

AI-Enabled Government Efficiency and Trusted Digital Public Services for Rwanda

An implementation framework for responsible AI, administrative automation, and accountable service delivery across Rwanda's digital government ecosystem

AUTHOR	Samuel Abinsinguza · Founder, CentPol
SERIES	CentPol Papers No. 1
EDITION	Published 12 May 2026 · Version 1.0
TOPIC	AI governance · Rwanda
ONLINE	www.centpol.com/proposals/rwanda-ai-government-efficiency

A governance concept note on responsible AI, administrative automation, and trusted digital public services.

SUGGESTED CITATION

Abinsinguza, S. (2026, May 12). AI-Enabled Government Efficiency and Trusted Digital Public Services for Rwanda (CentPol Papers No. 1). CentPol. <https://www.centpol.com/proposals/rwanda-ai-government-efficiency>

Executive Summary

Rwanda should not build its next governance innovation agenda around blockchain as the central idea. The evidence is too uneven, and the public-sector value proposition is too narrow. Blockchain has useful applications in government, especially for tamper-evident audit logs, verifiable credentials, and procurement integrity, but it does not by itself make services faster, fairer, or more citizen-responsive. The stronger opportunity is to build an **AI-enabled public-service operating layer** that improves how the government receives requests, reads documents, routes cases, supports frontline officers, detects bottlenecks, and reports performance.

This concept note proposes the **Rwanda AI-Enabled Government Efficiency Framework**, a practical national programme for applying responsible artificial intelligence to high-volume, low-to-medium-risk government workflows. The proposal is deliberately not a technology showcase. It is an execution framework for public administration. Its purpose is to help Rwanda move from digital access to digital performance: from services being available online to services being processed intelligently, consistently, and transparently.

Rwanda is unusually well positioned for this shift. The country already has a mature digital-government trajectory, including the IremboGov public-service portal, which began in 2015 and provides access to more than 90 government services according to UN DESA, while Irembo reports that more than 100 public services have been digitised through its platform.^{1 2} Rwanda also has a National Artificial Intelligence Policy that identifies AI as a tool for public-service improvement, economic transformation, and responsible innovation.³ This creates a rare alignment between existing digital infrastructure and future-oriented policy. The missing link is not another broad strategy. The missing link is a governed implementation engine.

Strategic proposition: Rwanda can become one of Africa's first AI-enabled public-service states by applying responsible AI to the administrative layer of government: document processing, case triage, citizen support, service analytics, fraud-pattern detection, workflow routing, and performance monitoring, while using tamper-evident audit technologies only where trust and record integrity require them.

The framework should be implemented through a dedicated **Government AI Efficiency Unit** operating under a clear public-interest mandate and in partnership with relevant ministries, agencies, regulators, academic institutions, and local technology firms. The first phase should focus on a controlled set of priority services where AI can reduce delays without replacing legal accountability. High-risk decisions, such as benefit denial, law enforcement action, immigration refusal, or punitive enforcement, should remain outside automated decision-making unless strong safeguards, human review, appeal mechanisms, and legal authorization are in place.

The proposed initial investment envelope is **US\$45–75 million over three years**, structured as a phased public-sector modernization facility rather than a single software procurement. The investment would fund secure AI infrastructure, data readiness, service redesign, model development, human-in-the-loop systems, cybersecurity, algorithmic accountability, civil-service training, local vendor participation, and a limited trusted-audit layer for procurement, credentials, and registry integrity. The ex-

pected results are not speculative “AI transformation” claims. They are measurable administrative outcomes: faster processing times, fewer manual handoffs, better citizen support, improved service-level compliance, stronger auditability, and greater managerial visibility across priority government services.

Why the Original Blockchain Idea Should Be Reframed

The starting question was whether Rwanda should pursue **blockchain-based governance systems** for transparent government services. The research does not support blockchain as the flagship concept. The most credible government blockchain examples show that the technology is useful in specific integrity functions, not as a general operating system for government.

Estonia is often presented as the model for blockchain government, but this is frequently misunderstood. Estonia’s KSI blockchain is not a cryptocurrency-style public blockchain replacing public institutions. It is a cryptographic integrity system that helps verify that government data and system logs have not been altered.⁴ That lesson is important for Rwanda. Estonia’s success came from strong digital government, secure registries, interoperable systems, and institutional trust; blockchain strengthened the integrity layer. It did not substitute for state capacity.

