



THE AI INTEGRITY INSTITUTE

Responsible, inclusive and practical AI governance

Submission to the public consultation on the

National AI Strategy of the Republic of Cyprus 2032

Comments and proposed amendments

Submitted to the Deputy Ministry of Research, Innovation and Digital Policy
through the e-Consultation platform of the Republic of Cyprus

Consultation open until 31 August 2026

Date of submission: 29 August 2026

Prepared and submitted by
The AI Integrity Institute Global
[Website](#) | [Email](#) | [LinkedIn](#)

Contents

1. About AI³ and the purpose of this submission

2. Executive summary

3. What the draft Strategy gets right

4. Part A. A citizen-centric approach

A1 Citizen AI Charter. A2 Public register of AI systems. A3 Human and non-digital routes.

A4 The Greek language. A5 Local government. A6 Civil society in governance.

5. Part B. Responsible AI in practice

B1 Regulation (EU) 2026/1744. B2 Authorities and mandates. B3 Regulatory sandbox.

B4 Impact assessment. B5 Incidents and redress. B6 Procurement.

B7 Private-sector conditionality. B8 Judicial AI and AIREG.

6. Part C. AI sovereignty across five layers

C1 Definition. C2 Infrastructure. C3 Data. C4 Models. C5 Assurance.

7. Part D. Delivery, sustainability and measurement

D1 Costed plan. D2 Institutional scale. D3 Environment. D4 Milestones and targets.

D5 International position. D6 AI literacy cooperation. D7 Internal consistency.

8. Conclusion and offer of contribution

Annex 1. Mapping of proposals to affected sections

Annex 2. International reference points

Note on structure. Each proposal states the sections of the draft Strategy that it affects, summarises the current position, sets out the proposed change and explains the reasoning. Annex 1 consolidates all proposals into a single mapping table.

1. About AI³ and the purpose of this submission

AI³, the AI Integrity Institute, is a non-profit organisation working to make artificial intelligence governance responsible, inclusive and practical. We connect researchers, policymakers, industry, public institutions and civil society, and we work across three areas of activity. The first is dialogue and governance, through expert groups, roundtables and international partnerships. The second is research and thought leadership, through policy papers, comparative analysis and evidence for decision-makers. The third is implementation and capacity building, through toolkits, guidance and training that help organisations move from principles to practice.

Cyprus is the international launchpad for our work. We have a direct interest in the success of this Strategy, and we are submitting these comments in that spirit. We regard the draft as a serious and unusually honest document. It names its own weaknesses, it sets measurable ambitions, and it treats governance as an enabler rather than an obstacle. Our comments are intended to strengthen it, not to replace its direction.

This submission is organised around three themes that we believe determine whether the Strategy will succeed in practice.

- A citizen-centric approach. The Strategy speaks about citizens frequently, but it gives citizens very few enforceable rights, few visible mechanisms and no seat at the governance table. We propose specific ways to close that gap.
- Responsible AI as an operating system rather than a principle. The ethical commitments in the draft are sound. What is missing is the machinery that turns them into repeatable practice, in particular assessment methodology, incident handling, procurement controls and public reporting.
- AI sovereignty defined in operational terms. The Strategy uses the word sovereignty in several different senses. We propose a single working definition across five layers, so that sovereignty can be planned, budgeted and measured rather than asserted.

Method. This submission draws on three sources. The first is the draft Strategy itself, read in full including all annexes. The second is the full set of comments published on the e-consultation platform up to the date of submission. The third is a comparative review of national AI strategies, governance instruments and legal frameworks from the European Union, the United States, the Council of Europe and other jurisdictions, selected for their relevance to a small state with a services-based economy. Working drafts of this submission were shared for comment with a small number of external reviewers, including from industry. All assessments, positions and editorial decisions in this document are AI³'s own.

A number of the proposals set out below converge with points raised independently by other respondents to this consultation. We regard that convergence as useful information for the drafting team rather than as duplication. Where several parties arrive at the same conclusion from different professional starting points, that is evidence about the draft rather than about the respondents. We have not repeated analysis that others have already set out well, and we have concentrated our own contribution on the areas where we believe the consultation record is thinnest.

We have written this document so that every proposal can be acted on directly. Each proposal names the sections it affects, states what the draft currently says, sets out the change we recommend, and explains the reasoning. Annex 1 collects all proposals into a single mapping table for the drafting team.

2. Executive summary

The Cyprus National AI Strategy 2032 is a strong foundation. It correctly identifies the two structural weaknesses that will determine outcomes, which are fragmented data and infrastructure, and limited execution capacity. It refuses to present pilot activity as achievement. It anchors national ambition in European law rather than in a parallel national rulebook. These are the right choices for a small, open, services-based economy inside the European Union.

Our review identifies five areas where the draft is, in our assessment, materially incomplete.

- First, citizens appear in the Strategy as beneficiaries rather than as participants or rights-holders. There is no public register of the AI systems used on citizens, no single statement of what a citizen can demand when an AI system affects them, no guaranteed non-digital route to a public service, and no standing role for civil society in the governance structure. Civil society is mentioned once in a document of over one hundred pages.
- Second, responsible AI is described but not operationalised. There is no single national methodology for assessing impact on fundamental rights, no published incident and redress mechanism, no standard procurement clauses, and no obligation to publish what goes wrong. Without these, ethical commitments depend on the goodwill of individual project teams.
- Third, sovereignty is asserted across compute, data and skills, but the model layer is largely absent and the application and assurance layer is underdeveloped. The Strategy also does not mention the Greek language at any point. For a country whose citizens will interact with AI in Greek, and whose public administration operates in Greek, this is a significant omission at both the citizen and sovereignty levels.
- Fourth, several provisions applying to the private sector are drafted in terms that would create legal uncertainty. The draft states that private-sector adoption is voluntary and then conditions the grant or renewal of operating licences in strategic sectors on disclosure of a company's AI strategy, using triggering terms that are nowhere defined. It also proposes an AI Judge Capability without stating that judicial authority remains with judges. Both provisions concern the exercise of public power over private parties and both require correction in the final text.
- Fifth, the delivery architecture is not proportionate to the scale of the country, the financial commitment is not stated, and the headline targets do not reconcile with each other or with the baseline data in the document. The Strategy proposes eleven or more new bodies for a country of roughly one million people, without a published budget, a responsibility matrix or an independent evaluation mechanism.

We also note one point of currency. The draft was prepared before Regulation (EU) 2026/1744, the Digital Omnibus on AI, entered into force on 27 July 2026. That Regulation amends the AI Act and changes the compliance calendar that several parts of the draft rely on. The final Strategy should reflect the amended dates.

This submission contains twenty-six proposals. They are grouped into four parts. Part A concerns the citizen. Part B concerns responsible AI in practice. Part C concerns AI sovereignty. Part D concerns delivery, sustainability and measurement. Fourteen of the twenty-six proposals require no new institutions and no significant new expenditure. We have marked these in Annex 1, because they are the fastest route to visible progress.

Summary of proposals

Part A. A citizen-centric approach

- A1. Publish a Citizen AI Charter with enforceable service commitments.
- A2. Establish a public register of AI systems used by public bodies.
- A3. Guarantee a human route, a non-digital route and accessible design for every citizen-facing system.
- A4. Treat the Greek language as both a citizen right and a sovereign capability.
- A5. Bring local government into the Strategy as a delivery partner.
- A6. Give civil society a defined institutional role in governance.

Part B. Responsible AI in practice

- B1. Update the Strategy to reflect Regulation (EU) 2026/1744 and restate the compliance calendar.
- B2. Clarify the relationship between the National AI Authority and the authorities already designated under the AI Act.
- B3. Commit to a dated, operational AI regulatory sandbox under Article 57 of the AI Act.
- B4. Adopt one national methodology for fundamental rights and ethical impact assessment.
- B5. Establish incident reporting, redress and public disclosure of failures.
- B6. Use procurement as the principal control point, replace the AI-first framing, and open tenders to Cypriot suppliers.
- B7. Remove the link between private AI policy disclosure and operating licences, and define the private-sector obligations precisely.
- B8. Rename and constrain the judicial AI provisions, and clarify the scope of the AIREG registry.

Part C. AI sovereignty across five layers

- C1. Adopt a single operational definition of sovereignty across five layers.
- C2. Define infrastructure sovereignty in testable terms and track concentration risk.
- C3. Sequence data sovereignty correctly and align health data with the European Health Data Space.
- C4. Add the missing model layer, model selection criteria, evaluation and Greek-language capability.
- C5. Build a national AI assurance and conformity capability, and treat it as an export sector.

Part D. Delivery, sustainability and measurement

- D1. Publish a costed delivery plan with lifecycle costs and named owners.
- D2. Right-size the institutional architecture to the scale of Cyprus.
- D3. Make environmental sustainability an operational requirement, not a principle.
- D4. Set dated milestones, fix the target arithmetic and report annually to the House of Representatives.
- D5. Strengthen the international position, including the Council of Europe Framework Convention.
- D6. Frame international cooperation on AI literacy as open and non-exhaustive.
- D7. Correct a small number of internal inconsistencies in the text.

3. What the draft Strategy gets right

We begin with what should be preserved. Consultation responses that list only deficiencies are easy to discount, and several strengths of this draft are unusual in national AI strategies and should not be lost in redrafting.

