INDEPENDENT STATE OF PAPUA NEW GUINEA

National Artificial Intelligence Policy

Consultation draft

9 November 2026

Five years from approval

$$DG \times DPI^{AI}$$

Digital government and shared public infrastructure accelerated by AI

Department of Information and Communications Technology

This policy sets the national direction for using AI to improve public services, reduce waste and misuse of public money, and help Papua New Guinean businesses grow. It explains how government, businesses, educators and communities will participate and how people will be protected.

Ministerial Foreword

People should be able to obtain public services without unnecessary delay, expense or confusion. Businesses should be able to work more efficiently and reach new markets. Public money should be easier to track and harder to misuse. These are the results we seek from artificial intelligence.

Our approach is $DG \times DPI^{AI}$: improve Digital Government through shared Digital Public Infrastructure, accelerated by AI. Government will lead through controlled pilots and phased adoption, expanding when benefits and safeguards are demonstrated. Other sectors will follow published schedules reflecting their needs and readiness.

People remain at the centre of this change. A person will remain responsible for decisions affecting rights, liberty, health and access to essential services. Human assistance will remain available. Our languages, cultural knowledge and public information must be respected and protected.

I call on public bodies, businesses, educators and communities to help shape and deliver this policy. Its success will be judged by improvements people can see in their daily lives.

HON ROBERT NAGURI, MP

Minister for Information and Communications Technology

Secretary Statement

The Department of Information and Communications Technology will coordinate this policy across government and all sectors. Our task is to translate national goals into costed plans, clear responsibilities and services that people can trust.

The first six months will establish baselines and costed plans connecting service reform, shared infrastructure and AI. We will coordinate these investments, support local skills and businesses, and select models for capability and cost, subject to security and law. Agencies will demonstrate how AI improves delivery through common foundations.

Agencies will report progress every six months. A comprehensive annual report and independent evaluations will show what is working, what costs more than expected and what must change. Public participation will help keep implementation grounded in the needs of Papua New Guineans.

STEVEN MATAINAHO

Secretary

Executive Summary

Artificial intelligence (AI) uses information to produce predictions, content, recommendations or decisions. It can make mistakes. People and organisations remain accountable for how it is used.

What this policy aims to achieve

- Faster, more reliable public services that people can access in their communities.
- Less corruption, waste and misuse of public money, supported by better records and lawful investigations.
- Stronger PNG businesses, better productivity and new employment opportunities.

$DG \times DPI^{AI}$ is the core concept. Digital Government (DG) improves services. Shared Digital Public Infrastructure (DPI), including digital identity, payments and secure data exchange, supports delivery. AI amplifies these capabilities to improve daily life and expand the Digital Economy.

The policy runs for five years from approval. Priorities are government administration; justice and public safety; health; agriculture, fisheries and forestry; education; and tourism and hospitality.

What will change

Government bodies and sector leads will submit costed plans within six months of approval, aligned with national planning and budget cycles. Ready agencies will run controlled pilots. Organisations must assess opportunities and adopt suitable AI with demonstrated benefits and safeguards, or justify an exemption or deferral. Non-government deadlines will follow consultation and readiness assessments. New duties require lawful instruments.

GovAoS, the Government Agentic Operating System, will connect official records, correspondence and AI assistance to shared infrastructure. Public servants will teach and test GovAI Agents using approved knowledge. Agencies will join as their staff, systems, funding and safeguards become ready.

What you can expect

- Human assistance and a non-digital route to essential services, without disadvantage because you cannot use AI.
- A human decision-maker for matters affecting your rights, liberty, health or access to essential services.
- Clear notice when you interact with AI, and a way to question an AI-supported decision and seek review.
- Public services designed for English, Tok Pisin and Hiri Motu from the first phase, with other PNG languages added progressively.
- Protection of personal information and informed consent from legitimate knowledge holders before traditional knowledge or community-held cultural material is used.

Who will be accountable

The ICT department will coordinate policy and is proposed to become the Department of Digital Development and Artificial Intelligence. NICTA will regulate AI under its proposed expanded mandate. SevisPNG Ltd is proposed to operate shared services under government ownership and control, subject to the required approvals.

The Government will maintain a public AI register and complaint routes, publish six-monthly progress and annual reports, and commission independent mid-term and final evaluations. This policy also guides AI legislation.

Contents

Ministerial Foreword 2
Secretary Statement 3
Executive Summary 4
1. Introduction 6
2. Strategic Direction 8
3. Policy Pillars 9
4. Institutional Arrangements and Accountability 22
5. Implementation and Reform Roadmap 24
6. Financing and Resource Mobilisation 27
7. Monitoring and Evaluation 28
Annex A Implementation Responsibilities 29
Annex B Results framework 30
Annex C Legal and institutional reform programme 32
Annex D Definitions 35
Annex E Proposed GovAoS Agents 39
Annex F Sources and references 40

1. Introduction

1.1 Purpose and coverage

This policy guides AI development, supply, adoption and use throughout Papua New Guinea and provides the policy basis for AI legislation. It covers national and subnational government, state enterprises, businesses, educational and research institutions, civil society and communities. It applies to AI obtained locally or from overseas when it is deployed in PNG or used to deliver services to people in PNG, within the reach of PNG law.

Mandatory adoption duties apply to organisations through phased, lawful arrangements. Citizens are not required to buy or use AI to obtain essential services. Research and small community activities will be treated proportionately to their risks and capacity.

Implementation must respect the constitutional independence of Parliament, the judiciary and oversight bodies. Applications in the Autonomous Region of Bougainville must respect applicable constitutional and legal arrangements and be coordinated with the Autonomous Bougainville Government. State enterprises and other bodies must be covered through the instruments that lawfully apply to them.

The policy runs for five years from National Executive Council approval. The indicative planning horizon is 2027–2031; the approval instrument will confirm the start and end dates. Month 0 means the approval date, and all month-based deadlines run from that date. Implementation plans will use the confirmed policy period.

1.2 How to read this policy

“Will” states a government commitment. “Must” identifies a required standard to be implemented through existing powers, lawful administrative directions, contracts or new legislation, as appropriate. This policy alone creates no offence, fine, regulatory power or spending authority. Existing legal protections continue to apply.

The coordinating department means the Department of Information and Communications Technology (DICT) and its lawful successor responsible for ICT and AI. This policy sets out the proposed change to the Department of Digital Development and Artificial Intelligence (DDDAI). References to NICTA’s new AI powers apply only after those powers lawfully commence.

1.3 Situation analysis and rationale

The Digital Government Plan 2023–2027 prioritises coordinated government cloud services, secure data exchange and a common portal. It identifies and pursues simpler processes, shared services, accessibility and reliable information as foundations for better delivery. These priorities provide the starting point for AI adoption. [4]

Government bodies, businesses and communities have different levels of connectivity, electricity, affordable access, reliable records and staff capability. The first six-month assessment will establish these differences by sector and location, alongside service waiting times, backlogs, user effort, costs and errors. Agencies must distinguish measured results from estimates and identify gaps in the evidence.

A phased national mandate is intended to make every covered organisation assess useful opportunities, plan capability and account for results. Shared requirements support reuse and consistent protection across services. Adoption must follow demonstrated benefit, readiness and lawful safeguards. An immediate uniform deadline would not reflect different capacities; voluntary uptake alone would not provide the same common planning obligation. Consultation and evaluation will test costs and burdens, especially for smaller organisations.

1.4 Policy and Legal Alignment

Implementation must respect the Constitution, its National Goals and Directive Principles, rights, institutional arrangements and access to remedies.

The Digital Government Act 2022 provides for public-sector digital coordination and the Public Service ICT Steering Committee. The National Information and Communications Technology Act 2009 establishes NICTA. The proposed changes to statutory functions and committee arrangements require the appropriate legislative reform. [2–3]

The policy supports Vision 2050, successive Medium Term Development Plans, the Digital Government Plan 2023–2027 and its successors, and the Digital Transformation Policy 2020. Implementation will align with the Government Cloud Policy 2023, National Cyber Security Policy 2021, Data Governance and Data Protection Policy 2024 and National Digital ID Policy 2025. [4–5]

$DG \times DPI^{AI}$ connects this policy to PNG's wider digital reform programme. Related strategies and institutional proposals must be reconciled through the appropriate government decisions. [6]

UNESCO's Recommendation on the Ethics of Artificial Intelligence and the OECD AI Principles inform responsible, inclusive and trustworthy use. Council of Europe assessment guidance and comparative criminal-justice resources inform the relevant safeguards and legislative review. These resources do not replace PNG law. Annex F identifies the sources and their status. [7–12]

1.5 Public Consultation and Finalisation

Citizens, businesses, workers, educators, professional bodies and communities are invited to comment on this consultation draft. The coordinating department will publish the submission opening and closing dates, designated email and postal contacts, public meeting arrangements and accessible submission options in an official notice on www.ict.gov.pg and through public-information channels. No deadline or submission channel is established by this draft alone.

Consultation will provide plain-language explanations and opportunities for written, oral and assisted submissions, including participation in Tok Pisin and Hiri Motu. Engagement will include provinces, districts, the Autonomous Bougainville Government, people with disability, women, young people, small businesses and cultural knowledge holders. People will be able to comment without an AI account or smartphone.

The coordinating department will publish a summary of the issues raised and explain how the final policy responds, while protecting personal and restricted information. It will identify material changes to the mandate, institutional proposals, safeguards, exemptions, costs and implementation arrangements for the Government's consideration before approval.

2. Strategic Direction

2.1 Vision and Mission

Our vision is a Papua New Guinea where digital government, shared infrastructure and trusted AI improve daily life. Our mission is beneficial adoption across all sectors, protecting people, public resources and national control.

2.2 Outcomes and Objectives

The first three years prioritise better public services, reduced corruption and waste, and stronger local businesses and employment. The five-year programme builds oversight, infrastructure, skills, responsible use, innovation and sector delivery.

Baselines are due within six months, followed by published sector targets with owners, budgets, deadlines and evidence, advancing Digital Government and the Digital Economy.

2.3 Principles

Decisions will follow public benefit, human responsibility, inclusion, privacy, cultural respect, safety, transparency, sovereign control and value. Plan reform, shared infrastructure and AI together; deploy with the required foundations and safeguards. Match controls to risk and assess fairness, energy, water and equipment impacts.

2.4 Core concept and delivery model

$$DG \times DPI^{AI}$$

Digital Government multiplied by Digital Public Infrastructure to the power of Artificial Intelligence.

Element	Role in national development	What implementation must show
DG Digital Government	The public purpose and reform agenda: accessible services, effective administration, public integrity and economic opportunity.	A service owner, simpler processes, capable staff and measurable improvements for citizens and businesses.
DPI Digital Public Infrastructure	Shared, secure and interoperable foundations for identity, payments, data exchange and core registers. [9]	Reuse of approved components, lawful access, reliable records and services that work across institutions.
AI Artificial Intelligence	Capabilities that amplify speed, reach, quality and responsiveness across services and their foundations.	Locally validated models and agents, human responsibility, controlled actions and verified additional benefit.