Other public-sector blockchain experiments show similar limits. Land-registry pilots in countries such as Georgia and Honduras demonstrate that a ledger cannot fix weak underlying records, contested legal rights, unclear institutional authority, or poor dispute-resolution mechanisms.⁵ Public procurement is more promising because it involves public money, multiple actors, and audit trails, but even here the World Economic Forum’s work emphasizes that blockchain must sit inside broader procurement reform, legal design, user incentives, and oversight.⁶

The conclusion is simple but important: blockchain is a **trust technology**, not a government-performance technology. Rwanda’s immediate opportunity is not to create a distributed ledger for government services. It is to make existing and future digital services work faster, more intelligently, and more transparently. That is an AI problem, a data-governance problem, a workflow problem, and an institutional-delivery problem.

PROPOSED TECHNOLOGY ROUTE	EVIDENCE FROM GLOBAL EXPERIENCE	PRACTICAL IMPLICATION FOR RWANDA
Broad blockchain governance platform	Weak evidence of scaled public-service transformation; high implementation complexity	Not recommended as the lead concept
Blockchain for registries and procurement audit trails	Plausible where tamper-evidence and multi-party record integrity are needed	Include only as a targeted trust module
AI for public-service efficiency	Stronger evidence from OECD and World Bank examples in service delivery, document review, process automation, and decision support	Recommended as the main framework
AI with auditability safeguards	Best fit where efficiency must be matched with public trust, human review, and traceability	Recommended architecture for Rwanda

The Strategic Problem: Rwanda Has Digitised Access; the Next Frontier Is Administrative Intelligence

Digital government often begins by moving forms, payments, and applications online. That is necessary, but it is not the end of transformation. Once services are digitised, the deeper question becomes: what happens behind the screen? Does the application move quickly? Is the document checked consistently? Is the citizen informed? Can managers see where delays are forming? Can ministries learn from service data? Can the system identify duplication, error, fraud patterns, or capacity bottlenecks before they become public frustration?

Rwanda's existing digital-service foundation gives the country a strong platform for this next step. UN DESA describes Irembo as Rwanda's national online service portal, created through a partnership between the Government of Rwanda and RwandaOnline Platform Ltd, and notes that it has grown since 2015 to provide access to more than 90 government services online.⁷ Irembo itself reports that IremboGov provides access to over 100 public services and has produced measurable service-access improvements, including reported reductions in service access time and millions of working hours saved by citizens.⁸ These achievements matter because they mean Rwanda is not starting from zero. It already has channels, users, transactions, agencies, and data flows.

The next constraint is different. Rwanda now needs to improve the **administrative layer** that sits behind digital portals. Many governments digitise the front door while leaving the back office manual, fragmented, document-heavy, and dependent on overburdened staff. AI can help close that gap if it is applied carefully. It can read and classify documents, flag missing information, route cases to the right officer, summarize applications, detect duplicate submissions, support citizen-service agents, translate common questions, monitor service-level agreements, and identify where workloads are becoming unmanageable.

This is not about replacing public servants. It is about giving them better tools. In a well-designed public-sector AI programme, the human officer remains responsible for judgment, discretion, fairness, and final decision-making. AI handles repetitive pattern work, information retrieval, triage, summarization, and early warning.

The Evidence for AI in Government Efficiency

The international evidence base is stronger for AI-enabled administrative efficiency than for broad blockchain governance. OECD's 2025 work on AI in public-service design and delivery reports that 67% of OECD countries use AI to improve public-service design and delivery, particularly for process automation, resource allocation, personalized services, and decision support.⁹ These are not abstract use cases. They are the ordinary mechanics of government made faster and more manageable.