The draft is honest about the starting position. It states that enterprise AI adoption in Cyprus stands at 9.27 per cent against an EU average close to 20 per cent. It states that institutional compute is fragmented and that data governance is immature. It names the risk of AI theatre, meaning many experiments and little impact. Very few national strategies publish figures that are inconvenient to their own narrative. This candour is an asset and should be retained in the final text.¹

The strategic positioning is correct for the scale of the country. Cyprus cannot compete on compute volume or capital. It can compete on regulatory credibility, speed of decision-making, and the ability to serve as a place where AI systems intended for regulated European markets are built, validated and certified. Treating compliance as a product rather than as a cost is the right bet, and it is consistent with what small jurisdictions have achieved in financial services, shipping and professional services.

The design principles and the control gate approach in Annex B are sound. The insistence that problems are defined before solutions are chosen, and that pilots scale only where value is demonstrated, addresses the most common failure mode in public sector AI programmes.

The treatment of data as foundational infrastructure rather than as a parallel workstream is correct. So is the decision to build on a federated architecture rather than to centralise sensitive information unnecessarily.

Finally, the Strategy places human oversight at the centre of high-impact systems and states that no high-risk system affecting individuals should operate without it. That commitment is the foundation on which most of our Part A and Part B proposals build.

¹ National AI Strategy of the Republic of Cyprus 2032, consultation draft, Section 2.2 and Annex D.

4. Part A. A citizen-centric approach

The Strategy is described as technology for citizens. The vision, the objectives and the design principles all refer to human-centred AI. Our concern is that the citizen appears throughout the document as a beneficiary of better services and almost never as a person with rights, choices or a voice in decisions.

Three tests are useful here. Can a citizen find out which AI systems are being used on them? Can a citizen understand what they are entitled to demand when an AI system affects a decision about them? Can a citizen or a representative organisation influence the Strategy as it is implemented?

On the first test, the draft provides a partial answer through Section 3.9.6, which mandates registration in the European database. That instrument is limited in three respects that matter here. It covers high-risk systems only, it is populated by providers rather than by the Cypriot public bodies deploying the systems, and it is a European compliance mechanism rather than a national tool designed for a citizen who wants to know what is being used on them. On the second and third tests the draft does not currently provide an answer at all. The proposals in this Part are designed to close all three gaps, using mechanisms that already work in other European jurisdictions.

Proposal A1. Publish a Citizen AI Charter with enforceable service commitments

Sections affected: 1.1 (Human Oversight and Social Contract Rights Protection), 1.4.8, 2.5.7, 3.1.2, 3.4.3, 3.9.3 to 3.9.6, Annex A (Ethics and Fundamental Rights), Annex F.

Current position. The draft already contains most of the substance of citizen rights, and Section 3.9.6 is where it comes closest to assembling them. That section embeds rights to explainability, contestation and human review in sensitive areas including health, credit, employment and social benefits. It commits to AI literacy for citizens, to transparent communication about deployments and to mechanisms for redress. It states that no high-risk or other system whose decisions may significantly affect individuals should operate without human oversight. Section 1.1 and the design principles in Section 3.1.2 make related commitments, and Annex A describes an aspiration to legislate for a right to human intervention and redress.

Our observation is about form rather than substance. Section 3.9.6 frames these as mandates for AI adoption, meaning obligations imposed on organisations. A mandate on an organisation and an entitlement of a person are not the same instrument, even where the underlying commitment is identical. A mandate is discharged through compliance documentation and is visible to a regulator. An entitlement is exercised by an individual and must therefore be findable, intelligible and enforceable by that individual. The draft has the first and not the second. There is no single statement of what a person may demand, no time limit attached to any of these commitments, no named point of contact, and no indication of where a citizen goes if the commitment is not honoured.

Proposal. Publish a single Citizen AI Charter as a short public document accompanying the Strategy, written in plain Greek and English, setting out what a person can expect and demand when a public body uses an AI system in a matter affecting them. The Charter should state the right to be told that AI is in use, the right to a meaningful explanation, the right to human review by a named responsible body, the right to challenge and obtain correction, the right to an alternative service route, and the right to know where to complain. Each right should carry a service standard, meaning a maximum response time and a named point of contact. The Charter should identify the source and legal status of each protection, distinguishing among rights established under Union or national law, administrative remedies that already exist, and service

commitments made through the Strategy. A policy commitment should not be presented as a new enforceable right unless a legal basis establishes it. Where the explanation right derives from Article 86 of the AI Act, the Charter should describe it accurately, meaning an explanation of the role of the AI system in the decision-making procedure and the main elements of the decision, rather than a general right to the internal workings of a model. The Charter should be referenced in Section 3.9 and reflected in the KPI framework in Annex F.²

Why this matters. The decisive element is the service standard. A right without a time limit and a named responsible body is difficult for an ordinary person to exercise, and it is difficult for a public body to plan around. International practice supports the instrument. The United States published a Blueprint for an AI Bill of Rights in 2022, which established the format of a short public-facing statement of expectations covering notice, explanation, human alternatives and protection from discriminatory outcomes. European practice has generally embedded equivalent rights in binding law rather than in a public statement, which is effective legally but leaves most citizens unaware that the rights exist. A charter closes that gap without changing the underlying legal position. Cyprus can move faster than either because it is smaller and because the underlying commitments already exist in the draft. A charter of one or two pages, published at the same time as the Strategy, would be the single most visible citizen-facing commitment in the document.³

Proposal A2. Establish a public register of AI systems used by public bodies

Sections affected: 3.3.2 (Control Framework), 3.4 (AI Adoption Across Government), 3.9.3 (Transparency and Explainability), 3.9.6, Section 5 item 4 (National AI application and certification registry, Development of an AI Registry), Annex B.

Current position. The draft proposes three registration instruments. Section 3.9.6 mandates risk classification of systems and registration in the European database, complemented by documented conformity assessments and an annual audit. Section 5 proposes a national AI application and certification registry, and separately an AI Registry cataloguing AI solutions available in the market. All three are compliance or market instruments. None is a national list, maintained by the Republic, of the systems that Cypriot public bodies are using on the people they serve.

The European is confined to high-risk systems as classified under the AI Act, which excludes a wide range of systems that materially affect a citizen without meeting that threshold. It is populated principally by providers rather than by the public bodies deploying the systems, so it does not answer the question of who in Cyprus is using what. It is structured for conformity purposes rather than for public comprehension. A citizen wanting to know whether an AI system was involved in a decision about their benefits would not find the answer there, and would not think to look.

Proposal. Create a single public register of AI systems and materially consequential algorithmic tools used by ministries, deputy ministries, public law bodies, local authorities and other public entities. The scope should be risk based. It should cover systems that are operational or approved for deployment and that materially affect rights, health, safety, eligibility or access to public services, together with pilots that involve live personal data or direct interaction with the public. Purely exploratory work need not be registered before it reaches that threshold. For each system, publish the purpose and scope of use, the

² Regulation (EU) 2024/1689 laying down harmonised rules on artificial intelligence ([Artificial Intelligence Act](#)), Article 86, on the right to explanation of individual decision-making.

³ The White House Office of Science and Technology Policy, Blueprint for an AI Bill of Rights: Making Automated Systems Work for the American People (October 2022).

responsible public body and contact point, the supplier or developer, the broad categories of data used, the risk classification and legal basis, the form of human oversight, the lifecycle stage, the impact assessment summary, the routes to challenge and human review, and any serious incidents, suspensions or withdrawals. Publication should be subject to justified exceptions for national security, cybersecurity and legitimate commercial confidentiality, and those exceptions should be narrow and stated. The register should be machine readable, searchable, and updated before a system enters operational use and after any substantial modification. It should be linked to the control gates in Annex B, so that registration is a condition of approval rather than an afterthought.

Why this matters. This is proven practice, not an experiment. The Netherlands has operated a national algorithm register since 2022 and has recorded approximately one thousand entries, with its data protection authority now calling for registration to be made mandatory for government organisations. The United Kingdom has made its Algorithmic Transparency Recording Standard mandatory for central government departments and for arm's length bodies that deliver public services or interact with the public. France, Norway and Scotland maintain national registers, and a group of European cities including Amsterdam, Helsinki, Barcelona and Rotterdam has developed a shared transparency data schema. Cyprus should adopt an existing schema rather than design its own, which would reduce cost and improve comparability.⁴

Practice outside Europe points the same way. In the United States, federal agencies are required under Office of Management and Budget Memorandum M-25-21 to compile and publish annual inventories of their AI use cases, with an expanded set of fields for uses that affect rights or safety. The instrument is different in legal form from the European registers, but the underlying commitment is the same, which is that the public is entitled to know where the state is using these systems.⁵

We would add one point of efficiency. The Strategy currently contemplates three separate registration instruments, and our proposal would be a fourth if implemented in isolation. We recommend instead that the instruments share a single data model and interoperable fields, published through a documented public interface in a machine readable schema, so that information is entered once and reused wherever it is needed. The obligations differ in purpose, scope and legal basis, so one record will not automatically discharge all of them, but a shared data model avoids parallel data entry and inconsistent descriptions of the same system. Building four unconnected registers for a public administration of this size would be a poor use of scarce capacity.

⁴ Autoriteit Persoonsgegevens, [Algorithm registration in the Netherlands needs improvement](#) (2025), recording that 1,000 algorithms have been entered in the [Algorithm Register of the Dutch government](#) since the beginning of 2023 and calling for registration to be made mandatory for government organisations; Government Digital Service, [Algorithmic Transparency Recording Standard Hub](#), GOV.UK, and the [Mandatory Scope and Exemptions Policy](#) (December 2024), under which the Standard is mandatory for all government departments and for arm's-length bodies delivering public or frontline services or interacting directly with the public.