Multiplication means reform and shared foundations advance together. The AI exponent expresses the ambition for exponential improvement through reusable capability. Benefits must be verified: this is a policy model, not a mathematical forecast or automatic savings claim.

GovAoS applies the concept through shared records and workflows using approved DPI interfaces. GovAI Agents provide controlled assistance. Sovereign training makes approved PNG knowledge and tested capability reusable, subject to permissions and revalidation. DPI also supports services outside GovAoS.

Connectivity, power, cloud, computing, cybersecurity and data governance support delivery and PNG control.

$DG \times DPI^{AI}$ explains public value. The five technology-control layers (section 3.2) identify national control requirements. Risk tiers (section 3.4) set safeguards and approvals; HUDERIA and COBRA support risk assessment.

3. Policy Pillars

3.1 Governance and Adoption

The intended result is coordinated adoption with clear accountability. Public bodies will adopt AI in phases, based on readiness, suitable uses and documented risk assessments. The same assess-and-adopt approach will extend across other sectors through lawful sector schedules.

Assess opportunities and justify investment

Each organisation covered by a commenced mandate must assess where AI can deliver a useful benefit. It must adopt suitable uses where benefits and safeguards are demonstrated, or justify an exemption or deferral. Buying a licence or increasing the number of AI tools is not, by itself, evidence of successful adoption.

Government assessments must connect a digital government service outcome, required DPI components and the additional contribution of AI. Compare AI with process reform and simpler digital options. Record purpose, affected people, data, cost, benefits and workforce effects. Agency heads must appoint owners for each deployment and its results.

Prepare costed adoption plans

Within six months of policy approval, government bodies and sector leads will submit initial costed adoption plans to the coordinating department. Plans must map DG outcomes, DPI dependencies and AI value, with baselines, targets, readiness gaps, safeguards, workforce changes, funding, owners and dates. Align them with national planning and budget cycles, identify reusable components and justify duplication. Consult businesses, educators and communities. Later deployment dates will be set through the applicable lawful schedules.

The Government will begin with controlled pilots in ready agencies and priority sectors. Expansion requires evidence of benefit, acceptable risk, competent staff, affordable ongoing operation and required approvals. A pilot that affects real people or live services is subject to the same legal and risk controls as other deployments.

Apply the mandate fairly

Non-government adoption deadlines may be set by sector after readiness assessments and consultation. Published schedules will explain who is covered, transition periods, support, exemptions, decision rights and review routes. The legal commencement of each duty must be clear before enforcement begins.

Persistent, unjustified non-compliance will first receive support and a time-bound corrective plan. Proportionate sanctions may follow continued non-compliance only where authorised by law, with notice, reasons and an opportunity for review.

GovAoS and GovAI Agents will be a common government delivery programme. Agency plans must address onboarding, approved knowledge, staff participation and benefits. Progress will track readiness and results. The coordinating department coordinates adoption; agency heads and sector leads remain accountable.

3.2 Infrastructure and Data

The intended result is reusable DPI and AI capability supporting better services under PNG control. Shared government platforms are the default; approved specialised systems may meet sector needs. Each use must identify and test the components it requires.

Shared services and national control

SevisPNG Ltd is the proposed operator of designated common DPI, GovAoS and shared government AI infrastructure, supported by contracted private providers. Its establishment or designation, ownership, service mandate, financing and assurance arrangements must meet section 4.5 before operations begin. Sector bodies retain responsibility for their authorised specialist systems and records. Interim delivery must remain within existing lawful arrangements.

The coordinating department will set common architecture and interoperability requirements. Approved connections may use SevisPass identity credentials, SevisWallet access, SevisDEx data exchange, core registers and authorised payment interfaces. Agencies retain record custody and data authority.

The PNG AI Technology Stack has five control layers. Cloud storage and computing support DPI and applications, while AI models may assist services across these layers. Government control may be secured through ownership, configuration, contracts and enforceable rights appropriate to the risk.

Control layer	Control government must retain
AI models	Choose and replace models; test suitability; control permitted data use and access.
Cloud storage	Approve locations, access, retention, backup, encryption and data export.
Computing resources	Manage capacity, access, availability, security and operating costs.
Digital public infrastructure	Govern identity, payments, data exchange, core registers and shared interfaces under lawful access rules.
Applications	Control workflows, permissions, custom software rights, updates and exit arrangements.

Data protection and cultural knowledge

All organisations using AI must respect applicable privacy, data-protection and cultural rights. Government data may be processed only for a defined, lawful purpose, with permission, access controls and records. Government bodies must collect only necessary data, check its quality, correct errors and apply lawful retention and deletion rules. Consent must be obtained where law requires it. This policy creates no alternative legal basis for access.

Government data will be stored and processed within PNG and in specifically approved overseas locations. Approval must cover the actual processing, backups, support access and subcontractors, not just the provider’s headquarters. Sensitive and classified data will face stricter controls. Existing statutory restrictions, required permissions and security classifications prevail.

The coordinating department will coordinate the approval process with source agencies and competent security and data authorities, subject to their lawful powers. Decisions must identify the authorised approver and assess foreign legal access, confidentiality, encryption and key control, incident response, continuity and enforceability. Approval for ordinary data does not authorise transfer of sensitive or classified data.

Suppliers must not use non-public government data to train or improve models used by other customers. An exception requires a defined lawful purpose, explicit written government approval from authorised decision-makers and enforceable safeguards. Where required, affected persons' consent or other legal permissions must also be obtained. A supplier's standard terms do not constitute government approval.

Contracts must cover prompts, uploaded documents, outputs, logs and derived data; access by staff and subcontractors; retention; incident reporting; audit evidence; and secure return or deletion. Agencies must not enter restricted data into unapproved public AI services.

Publicly funded datasets will be openly accessible when they can lawfully be published and contain no protected or sensitive material. Publication must address re-identification risk, licences, data quality, privacy and cultural rights. Public funding does not extinguish another person's or community's rights.

Use of traditional knowledge and community-held cultural material requires informed consent from legitimate knowledge holders, respect for customary restrictions and agreed benefit-sharing for commercial use. Consent discussions must explain intended uses, risks, future reuse, access, withdrawal arrangements and practical limits on removing knowledge from a trained model. Restricted or sacred knowledge must not be made open merely because it has been collected digitally.

Local-language datasets will be developed with participating communities and competent language experts. Data stewards must document provenance, permissions and permitted uses. The coordinating department will coordinate standards; agencies and knowledge holders retain their respective lawful responsibilities. Indicators include data quality, approved hosting, access compliance and the number of properly governed local datasets.

Build capacity in stages

Investment will strengthen domestic data centres and computing where demand, power supply, security and whole-life costs support the case. Approved overseas services may support delivery under the data controls in this section. Procurement must provide tested backup, recovery and a workable exit from suppliers.

The Government will use existing models and remain LLM agnostic, allowing changes of language-model provider. Systems may combine generative, predictive and analytical models with retrieval, memory and tools. Selection will assess the complete system in its intended PNG setting, including capability, cost, security and licensing. Local hosting, open-source code or model ownership alone does not prove safety or accuracy. PNG capability will grow through suitable model reuse, local data, applications, adaptation and testing.

3.3 Skills and workforce

The intended result is a workforce and population able to use AI safely and productively. Training must connect service design, records, DPI permissions and AI verification. Adoption will fund skills and job redesign. Accessible community information will explain safe use, scams and complaints.

Public service capability

All public servants will receive foundational AI literacy. It will cover safe use, privacy, checking outputs, fraud and deception, records management and personal responsibility. Leaders will receive training in investment and accountability. Technical staff and people in high-risk roles will receive practical assessment. Providers and deployers must also ensure suitable staff AI literacy. Justice-sector training will cover AI-related crime, electronic evidence, bias, disclosure and victims' needs.

The department responsible for personnel management will lead workforce planning with the coordinating department, agency heads and relevant training institutions. Agencies must identify the skills needed before deployment, provide time for learning and ensure people can seek help when AI performs poorly.

AI literacy

AI literacy will be introduced across primary, secondary, technical and vocational education and training, and higher education. Education authorities will sequence delivery using age-appropriate content, trained teachers and suitable learning materials. Teaching will include critical thinking, recognising unreliable outputs, protecting information and using AI honestly.

Children's safety and learning needs must guide design. Schools must avoid unnecessary collection of children's data and commercial exploitation of learners. Student work and assessment information must not be reused for supplier training without lawful authority and required permissions. AI must support teachers' professional responsibilities.

Universities and training providers will strengthen specialist AI education, local research, testing and cybersecurity skills. Partnerships with industry will connect learning with PNG service and business needs. Programmes must include women, people with disability and learners outside major centres.

Manage changes to work fairly

The Government may actively restructure public service functions, teams and staffing where evidence demonstrates better delivery and value. This may include redundancies where justified, subject to law, required approvals, consultation, employment entitlements, retraining and transition support.

Each restructuring proposal must distinguish work that disappears, work that changes and new work that is needed. It must assess service continuity, workload, staff capability, transition costs and the effect on communities. AI output alone must not determine dismissal, discipline or appointment.

Time saved does not automatically become a budget saving or a reduction in staff. Agencies must show how released time improves services or changes actual spending. Targets will track training competence, successful redeployment and transition support, service quality, and verified workforce and financial effects.

3.4 Responsible AI

The intended result is AI that is safe, fair and accountable throughout its use. The government will establish national risk tiers, with duties proportionate to the purpose, scale and likely harm of a system. An application using a general-purpose model may be low risk or high risk depending on its use. Risk classification governs safeguards; it does not by itself establish criminal guilt.

Risk tier	Policy treatment
Low risk	Basic security, privacy, competent oversight and a named accountable owner. No general prior AI approval requirement, unless another law requires it.
High risk	Documented impact assessment, stronger testing and monitoring, effective human control, and prior NICTA approval coordinated with sector regulators.
Unacceptable risk	Prohibited uses, clearly defined through lawful instruments. An adoption mandate or testing programme cannot authorise them.

High-risk decisions and approval

The high-risk framework will cover AI with material potential to affect rights, liberty, health, safety, livelihoods or access to essential services. Relevant contexts include policing and justice, clinical care, employment, education assessment, credit and critical infrastructure. Classification will depend on the actual use and consequences, not a sector label alone.

Before live deployment, the organisation must demonstrate a lawful purpose, data quality, local suitability, security, fairness, human oversight, explainability, complaint arrangements and a safe fallback. Tests must examine performance across affected groups and relevant languages. High-risk approval does not replace medical, judicial, financial or other statutory requirements.

NICTA will set transparent assessment criteria, service standards, conditions and review procedures once legally empowered. Material changes to a model, purpose, data or operating context require reassessment and renewed approval where required. Approval is not a guarantee of safety or a transfer of responsibility from the deploying organisation.