Three examples are especially relevant for Rwanda. Austria's BRISE-Vienna project uses AI-assisted verification to support building-permit processes, helping reviewers check applications against regulatory requirements.¹⁰ Catalunya applies AI, robotic process automation, and business-process management to

reduce administrative burden in energy-poverty reporting.¹¹ Greece's Hellenic Cadastre uses AI to read, categorize, and legally assess property contracts; OECD reports that the system reduced assessment times from several hours to less than ten minutes and lowered costs from EUR 15 to EUR 0.11 per case.¹²

The World Bank's work on trustworthy AI reinforces the opportunity while warning against careless deployment. It notes that AI can improve public governance, service delivery, administrative efficiency, and data-driven decision-making, but that governments must manage risks such as bias, privacy violations, weak accountability, opacity, and regulatory fragmentation.¹³ The Open Government Partnership, Ada Lovelace Institute, and AI Now Institute reach a similar conclusion from an accountability perspective: algorithmic systems can improve efficiency and reduce costs, but they can also cause harm unless governments provide transparency, impact assessments, oversight, and public accountability mechanisms.¹⁴

The implication for Rwanda is not “automate government.” That would be a dangerous simplification. The correct implication is more disciplined: Rwanda should deploy **governed AI in administrative workflows where benefits are clear, risks are manageable, and human accountability remains intact.**

Rwanda's Advantage: Policy Alignment, Digital Foundations, and Institutional Readiness

Rwanda's National AI Policy gives the concept note a strong policy anchor. It positions AI as a tool for national development and sets out the need for skills, data, infrastructure, governance, and responsible adoption.¹⁵ This matters because many AI proposals fail by arriving before the institutional conversation is ready. Rwanda has already begun that conversation. The concept note should therefore be written not as a speculative idea, but as an implementation instrument for a policy direction the country has already endorsed.

The Smart Rwanda Master Plan also supports the framework. It treats ICT as an enabler of productivity, service delivery, and national transformation, and it emphasizes shared infrastructure, e-government, and digital transformation as part of Rwanda's development model.¹⁶ The new framework would extend that logic. Smart Rwanda built the foundations for digital services; the AI-enabled efficiency framework would improve how those services perform.

Rwanda also has a practical advantage in scale. Its governance system is small enough to coordinate and ambitious enough to move quickly, yet complex enough that service bottlenecks still matter. A phased AI programme can begin with a few priority service families, prove measurable gains, and then expand through reusable tools, shared standards, and trained civil-service teams.

What Is Proposed

The proposal is to establish a **Rwanda AI-Enabled Government Efficiency Framework** with five integrated components: a Government AI Efficiency Unit, a priority-service transformation pipeline, a responsible AI and data-governance layer, a trusted auditability module, and a public-sector capability programme.

This framework should not be procured as one large platform. That would be risky, expensive, and likely to lock Rwanda into a vendor-driven architecture. It should be built as a **modular public-service operating layer**: shared standards, reusable tools, secure infrastructure, reference models, approved vendors, agency-level implementation teams, and a central delivery unit that measures performance.

COMPONENT	FUNCTION	WHY IT MATTERS
Government AI Efficiency Unit	Coordinates AI use cases, standards, procurement support, technical review, and performance monitoring across agencies	Prevents fragmented pilots and turns AI into a managed public-administration programme
Priority-service transformation pipeline	Selects high-volume workflows for AI-assisted redesign, automation, and human-in-the-loop processing	Keeps the programme focused on measurable service gains rather than technology experimentation
Responsible AI and data-governance layer	Provides impact assessments, privacy rules, model documentation, bias testing, human review, and appeal mechanisms	Protects citizens and makes the programme credible to funders, regulators, and public institutions
Trusted auditability module	Uses digital signatures, tamper-evident logs, verifiable credentials, and limited blockchain-style integrity proofs where needed	Preserves the useful part of the blockchain idea without making blockchain the main architecture
Public-sector capability programme	Trains civil servants, service managers, procurement officers, data stewards, and legal teams	Ensures that AI is governed and operated by institutions, not only by vendors

Where Rwanda Should Start

The first phase should avoid high-stakes automated decision-making. Rwanda should begin where AI can assist, not decide. This distinction is essential. The safest early use cases are those in which AI improves speed, classification, communication, or analysis while a human officer retains responsibility for final decisions.

