Compliance review:** an agent that analyses submitted drawings and documents against the Building Code, flags missing or non-compliant elements with the rule cited and the drawing location, gives applicants feedback before human review, and gives the reviewing officer a compliance summary and checklist. Building Code and zoning rules are to be expressed as versioned, machine-readable rules and linked to zoning, cadastral and parcel data where those sources can be accessed.
- **Intelligent support:** improvements to the existing support agent in routing, escalation and hand-off quality for citizen and engineer queries.

The permit decision stays with an authorised official. Where the AI is uncertain or evidence is incomplete, the case goes to the reviewing officer, and nothing is sent to an applicant or approved without the officer's action.

Work done so far. The in-house team has built a compliance agent that reads submitted PDFs and drawings (OCR and computer vision) and checks them against the Building Code. It has analysed close to 10,000 documents. It is optional and in pilot, with reported accuracy of about 75% against a target of 90%; the main gaps are better rubrics, agreed with RHA, and annotated examples to improve how the model reads drawings.

What the proof of concept must add. A validated rule catalogue and a labelled, held-out evaluation set built with RHA reviewers; higher accuracy and faster processing; reviewer checklists and officer-approved corrective guidance for applicants; secure integration through governed APIs; and benchmarks against an agreed baseline.

Expected outcomes. More applications processed per reviewer; fewer resubmissions and shorter, more predictable approval times; consistent, transparent compliance checks with the rule cited; faster page analysis (the target is at least 60% faster, subject to benchmarking); and higher first-contact resolution of support cases.

Expectations common to all lots

- **Data stays in Rwanda.** Identifiable citizen, business and permit data must not leave the country. Where an external model is considered, data must first be anonymised, subject to approval from and registration with the national cybersecurity authority. Firms should plan for hosting on the national compute platform.
- **Human accountability.** Agents act within delegated authority; citizens and businesses confirm submissions and payments; officials make statutory and case decisions. Escalation paths to a person are part of the design.
- **Build on what exists.** Firms work alongside the owners' technical teams, reuse existing APIs, models and data, and do not replace working modules without justification.
- **Reusable components.** Language and voice, identity and delegated authority, API wrappers and tool registry, rules as code, and an evaluation harness are built so that other government services can reuse them, and are handed over as Government-owned assets.
- **Evidence for scale-up.** Results are measured against a baseline agreed at inception, by language and channel where relevant, to show whether the approach is ready to scale.