⁵ Office of Management and Budget, Memorandum M-25-21, Accelerating Federal Use of AI through Innovation, Governance, and Public Trust (3 April 2025), which replaced Memorandum M-24-10 (28 March 2024) and retains the requirement that agencies inventory their AI use cases at least annually and publish the releasable inventory. For an example of a published inventory see United States Department of Justice, [AI Use Case Inventory](#).

Proposal A3. Guarantee a human route, a non-digital route and accessible design

Sections affected: 2.5.4, 3.4.1, 3.4.3, 3.1.1 (Personalisation at scale), Annex A (Inclusive growth), Annex C (Government and Public Sector).

Current position. The draft commits to human oversight of high-risk systems and refers to multilingual support and inclusive access tools so that people with disabilities can interact with services. Digital inclusion is mentioned in the Executive Summary in relation to people with disabilities and pensioners. There is no commitment that a citizen who cannot or does not wish to use an AI-mediated channel retains an equivalent route to the service.

Proposal. Add an explicit commitment structured by the significance of the service. For essential public services, and for any service in which an AI system makes or materially informs a decision affecting a person's rights, eligibility, entitlements or effective access, the responsible public body should guarantee timely and meaningful access to human assistance, and where appropriate human review or an alternative service route, for people who cannot or do not wish to use the AI-enabled channel. The alternative arrangements should carry no additional charge, should not place the person at an unjustified disadvantage in timeliness or completeness, and may be provided through digital, telephone or in-person channels as appropriate. They need not require a wholly separate end-to-end service infrastructure where AI performs only a limited back-office function. Add an explicit commitment that accessibility requirements under the European Accessibility Act apply to AI-enabled interfaces from design stage rather than as remediation. Add a requirement that AI-enabled citizen services are tested with older users, users with disabilities and users with low digital confidence before deployment, and that the results of that testing are published in summary form.⁶

Why this matters. Cyprus has high generative AI use overall, with roughly 44 per cent of people aged 16 to 74 reporting recent use, but that figure conceals a widespread by age and by circumstance. A strategy that improves service for confident digital users while degrading it for everyone else will reduce trust even as it raises efficiency. The commitment costs little at design stage and is expensive to retrofit. It also protects the Strategy politically, because the most damaging early failure in a national AI programme is usually a vulnerable person unable to reach a service.⁷

Proposal A4. Treat the Greek language as both a citizen right and a sovereign capability

Sections affected: 2.5.6, 3.2 (Data as a Strategic National Asset), 3.5 (Legal Services), 3.7 (Infrastructure, Compute and Digital Sovereignty), 3.8.6 (Pharos-CY), Annex C (Government, Tourism, Legal, Healthcare).

Current position. The draft refers to multilingual services in at least a dozen places. It proposes a legal language model trained on Cyprus law and a language model capability able to handle Cyprus case law. It does not name the Greek language anywhere in the document. It does not refer to Cypriot Greek, and it does not refer to Turkish, which is an official language of the Republic.

Proposal. State the language position explicitly. Establish a national language resources programme, coordinated through Pharos-CY and the research ecosystem, with four components. First, the assembly and open publication of Greek-language corpora and terminology resources for public administration,

⁶ Directive (EU) 2019/882 on the accessibility requirements for products and services ([European Accessibility Act](#)).

⁷ National AI Strategy of the Republic of Cyprus 2032, consultation draft, Section 2.2, recording that 44.2 per cent of people aged 16 to 74 in Cyprus used generative AI tools in the preceding three months.

health, law and tourism, with clear licensing and in compliance with applicable data protection, copyright and confidentiality rules. Second, the creation and open publication of Greek-language evaluation datasets and benchmarks, including Cypriot Greek usage, so that procurement can distinguish between a system that genuinely performs in Greek and one that is machine translated. Third, participation in European multilingual model initiatives rather than an attempt to build a national foundation model from scratch. Fourth, a stated commitment on the official languages of the Republic and on the practical language needs of residents and visitors.

Why this matters. This is where the citizen agenda and the sovereignty agenda meet. If public services are delivered through models that perform poorly in Greek, the citizen experience degrades and the state becomes dependent on suppliers whose language performance it cannot verify. European practice is now well established. Spain has invested public funding in the ALIA family of models at the Barcelona Supercomputing Centre, with a strong focus on Spanish and the co-official languages, released openly. The Netherlands is building GPT-NL as a public interest model for the Dutch language and context. Germany produced Teuken-7B through a publicly funded consortium. The OpenEuroLLM project aims at model families covering official Union languages with access to EuroHPC resources. Estonia is building an Estonian-adapted model as part of its public services programme.⁸

Cyprus should not attempt to replicate these projects at national scale. The proportionate strategy for a country of this size is to own the data and the evaluation layer, which is where language capability is actually determined, and to participate in European model development rather than to finance a national model alone. Evaluation infrastructure in particular is a classic public good. It is too expensive for any single supplier to build, it is inexpensive at national scale, and it disproportionately benefits suppliers who have genuine Greek-language capability to demonstrate. Without it, the phrase supports Greek is a claim on a sales document that procurement has no means to test.

Proposal A5. Bring local government into the Strategy as a delivery partner

Sections affected: 2.5.4, 3.3.1 (Governance Structure), 3.4 (AI Adoption Across Government), 3.4.4 (Shared Capabilities and Reuse), Annex C (Government and Public Sector).

Current position. The Strategy is written almost entirely for central government. Ministries and deputy ministries carry the obligations, AI Officers sit inside ministries, and the Applied AI Strategy requirement is directed at ministries and public bodies. Local authorities appear only marginally.

Proposal. Include local authorities explicitly in Section 3.4 and in the governance structure. Provide a proportionate route for them, since a municipality cannot reasonably be expected to produce an Applied AI Strategy on the same terms as a ministry. We suggest three elements. First, a shared services model in which local authorities consume national capabilities such as the citizen assistant, document intelligence and translation rather than procuring their own. Second, a simplified adoption template and a single point of support within the National AI Authority. Third, representation of local government in the governance structure, because local authorities deliver the services citizens interact with most often.

Why this matters. In most European countries, local government is where public sector AI meets the citizen most frequently, in areas such as permits, waste, local planning, social support and local information services. It is also where capacity is weakest. A strategy that assumes ministry-level capability throughout

⁸ Barcelona Supercomputing Centre, ALIA and the Salamandra model family; TNO, SURF and the Netherlands Forensic Institute, GPT-NL; OpenGPT-X consortium, Teuken-7B; and the OpenEuroLLM consortium, which develops open model families covering the official languages of the Union with access to EuroHPC resources.

will produce a two-speed public sector, in which the services citizens use most often are the least well governed. The municipal algorithm registers developed in Amsterdam, Helsinki, Rotterdam and other European cities show that local government can lead on transparency when it is given a framework to work with rather than an obligation to invent one.

Proposal A6. Give civil society a defined institutional role in governance

Sections affected: 3.3.1 (Governance Structure), 3.3.1.3 (National AI Taskforce), 3.3.1.5 (National Ethics and Values Committee), 3.3.1.6, 3.11.3, 4.1.6 (Continuous Feedback Loop).

Current position. The Strategy places considerable weight on public trust, human-centred AI, ethics and societal benefit. The governance architecture it proposes does not, however, provide any explicit institutional role for civil society. The structure comprises governing bodies, advisory bodies, monitoring and oversight bodies, innovation bodies and ministry level roles, and every one of them is drawn from government, academia, research and industry. There is no route through which organisations working on fundamental rights, consumer protection, accessibility and inclusion, or organisations representing the groups most affected by AI systems, can contribute to implementation.

Proposal. We recommend three measures, which are complementary rather than alternative.

- **Representation on the National Ethics and Values Committee.** The Strategy states at Section 3.3.1.5 that this Committee shall oversee the promotion of ethical principles and public trust across the implementation and use of AI systems. That mandate is inseparable from questions of societal values, fundamental rights, inclusion and public confidence, and it cannot be discharged credibly by a body composed entirely of state, academic and industry representatives. Civil society representation on this Committee should be provided for expressly.
- **Representation on the National AI Taskforce.** Section 3.3.1.5 provides that all recommendations and guidelines of the Ethics and Values Committee are submitted to the National AI Taskforce. Representation at Committee level therefore has limited effect unless the Taskforce, which receives and acts on those recommendations, also includes a civil society voice.
- **A dedicated Civil Society and Stakeholder Engagement Mechanism.** Representation within specific bodies is necessary but not sufficient, because it depends on a small number of individuals and on the frequency with which those bodies meet. We recommend a standing mechanism comprising transparent and open calls for civil society representatives, periodic structured consultation, and a permanent channel through which organisations may submit observations and recommendations at any point during implementation. The mechanism should include representation from consumer and human rights bodies, trade unions and professional associations, organisations representing people with disabilities and older people, educators and parents, journalists, small business, and citizens selected through structured deliberative processes.

We further recommend that the mechanism be tied to three fixed decision points. It should comment publicly on the annual implementation report. It should be consulted before deployment of high-impact citizen-facing systems. It should publish an independent opinion before any substantial revision of the Strategy. Transparency and conflict of interest rules should apply to composition and to the selection process. Participation should complement, and not displace, accountable decision-making by the responsible public bodies and the independent exercise of statutory functions.