The most promising initial use cases are **document-heavy, high-volume, rules-based, and service-facing**. These include permit pre-screening, certificate applications, visa and travel-service support, civil-status service routing, business-registration assistance, land-document intake, complaint triage, call-centre support, procurement milestone tracking, and internal correspondence classification. These are places where government spends time reading, sorting, checking, routing, responding, and escalating.

USE-CASE FAMILY	AI ROLE	HUMAN ROLE	EARLY SUCCESS METRIC
Document intake and pre-screening	Classify documents, detect missing fields, flag inconsistencies, summarize submissions	Review flagged cases, approve completeness, make final administrative judgment	Reduction in incomplete applications reaching officers
Citizen-service support	Provide multilingual answers, guide users through service requirements, escalate complex cases	Handle appeals, sensitive cases, and unresolved questions	Faster response times and higher first-contact resolution
Workflow routing	Direct cases to the right unit based on service type, urgency, geography, and required expertise	Override routing and resolve exceptions	Fewer misrouted cases and shorter internal handoff times
Service analytics	Detect bottlenecks, workload imbalances, repeated errors, and service-level breaches	Redesign processes and allocate staff	Improved compliance with service-level targets
Procurement and audit support	Track procurement milestones, compare documents, flag missing approvals, preserve tamper-evident event logs	Conduct procurement decisions, investigations, and legal review	Stronger audit trails and fewer documentation gaps
Registry and credential verification	Support verification of certificates, licenses, and registry extracts through digital signatures or tamper-evident proofs	Resolve disputes and authorize corrections	Faster authenticity checks and reduced manual verification burden

This first wave should be selected through evidence, not enthusiasm. Each candidate service should pass a readiness test: enough volume to matter, enough digitized data to support AI, clear process rules, manageable legal risk, agency leadership commitment, and a measurable baseline. If those conditions are absent, AI should wait until the service is redesigned and the data is improved.

The Trust Architecture: Responsible AI First, Blockchain Only Where It Adds Value

The central risk of AI in government is not technical failure alone. It is administrative overreach: a system that makes decisions people cannot understand, contest, or correct. Rwanda's framework should therefore make **trust architecture** part of the design from the beginning.

Every AI use case should be classified by risk. Low-risk tools, such as internal document search or service FAQs, can move faster. Medium-risk tools, such as case triage or fraud-pattern detection, should require human review and model monitoring. High-risk tools, such as automated denial of benefits, enforcement action, policing, surveillance, immigration refusal, or eligibility determination, should be excluded from the first phase unless a specific law, appeal process, independent review, and impact assessment are in place.

Blockchain enters here only as a narrow safeguard. For procurement, registries, and credentials, Rwanda may use tamper-evident logs or distributed-ledger-style integrity proofs to show that a record, approval, or milestone has not been altered after the fact. But the primary governance tools should re-

main simpler and more familiar: digital signatures, public audit trails, role-based access controls, secure APIs, privacy-by-design, logs, dashboards, and independent oversight.

SAFEGUARD	WHAT IT DOES	REQUIRED IN FIRST PHASE
Algorithmic impact assessment	Evaluates risk, affected groups, data quality, legal basis, and mitigation before deployment	Yes
Human-in-the-loop review	Keeps accountable public officers responsible for final decisions	Yes
Model documentation	Records data sources, model purpose, performance limits, and known risks	Yes
Citizen explanation and appeal pathway	Gives affected users a route to understand and challenge outcomes	Yes for service-impacting use cases
Public AI use-case registry	Lists approved public-sector AI systems and their purpose	Yes
Privacy and data-minimization controls	Limits collection, retention, and unnecessary sharing of personal data	Yes
Tamper-evident audit logs	Protects selected high-value records from undetected alteration	Yes, where justified
Independent review	Provides external scrutiny for high-risk or sensitive systems	Required before high-risk expansion

Institutional Model

The framework should be anchored by a **Government AI Efficiency Unit** with a mandate to coordinate, not centralize every implementation. It should operate as a delivery and assurance body. Its work would include selecting use cases, setting standards, supporting agencies, reviewing procurement, managing shared infrastructure, measuring outcomes, and ensuring that AI systems comply with Rwanda's responsible AI principles.