Unacceptable uses

The legal programme will define prohibited uses with public consultation. Proposed categories include AI used for unlawful surveillance or discrimination; social scoring that unfairly denies rights or essential services; and manipulation exploiting vulnerability that causes or is likely to cause serious harm. The final definitions must be precise, proportionate and consistent with constitutional protections.

The Government must not deploy AI that removes the human final decision required by this policy, or a system whose serious risks cannot be controlled. Research and supervised testing do not waive citizens' rights. Until the new approval regime has lawful effect, new government high-risk systems must remain in safe testing that does not determine live outcomes for people.

Assess risk throughout the AI lifecycle

A common assessment method will be developed that supports PNG's risk tiers. The method must consider the whole system, its users and institutional setting, from design and procurement through operation, change and retirement. A model's technical category does not determine its legal risk. The same technology can have very different consequences in different uses.

The assessment method will draw on the Council of Europe's HUDERIA guidance, adapted to PNG law and capacity. Its four elements are context-based risk analysis, stakeholder engagement, risk and impact

assessment, and a mitigation plan. COBRA supports analysis of an application's context, design and deployment. These are non-binding assessment resources. PNG's risk tiers and lawful approval requirements remain separate. [10]

- Describe the problem, intended use, affected people, data, model, tools and institutional decisions. Compare practical alternatives, including approaches that do not use AI.
- Engage affected people before major procurement and deployment choices. Include relevant language, disability, rural and cultural perspectives, while protecting sensitive information.
- Assess the likelihood, seriousness, scale and reversibility of harm. Test data quality, unequal errors, privacy, security and cascading failures across shared DPI, connected services and agent actions.
- Assign safeguards, responsible owners, resources, review dates and stop conditions. Record remaining risk and the evidence supporting the decision to proceed, restrict or reject the use.

Assessment must continue during use. Revisit it when a model, dataset, purpose, permission, supplier or operating environment changes, and after serious incidents. Check whether new data merely repeat earlier errors or patterns of unequal treatment. Agencies sharing information must assess their combined effects, including how an initial AI flag can influence later decisions. Retirement must preserve required records and remove obsolete access safely.

High-risk systems require validation independent of the delivery team and meaningful oversight by competent people. Human reviewers need training, time, access to relevant evidence and authority to reject an output. Record material reliance on AI, corrections and the reasons for consequential decisions.

An explanation must address the relevant model output, the data and system that shaped it, how the official used it, and the lawful justification for the action. Technical explanations help scrutiny but do not prove that an output is correct or an intervention justified. Affected people must have a practical opportunity to correct errors and challenge decisions.

Human authority and AI agents

AI may automate routine tasks. A competent person must make the final decision where it affects rights, liberty, health or access to essential services. That person must have authority, relevant information and enough time to question or reject the AI output. A routine human sign-off without real assessment is not adequate oversight.

A person affected by an AI-supported decision must be able to understand the main reasons, correct relevant errors and seek human review through the proper process. Explanations must be useful to the person. Courts must be able to require disclosure needed for a fair hearing, using lawful confidentiality measures for protected information and trade secrets.

Government AI agents may update records, send messages or trigger routine transactions within lawful, pre-approved access and transaction limits. Each agent must have a named human owner, a defined purpose and a documented operating permission. AI cannot create its own legal delegation or expand its own powers.

- Use identifiable agents with limited access. Bound memory, tool use and delegation to other agents; delegated actions must remain within the original authority.
- Set transaction, time and resource limits, recipients and escalation rules. Require human approval for high-impact actions and any action outside those limits.
- Keep auditable records of instructions, data access, actions, approvals and material model changes, subject to lawful privacy and retention controls.
- Provide an immediate stop mechanism covering connected agents and tools; revoke access when necessary and maintain a tested fallback.

- Test for malicious instructions in documents or messages, unauthorised data disclosure and unwanted actions before connection to live systems.

Agents must not override financial delegations, procurement approvals, evidence rules or other statutory responsibilities. An agent's action remains attributable to the responsible organisation. Officials must be able to correct or reverse actions where possible and promptly address any harm.

Essential services must retain human assistance and a non-digital route. Citizens who cannot use AI must not receive worse access or treatment for that reason. Services must accommodate disability, low literacy, weak connectivity and differing language needs.

English, Tok Pisin and Hiri Motu will be supported from the first implementation phase, with other PNG languages progressively added. Language performance must be tested with competent speakers. Where automated language quality is insufficient, trained human assistance must provide the required support.

People must receive clear notice when interacting with AI and when it materially supports a consequential decision. Public information must state its limitations and how to reach a person. AI-generated official advice must be checked against authoritative, current sources before it is relied upon.

Transparency, complaints and incident response

A public register will be maintained covering all government AI systems, including embedded and purchased systems. Each agency must supply and update its entries. The register will show purpose, responsible agency, risk level, deployment status, main safeguards and complaint route. High-risk entries will include approval status and a public summary of the impact assessment.

Sensitive operational details may be withheld where law permits, with the basis recorded and disclosure limited only as necessary. Full information must remain available to authorised oversight bodies. The register must be accessible in plain language, with assisted access for people who cannot use the internet.

A person may complain to the organisation using AI or directly to NICTA under the proposed AI complaints mandate. Complaints must be possible through accessible online and non-digital channels, without requiring technical knowledge. Organisations must provide human contact, acknowledge complaints, explain the process and report the outcome.

NICTA will publish complaint service standards and coordinate referrals to the authority legally responsible. A person should not have to repeatedly submit the same complaint because agencies disagree about responsibility. The receiving body must explain the referral and maintain continuity of assistance.

Remedies may include correcting records, reconsidering a decision, stopping an unsafe use or compensation where law provides it. NICTA will not replace the courts, clinical decision-makers or other statutory appeal bodies. Existing review rights and deadlines remain available. The legal programme must establish an appropriate route to challenge NICTA's own AI decisions.

All organisations will be required, through applicable law, to promptly report AI incidents causing serious harm, or creating an immediate risk of serious harm, to NICTA and relevant authorities. They must act to contain harm and notify affected people where appropriate. Standards will define thresholds, reporting times and the information required, aligned with other reporting duties.

The response must preserve evidence, protect people, stop or restrict unsafe operation and restore essential services safely. NICTA will coordinate regulatory action; the National Cyber Security Centre and other competent bodies will handle incidents within their mandates. Notifications to people must provide practical advice and respect lawful restrictions, including protection of victims and investigations.

Safety action must not wait for the completion of an adoption corrective plan. Serious incidents will be investigated, lessons shared in a form that protects sensitive information, and approvals revisited when

necessary. Measures include incident response time, recurrence, complaint resolution and whether safeguards work in practice.

3.5 Innovation and Industry

The Government will build PNG's AI industry and Digital Economy through research partnerships, supervised testing and procurement. Firms may use approved DPI interfaces and lawfully available data under published access rules and proportionate charges. Private adoption need not use government systems.

Domestic capability

Priority investment will support local-language datasets and tools; PNG-developed applications; domestic data centres and computing; university research and specialist training; locally adapted general-purpose AI capability using existing models; and local testing, safety and cybersecurity expertise. Projects will be phased and assessed for demand, affordability and measurable results.

University–industry research partnerships must agree on intellectual property, data access, community permissions, publication and commercial benefit arrangements before work begins. International partnerships should transfer skills, strengthen local institutions and provide access to research and testing while preserving PNG control over government data and decisions.

Supervised testing

NICTA, under its proposed AI mandate, will coordinate supervised testing with relevant sector regulators. Before that mandate commences, competent authorities may support testing only within their existing powers. Each test must have a defined purpose, duration, participants, risk controls and exit criteria. Participation does not authorise disregard of law or uncontrolled exposure to risk. Live high-risk testing requires the necessary approvals.

Useful findings will be published without exposing protected information. Successful testing may support an application for deployment, but does not guarantee approval or government procurement.

Procurement that builds local industry

The Government will prefer qualified PNG-owned suppliers where quality, security and overall cost are comparable, consistently with procurement law. Published evaluation rules must define qualification and verify beneficial ownership. Local preference must not conceal conflicts of interest or weaken competition.

Contracts will be assessed over their full life, including usage, integration, training, audits, hosting, replacement and exit. Buyers must secure interoperability, data portability, validation records, audit access and notice of material changes. Supplier claims must be tested in PNG conditions. Justice applications must provide the evidence needed for lawful disclosure and challenge, including where third-party models are used.

Ownership and continuity

The Government will own newly developed software and source code that it commissions, with rights assigned in the contract. Pre-existing and third-party components remain under their licences. Contracts must identify those components and secure the rights needed to operate, maintain, modify and transfer the commissioned solution lawfully.

Handover must include source code, build instructions, documentation, configuration and operational knowledge. Rights in models, weights and datasets must be explicit and respect licences, privacy and cultural rights. Owning custom software does not imply ownership of its underlying model or data.

3.6 Adoption in Priority Sectors

All five sectors below will participate alongside government administration, according to readiness, risk and resources. Plans due within six months must connect the sector outcome, reusable DPI and proposed AI contribution, with owners, targets and safeguards. Use only the identity, payment, register and exchange functions each service needs. Annex B sets starting indicators.

Priority Sector 1: Government administration

Improve the speed, accuracy and traceability of correspondence, records, approvals and expenditure management. AI may help staff find information and identify transactions needing review. Flags are leads for authorised checking, not findings of corruption. Finance, procurement, audit and investigation responsibilities remain with their lawful owners.

Priority Sector 2: Justice, policing and public safety

Justice, policing and public-safety agencies will improve case preparation, records, lawful analysis and service timeliness while protecting liberty, fair process and institutional independence. Humans retain decisions on investigative powers, arrest, prosecution, bail, judgment and sentencing. The safeguards below apply throughout the justice process; Annex C sets the related legislative requirements.

Justice agencies must distinguish original records, AI-assisted selection or analysis, and generated or reconstructed material. Record material AI use where relevant to disclosure and fairness, including relevant versions, inputs, outputs, transformations and human checks. Preserve original evidence and its chain of custody. Check generated legal references, translations, summaries and factual claims against reliable sources. An AI reconstruction or opinion must never be presented as an original record or eyewitness account.

Forensic use requires a sound scientific basis, validation in conditions resembling the intended use, documented uncertainty and error rates, and an independently examinable result in the individual case. Preserve the relevant configuration, version, inputs and transformations. Record any limits on reproducing an output. A confidence score must not be presented as proof of guilt. Where reliability cannot be established, restrict use to suitable investigative leads or do not use it.

Courts retain control of admissibility, evidential weight and the disclosure needed for a fair challenge. Explainability or source code alone may not establish reliability. Procurement must secure necessary documentation, validation evidence and lawful audit access. The programme must fund independent technical expertise and appropriate legal-aid support so the defence can use disclosed information. Evidence must not be rejected solely because a victim submitted it online; authenticity and other legal requirements still apply.