Why this matters. It gives effect to objectives the Strategy has already adopted, namely protection of the public interest, public trust in AI, transparency, accountability, and the ethical and responsible

development and use of AI systems. Structured civil society participation reinforces those commitments and helps translate the Strategy's stated vision into practice. It also strengthens pluralism, independent scrutiny and societal legitimacy, which are precisely the qualities that determine whether a national AI programme survives its first public controversy.

The approach is consistent with established international practice. Multi-stakeholder governance models at European Union, Council of Europe and UNESCO level have expressly involved civil society organisations alongside academia, industry and public authorities in the development of AI policy and governance frameworks. This is not an unusual request and Cyprus would not be an outlier in granting it.⁹

The design details are what make this proposal effective. Participation mechanisms often fail when they remain consultative in the abstract, without fixed decision points to which they are attached. They can then deteriorate into an annual meeting that neither side finds particularly valuable. Tying participation to the annual report, high-impact deployments, and revisions of the Strategy gives it a concrete function rather than merely a formal status. This is also among the least expensive proposals in this submission. It requires a secretariat function, published selection criteria, and dedicated meeting time, but does not require the establishment of a new institution.

⁹ European Commission, [General-Purpose AI Code of Practice](#), noting the multi-stakeholder process and consultation involving a broad range of stakeholders in the development of the Code; Council of Europe, [Framework Convention on Artificial Intelligence and Human Rights, Democracy and the Rule of Law](#), noting that 68 international representatives from civil society, academia and industry were actively involved in the development of the Convention; UNESCO, [Global Civil Society Organizations and Academic Network on AI Ethics and Governance](#), noting that the 2021 Recommendation on the Ethics of Artificial Intelligence recognises the essential role of civil society organisations in shaping responsible and inclusive AI development and governance.

5. Part B. Responsible AI in practice

The ethical content of the draft is strong at the level of principle. Section 3.9 covers ethical foundations, human oversight, transparency, fairness and responsible deployment. Annex A restates the OECD principles. The gap is between principle and operation. A principle becomes an operating system when it has a method, an owner, a trigger point in the delivery lifecycle, and a consequence when it is not followed. The first six proposals in this Part supply those four elements.

The final two proposals address a different problem. Responsible AI governance requires not only that the state act well, but that the obligations it places on others are lawful, defined and proportionate. Two sets of provisions in the draft fall short of that standard, and we address them at B7 and B8.

Proposal B1. Update the Strategy to reflect Regulation (EU) 2026/1744 and restate the compliance calendar

Sections affected: 1.4.10, 2.3.3, 3.11.2, 3.12.1, Annex A (EU AI Act), Section 5.

Current position. The Strategy cites the AI Act as Regulation 2024/1689 in Section 1.4.10 and in its footnotes. It gives one compliance date, in the Section 5 table, which asks for risk management, auditability and governance maturity ahead of 2 August 2026. That date has passed and the obligations behind it have changed. The Strategy makes no mention of Regulation (EU) 2026/1744, which amended the AI Act on 27 July 2026, six days before that deadline. There is no calendar of dates a ministry or supplier could plan against.

Proposal. Three changes. Update the citation in Section 1.4.10 and the footnotes to read Regulation (EU) 2024/1689 as amended by Regulation (EU) 2026/1744. Replace the 2 August 2026 milestone in Section 5 with the dates that now apply. Add a short calendar table to Annex A setting them out.

Some obligations are already in force. The prohibitions under Article 5 have applied since 2 February 2025, and the new provisions added by the amending Regulation apply from 2 December 2026. Obligations for general purpose AI models have applied since 2 August 2025, with models placed on the market before that date required to comply by 2 August 2027. The transparency obligations under Article 50 have applied since 2 August 2026, with a transition allowing systems already on the market until 2 December 2026 to mark synthetic content in machine readable form.¹⁰

The remaining dates fall in the years ahead. At least one national AI regulatory sandbox must be operational by 2 August 2027. Stand-alone high-risk systems listed in Annex III come into scope on 2 December 2027, and high-risk AI built into products regulated under Annex I on 2 August 2028. High-risk systems already in use by public authorities have until 2 August 2030. The final text should note that these dates remain subject to further amendment of Union law.

¹⁰ [Regulation \(EU\) 2026/1744](#) amending Regulations (EU) 2024/1689, (EU) 2018/1139 and (EU) 2023/1230 as regards the simplification of the implementation of harmonised rules on artificial intelligence (Digital Omnibus on AI), OJ L, 2026/1744, 24 July 2026, in force 27 July 2026; and Regulation (EU) 2024/1689 ([Artificial Intelligence Act](#)), Articles 5, 50, 111 and 113.

Proposal B2. Clarify the relationship between the National AI Authority and the designated competent authorities

Sections affected: 3.3.1 (Governance Structure and Figure 2), 3.3.1.1 (National AI Authority), 3.3.1.5, 3.9.6, Annex A (EU AI Act).

Current position. The Strategy contains two institutional maps and they do not reference each other. Section 3.3.1 and Figure 2 set out the governance structure for the Strategy, comprising governing, advisory, monitoring, innovation and ministry level bodies. Annex A separately states that the Communications Commissioner and the Commissioner for Personal Data Protection are appointed as competent authorities to enforce the AI Act, and that the Strategy aims to support these authorities. Neither Commissioner appears in Figure 2 or in the Section 3.3.1 list.

The consequence is a set of unresolved questions. Section 3.3.1.1 proposes a National AI Authority acting as policy coordination body, execution body and governance and control gatekeeper, whose designation is to be determined by the President and ratified by the Council of Ministers, but does not state how that role relates to authorities that already hold statutory enforcement powers. Figure 2 describes the National Ethics and Values Committee as overseeing compliance with the EU AI Act, which is a function of the designated authorities rather than of an advisory committee. Section 3.9.6 refers to an annual audit by a Cyprus AI Security and Certification Authority, a body that appears nowhere else in the document and is not defined.

Proposal. Merge the two maps. Annex A should carry a single institutional map covering both the bodies created by the Strategy and the authorities designated under Union law, and Figure 2 should be redrawn to include the designated authorities and to show the interface between them and the Strategy bodies. State that the Deputy Ministry of Research, Innovation and Digital Policy coordinates implementation of the AI Act and represents Cyprus on the European Artificial Intelligence Board. State that the Commissioner of Communications acts as notifying authority, market surveillance authority and national single point of contact. State the role of the Commissioner for Personal Data Protection. Identify the public authorities responsible under Article 77 for supervising fundamental rights obligations in relation to high-risk systems listed in Annex III of the AI Act. Then state clearly what the National AI Authority is and is not. The draft is currently inconsistent on this point. Section 3.3.1.1 states that the Authority does not assume operational control over AI initiatives, while Section 5 assigns it an executive mandate covering standards, enablement, compliance oversight and platform operations. Those two descriptions cannot both be accurate. In our reading, the Authority should be a strategy delivery and capability body for the public sector, its enablement and platform functions should be institutionally and procedurally separated from any compliance assessment role, and it should not exercise or appear to exercise regulatory enforcement powers that are vested in independent authorities. Regulated entities should have a single point of contact, formal coordination and escalation arrangements among the authorities, and predictable decision timelines. Finally, either define the Cyprus AI Security and Certification Authority properly or remove the reference and reassign the audit function to an identified body.

Why this matters. Regulatory clarity is the product Cyprus is selling. A strategy that creates apparent ambiguity about who supervises AI in Cyprus undermines the trusted jurisdiction claim at its foundation. There is also a point of constitutional form that applies across the whole governance chapter. A strategy can establish advisory, assurance and control gate arrangements binding on public bodies. It cannot confer licensing, enforcement or adjudicative powers on any body unless legislation provides them. The final text should be precise about which of its proposed arrangements require primary or secondary legislation, which require a Council of Ministers decision, and which can be implemented administratively. Without

that distinction the implementation timetable is not reliable, because the legislative items have a different lead time from the administrative ones.

Proposal B3. Commit to a dated, operational AI regulatory sandbox

Sections affected: 3.8.5 (National Testbeds), Section 5 item 2 (Industry sandboxes programme), Annex A (Leverage Regulatory Sandboxes).

Current position. The draft refers to sandboxes in many places, including industry sandboxes, sectoral sandboxes, a maritime sandbox and regulatory sandboxes generally. Annex A describes establishing AI regulatory sandboxes as an action that will safeguard competitiveness. No date, no responsible authority and no funding are attached.

Proposal. State the legal position and commit to a date. Article 57 of the AI Act requires each Member State to ensure that its competent authorities establish at least one AI regulatory sandbox at national level. The original deadline of 2 August 2026 was moved by Regulation (EU) 2026/1744, and the obligation now stands on 2 August 2027. The obligation may be satisfied jointly with another Member State, or by participating in an existing sandbox that provides equivalent national coverage. The final Strategy should name the responsible authority, state the operating model, commit to an operational date within the deadline, and state whether Cyprus intends to meet the obligation alone or jointly. Participation should be voluntary and should support testing and compliance readiness. It should not become a precertification mechanism or an additional prerequisite for procurement, deployment or market access. The Strategy should also confirm the priority access for small and medium enterprises and start-ups required under Articles 58 and 62.¹¹

Why this matters. Cyprus has a live legal deadline eleven months away and no named authority, no operating model and no date in its national strategy. Many Member States were slow enough to establish sandboxes that the Union moved the deadline for everyone. A Cypriot sandbox that is genuinely operational ahead of August 2027, well documented and open to foreign small and medium enterprises would be a concrete demonstration of the trusted jurisdiction proposition, and it is achievable within the resources the Strategy already contemplates. A joint sandbox with another Member State would reduce cost further and is expressly permitted.