The unit should work closely with the ministry responsible for ICT and innovation, public-service institutions, data-protection authorities, service-delivery agencies, cybersecurity bodies, universities, and the private sector. Rwanda should avoid building a programme that depends entirely on foreign vendors. International partners can provide models, tools, cloud infrastructure, technical assistance, and financing, but Rwanda should deliberately build domestic capacity among civil servants, local firms, universities, and GovTech entrepreneurs.

INSTITUTION OR ACTOR	PROPOSED ROLE
Central ICT and innovation leadership	Policy alignment, national coordination, approval of AI governance standards
Public-service and civil-service institutions	Workflow redesign, staff training, adoption management, performance accountability

INSTITUTION OR ACTOR	PROPOSED ROLE
Service-delivery agencies	Identification of priority use cases, business ownership, human review, operational deployment
Data-protection and cybersecurity authorities	Privacy, security, compliance review, incident response, audit standards
Universities and research centres	Model evaluation, skills development, local language resources, independent research
Local technology firms	Build and adapt workflow tools, user interfaces, analytics, and service integrations
Development partners	Finance, technical assistance, procurement support, international benchmarking
Citizens and civil society	Feedback, transparency, usability testing, accountability, and grievance pathways

Financing and Investment Architecture

The proposed first-phase facility should be sized between **US\$45 million and US\$75 million over three years**. This is an indicative planning envelope, not a final budget. The amount is designed to be large enough to build shared capability and deliver measurable transformation across priority services, but disciplined enough to avoid a mega-platform procurement.

The financing should be blended across government budget allocations, development-partner support, concessional financing, technical assistance, and controlled private-sector participation. The goal is not simply to buy AI tools. The goal is to create public-sector capacity to identify problems, govern technology, improve workflows, and measure results.

INVESTMENT LINE	INDICATIVE ALLOCATION	PURPOSE
Shared AI infrastructure and secure integration layer	US\$10–15 million	Hosting, APIs, secure data access, model operations, monitoring, and cybersecurity
Priority-service redesign and implementation	US\$12–20 million	Workflow mapping, automation, pilots, agency deployments, user testing, and integration with existing platforms
Responsible AI governance and assurance	US\$5–8 million	Impact assessments, audits, model documentation, legal review, privacy controls, public registry, and oversight mechanisms
Public-sector capability and change management	US\$6–10 million	Training for civil servants, managers, procurement teams, lawyers, data stewards, and frontline service staff
Local innovation and vendor participation fund	US\$5–10 million	Competitive grants and contracts for local firms and universities to build Rwanda-specific solutions
Trusted auditability and verification module	US\$3–6 million	Tamper-evident logs, digital signatures, verifiable credentials, and limited ledger-based integrity proofs where justified
Monitoring, evaluation, and citizen feedback	US\$4–6 million	Baselines, service dashboards, independent evaluations, citizen surveys, and public reporting

The financing model should reward measurable service improvement. Vendors and implementing agencies should be evaluated not only on system delivery, but on reductions in processing time, improved service-level compliance, fewer manual handoffs, better user satisfaction, reduced error rates, and documented safeguards.

Twelve-Month Implementation Plan

The first year should prove that the framework can deliver real service improvements without creating new governance risks. It should not attempt to automate everything. The right first year is narrow, visible, and disciplined.

PERIOD	PRIORITY ACTION	OUTPUT
Months 1–2	Establish the Government AI Efficiency Unit and approve responsible AI operating principles	Mandate, governance charter, risk-classification model, and initial staffing
Months 2–3	Select 8–12 candidate services using readiness criteria	Prioritized service pipeline with baselines, data-readiness review, and legal-risk screening
Months 3–4	Conduct workflow mapping and data-quality assessment	Service maps, bottleneck analysis, data inventories, and implementation requirements
Months 4–6	Build the first shared AI assistance tools	Document classification, service chatbot prototype, case-routing support, and analytics dashboard
Months 6–8	Launch 3–5 controlled pilots with human review	Pilot deployments in selected services with monitoring and citizen feedback
Months 8–10	Establish public AI use-case registry and audit protocols	Registry, model cards, impact-assessment summaries, and audit-log standards
Months 10–12	Evaluate pilots and approve scale-up package	Independent review, performance report, revised safeguards, and year-two expansion plan

By the end of twelve months, Rwanda should be able to answer five questions with evidence: which services improved, how much time was saved, whether users experienced better service, whether staff trust the tools, and whether safeguards worked. If the answer is unclear, the programme should slow down and fix the design before scaling.