Victims must have accessible online and non-digital reporting, privacy protection, information about progress, support and referral to legal assistance. The programme will provide free initial assistance and specialised support for children and people at greater risk of exploitation or repeated harm. Procedures must help victims identify or preserve evidence safely without requiring them to obtain or circulate unlawful material. Legislation will provide timely protective measures, lawful content restriction and effective routes to compensation and other remedies.

AI may help competent authorities draft and translate assistance requests, identify preservation needs, track deadlines and route matters through authorised cooperation channels, including the 24/7 contact point. A competent official must verify the legal basis, accuracy, data sought and any emergency justification before an official request or compulsory measure is approved. The policy’s agent limits also apply to cross-border workflows.

Official requests must identify material AI processing and any special procedures needed to preserve evidential value. Transferred case data retain their purpose, access, retention and onward-transfer restrictions and must not be reused for model training without separate lawful authority. Foreign-hosting approval does not confer a search power. Treaty and protocol mechanisms may be used only when legally applicable. Annex C distinguishes the relevant legal powers.

Priority Sector 3: Health services

Improve service access, continuity of care, administration and the availability of accurate information. Clinical AI must be validated for its intended setting, protect patient confidentiality and be supervised by competent professionals. A health professional remains responsible for final clinical decisions. Safe referral and fallback arrangements are required.

Priority Sector 4: Agriculture, fisheries and forestry

Improve access to reliable advice, markets and resource-management information. Tools must be tested for local conditions and designed for producers with limited connectivity. Agencies must protect customary knowledge and land-related information. AI advice must clearly state uncertainty where livelihoods, food safety or natural resources may be affected.

Priority Sector 5: Education and training

Improve learning support, teacher preparation and education administration while expanding AI literacy. Education authorities will protect children's data, maintain fair assessment and evaluate learning outcomes. Human educators retain responsibility for consequential decisions about learners. Adoption must not exclude schools because they lack connectivity or equipment.

Priority Sector 6: Tourism and hospitality

Help PNG operators improve customer service, language access, booking administration, marketing and productivity. Information about destinations, safety and communities must be accurate. Cultural material must be used with permission, and businesses must protect customer data and avoid misleading AI-generated promotion.

3.7 Government AI adoption and shared platforms

GovAoS and government records

The Government will develop and adopt GovAoS, the Government Agentic Operating System, as the common platform for official document management, correspondence and AI-assisted work. A governed government document repository will support both people and GovAI Agents. It implements the $DG \times DPI^{AI}$ approach through shared workflows and approved DPI connections, phased by readiness, value, lawful authority and safeguards.

GovAoS will support capture and filing, document identifiers, metadata, search, editing, version history, permissions, retention and audit records. The repository may connect existing agency stores; a shared workspace does not require every record to move into one physical database. Agencies retain custody and approve access to their records.

The correspondence workflow will cover letters between agencies and internal memos. An authorised user or agent prepares a draft from permitted records. The responsible official reviews and signs off the final version. Approved dispatch, receipt, assignment, due dates, responses and closure are recorded against the document. Electronic signatures must meet applicable legal and government requirements.

Official work context will pass to authorised successors when roles change. Personal preferences must remain separate. Agent memory supplements the official record and is subject to retention, deletion, legal

holds and access controls. A change of office holder must neither erase required records nor expose unrelated information.

SevisDEx will provide authorised exchanges between GovAoS and connected systems, including Microsoft 365, the Integrated Financial Management System (IFMS), electronic procurement and recruitment, government payroll and SevisWallet. SevisWallet supports identity and sign-in; the responsible agency separately grants role, document and case permissions. Source-system owners must approve each connection and its uses. Verify interfaces, licences and capacity before connection. Test the DPI components required for the service; completion of every national component is not a precondition.

The AI Technology Stack will apply to the platform, every model provider and each processing route. Government hosting of an application does not itself approve an external model service. Approval must cover storage, retrieval, inference, optional model training, support access, backups and derived records. Enforce source permissions in search indexes, caches, logs and agent tools.

The GovAoS will provide role-based functions for public servants, agency heads, MPs, Ministers and the Prime Minister, respecting institutional independence. A senior role or subscription grants no automatic access to another agency's records or protected cases. Agencies will retain tested fallback, recovery and export arrangements.

GovAI Agents and government adoption

GovAI Agents are approved government AI assistants that can carry out sequences of tasks within defined permissions. They combine models, instructions, permitted knowledge and controlled tools within GovAoS or approved specialised systems. They use authorised DPI interfaces for shared workflows. Every agent must have a named functional owner, approved tasks and sources, a recorded risk assessment and an identifiable released version.

The initial core portfolio comprises GovAsk, GovWrite, GovMeet, GovAssist, GovPolicy, GovPulse, GovAudit, GovContract, GovInstruct and GovDraft. The first six are the proposed starting pilot set. The remaining four will be activated after professional, technical and security acceptance. Annex E identifies each function and proposed accountable lead. Approval applies to the actual deployment.

GovInvestigate and GovProsecute form a separate specialist cybercrime track. Police and forensic specialists will own investigation support; the Office of the Public Prosecutor and authorised prosecutors will own prosecution support. They retain their distinct legal functions and independent decisions. Neither agent may exercise coercive powers, determine guilt or autonomously file case documents.

Specialist agents require segregated case workspaces, lawful access and protected evidence-system connections through SevisDEx. Original evidence and chain-of-custody records remain in the authorised evidence system. Generated analysis must remain distinguishable from evidence. Live case files must stay outside model training.

Public bodies will use the approved common platform and agent catalogue for suitable shared functions under lawful adoption schedules, with justified exemptions or approved specialised alternatives. Readiness assessments must cover records, connectivity, staff, language needs, delegations and operating funds. Agencies will reuse approved capabilities and teaching resources where permissions allow, rather than repeatedly commissioning the same functions.

The initial GovAoS pilot will require explicit user approval before each external action or change to an official record. Approval must identify the final content, recipient or change and be logged with execution. Finance, procurement, HR and payroll connections will remain read-only during the pilot. Approval cannot override that restriction.

The pilot will begin with ready agencies, followed by Central Agencies and further government waves after verified results. Internal pilot language and feature priorities do not defer the public-facing language and access guarantees.

Sovereign Agent Training

The Government will establish a whole-of-government sovereign training programme for GovAI Agents. Public servants and authorised professionals will lead the teaching, testing and acceptance of agents used for their functions. It will turn institutional knowledge into reusable, tested assistance for approved digital government workflows and DPI interfaces. It complements staff literacy; responsible institutions supply the knowledge, procedures and professional judgement.

Sovereign agent training means that the government controls the approved tasks, knowledge, instructions, tools, data permissions, testing and release of its agents. It does not require PNG to build a foundation model from the beginning. The programme will remain model agnostic and retain the ability to change providers, subject to fresh validation.

Training will normally begin with role instructions, approved knowledge retrieval, worked examples, templates and tool configuration. Retrieval provides permitted information at the time of use and does not itself change model weights. Fine-tuning, which changes model parameters, is optional and requires a separate business case and lawful data, privacy and security approval. Ordinary conversations and feedback must not automatically become training data.

Source custodians must approve indexing and each proposed reuse separately. Teaching material will include current laws, policies, procedures, forms and carefully checked examples. Record its origin, permissions, version and review date. Distinguish draft laws from law in force. Use synthetic or lawfully redacted cases where appropriate, minimise personal data and respect cultural restrictions. Secure access for one task does not authorise general reuse.

Functional owners will nominate leads and specialists to teach tasks and maintain sources. Each agency joining GovAoS must appoint a knowledge custodian and training liaison, contribute approved examples and test its workflows. Agency heads must protect staff time and fund support. Separate validators assess unseen cases; high-risk validation must be independent of the delivery team.

The Government will maintain a controlled catalogue of agent instructions, approved knowledge collections, teaching examples, assessment cases, tool permissions and release records. Assets may be shared across agencies only with the necessary source-owner permissions. Specialist, personal and classified material must remain segregated. Local-language and subnational expertise will inform relevant agents and tests.

Contracts must secure government ownership of newly commissioned agent software and project-specific training assets to the extent lawfully assignable, with enduring rights to use, maintain, export and transfer the complete solution. Pre-existing software, model weights and third-party material remain subject to their licences. The Government must control approved configurations and receive the documentation and skills needed for continuity.

Training, validation and acceptance

- Define the task, owner, legal basis, sources, access and performance criteria before teaching begins.
- Prepare approved examples and configure instructions, retrieval and tools in a controlled environment. Keep assessment cases separate from teaching material.
- Test professional quality, source citations, unequal errors, unsupported conclusions, prompt injection, data leakage, permission boundaries and stop controls. Set risk-appropriate thresholds before testing.

- Run supervised trials with trained officers. Record corrections and compare quality, time and cost with the baseline. Failed agents remain in testing.

Release requires functional-owner acceptance, independent validation evidence and technical and security clearance. The coordinating department will coordinate technical assurance, supported by NCSC within its mandate. The Working Group recommends readiness and the Steering Committee endorses programme entry within its remit. These steps do not replace NICTA's required high-risk approval, sector approvals or agency delegations. No unresolved high or critical defect may be waived merely to meet a launch date.

Continuing Assurance

Each release must identify the model, instructions, knowledge versions, tests, permissions, owner and review date. Reassess material changes to models, law, sources, tools or duties before release. Monitor performance and incidents, provide immediate suspension and rollback, and restore service only after documented clearance and owner approval. Report readiness and defects monthly during training and pilot delivery, alongside the policy's public reporting cycle.

Annex E records the agent portfolio and delivery stages. The supporting GovAoS Implementation and Training Plan holds provisional pilot dates, staffing assumptions, teaching volumes and testing thresholds. The Working Group will maintain that plan, with functional, security and regulatory approvals where required. Changes to the plan cannot weaken this policy's safeguards.

4. Institutional Arrangements and Accountability

4.1 National Strategic Oversight

The Government will make the legal and administrative changes needed to rename the Public Service ICT Steering Committee as the National Digital Development and AI Steering Committee. Its remit will cover digital development and AI across all sectors. As a subcommittee of the Central Agencies Coordination Committee, it will report on priorities, delivery and matters requiring National Executive Council decisions.

The Secretary of the coordinating department will chair the Committee. Government agency heads will be its decision-making members, including central agencies, NICTA and relevant sector heads, with subnational participation. The necessary legal instruments will establish membership and functions.

The Steering Committee will oversee one coordinated portfolio of DG reform, DPI investment and AI adoption. It will resolve dependencies and duplication, review results and endorse sector targets. Ministers, agency heads, regulators and independent institutions retain their lawful powers.

4.2 Policy Coordination and Regulation

The Government proposes to rename the Department of Information and Communications Technology (DICT) as the Department of Digital Development and Artificial Intelligence (DDDAI). Until the change lawfully takes effect, DICT continues within its existing mandate. The term “coordinating department” in this policy means DICT and, after the change takes effect, its lawful successor responsible for ICT and AI.