Proposal B4. Adopt one national methodology for fundamental rights and ethical impact assessment

Sections affected: 3.9 (new subsection), 3.3.2 (Control Framework), 3.4.2, Annex B (Stage 2 and Control Gate 2), Annex F.

Current position. The draft requires mandatory legal, ethical and data protection assessments for high-impact systems and approval gates before procurement, deployment and scaling. It does not specify a methodology, a template, a publication requirement or a consultation requirement.

Proposal. Adopt a single national methodology for an integrated impact assessment covering fundamental rights, ethics, data protection and, where material, environmental and workforce effects. Require it for high-risk systems and for any system that materially affects rights or access to public services, whether or not it is classified as high risk under the AI Act. Require that it be completed before Control Gate 2 in Annex

¹¹ Regulation (EU) 2024/1689, [Articles 57, 58 and 62](#), as amended by [Regulation \(EU\) 2026/1744](#), which moved the date by which at least one national AI regulatory sandbox must be operational from 2 August 2026 to 2 August 2027.

B, that a non-confidential summary be published in the public register proposed at A2, and that it be repeated after substantial modification, material model or data drift, a serious incident, or a significant change of context. Require proportionate consultation with affected groups and frontline professionals where the likely impact is material. Publish the template and worked examples, so that a small public body can complete an assessment without external consultants. The methodology governs the public sector's own use of AI systems. It should not impose obligations on private organisations beyond those arising under Union and applicable sectoral law, and it should reflect the respective roles and degree of control of actors across the AI value chain.

Why this matters. The Chair of the Cyprus National Bioethics Committee has proposed an Ethical and Fundamental Rights Impact Assessment in this consultation, with detailed suggested wording. We support that proposal and add three operational points. First, there should be one methodology rather than several, so that a public body completes a single assessment rather than three overlapping ones. Second, the assessment must attach to a gate in the delivery lifecycle, or it becomes documentation produced after decisions are already made. Third, the template must be usable without specialist support, or small bodies and local authorities will not complete it. The Netherlands has established practice here with its Impact Assessment for Human Rights and Algorithms, applied to impactful public sector systems, and Canada has operated a mandatory Algorithmic Impact Assessment for automated decision systems in federal government for several years. Cyprus should adapt an existing instrument rather than write one from first principles.¹²

Proposal B5. Establish incident reporting, redress and public disclosure of failures

Sections affected: 3.3.2 (Control Framework), 3.9.5, 3.11.3, 3.13, Annex B, Annex F.

Current position. The commitment exists in three places. The design principles in Section 1.1 refer to clear redress mechanisms. Section 3.9.6 commits to mechanisms for redress as part of the social contract underpinning AI adoption. Annex A describes an aspiration to establish a robust and accelerated legal mechanism for human intervention and redress. The control framework adds that governance arrangements will include clear procedures for incident reporting, escalation and remediation. None of these passages states to whom an incident is reported, how a citizen reports harm, what is published, or within what period. Redress is named four times and defined once, in the future tense.

Proposal. Define the incident pathway in operational terms. Establish a single reporting channel for citizens, public servants and suppliers. Define what constitutes a serious incident, aligned with Article 73 of the AI Act. State the internal reporting timeline and the escalation route. Publish an annual account of serious incidents, systems suspended or withdrawn, corrective actions taken and time to remediation, in aggregated and anonymised form. Include reversal and override rates, and the number and resolution time of requests for human review, in the KPI framework in Annex F. Definitions, responsibilities and timelines should follow Article 73 rather than creating a parallel national scheme, and a single report should discharge both national and Union obligations wherever the law permits. Published material should be aggregated and anonymised and should protect personal data, trade secrets and security-sensitive information. State explicitly that transparent incident reporting will be treated as evidence of good governance rather than as failure, so that reporting is not suppressed.¹³

¹² Ministry of the Interior and Kingdom Relations (Netherlands), Impact Assessment Fundamental Rights and Algorithms (IAMA); Treasury Board of Canada Secretariat, Directive on Automated Decision-Making and the accompanying Algorithmic Impact Assessment tool.

¹³ Regulation (EU) 2024/1689, [Article 73](#), on the reporting of serious incidents.

Why this matters. Public trust in a national AI programme depends less on the absence of failure than on the visible handling of failure. A programme that publishes only successes will be assumed to be concealing problems, and the assumption will usually be correct. The point about not penalising disclosure is essential rather than decorative. If reporting an incident is treated as an admission of fault, incidents will be resolved quietly and the national record will show a safety performance that does not exist. There is also a practical benefit. Without an incident record the Strategy cannot learn, and Section 3.11.4 commits it to learning. The NIST AI Risk Management Framework, which is widely used in the United States and increasingly referenced by European organisations, treats measurement and incident feedback as core functions rather than as reporting overhead, and offers a usable structure for this.¹⁴

Proposal B6. Use procurement as the principal control point, replace the AI-first framing, and open tenders to Cypriot suppliers

Sections affected: 3.4.1 (AI-first procurement), 2.5.8, 3.1.2 (Design Principles), 3.3.2, 3.6.3, Annex B, Annex F.

Current position. Section 3.4.1 states that the government will work towards establishing principles and prerequisites to support the transition towards AI-first procurement models by 2032. The control framework refers to approval gates and to the use of approved and certified delivery partners. There are no standard contractual requirements.

Proposal. Two changes. First, replace the AI-first procurement framing with a problem-first framing, for example a commitment that public procurement will systematically assess whether AI is the appropriate solution to a defined problem, and that AI-enabled options will be evaluated on outcomes rather than on technology category.

Second, mandate a standard set of contractual requirements for AI procurement, issued by the National AI Authority in cooperation with the Treasury, the Law Office and the relevant supervisory authorities. These should cover access to technical documentation sufficient for oversight, the right to independent technical, legal, ethical and security audit, ownership and deletion of data, portability of exportable data and digital assets in documented and interoperable formats, substitutability of models and components where technically and legally feasible, avoidance of lock-in and unjustified exclusivity, open standards and interoperability where feasible, documented accuracy and fitness for the stated purpose, continuous monitoring of failure and bias, liability and remediation clauses, safe termination and tested exit, complete audit logs, publication of lifecycle cost, and evaluation of environmental impact where material. Audit and documentation rights should be scoped to the oversight objective and should protect trade secrets and security-sensitive information, with recognised certifications and independent third-party audit reports accepted where they provide sufficient evidence, and direct access reserved for cases where it is necessary and proportionate. Obligations should be allocated according to each party's role and degree of control in the value chain. We would note that model weights and capabilities provided through a managed service may not be technically transferable or legally available, which is why the requirement is framed as portability of data and assets and substitutability of components rather than as a general right to take a model elsewhere.

Why this matters. Procurement is where sovereignty, responsible AI and industrial policy are actually decided. Everything else in the Strategy is a statement of intent that the contract either implements or discards. A strategy can commit to portability, to auditability and to human oversight, but if the contract

¹⁴ National Institute of Standards and Technology, Artificial Intelligence Risk Management Framework (AI RMF 1.0), NIST AI 100-1 (January 2023).

does not require them, the state will not have them and will discover this at the point when it most needs them. On the AI-first point, we note a tension within the draft itself. Section 3.1.2 states that AI initiatives will be selected on measurable outcomes and not on technological novelty. The AI-first formulation states the opposite, and it appears twice, once in Section 3.4.1 and once in the description of the Government Innovation Hub, where ministries are to be supported in an AI-first and cloud-first approach. The final text should resolve this in favour of Section 3.1.2.

We would also add one point on lot sizing. Large public tenders normally require bidders to show a minimum annual turnover and evidence of comparable prior contracts. Applied to national AI programmes, that test is circular. A company can only qualify by having already delivered a programme of similar scale, and since Cyprus has never run one, no Cypriot company can meet the condition. Every domestic supplier is excluded before its capabilities are examined, and every flagship contract goes abroad by default. If the Strategy intends to build a Cypriot AI industry, dividing flagship programmes into lots that a small domestic supplier can bid for is the instrument that does it. It requires no additional funding, no new institution and no legislation, only a decision taken when the tender is drafted. Lot design should broaden participation without fragmenting accountability, security or coherent end-to-end delivery.

Proposal B7. Remove the link between private AI policy disclosure and operating licences, and define the private-sector obligations precisely

Sections affected: 3.4.2 (Institutional AI Strategies), 3.5, Annex C (Financial Services and related governance provisions), Annex A.

Current position. The Strategy states that private-sector AI adoption is voluntary. Other provisions qualify that statement significantly. Section 3.4.2 provides that every sizeable organisation must formalise an AI adoption plan. Annex C provides that in regulated sectors, institutional AI strategies become part of supervisory reporting to the respective regulators with the same legal weight as risk-management frameworks. It further provides that private enterprises dealing with critical and high-risk systems will be required to share their AI strategy and acceptable AI use policy in order to obtain or renew participation in state-subsidised AI programmes, access to the national GPU cloud or datasets, and new or renewed operating licences in strategic sectors. The terms sizeable organisation, critical and high-risk systems, strategic sectors and basic compliance are not defined.