Performance Dashboard

The framework should be judged through public-service outcomes, not technical novelty. A dashboard should be published quarterly for approved service areas, with internal detail available to responsible agencies and oversight institutions.

INDICATOR	WHY IT MATTERS	FREQUENCY
Average processing time for pilot services	Measures whether AI is improving service delivery	Monthly

INDICATOR	WHY IT MATTERS	FREQUENCY
Percentage of applications flagged as incomplete before officer review	Shows whether AI reduces wasted administrative time	Monthly
Service-level compliance rate	Tracks whether agencies meet promised timelines	Monthly
Citizen satisfaction score for AI-supported services	Ensures efficiency does not come at the cost of user experience	Quarterly
Human override rate	Shows whether AI recommendations are useful or frequently corrected	Monthly
Appeal or complaint rate	Detects possible harm, confusion, or unfair outcomes	Monthly
Model performance and drift	Monitors whether tools remain reliable over time	Quarterly
Number of trained public servants	Measures institutional capacity rather than vendor dependence	Quarterly
Number of local firms or universities engaged	Tracks ecosystem development	Quarterly
Number of approved use cases in public AI registry	Supports transparency and accountability	Quarterly
Audit-log completeness for selected procurement or registry workflows	Measures integrity of high-value records	Quarterly

Risk and Mitigation

The largest risk is not that Rwanda tries AI. The larger risk is that it tries AI in a way that is either too timid to matter or too aggressive to trust. The framework therefore needs practical discipline.

RISK	CONSEQUENCE	MITIGATION
Technology-first implementation	Tools are deployed without solving real service bottlenecks	Select use cases through workflow evidence, volume data, and agency readiness
Bias or unfair treatment	Citizens may receive inconsistent or harmful outcomes	Require impact assessments, bias testing, human review, and appeal pathways
Privacy and data misuse	Sensitive citizen data may be over-collected or misused	Apply data minimization, access controls, retention limits, and privacy review
Vendor lock-in	Government becomes dependent on one provider or closed platform	Use open standards, modular procurement, data portability, and local capacity building
Weak civil-service adoption	Tools are ignored or resisted by frontline staff	Train staff early, involve them in design, and measure workload reduction
Over-automation of high-risk decisions	Public trust may be damaged by opaque or unfair automated decisions	Exclude high-risk automated decisions in phase one and require legal authorization for expansion
Poor data quality	AI tools produce unreliable outputs	Invest in data readiness before deployment and use AI only where records are sufficiently structured

RISK	CONSEQUENCE	MITIGATION
Blockchain overreach	A narrow audit tool becomes an expensive distraction	Limit ledger-style tools to specific integrity cases with clear justification
Cybersecurity exposure	Integrated systems create new attack surfaces	Build security architecture, monitoring, incident response, and independent testing from the start

Expected Outcomes

The framework's outcomes should be expressed in administrative performance rather than vague transformation language. Within three years, Rwanda should aim to deploy AI assistance across **30–50 priority public-service workflows**, train **8,000–12,000 public officials and service staff** in responsible AI use and data-enabled service management, establish a public AI use-case registry, and produce quarterly performance reports for AI-supported services. These are proposed programme targets, not claims of existing performance.

The deeper outcome is institutional. Rwanda would develop a repeatable method for improving services: identify bottlenecks, prepare data, classify risk, deploy AI assistance, preserve human accountability, measure performance, and scale only what works. That method would be more valuable than any single model or platform.

Conclusion

Rwanda does not need to prove that it can follow a global technology trend. It has already shown that it can build digital public-service systems. The next step is more demanding: making those systems intelligent, responsive, accountable, and useful in the everyday work of government.

A blockchain-led concept would be narrower than the problem. It would speak the language of transparency, but not necessarily the language of service performance. An AI-enabled government efficiency framework is more honest about the work ahead. It recognizes that citizens experience government through waiting times, forms, approvals, explanations, corrections, and follow-up. If AI can reduce those frictions while preserving human judgment and public accountability, it becomes more than a technology project. It becomes public-sector capacity.