The coordinating department will develop, maintain and coordinate this policy, provide the Steering Committee secretariat, consolidate costed plans, set lawful architecture and implementation guidance, maintain the government AI register and publish progress reports.

The Government will expand NICTA’s mandate to regulate AI across sectors. Its proposed functions include risk classification, high-risk approvals, compliance monitoring, serious-incident reporting, complaints and enforcement. NICTA will work with specialist regulators through clear referral and assessment arrangements. Suspected crimes will be referred to the police and competent prosecution authorities. NICTA’s proposed regulatory role does not transfer the functions of police, prosecutors or courts. Safety requirements apply to government and private providers.

4.3 Whole of Government Agent Training Coordination

A Whole of Government AI Agent Training and Validation Working Group will coordinate the programme under the National Digital Development and AI Steering Committee. The coordinating department will chair and provide the secretariat, with the Department of Personnel Management as deputy chair. Until the Committee’s reforms take effect, the existing Public Service ICT Steering Committee may oversee preparation within its lawful powers.

The Working Group will include the coordinating department; the departments responsible for Prime Minister and NEC, personnel management, finance, treasury, national planning, and justice; the Office of Legislative Counsel; and the National Procurement Commission. The Auditor-General may advise independently. Police and prosecution authorities will lead their specialist workstreams while retaining case and prosecutorial independence.

Three technical groups will support the core portfolio: Executive and Administration led by Personnel Management; Policy and Performance led by Prime Minister and NEC; and Assurance and Legal led by Justice and Attorney General. They will coordinate teaching sessions, resolve procedural differences, maintain versioned examples and report readiness. Agency owners retain responsibility for their functions.

Contracted technical partners support configuration and skills transfer, without voting on government acceptance.

The Working Group coordinates teaching, validation evidence and readiness recommendations. Functional owners retain professional acceptance; competent technical and security authorities provide the necessary assurance; NICTA and sector regulators retain their lawful approval functions. Section 3.7 sets the training and release requirements. Committee endorsement does not substitute for a legal approval or delegation.

4.4 Independent advice and public participation

An independent expert panel will be recruited in consultation with the Vice Chancellor's committee and will advise NICTA and the Steering Committee on ethics and citizens' rights. Transparent selection will consider expertise, independence and diverse PNG language, cultural, disability, gender and community perspectives. Members must disclose interests and manage conflicts. Recommendations will be published, with lawful redactions.

A formal advisory forum for business, academia and civil society will inform strategy and implementation. The forum and expert panel are advisory bodies. Neither grants regulatory approvals nor replaces government decision-making accountability.

4.5 Delivery and Assurance

SevisPNG Ltd is the proposed shared-service operator for designated DPI, GovAoS and common government AI infrastructure. Before designation or commencement, the Government must confirm its legal status, mandate and ownership arrangements against this policy's government-owned implementation model. Required approvals must establish government control, financial accountability, service obligations and independent financial and security assurance. This policy does not itself incorporate a company or alter existing ownership.

Any joint venture or departure from the government-owned implementation model requires a separate explicit government decision, lawful approvals and arrangements that preserve sovereign control. Private participation must not transfer policy or regulatory powers to the operator. Contracts must protect government data, newly commissioned software and training assets, continuity and supplier exit. This policy awards no contract.

Agency heads retain accountability for records, benefits and lawful use. The Working Group coordinates readiness and functional owners accept agents. Sector leads coordinate plans, while subnational administrations support delivery and access. SevisPNG Ltd's operating mandate must identify the common services assigned to it and its responsibilities to source agencies and specialist service operators.

Central agencies retain responsibility for public resources and workforce reform. Police investigate crimes; prosecutors and courts retain their statutory functions. The justice department leads legal coherence with the coordinating department. NCSC provides technical support within its mandate. Audit and integrity bodies remain independent. Policy, regulation, service operation and assurance must remain separate.

5. Implementation and Reform Roadmap

Implementation will proceed in phases. The six-month planning and baseline deadlines are firm policy commitments. Other milestones below are proposed delivery dates for confirmation in the costed national programme. No date overrides a required legal, safety or funding approval.

Phase	Main deliverables	Condition for proceeding
Months 0 to 6	Map DG outcomes, DPI dependencies and AI uses; establish coordination and training; complete baselines, costed plans and reform proposals; prepare pilots.	Use existing lawful authority. Identify funding and approvals before live activity.
Months 7 to 12	Publish targets; prepare required shared components; run integrated service pilots; release validated agents; begin reporting and regulatory capacity.	Documented benefit, risk assessment, accountable owner, staff readiness and required approvals.
Years 2 and 3	Reuse proven services, DPI and agent capabilities across sectors; apply lawful schedules; deepen infrastructure and skills; evaluate independently around Month 30.	Evidence of improved outcomes and affordable operation; unresolved risks addressed.
Years 4 and 5	Extend proven services and language coverage; improve weak programmes; replace poor-performing systems; complete final evaluation and prepare the next policy.	Review value, safety, inclusion and sustainability before renewal or expansion.

5.1 Reform and Commencement

The coordinating department and the justice department will prepare a coordinated reform programme with NICTA, sector regulators and central agencies. It will address institutional mandates, the adoption duty, risk tiers, approvals, prohibited uses, incident reporting, complaints, sanctions and review rights. Annex C sets out the main legal tasks.

New mandatory duties outside existing authority will commence only after the required law or instrument is in force and implementation arrangements are published. Adoption notices must specify coverage, dates, evidence required, support and exemptions. NICTA must have the staff and processes to perform its new duties before the corresponding regime begins.

Until reforms take effect, the existing committee and institutions will continue within their lawful mandates. Interim working groups may support preparation. They must not claim new statutory powers. Existing AI uses will be inventoried and reviewed, with lawful transition and service-continuity arrangements for systems requiring change.

5.2 Readiness before Scaled Adoption

GovAoS and sector pilots must compare results with baselines, test safeguards and seek affected people’s feedback. Existing lawful pilots may proceed earlier than the indicative phases. Annex E gives the GovAoS sequence. Expansion requires written acceptance of benefits, risks, costs and accountability; unsafe or ineffective systems must be corrected, restricted or retired.

5.3 Exemptions, Support and Delivery Risks

Proposed exemption criteria for consultation

The following exemption grounds are proposed for public consultation. The final framework will specify evidence, decision-makers, conditions and review arrangements before the relevant adoption duties commence:

- No suitable use has demonstrated a worthwhile benefit after a reasonable assessment.
- Necessary connectivity, power, reliable data, skills or support are unavailable, with a plan to address remediable gaps.
- Total costs are disproportionate to the expected benefits or the organisation's capacity.
- Privacy, safety, cultural or other rights risks cannot be reduced to an acceptable level.
- A proportionate threshold or simplified obligation is justified for a microbusiness or small community organisation.

The adopted framework must distinguish a temporary deferral from an exemption. Applicants must provide proportionate evidence; decisions must give reasons, scope, conditions and review dates. Deferrals should include a practical improvement plan. Exemptions must not become permanent simply through inaction.

The legal programme will specify decision-makers and appeal routes. For the Government, the proposed route is a reasoned decision by the authorised official of the coordinating department after sector advice; for other organisations, NICTA would decide under the new mandate in coordination with sector regulators. Sensitive details may be protected while aggregate exemption information is published.

An adoption exemption cannot waive existing law, permit a prohibited AI use or remove safeguards from a system that is used. No organisation should be penalised for a duty that has not lawfully commenced or for failing to deploy a system whose risks cannot be controlled.

Support and enforcement

Support will include clear guidance, common assessment tools, relevant training and access to the approved programmes in this policy. Persistent, unjustified non-compliance will move from support to a time-bound corrective plan, then to proportionate sanctions only under law. Enforcement decisions must consider actual capacity and evidence, avoid conflicts of interest and preserve review rights.

Delivery risks

Each programme must maintain a risk register with a named owner and action dates. National monitoring will focus on weak infrastructure and skills; unreliable or unauthorised data; financial pressure; supplier dependence; cyber incidents; discrimination or exclusion; workforce disruption; and insufficient regulatory capacity.

Responses will include phased investment, independent assurance, interoperable systems and exit rights, service fallback plans, workforce support, targeted accessibility measures and transparent reporting of delays. Failed pilots must inform redesign rather than create pressure to scale a poor investment.

5.4 Policy basis for AI legislation

The coordinating department and the Department of Justice and Attorney General (DJAG) will translate this policy into legislative instructions with NICTA, sector authorities and legislative counsel. They must define duties, powers, safeguards, resources and commencement. Annex C sets out the legal reform programme and detailed drafting requirements; Annex F identifies the comparative sources.

A coherent legislative package

The package will assess a dedicated AI Bill and amendments to digital government, ICT, cybercrime, procedure, evidence and sector laws. Cabinet will choose the legislative vehicle after gap analysis and consultation. Reconcile the related digital reform programme, identify which instruments are in force, provide complete coverage and avoid conflicting powers. [14]

The AI governance component must establish common definitions; scope; responsibilities across the supply chain; risk tiers and prohibited uses; NICTA's mandate; prior approval of high-risk uses; human oversight; transparency; data and cultural safeguards; complaints and review; incident reporting; inspections and proportionate enforcement; and phased adoption with exemptions. It must also establish the proposed Steering Committee arrangements and preserve institutional independence.

The criminal justice component must address demonstrated gaps in offences, responsibility, evidence, victim protection, investigative powers, jurisdiction and cooperation, separately from administrative regulation. Annex C provides the drafting requirements. Technology-neutral provisions must connect DG service duties, DPI interoperability and lawful data access with AI accountability, records, signatures, agent permissions and training-data controls. The core concept creates no new access power. Programme specifications belong in implementation instruments.

Matters for Parliament and delegated rules

Primary legislation must establish core rights, duties, offences, penalties, intrusive powers, appeal routes and the boundaries of delegated authority. Regulations and standards may address specified technical details within those limits, with consultation, publication and review requirements. Guidance or a technical standard must not create an offence or extend a power beyond its lawful basis.

Every proposed offence must state conduct, fault, jurisdiction, applicable defences and a proportionate penalty. Every investigative or regulatory power must identify who may use it, on what evidence, for what purpose and duration, with what authorisation, safeguards and remedy. Supplier protections must be reconciled with safety duties and liability for the supplier's own conduct. Courts retain control of criminal adjudication and fair-trial safeguards.

Drafting instructions and commencement

The proposed Month 6 reform submission will include an existing-law gap analysis, common definitions, a policy-to-clause matrix, Cabinet options, institutional responsibilities, financial and regulatory impact assessments, and a commencement plan. The matrix must record the treatment of an appropriate Model Law, distinguish binding law from guidance, and justify each adaptation. It must address lifecycle duties, provider categories, evidence powers, liability and remedies.