Proposal. Four changes. First, preserve the voluntary character of private-sector AI adoption, and recast the statement that every sizeable organisation must formalise an AI adoption plan as guidance rather than obligation. Second, remove the connection between disclosure of an enterprise AI strategy or acceptable use policy and the grant or renewal of operating licences. Any condition capable of affecting an operating licence should arise from clearly defined requirements under applicable law, be administered by the competent sectoral authority, be proportionate, and carry procedural safeguards and rights of review. Third, do not give institutional AI strategies the same legal weight as risk-management frameworks by way of a strategy document. Supervisory reporting obligations should be established under Union and sectoral law, within the mandate of the competent regulator, and limited to information necessary for a defined supervisory purpose. Fourth, where eligibility conditions are attached to state-supported programmes, national compute or public datasets, they should be transparent, risk based, proportionate and non-discriminatory, should require only the information necessary to assess eligibility, and should accept targeted attestations or summaries rather than requiring disclosure of a full institutional AI strategy. Where high risk is intended to carry its AI Act meaning, the defined term should be used, and each of the undefined terms above should be defined or removed.

Why this matters. This set of provisions is, in our assessment, the most legally exposed part of the draft. A corporate AI strategy contains forward-looking commercial plans, technical architecture, security controls and supplier relationships. Conditioning market access on its disclosure, through a strategy document rather than legislation, with undefined triggering terms, creates legal uncertainty for every company operating in or considering Cyprus, and sits uneasily with the trusted jurisdiction positioning of the whole Strategy. It also risks creating a Cyprus-specific compliance layer diverging from the harmonised Union framework, which is the outcome Annex A is otherwise careful to avoid. The underlying objective, which is that significant AI users take governance seriously, is legitimate. The instrument should be Union and sectoral law for obligations, and proportionate eligibility criteria for access to public resources.

Proposal B8. Rename and constrain the judicial AI provisions, and clarify the scope of the AIREG registry

Sections affected: 3.5 (Legal Services), Annex C (Legal Services and Dispute Resolution).

Current position. Section 3.5 and Annex C propose an AI Judge Capability for very basic, factually uncontested cases valued under five thousand euro, subject to strict legal constraints and human oversight. They also propose AIREG, a blockchain-based registry providing immutable certification of AI assets.

Proposal. On the judicial provision, replace the term AI Judge Capability with AI-assisted judicial administration and decision support, and state expressly that AI may support administration, research and analysis but does not determine case outcomes, evaluation of evidence, interpretation of law, issue judgments or exercise judicial authority. The competent judge should retain independent judgment, decisional authority, responsibility for the reasoning and accountability for the outcome, and any deployment should protect the right to a fair trial, procedural fairness, contestability and effective human oversight. We note that judicial and case-outcome uses fall within the high-risk categories of the AI Act, with the obligations that follow. On AIREG, define its scope, legal effect and disclosure requirements. Registration should be voluntary and evidentiary, supporting proof that a record existed at a point in time and has not been altered. It should not create or determine ownership, alter copyright or intellectual property rules, or be presented as establishing Union-wide presumptions, and it should not require disclosure of model weights, training data, source code or trade secrets.

Why this matters. A national strategy that announces an AI Judge, whatever the safeguards in the surrounding text, invites every reader to conclude that Cyprus intends machines to decide cases, and that impression will outlive the caveats. The substantive design may well be defensible as decision support for uncontested small claims. The name is not, and the constitutional principle at stake, which is that judicial authority is exercised by judges, should be stated in the Strategy rather than assumed. On AIREG, a registry that certifies existence and integrity is useful and modest. A registry read as certifying ownership or conferring legal status across the Union would promise more than national law can deliver, and the final text should make the distinction explicit.

6. Part C. AI sovereignty across five layers

Sovereignty is central to the draft. It appears in the vision, in Objective 6, in Objective 8, in the infrastructure chapter and in the description of the Government Innovation Hub. It is used in at least three different senses across the document. Sometimes it means physical control of compute. Sometimes it means legal and regulatory autonomy. Sometimes it means the ability to make independent decisions while relying on partners for delivery.

This ambiguity has a practical cost. Sovereignty that is not defined cannot be planned, budgeted, procured or measured. It becomes a word that every stakeholder interprets in their own favour, and a supplier can claim to satisfy it without changing anything material. We therefore propose a single operational definition.

Proposal C1. Adopt a single operational definition of sovereignty across five layers

Sections affected: 1.1 (Vision), 2.4 (Vision), 2.5.6 (Objective 6), 2.5.8 (Objective 8), 3.7.3 (Digital Sovereignty and Strategic Autonomy), Annex G (Glossary).

Proposal. Define AI sovereignty in the Strategy as meaningful national control over five layers, each with its own test, owner and indicator.

- Infrastructure sovereignty. Where computation runs, under whose jurisdiction, and whether Cyprus can continue operating critical services if a supplier relationship ends.
- Data sovereignty. Who decides access to national datasets, under what legal basis, with what auditability, and whether that decision can be exercised independently.
- Model sovereignty. Whether Cyprus can choose, inspect, adapt, evaluate and replace the models on which public services depend.
- Application and assurance sovereignty. Whether Cyprus can independently verify that a system does what it claims, safely and lawfully, without relying on the supplier's own assurance.
- Talent sovereignty. Whether the knowledge required to specify, commission, operate, supervise and, if necessary, rebuild these systems resides inside Cypriot institutions.

Add the definition to the glossary in Annex G, use it consistently across the document, and add a small number of sovereignty indicators to Annex F. The Strategy should say where Cyprus intends to be on each layer by 2032.

Three points of interpretation keep the framework practical. Each layer is measured through verifiable outcomes and controls, meaning what Cyprus can actually do, verify and reverse, and each layer can be satisfied while working with international partners. In the model layer, inspection means access to information and evidence proportionate to the use case, sufficient to evaluate performance, limitations, safety, security and compliance. Talent sovereignty means the capability to specify, commission, evaluate and govern AI systems, which is the layer where the Strategy's existing education and skills provisions already do most of the work.

Why this matters. The five-layer framing is useful precisely because it shows where the current draft is strong and where it is thin. Infrastructure, data and talent are well covered. The model layer is barely addressed. The assurance layer is present in ambition but has no delivery mechanism. Naming the layers makes the gaps visible and makes trade-offs explicit. It also gives Cyprus a distinctive and exportable governance contribution, since most national strategies still treat sovereignty as a single undifferentiated concept.

Proposal C2. Define infrastructure sovereignty in testable terms and track concentration risk

Sections affected: 3.7.1, 3.7.2, 3.7.3, 3.7.4, 3.7.5, 3.10.2, Annex F.

Current position. Section 3.7.3 states that infrastructure choices should be guided by interoperability, portability, vendor diversity, supply chain resilience, security and long-term sustainability. This is the right list. It is expressed as guidance rather than as requirement, and none of it is measured.

Proposal. Convert the guidance into testable criteria applied at procurement and reviewed annually. We suggest five. First, legal and regulatory exposure, meaning the extent to which applicable legal regimes may compel access to data or affect service continuity, assessed for the relevant workload and alongside the technical, contractual and organisational safeguards available to the customer. Second, key control, meaning who holds encryption keys and administrative access. Third, exit, meaning a documented and tested plan to move a workload to another provider within a stated period. Fourth, dependency, meaning a published register of critical dependencies on single suppliers across compute, models and platforms. Fifth, continuity, meaning what happens to critical services if a supplier relationship ends at short notice. Add two indicators to Annex F. The first should measure material dependency on any single provider across the compute, model and platform layers, regardless of the provider's nationality, ownership or place of establishment, assessed through workload criticality, substitutability, interoperability, tested exit arrangements and service continuity. The second should measure the share of critical public workloads subject to legal regimes that permit compelled access from outside the Union without effective technical or contractual safeguards. The first indicator measures resilience. The second measures the specific exposure that motivates the sovereignty objective, and neither can substitute for the other.

We recommend that Cyprus adopt the European Commission's Cloud Sovereignty Framework rather than design national criteria of its own. The framework assesses providers against eight sovereignty objectives, including whether the bodies with decisive authority over the provider sit within Union jurisdiction and the extent of exposure to non-Union laws with extraterritorial reach. It assigns assurance levels, and the Commission applied it in April 2026 to award a sovereign cloud contract worth up to one hundred and eighty million euro, setting a minimum assurance level as a condition of eligibility. Adopting it gives Cyprus a tested methodology at no development cost, aligns national procurement with Union practice rather than diverging from it, and answers in advance the objection that sovereignty criteria fragment the internal market.¹⁵

Why this matters. The Strategy accepts a hybrid model of national compute combined with European and commercial partnership, which is the correct and realistic choice for a country of this size. The risk in a hybrid model is not the partnership itself. It is the absence of an exit path. The distinction between sovereignty as ownership and sovereignty as reversibility is the one that matters for a small state, and reversibility, unlike ownership, is testable. We would add that Cyprus does not need to design this discipline from scratch. The third-party risk regime applied to the financial sector under the Digital Operational Resilience Act already addresses concentration risk, exit strategy, documentation access and continuous monitoring of critical providers. Cypriot financial institutions and their supervisors are already operating it.

¹⁵ European Commission, [Sovereign Cloud Framework explained](#) (1 June 2026), describing the eight Sovereignty Objectives, the Sovereignty Effectiveness Assurance Level and the award in April 2026 of a sovereign cloud contract worth up to EUR 180 million to four providers, with SEAL-2 set as the minimum eligibility threshold.