The right proposition for Rwanda is therefore not “government on blockchain.” It is **trusted, intelligent government services**: AI where it improves work, humans where judgment matters, and tamper-evident records where public trust requires proof.

References

1. United Nations Department of Economic and Social Affairs, “Irembo: Rwandan Government E-Service Portal,” <https://publicadministration.desa.un.org/good-practices-for-digital-government/compendium/irembo-rwandan-government-e-service-portal>
2. Irembo, “IremboGov — All the public services you need in one place,” <https://irembo.com/>

3. Government of Rwanda, Ministry of ICT and Innovation, *The National AI Policy*, 2023, https://www.ictworks.org/wp-content/uploads/2023/12/Rwanda_Artificial_Intelligence_Policy.pdf
4. e-Estonia, “KSI Blockchain,” <https://e-estonia.com/solutions/cyber-security/ksi-blockchain/>
5. OECD, *Blockchains Unchained: Blockchain Technology and its Use in the Public Sector*, OECD Working Papers on Public Governance, 2018, <http://uatpfmkin.cnkonline.in/sites/default/files/2020-02/185.%20Blockchains%20Unchained.pdf>
6. World Economic Forum, *Exploring Blockchain Technology for Government Transparency: Blockchain-Based Public Procurement to Reduce Corruption*, Supplementary Research, https://www3.weforum.org/docs/WEF_Blockchain_Government_Transparency_Report_Supplementary%20Research.pdf
7. United Nations Department of Economic and Social Affairs, “Irembo: Rwandan Government E-Service Portal,” <https://publicadministration.desa.un.org/good-practices-for-digital-government/compendium/irembo-rwandan-government-e-service-portal>
8. Irembo, “IremboGov — All the public services you need in one place,” <https://irembo.com/>
9. OECD, *Governing with Artificial Intelligence: The State of Play and Way Forward in Core Government Functions*, “AI in public service design and delivery,” 2025, https://www.oecd.org/en/publications/governing-with-artificial-intelligence_795de142-en/full-report/ai-in-public-service-design-and-delivery_09704c1a.html
10. OECD, *Governing with Artificial Intelligence: The State of Play and Way Forward in Core Government Functions*, “AI in public service design and delivery,” 2025, https://www.oecd.org/en/publications/governing-with-artificial-intelligence_795de142-en/full-report/ai-in-public-service-design-and-delivery_09704c1a.html
11. OECD, *Governing with Artificial Intelligence: The State of Play and Way Forward in Core Government Functions*, “AI in public service design and delivery,” 2025, https://www.oecd.org/en/publications/governing-with-artificial-intelligence_795de142-en/full-report/ai-in-public-service-design-and-delivery_09704c1a.html
12. OECD, *Governing with Artificial Intelligence: The State of Play and Way Forward in Core Government Functions*, “AI in public service design and delivery,” 2025, https://www.oecd.org/en/publications/governing-with-artificial-intelligence_795de142-en/full-report/ai-in-public-service-design-and-delivery_09704c1a.html
13. World Bank, *Building Trustworthy Artificial Intelligence: Frameworks, Applications, and Self-Assessment for Readiness*, 2025, <https://documents1.worldbank.org/curated/en/099805309022518222/pdf/IDU-1e9a05ec-ab52-425d-a18a-c9c91ed04a37.pdf>
14. Ada Lovelace Institute, AI Now Institute, and Open Government Partnership, *Algorithmic Accountability for the Public Sector*, 2021, <https://www.opengovpartnership.org/wp-content/uploads/2021/08/executive-summary-algorithmic-accountability.pdf>
15. Government of Rwanda, Ministry of ICT and Innovation, *The National AI Policy*, 2023, https://www.ictworks.org/wp-content/uploads/2023/12/Rwanda_Artificial_Intelligence_Policy.pdf
16. Government of Rwanda, Ministry of ICT and Innovation, *Smart Rwanda 2020 Master Plan*, 2015, https://www.minict.gov.rw/fileadmin/user_upload/minict_user_upload/Documents/Policies/SMART_RWANDA_MASTERPLAN.pdf