Commencement requires staff, funding, accessible procedures, guidance and review channels. Consultation must test effects on small organisations, workers, children, cultural knowledge holders and users of non-digital services. These legislative measures remain proposals until lawfully enacted.

6. Financing and Resource Mobilisation

The government-funded programme will be phased and costed for budget consideration. This policy sets no predetermined expenditure ceiling and commits no unappropriated funds. Funding must follow normal budget, procurement and financial-control processes.

6.1 Prepare a credible costed programme

Agency and sector plans due within six months must distinguish initial investment from annual operating costs. They must cost data preparation, connectivity, hosting, model use, integrations, licences, security, testing, training, workforce transition, accessibility, regulation and supplier exit. GovAoS estimates must also cover records onboarding, protected staff teaching time, validation, training-asset maintenance and specialist case controls.

The coordinating department will consolidate the programme with Treasury, Finance and National Planning and Monitoring. A proposed Month 9 submission will present phased options, costs, benefits, dependencies and funding gaps. Estimates must state quantities, unit costs, assumptions, exchange-rate exposure, contingency and confidence ranges.

Budget proposals must cost service reform, shared DPI and AI together, allocating shared costs once and justifying duplication. Baselines, validated demand and procurement evidence determine funding. Shortfalls require transparent changes to scope or timing while maintaining safeguards.

6.2 Funding Sources

- National budget allocations for agreed public priorities, shared capability and regulation.
- Development partner grants and technical assistance aligned with the national programme.
- Private-sector co-investment under arrangements that preserve government control and lawful public accountability.
- Cost-recovery charges to agencies and businesses using DPIs and shared services, where authorised and transparently priced.

Cost recovery must be based on an approved service model and reasonable costs, with clear service levels and financial reporting. It must not disadvantage citizens who need human or non-digital access to essential services. Any public subsidy or separate fee proposal requires its own lawful budget or charging decision.

6.3 Sustainability and Value

Shared services must have funded maintenance, staff, security, recovery and replacement arrangements. Co-investment agreements must disclose public liabilities, protect data and avoid exclusive arrangements that prevent future competition or supplier replacement.

Benefits will be assessed after full costs. Agencies must separately report actual expenditure reductions, additional revenue, staff time released and service improvements. Avoided losses and suspected irregularities must not be counted as cash savings without evidence. Independent assurance will examine material claims.

7. Monitoring and Evaluation

7.1 Baselines and Targets

Within six months, government bodies and sector leads will establish baselines. Each indicator must have a definition, starting value, data source, accountable owner and measurement method. Record weaknesses in the data and how they will be addressed.

Sector targets will be approved through national coordination and relevant sector processes and published, with Month 9 proposed. Targets must cover quality, safety, inclusion, speed and cost. Annex B provides starting measures; numerical improvement targets will follow the evidence.

Evaluation must trace DG reform, DPI reuse and AI outputs through decisions to outcomes. Where feasible, compare process and DPI improvements alone with added AI to estimate its contribution. Speed and accuracy do not prove reduced crime or corruption. Test outcomes and unequal effects against appropriate alternatives.

7.2 Public Reporting

The coordinating department will publish a progress summary every six months and a comprehensive annual report. Agency heads and sector leads will certify their contributions. The annual report may incorporate the second six-month summary while retaining the required level of detail.

Reports will link DG outcomes, DPI reuse and AI contribution. Show targets, spending, funding gaps, verified benefits, inclusion, workforce effects, complaints, incidents, exemptions and corrective action. Explain shortfalls and scope changes. Protect sensitive information without concealing performance.

The Steering Committee will review results and report through CACC to the National Executive Council. Reports will be made available through government websites and accessible public-information channels. Any submission to Parliament will follow the applicable procedures.

7.3 Independent Evaluation and Assurance

Independent evaluators will be selected through a transparent process for the mid-term and final evaluations. The Steering Committee will oversee commissioning through lawful procurement, with conflict controls and access to the evidence needed. Delivery providers must not evaluate their own performance.

The mid-term evaluation, around Month 30, will examine whether the policy is improving services, public integrity and economic opportunity, at what cost and for whom. It will assess burdens created by the mandate, safeguards, exemptions, institutional performance and whether priorities should change. The final evaluation will be completed by Month 60 and inform the next policy period.

Evaluation findings and the government’s response will be published. Where justified, protected material may be withheld with reasons. Existing audit and integrity institutions retain their powers; independent policy evaluation does not replace statutory audit or investigation.

7.4 Adaptation

The coordinating department will maintain a public record of material policy amendments and implementation changes. Major changes to rights, mandates, institutional responsibilities or funding commitments require the appropriate consultation and approval. Standards may be updated more frequently under lawful authority, with notice and reasonable transition arrangements.

Annex A Implementation Responsibilities

Accountable owners are responsible for completion, with support from other institutions within their mandates. Milestones marked “proposed” are to be confirmed through the costed programme. “Commencement” means the relevant duty has lawfully taken effect.

Deliverable	Accountable owner	Milestone
National coordination and interim arrangements	Coordinating department with CACC	Proposed Month 3
AI inventories, readiness assessments and baselines	Agency heads and sector leads	Month 6
Costed plans linking DG outcomes, DPI reuse, AI value and safeguards	Agency heads and sector leads	Month 6
AI drafting instructions and coordinated reform package under section 5.4	Coordinating department and justice department with NICTA and legislative counsel	Proposed Month 6 submission; lawful commencement later
Sector targets and joint DG, DPI and AI costed programme	Coordinating department and sector leads with central finance and planning agencies	After baselines; proposed Month 9
Public AI register and service notices	Coordinating department and agency heads	Proposed Month 6; ongoing updates
GovAoS, shared platform and SevisPNG Ltd implementation arrangements	Coordinating department with ownership and finance authorities	Proposed Month 6 design; establishment after approval
Approval, incident, complaint and exemption procedures	NICTA with sector regulators	Before relevant duties commence
Controlled pilots and expansion decisions	Agency heads and sector leads	Year 1; expansion after evidence and approvals
Staff literacy and sovereign agent training programme	Coordinating department, Personnel Management and functional owners	Proposed Month 6 plan; validation before release
Education delivery plan and teacher preparation	Education and higher education authorities	Proposed Year 1 plan; phased delivery
Progress summaries and comprehensive annual report	Coordinating department; agencies certify inputs	Every six months and annually
Independent evaluations and published government response	Steering Committee commissions evaluations	Around Month 30; final by Month 60

Annex B Results framework

Baselines are due within six months; approved targets follow, with Month 9 proposed. Each measure must link a DG outcome, DPI contribution and AI contribution, with an owner, definition, denominator, source and year. Report additional value without counting shared benefits twice.

National outcome measures

Measure	Definition and evidence	Owner
Service timeliness	Median completion time and the time within which 90% of cases finish, using comparable case records. Report backlogs and completion rates alongside speed.	Agency heads and sector leads
Citizen effort and access	Cost, travel and repeat visits needed to complete a service; successful access through human and non-digital routes. Use service data and user surveys.	Service providers and sector leads
Quality and fairness	Verified error and rework rates; decisions corrected on review; performance differences across relevant groups and languages.	Service providers, checked by sector assurance bodies
Net financial benefits	Verified expenditure reductions and additional revenue, reported separately after implementation and operating costs. Compare with an agreed baseline and explain other causes.	Agencies with Finance and Treasury assurance
Public integrity	Confirmed irregularities, action taken and money actually recovered or protected with evidence. Report false alerts separately. Do not equate AI flags with proven corruption.	Finance and responsible audit or integrity bodies within their mandates
Business and employment	Productivity and revenue changes in participating PNG businesses; local procurement participation; jobs created, changed and lost, separately identified.	Sector leads, business partners and relevant economic agencies
Trust and inclusion	User satisfaction, understanding of AI use, and access and outcomes by location, language and other relevant characteristics, collected lawfully.	Coordinating department and service providers

Staff time released must be reported separately from cash savings. Broad economic effects require a credible evaluation method and must identify uncertainty and other influences. Baselines and results must use comparable services and populations; changed definitions must be disclosed.

Crime-prevention evaluations must examine false alerts, unequal effects, displacement of harm and community trust. Increased detection may increase recorded crime while underlying offending remains unchanged. Follow the intervention and its outcomes, rather than relying on prediction accuracy alone.

Starting indicators for priority sectors

Sector	Initial outcome indicators
Government administration	GovAoS retrieval and correspondence turnaround; overdue actions; complete record and approval history; quality, rework and cost per completed task.
Justice, policing and public safety	Case-processing time; validated accuracy and errors across groups; evidence failures and review outcomes. For preventive uses, verified changes in harm and trust against an appropriate comparison.
Health	Waiting or processing time for selected services; continuity of care; validated clinical performance and adverse events where clinical AI is used.
Agriculture, fisheries and forestry	Reach and accuracy of locally relevant advice; producers' measured productivity or income changes, with seasonal and market influences explained.
Education and training	Teacher administration time; assessed learner outcomes and AI competence; access gaps and incidents affecting children.
Tourism and hospitality	Time and cost of customer-service and booking tasks; accuracy of visitor information; productivity changes in participating PNG businesses.

Delivery and safety measures

Area	Required evidence
Adoption and skills	Agency readiness and active GovAoS use; staff competence; agent owners appointed; approved teaching sources; unseen tests and validation passed; workforce transition outcomes.
Infrastructure and data	DPI reuse across services; availability, recovery and duplicate components retired; approved model routes, lawful exchanges, record completeness and export. Measure AI gains net of checking and correction.
Risk and accountability	High-risk systems approved before live deployment; current assessments and named owners; completeness of government register entries.
Incidents, complaints and justice	Serious incidents and recurrence; reporting and containment times; complaint resolution and victim assistance; evidence errors and lawful investigation response times, with privacy protected.
Local capability and culture	PNG suppliers reaching sustained use; research applied; local-language tools tested; cultural-use agreements and benefit-sharing obligations met.

The requirement for prior approval applies to every high-risk deployment once the relevant regime commences. Similarly, all government systems must be represented in the register. Numerical benefit targets will be set from evidence; these universal safeguards are compliance duties, not optional performance ambitions.

Annex C Legal and institutional reform programme

This annex guides the preparation of legislative instructions and implementation instruments. Before approval and enforcement, the justice department and relevant legal advisers must confirm current legislation, amendments, commencement notices, delegations and constitutional compatibility. Comparative model provisions and guidance must be adapted to PNG law; they are not binding merely because this policy cites them.