Extending the same discipline to government AI procurement uses existing national expertise rather than creating a new competence.¹⁶

Proposal C3. Sequence data sovereignty correctly and align health data with European infrastructure

Sections affected: 3.2 (all subsections), 2.5.6, Annex C (Healthcare and Life Sciences), Annex D (Focus Area 3).

Current position. The Strategy treats data modernisation and the National Data Policy as a prerequisite for AI adoption rather than as a parallel activity, which is correct. It proposes a federated architecture, a national data layer and a national intelligent digital API fabric. It proposes a National Health Data Repository. Annex C states, under international alignment for healthcare, that data and interoperability will follow established standards including HL7 FHIR, SNOMED CT and ICD, and will reflect European Health Data Space principles.

Proposal. Three additions. First, state the sequencing explicitly and accept that digital maturity differs across public bodies, so that AI adoption is staged according to the readiness of each organisation rather than applied uniformly. Second, establish a formal data access framework, including data access committees with published decision criteria and published decisions, so that access to high-value national datasets is governed by a visible process rather than by negotiation. Third, convert the existing reference to European Health Data Space principles into an operational commitment. Reflecting principles is not the same as implementing a Regulation. The final text should identify the national bodies responsible for primary and secondary use of health data, state how Cyprus will connect to the cross-border infrastructures and set a date. The semantic and technical standards already named in Annex C should be carried into the main body of the Strategy rather than left in a sectoral annex, because they apply well beyond healthcare.¹⁷

Why this matters. The Ministry of the Interior and the National eHealth Authority have both made these points in this consultation, from the operational side. We support them and add the governance point. Data sovereignty is not only achieved by holding data nationally. It is achieved by controlling, documenting and being able to justify decisions about access. A federated architecture without published access rules simply relocates the discretion rather than governing it. Published access decisions also protect the state, because they demonstrate lawfulness in a way that is difficult to challenge after the fact.

Proposal C4. Add the missing model layer

Sections affected: 2.5.8 (Objective 8), 3.7 (Infrastructure, Compute and Digital Sovereignty), 3.8 (Research, Innovation and Ecosystem Development), 3.8.6 (Pharos-CY), Annex D (Focus Area 6).

Current position. The draft addresses compute, data, skills and governance. On models it says relatively little. The Government Innovation Hub is to guide and support developments for sovereign specialised learning models. A legal language model trained on Cyprus law is proposed under Legal Services. There is no national position on which models the state will use, on what basis it will choose them, whether it will require open weights for certain uses, how it will evaluate them, or how it will change them.

Proposal. Add a short model layer subsection to Section 3.7 or Section 3.8, covering five points. First, a model selection policy for the public sector enabling choice among commercial, open weight and nationally

¹⁶ Regulation (EU) 2022/2554 on digital operational resilience for the financial sector ([DORA](#)), in particular Articles 28 to 30 on the management of third-party risk, concentration risk and exit strategies.

¹⁷ Regulation (EU) 2025/327 on the [European Health Data Space](#), OJ L, 2025/327, 5 March 2025, in force 26 March 2025. See also European Commission, [Directorate-General for Health and Food Safety](#).

adapted models on the functional requirements and risk profile of each use case, with stated criteria covering performance and accuracy including Greek-language performance, security, privacy, safety, transparency, legal compliance and licensing terms, interoperability, operational support, total cost of ownership, and the ability to substitute components or switch solutions. Second, a scoped weighting for open weight models in defined use categories where inspection, independent evaluation or long-term availability is material to the public interest, applied through the selection criteria of the specific use case rather than as a general preference or a mandatory procurement condition, and without excluding commercial models where they better meet the requirement. Third, national evaluation capability, meaning the ability to test a model against Cypriot requirements, including Greek language performance, domain accuracy, refusal and escalation behaviour, and known failure modes. Fourth, adaptation rather than creation, meaning that Cyprus fine-tunes and adapts existing models for national purposes and participates in European model initiatives rather than financing a national foundation model. Fifth, portability, meaning that public sector systems are architected so that the underlying model can be replaced without rebuilding the service.

Why this matters. Model choice is where dependency is created and where it is hardest to reverse. The point is frequently misunderstood. In agentic and assistant-based systems the model itself is comparatively easy to substitute. The lock-in accumulates in the layer above it, in orchestration, tool integration, evaluation harnesses and capability description. It follows that the Strategy should commit to open, vendor-neutral interoperability standards at that layer, adopted from recognised international bodies rather than defined nationally. Cyprus is too small a market to set a standard and large enough to be harmed by adopting the wrong one.

National practice in Europe now clusters around the approach we recommend, which is public investment in data, evaluation and adaptation, combined with participation in shared European model development. Spain, the Netherlands, Germany and Estonia have all taken versions of this route. Portugal is the clearest recent illustration. Its national model for European Portuguese, released openly in July 2026, was produced for an initial public investment of approximately five and a half million euro by a consortium of universities, and it was built by extending an existing open European model rather than by training one from scratch. That is the cost envelope and the method a country of Cyprus's size should be planning against. Cyprus can be a serious participant in this space at a cost proportionate to its size, provided the objective is capability rather than a headline model.¹⁸

Proposal C5. Build a national AI assurance and conformity capability and treat it as an export sector

Sections affected: 2.5.1 (Objective 1), 3.8 (Research, Innovation and Ecosystem Development), 4.3 (NAICF), Section 5 item 4, Annex C (Entrepreneurship and Innovation).

Current position. The Strategy proposes a National AI Compliance Framework, a certification registry, an annual audit function and centres of excellence. It positions Cyprus as a compliance and testing hub. It does not describe the capability required to perform assurance, and it does not identify who will be qualified to do it.

¹⁸ Government of Portugal, [LLM AMALIA shows Portugal's potential](#) (1 July 2026); and European Commission Interoperable Europe Portal, [Portugal launches open source AI to drive digital sovereignty](#) (July 2026), recording an initial investment of EUR 5.5 million in European recovery funds and development from the open EuroLLM-9B model.

Proposal. Set out an assurance capability plan with three elements. First, technical evaluation capacity, meaning people and facilities able to test AI systems for accuracy, robustness, bias, security and compliance, hosted in the research ecosystem and available to the public sector and to industry. Second, a route to conformity assessment infrastructure, meaning an explicit assessment of whether Cyprus should host one or more notified bodies under the AI Act, and if so, what accreditation and staffing that requires and by when. Third, a services proposition, meaning the deliberate development of AI assurance, audit, evaluation and compliance services as an export industry building on the existing professional services base. Where the capability performs conformity assessment under the AI Act, it should operate within the applicable accreditation and notified body framework. It should not create additional national certification, registration or audit requirements, and it should not extend to the supervision of general purpose AI models as such, which is the role of the European AI Office. Voluntary evaluation and compliance support services should complement the Union framework and should not become a prerequisite for procurement, deployment or market access.

Why this matters. This is the most underdeveloped commercial opportunity in the draft, and it is the one best matched to what Cyprus already has. The country has an established professional services sector in audit, law and compliance. The AI Act creates sustained demand for conformity assessment, technical documentation, evaluation and audit across the Union, and the deferral of the high-risk deadlines to December 2027 and August 2028 creates a preparation window rather than removing the demand. Assurance is also the layer at which sovereignty becomes economically productive. A country that can independently verify AI systems does not need to own them in order to trust them.

7. Part D. Delivery, sustainability and measurement

The final Part concerns whether the Strategy can be delivered by a state of this size, whether the environmental cost has been accounted for, and whether progress can be measured honestly. These are the questions a reader in 2032 will ask.

Proposal D1. Publish a costed delivery plan with lifecycle costs and named owners

Sections affected: 1.5 (Implementation Timeline and Success Metrics), Section 5 (Immediate Steps), 3.10.3 (Funding and Renewal Model), Annex F.

Current position. The Strategy sets an implementation timeline in phases and lists immediate steps. It proposes a National AI Research Fund, an AI Innovation Fund, a national adoption acceleration programme and a compliance funding programme. It is not silent on money. Section 3.10.3 states that funding may combine public investment, European co-funding through EuroHPC and Digital Europe, and structured cost-sharing with major private-sector users. Section 3.3.1 states that budget allocations will be aligned with strategic performance goals agreed between the Ministry of Finance and the Deputy Ministry. The AI Sovereign Investment Matching Fund carries a stated cap of one million euro per round per company.

What is absent is scale and attribution. There is no total investment envelope for the Strategy, no figure attached to any of the four proposed funds, no capital or operating cost for the national compute infrastructure, and no budget or named delivery owner against any individual action. The funding sources are named in the conditional. Funding may combine those elements, which is a description of the options available rather than a decision about which will be used.

Proposal. Publish a costed delivery plan for the period to 2032 as a companion document, updated annually. For each major action, state the responsible organisation, the delivery lead, the timetable and interim milestones, the expected deliverables, the estimated budget and funding source, the staffing requirement and the principal dependencies. Assess cost across the full lifecycle rather than at procurement, including operation, maintenance, security, audit, data management, model monitoring, licensing, support, upgrade and decommissioning. State the total public investment envelope for the Strategy and its expected sources, including national funds, European instruments and any cost sharing.