C1 Institutional and enabling reforms

Area	Current foundation or issue	Required action
Constitution and institutional independence	Constitutional rights and the powers of independent institutions govern all implementation.	Review proportionality, fair process, appeal rights, subnational application and Bougainville arrangements.
Steering Committee	Digital Government Act 2022, sections 11–13, provides for the current committee.	Amend its name, remit and membership as needed; lawfully establish the proposed CACC subcommittee relationship.
Coverage of public bodies	The Act's definition of public body excludes majority-owned State enterprises. Its coverage cannot be assumed to match this broader policy.	Identify the lawful instrument for each category, including State enterprises and independent bodies.
NICTA AI mandate	The NICT Act 2009 establishes NICTA and specifies its functions and powers.	Expressly authorise cross-sector AI functions, prior approval, coordination, complaints, enforcement and review rights.
Adoption and risk controls	New economy-wide duties and AI-specific sanctions need a sufficient legal basis.	Legislate or amend appropriate instruments; define coverage, risk tiers, prohibited uses, exemptions, transitions and proportionate sanctions.
Government data and hosting	The Digital Government Act contains cloud, classification and data-access requirements.	Check relevant approvals and commencement conditions. Reconcile overseas processing, supplier reuse and open-data rules with law.
Digital reform bills	Coordinating department has published a draft Data Governance and Data Protection Bill and other digital reforms. Draft publication is not enactment.	Align DG service duties, DPI access and interoperability, AI safeguards and NICTA functions with enacted instruments; avoid conflicting powers.
Operator finance and workforce	A policy does not itself establish a company, assign assets, appropriate funds or authorise redundancies.	Obtain required ownership, financial, procurement and employment approvals; preserve accountability and service continuity.

The policy supports preparation of a coherent AI legislative package, including consideration of a dedicated AI Bill and consequential amendments. The final legislative vehicle requires Cabinet decision.

C2 Responsibilities, offences and liability

The Department of Justice and Attorney General and the coordinating department will coordinate AI-related crime measures with NICTA, police, prosecutors and sector authorities. The review will

distinguish offences against AI systems and their data, offences committed by means of AI, and AI used by justice agencies. Existing offences must be mapped before new ones are proposed; the absence of the term AI in a law does not itself establish a gap.

Legislation will assign duties to providers, deployers, importers, distributors, integrators and other relevant operators according to their roles, control and risk. A person may have several roles. Contracts cannot remove a statutory duty. Suppliers must provide the information and cooperation needed for safe use, oversight and lawful investigations, with protection for confidential information.

People and legal entities remain accountable for their own conduct when using AI. Autonomy does not by itself excuse an offence, but harm does not automatically prove someone's criminal liability. The required conduct, fault and causal link must be established. Legislative review must examine intent, recklessness and any applicable criminal-negligence duties, and distinguish these from civil liability. A supply-chain role alone does not prove guilt. This policy gives AI neither legal personhood nor independent legal authority.

The legislative review will address fraud, extortion, impersonation and identity theft; attacks on critical infrastructure; interference with AI systems or evidence; intentional supply of tools for criminal use; child sexual exploitation and abuse material, including synthetic material; and non-consensual intimate imagery and harmful deepfakes. It will also examine organised crime, incitement, aiding, abetting and attempts. Existing offences should be used where they adequately cover the conduct; new offences must address an identified gap.

Offences must clearly state the prohibited conduct, required intent or other fault, consent requirements, harm threshold and applicable defences. Child-protection provisions must define who is a child and how real, synthetic and apparently child-like depictions are covered; sexual consent rules must not create a gap in protection. A deepfake is a type of content, not automatically an offence. Safeguards must protect lawful research, authorised security testing, journalism, art, satire and legitimate expression.

The sentencing review will consider AI as an aggravating factor where its use materially increases the scale, concealment, exploitation or harm of an offence. Courts must assess evidence and avoid counting the same harm twice. Penalties must match culpability and harm. Administrative breaches, civil compensation and criminal offences require distinct legal tests and procedures.

Legislation will define when intermediaries must respond to a sufficiently specific notice of unlawful content or a valid legal order. Any liability protection for mere transmission, access, storage or hosting must have clear conditions and must not protect a provider's own unlawful conduct. There must be no general obligation to monitor all user communications. Proportionate AI safety monitoring, incident reporting and cooperation with lawful orders will remain required.

C3 Evidence, powers and international cooperation

Legislative drafting must distinguish an AI provider from a service provider under cybercrime law. An AI provider qualifies for a service-provider rule only where it meets that rule's definition; other powers may apply to a wider category of persons. A recipient's location, services offered and possession or control of the specified data matter. Prompts, outputs, account records and logs must be classified by their content and function, rather than all being treated as subscriber information.

The drafting instructions must separately test preservation, production of stored data, production of subscriber information, search and seizure, real-time traffic collection and interception of content. A power to obtain subscriber information does not authorise access to all content. Stored-data powers do not automatically authorise real-time collection. Each requires its own statutory conditions, limits and safeguards.

Legislation will distinguish preservation from production, disclosure, search, seizure and interception. Preserving specified data prevents its loss; it does not itself authorise access. Each power must identify the competent authority, legal threshold, data concerned, scope, duration, oversight and review route. Intrusive measures require judicial authorisation appropriate to the power; any narrowly defined emergency exception must provide prompt independent review. Protect legal privilege, defence rights and unrelated persons' data.

Orders to restrict an AI system, remove unlawful content, freeze assets or confiscate proceeds must be lawful, necessary and proportionate. Preserve evidence securely while containing harm; protect essential services and innocent third parties. Confiscation and any exceptional measure without a conviction require their own statutory conditions, court oversight and challenge rights. A regulatory AI approval does not confer investigative powers.

The legal programme will define territorial and cross-border jurisdiction and resolve overlapping cases through lawful consultation. Mutual legal assistance, extradition, expedited preservation and any joint investigations must use applicable domestic law and international arrangements, with rights protections and lawful data transfers. Cooperation agreements do not by themselves authorise foreign access to government systems or bypass required legal assistance procedures.

PNG acceded to the Budapest Convention on 6 May 2026, and the Convention entered into force for PNG on 1 September 2026. Implementation will align with applicable obligations; the status of other treaties and protocols must be checked separately. Competent authorities will maintain operational cooperation channels and develop judicial, prosecutorial, police, forensic and provider capability. These activities and AI literacy will form part of the costed programme. [13]

C4 Comparative sources and drafting record

The EL PACCTO 2.0 Model Law on Artificial Intelligence and Crime is a comparative criminal-justice resource developed for Latin America and the Caribbean. PNG's economy-wide adoption mandate, risk tiers, approval regime and institutional choices require their own policy and legal basis. The T-CY mapping study informs the treatment of AI systems, electronic evidence and cooperation under cybercrime frameworks. Neither document is enacted PNG law. [11–12]

The Month 6 drafting submission must include a source-to-proposal record covering all relevant Model Law articles and T-CY recommendations. For each, identify the existing PNG provision, the demonstrated gap, the proposed response, the responsible authority and the necessary safeguards. Record departures and reasons. Claims of complete article-by-article alignment must be supported by that completed record.

Annex D Definitions

These definitions guide the policy and legislative instructions. Legislation must reconcile them with PNG law. Citizen explanations do not narrow these definitions. Technical categories describe capabilities; risk depends on actual use. [11–12]

Artificial intelligence system. A machine-based system that, for explicit or implicit objectives, infers from the input it receives how to generate predictions, content, recommendations or decisions that can influence physical or virtual environments. Systems may differ in their autonomy and ability to adapt after deployment. An output called a decision does not confer legal authority on the system.

AI model. A core component of an AI system that represents patterns, relationships or knowledge used to produce outputs. It may be built using expert knowledge, machine learning or both. A complete system also includes data, tools, interfaces and operating arrangements.

AI agent and agentic system. An AI system that can pursue an objective through a sequence of actions, using tools and, where designed, memory and planning. It may act alone or coordinate with other agents. Under this policy its actions remain bounded by human-set permissions and oversight.

AI provider. A person, legal entity or public body that develops, or has developed, an AI system or general-purpose model and places it on the market or puts the system into service under its own name or trademark, for payment or free of charge.

AI deployer. A person, legal entity or public body using an AI system under its authority, excluding purely personal non-professional use for this organisational role. That exclusion does not exempt anyone from criminal law or duties that expressly cover personal use.

Importer. A person or legal entity established in PNG that places on the PNG market an AI system bearing the name or trademark of a person or entity established overseas.

Distributor. A supply-chain person or legal entity, other than the provider or importer, that makes an AI system available on the PNG market.

AI operator. A collective term covering a provider, product manufacturer, deployer, authorised representative, importer or distributor. It is broader than the shared-service operator that runs government platforms.

Integrator and authorised representative. An integrator connects AI components into a working system and may also become a provider or deployer. An authorised representative acts under a written mandate for specified legal duties. Legislation must define each role and its limits.

Deepfake. AI-generated or manipulated image, audio, video or other content resembling existing persons, objects, places, entities or events that would falsely appear authentic or truthful. Whether it is unlawful depends on the conduct and applicable law.

Personal data. Information relating to an identified or identifiable natural person. Applicable privacy rules determine how it may be used.

Critical infrastructure. Assets, facilities, equipment, networks or systems needed to provide an essential service. Formal designation and associated duties must follow applicable PNG law.

High-risk AI. An AI use with material potential to harm rights, liberty, health, safety, livelihoods or access to essential services.

Impact assessment. A documented examination of an AI system’s purpose, affected people, benefits, alternatives, risks and safeguards throughout its lifecycle, including effects on human rights, democracy and the rule of law.

Criminal justice and implementation terms

AI offences. An analytical grouping of offences against AI systems or their data and offences committed by or by means of AI systems. It does not itself create an offence or make AI a legally responsible person. AI used in evidence gathering is a separate function.

Malicious purposes. Intentional use of an AI system to cause unlawful harm or obtain a benefit without right. Each offence must still specify its own required fault and other elements.

Organised criminal group. A structured group of three or more persons, existing for a period and acting together to commit serious crimes or offences for direct or indirect financial or other material benefit. AI offences may also be committed by individuals outside such groups.

Serious crime. For the Model Law comparison, an offence carrying a maximum deprivation of liberty of at least four years or a more serious penalty. PNG legislation and applicable treaties must determine the operative threshold. Serious AI harm is a separate incident-reporting concept.

Preservation and disclosure. Preservation protects specified data against loss or alteration. Disclosure supplies data to another party. Each requires its own lawful basis; preservation alone does not permit access.

People at particular risk. Children, people with disability, women facing gender-based abuse and others at increased risk of exploitation or repeated harm. Protection must respond to the person's circumstances.

Dual-use technology. A tool capable of lawful and unlawful uses. Its legitimate development, supply or possession is not by itself proof of criminal purpose.

LLM agnostic. Able to change between language models from different providers, subject to suitability, cost and safeguards.

Baseline and target. The measured starting position and the improvement to be achieved by a stated date.

Data governance. Rules for collecting, accessing, using, sharing, correcting, keeping or deleting information, with named accountability.

Exemption and deferral. An exemption removes a specified adoption duty under stated conditions. A deferral moves its deadline.

Interoperability and sovereign control. Interoperability allows systems to work together. Sovereign control is PNG's practical and lawful ability to govern public systems and data, including suppliers, access, changes and exit.

Supervised testing. A limited trial overseen by relevant authorities, with safeguards, a defined duration and conditions for stopping or expanding.

Whole-life cost and cost recovery. Whole-life cost covers acquisition, operation, security, support, change and retirement. Cost recovery charges authorised users an approved share of service costs.

Abbreviations: ICT — information and communications technology; DICT — Department of Information and Communications Technology; DDDAI — proposed Department of Digital Development and Artificial Intelligence; DJAG — Department of Justice and Attorney General; NICTA — National Information and Communications Technology Authority; CACC — Central Agencies Coordination Committee; NEC — National Executive Council; PMNEC — Department of Prime Minister and National Executive Council; DPM — Department of Personnel Management; RPNGC — Royal Papua New Guinea Constabulary; NCSC — National Cyber Security Centre; TVET — technical and vocational education and training; IFMS — Integrated Financial Management System; HR — human resources.

AI capabilities and evidence terms

Generative AI. Models and systems that produce content such as text, images, audio, video or code. Generated material can be plausible but inaccurate.

Large language model. An AI model trained on substantial language data to process or generate language. An LLM is one model type and does not describe all AI.

Predictive and analytical AI. Systems that estimate outcomes, classify information or identify patterns and relationships. A prediction or detected correlation does not establish a cause or a legal finding.

General-purpose and foundation models. Broadly capable models that can support a range of tasks and downstream systems. Foundation model commonly describes a model trained on broad data for later adaptation. These terms do not imply proven suitability for every use; legislation must specify the scope of any model-level duties.

Multimodal and combined systems. Multimodal systems process or produce more than one form of data, such as text, images and audio. A combined system may use several models, retrieval, memory and tools. These features can overlap with generative, predictive and agentic capabilities.

Training, inference and retrieval. Training builds or changes a model. Inference applies it to inputs. Retrieval finds information for use by a system and does not itself mean retraining. Each activity requires appropriate data permissions and safeguards.

Computer system and ICT system. A device or related group of devices that automatically processes data under a program. The T-CY study treats AI systems as a subset of computer systems. Legislative drafting must reconcile these terms with PNG law without treating every computer system as AI.

Computer data. Facts, information or concepts represented in a form that a computer system can process, including programs. AI inputs and outputs can be computer data.

Service provider under cybercrime law. Under the Budapest Convention, an entity enabling computer-based communication, or another entity processing or storing computer data for such a service or its users. This legal category is distinct from AI provider and the government shared-service operator. Eligibility and powers must be assessed under applicable law.

Content, traffic and subscriber information. Content conveys the substance of a communication. Traffic data describe its transmission, such as origin, destination, route and time. Subscriber information identifies a subscriber, service or account arrangements and excludes traffic and content data. Classify AI records by their actual function.

Hallucination and automation bias. A hallucination is generated information that appears credible but is false or unsupported. Automation bias is undue reliance on an automated output. Both require verification and meaningful human judgement.

HUDERIA and COBRA. HUDERIA is guidance for assessing effects on human rights, democracy and the rule of law. COBRA supports context-based risk analysis. HUDERIA's Model is an assessment resource, not a computational AI model.

Digital government and agent training terms

Digital Government or DG. The national reform agenda for accessible services and effective, accountable administration, using digital technology and data alongside changes to institutions, processes and skills.

Digital Public Infrastructure or DPI. Shared, secure and interoperable systems that support public and private services, including identity, payments, trusted data exchange and core registers, with governance and safeguards.

Core concept. $DG \times DPI^{AI}$ expresses integrated service reform and shared foundations, accelerated by AI. The exponent signals the ambition for exponential improvement; it is not a numerical forecast. Section 2.4 defines the concept and its relationship with the technology-control layers and risk framework.

Digital Economy. Economic activity enabled or improved by digital technologies, data and services. It is an intended beneficiary of this approach; DG continues to mean Digital Government.

GovAoS. The Government Agentic Operating System: a shared application platform for PNG government records, correspondence and AI-assisted workflows, connected to DPI through approved interfaces.

GovAI Agents. Government AI assistants approved for defined public duties. Each has a functional owner, authorised knowledge and tools, tested limits and a controlled release. The name does not confer legal authority.

Sovereign agent training. Government-led teaching, configuration and testing for approved workflows and DPI interfaces, with control of knowledge, permissions, tools and release. It can use existing models without foundation-model training.

Functional owner. The public body and accountable official responsible for an agent's professional purpose, lawful scope, approved sources and acceptance. This differs from owning or operating the platform.

Knowledge retrieval and fine-tuning. Retrieval supplies permitted source material when an agent performs a task. Fine-tuning changes a model's parameters through further training. Approval for retrieval does not automatically permit fine-tuning or other reuse.

Teaching and assessment cases. Teaching cases demonstrate approved work. Assessment cases are kept separate and used to test whether an agent performs reliably on work it was not taught. Validators must be independent of the work they assess.

Training assets. Approved task instructions, examples, labelled data, knowledge collections, evaluation cases and related configuration or documentation. These assets retain their access, licensing, retention and cultural restrictions.

Official work memory. Retained context that helps continue authorised government work. It supplements official records, follows role permissions and retention rules, and is separate from personal preferences.

SevisDEx and SevisWallet. SevisDEx provides authorised exchanges between connected systems and is the proposed integration platform for GovAoS. SevisWallet supports identity and sign-in. Neither integration nor identity alone grants access to government records.

Coordinating department. DICT and its lawful successor responsible for ICT and AI after the proposed departmental reform takes effect. Section 4.2 sets out its policy role.

SevisPNG Ltd. The proposed operator of designated common DPI, GovAoS and shared government AI infrastructure, subject to the ownership, mandate and approval conditions in section 4.5. Naming it in this policy does not establish or change its legal status.

Policy commencement. The date of National Executive Council approval, called Month 0. The policy lasts five years from that date. A statutory duty may begin later, when its enabling instrument and implementation arrangements take effect.

Annex E Proposed GovAoS Agents

Agent portfolio and proposed functional owners

Agent	Intended function	Proposed lead and contributors
GovAsk	Find and explain permitted official information.	Coordinating department; knowledge officers and source custodians.
GovWrite	Draft letters, memos, briefs and reports.	PMNEC; correspondence, policy and records officers.
GovMeet	Prepare agendas, draft minutes and action lists.	PMNEC; secretariat and records officers.
GovAssist	Support tasks, scheduling and routine administration.	DPM; HR and administration officers.
GovPolicy	Develop evidence, options and policy advice.	PMNEC; Treasury, National Planning and policy sponsors.
GovPulse	Track projects, budgets and commitments.	National Planning and Monitoring; Finance, Treasury and project owners.
GovAudit	Prepare internal audit working papers and control checks.	Finance internal audit; Auditor-General advises independently.
GovContract	Review contractual obligations and risks.	Justice and Attorney General; government lawyers and National Procurement Commission.
GovInstruct	Prepare legislative drafting instructions.	Justice and Attorney General; policy sponsors and Legislative Counsel.
GovDraft	Prepare draft laws and amendments.	Office of Legislative Counsel; authorised counsel and instructing officers.
GovInvestigate	Organise lawfully acquired records, timelines and draft investigation briefs.	RPNGC and forensic specialists; preserve evidence originals and professional verification.
GovProsecute	Map evidence to offences, research law and draft submissions or disclosure schedules.	Office of the Public Prosecutor and authorised prosecutors; independent decisions and fair disclosure.

Specialist preparation may run alongside the core programme. Live-case deployment requires separate scope, funding, trained staff and agency, legal and security acceptance. Conviction rates must not be agent performance targets. All professional approvals remain with the authorised institutions.

GovAoS delivery stages

Delivery proceeds through records onboarding, sovereign teaching, independent validation, a supervised pilot, and verified expansion and renewal. Apply section 3.7 at every stage. The supporting GovAoS Implementation and Training Plan holds provisional dates, staffing and testing assumptions; it remains subject to this policy, law and required approvals.

Measure records, turnaround, overdue actions, quality, citations, rework, active use, security and cost per validated task. Deduct checking and correction time, avoid duplicate benefits across agents and verify cash savings separately. Illustrative productivity and revenue scenarios are not policy targets.

Annex F Sources and references

Legal instruments apply according to their scope and commencement. Policy documents, model provisions and guidance inform the approach without replacing PNG law. References identify the editions used for this consultation draft; legal advisers must verify current law before approval, drafting and enforcement.

[1] Constitution of the Independent State of Papua New Guinea. National Goals and Directive Principles, rights and institutional arrangements. Use the current authorised text and applicable constitutional instruments.

[2] Independent State of Papua New Guinea. Digital Government Act 2022, No. 41 of 2022. In particular, sections 2, 5–13, 20–31, 33–37 and 44–53. Official source

[3] Independent State of Papua New Guinea. National Information and Communications Technology Act 2009. NICTA publication. Official source

[4] Department of Information and Communications Technology. Digital Government Plan 2023–2027, published December 2022. Sections on the national programme, service principles, delivery and monitoring. Official source

[5] Department of Information and Communications Technology. Endorsed policy collection: Digital Transformation Policy 2020; National Cyber Security Policy 2021; Government Cloud Policy 2023; Data Governance and Data Protection Policy 2024; National Digital ID Policy 2025. Official source

[6] Ministry of ICT. Acting ICT Minister announces release of the Draft AI Adoption Framework for Public Consultation, 24 February 2026. Explanation of $DG \times DPI^{AI}$. The framework is a consultation document. Official source

[7] UNESCO. Recommendation on the Ethics of Artificial Intelligence, adopted 23 November 2021. Non-binding international recommendation. Official source

[8] OECD. OECD AI Principles and Recommendation of the Council on Artificial Intelligence, OECD/LEGAL/0449. Responsible stewardship and national policy guidance. Official source

[9] OECD. Digital public infrastructure for digital governments, OECD Public Governance Policy Papers, No. 68, 2024. DOI 10.1787/ff525dc8-en. Official source

[10] Council of Europe. HUDERIA Methodology and Model, 2026. Methodology approved 26 February 2025; Context-Based Risk Analysis resources approved 25 February 2026. Non-binding risk and impact assessment guidance. Official source

[11] EL PACCTO 2.0. Model Law on Artificial Intelligence and Crime. Madrid, August 2025. Comparative model for criminal-justice legislation; not binding PNG law. Official source

[12] Council of Europe Cybercrime Convention Committee. Cybercrime, electronic evidence and artificial intelligence: Mapping study, T-CY(2025)1, adopted 12 October 2026. Official source

[13] Council of Europe. Convention on Cybercrime, ETS No. 185, and PNG accession notice of 6 May 2026. Entry into force for PNG: 1 September 2026. Additional protocols have separate participation requirements. Official source

[14] Department of Information and Communications Technology. Public Consultation Opens on Priority Digital Legislation, August 2026. The announcement identifies draft reforms; subsequent enactment and commencement must be verified. Official source