Why this matters. Comparison is instructive. Estonia, a country of comparable population with a stronger digital baseline, backed its AI and Data Action Plan for 2024 to 2026 with a stated package of approximately eighty-five million euro across public sector transformation, private sector innovation and skills. That is the order of magnitude at which a small European state operates when it is serious. A strategy that names no figure at all invites the assumption that no figure has been agreed, and that assumption will affect how the document is received by suppliers, by the research community and by international partners. Publishing the envelope also protects the Strategy, because it converts any future underfunding into a visible departure from a stated plan rather than an unremarked absence.¹⁹

Proposal D2. Right-size the institutional architecture to the scale of Cyprus

Sections affected: 3.3.1 (Governance Structure, Figure 2 and all subsections), 3.3.1.4, 4.2, Section 5 item 1 and item 3.

¹⁹ ERR News, [Estonia to invest €85M in boosting AI uptake across public, private sectors](#) (February 2024), on the AI and Data Action Plan (Kratt) 2024 to 2026 and the accompanying White Paper on Data and Artificial Intelligence.

Current position. The draft proposes a National AI Authority including a Government AI Accelerator, an Interministerial AI Council, a National AI Taskforce, a National AI Infrastructure Committee, a National Ethics and Values Committee, a National AI Misinformation and Security Council, a National AI Skills Observatory, a Government Innovation Hub, an Industrial Centre of Excellence, a National AI Cybersecurity Centre of Excellence, AI Officers in every ministry and AI Champions in every ministry. Additional bodies appear elsewhere in the text. Figure 2 presents these across four tiers alongside a control framework of seven lifecycle controls.

Figure 2 is a useful diagram and it is not an accountability instrument. It shows which bodies exist and how they are grouped. It does not show, for any given decision, which body decides, which approves, which is consulted and which remains answerable for the outcome. Those are different questions, and the second set is the one that determines whether a governance structure functions under pressure.

Proposal. Two changes. First, publish a governance and accountability matrix identifying, for every function, which body sets policy, which approves, which delivers, which supervises and which remains accountable for outcomes, together with escalation and dispute resolution procedures. Second, consolidate. We suggest that the monitoring and advisory functions currently distributed across four separate bodies could be delivered by a smaller number of standing committees reporting through the National AI Taskforce, and that the three centres of excellence should share a single governance and administrative structure even if they maintain distinct technical missions. Where an existing institution can perform a function, the Strategy should say so rather than create a new body.

Why this matters. The obvious objection to the proposed architecture is the risk of overlapping mandates and fragmented decision-making. The more serious objection is arithmetic. The proposed architecture requires a substantial number of senior, scarce and specialised people to staff committees. Cyprus does not have a deep pool of such people, and the Strategy itself identifies this as one of the two structural weaknesses it must overcome. Every additional body consumes the same scarce individuals who are needed to deliver actual systems. Small states that have succeeded in digital government have generally concentrated authority rather than distributed it. The risk here is not that these bodies will conflict. It is that several of them will exist only on paper, which damages credibility more than not creating them at all.

Proposal D3. Make environmental sustainability an operational requirement

Sections affected: 3.10.1 (Energy and Environmental Sustainability), 3.7 (Infrastructure), 3.4.1 (procurement), Annex D (Focus Area 1), Annex F.

Current position. Section 3.10.1 commits to green computing principles, renewable energy integration where feasible, efficient cooling and power utilisation, and shared services. This is one short subsection. There is no measurement, no reporting obligation, no procurement weighting and no target.

Proposal. Make the environmental provisions operational in four ways. First, state that data centres supporting national AI infrastructure will report energy and water key performance indicators in line with Article 12 of the recast Energy Efficiency Directive and Commission Delegated Regulation (EU) 2024/1364, which apply to facilities with installed information technology power demand of at least 500 kilowatts and cover energy consumption, capacity utilisation, waste heat, renewable energy and both total and drinking water consumption. Second, publish national aggregate figures annually, reusing information already reported under Union and national mechanisms rather than creating a duplicative Cyprus-specific disclosure channel, with aggregation that protects trade secrets and security-sensitive information. Third, include energy and water performance as weighted criteria in the procurement of AI infrastructure and

services, and require environmental impact assessment for material workloads. Fourth, add environmental indicators to Annex F, covering energy intensity of national AI infrastructure, water consumption, share of renewable supply and waste heat recovery.²⁰

Why this matters. This is a matter of national self-interest for Cyprus, not a general environmental preference. Cyprus operates an isolated electricity system with no interconnection to a continental grid, which makes additional large and inflexible loads more consequential than in most Member States. Cyprus also faces severe and worsening water stress, and the draft itself identifies water management as a research priority and refers to water scarcity prediction. Data centre cooling is water intensive. A national AI infrastructure programme that does not measure its own energy and water use is inconsistent with the Strategy's own research priorities and with the country's physical constraints.

There is a further point. The European Commission is developing a Union rating scheme for data centres, with electronic labels expected from 2027. Facilities that already measure and report will be able to demonstrate performance from the outset. Building measurement into the national infrastructure now is cheaper than retrofitting it, and it supports the Strategy's positioning on trust and verifiability.

Proposal D4. Set dated milestones, fix the target arithmetic and report annually to Parliament

Sections affected: 1.5, 2.5.2, 3.4, 3.5, 3.11 (all subsections), Section 5 item 4, Annex C, Annex F.

Current position. The Strategy has targets but very few milestones. The distinction matters. A target is an endpoint. A milestone is a dated checkpoint with an owner and published evidence.

Quantified endpoints appear throughout Section 1.5 and Annex C. They include a recommended productivity improvement of 15 per cent, a recommended GDP expansion of 12 per cent, more than three thousand AI professionals, a 40 per cent reduction in citizen waiting times, 50 per cent of the medical workforce trained in AI-enabled practice, movement in education and labour adoption from 13 per cent to 75 per cent, and sector figures for shipping and financial services. Almost all of these are dated 2030 or 2032 with nothing in between.

Dated commitments appear in only one place. Section 5 provides for procurement of the first six transformational AI solutions within eight months and the identification of five priority use cases per ministry within twelve months. These are genuine milestones and they are the exception.

The measurement framework itself is not yet operational. Annex F is expressly described as a sample and as indicative. It lists categories of indicator rather than indicators. It contains no baseline value, no numerical target, no data source, no measurement method, no reporting frequency and no organisation named as responsible for collection or verification. The implementation timeline in Section 1.5 uses relative bands of 0 to 8 months, 6 to 12 months and 12 to 24 months, with no anchor date, followed by a band labelled 2026 to 2032 that overlaps the earlier bands and is already out of date. The headline figures are prefixed with the word recommended, which is not the language of commitment.

Proposal. Six changes, of which the first three are the priority.

²⁰ Directive (EU) 2023/1791 on [energy efficiency](#), Article 12; and [Commission Delegated Regulation \(EU\) 2024/1364](#) on the first phase of the establishment of a common Union rating scheme for data centres, which applies to data centres with an installed information technology power demand of at least 500 kW.

- Publish a milestone schedule with absolute dates. Replace the relative time bands in Section 1.5 with calendar dates rebased to the date on which the Strategy is adopted. Each milestone should name the responsible body and the evidence that will be published to demonstrate completion.
- Complete Annex F. Attach to every principal indicator a baseline value, a numerical target, a target date, a data source, a measurement method, a reporting frequency and a named organisation responsible for collection and verification. Baseline measurement should be completed within twelve months of adoption, because an indicator without a baseline cannot be reported against.
- Provide for one mandatory mid-point review in 2029. The review should be conducted by an independent external evaluator, published in full, and used to restate, revise or retire targets. It should be accompanied by a corrective action plan where performance is behind trajectory. A single mid-point review is the minimum needed to make a 2032 horizon manageable, and it is the point at which course correction is still possible.
- Reconcile the adoption targets to a single headline definition with an explicit base, scope and date, and label sectoral targets distinctly.
- Publish the assumptions and methodology behind the productivity and GDP projections, including the baseline year, the counterfactual and the external variables, or restate them as scenarios rather than as targets and remove the word recommended.
- Report annually to the House of Representatives with public debate, alongside a published annual progress report in an open and machine readable format.

Why this matters. On milestones, the practical effect of the current drafting is that the Strategy can be judged in 2032 and almost nowhere before then. That is the opposite of what Section 3.11.4 states, which is that measurement should function as a learning tool rather than as a compliance exercise. It also creates a specific institutional risk. The Strategy will span at least one parliamentary term. A document whose first meaningful checkpoint falls near its end date offers no protection against quiet drift, and no evidence base for a future government deciding whether to continue, adjust or fund it further. Dated milestones with published evidence are what convert a strategy into a programme.

Why this matters. On the arithmetic, the figures in the draft do not reconcile with each other. The Strategy records enterprise AI adoption at 9.27 per cent in 2025 following annual progression of 17.3 per cent, and it sets a target of 75 per cent industry adoption by 2030. Held at the rate the document itself observes, 9.27 per cent compounds to approximately 21 per cent by 2030, which is short of the target by a factor of more than three. Reaching 75 per cent from that base within five years requires sustained annual growth of roughly 52 per cent, which is about three times the observed rate. Extending the horizon to 2032 does not resolve the problem, because it still requires roughly 35 per cent annual growth. This should be corrected before Annex F is finalised. A target that cannot be reached will not be managed, and its presence undermines confidence in the targets that are achievable.

On parliamentary reporting, the Strategy currently proposes a governance structure that is entirely executive and technocratic, with no role for the legislature in ongoing supervision. AI affects the exercise of public authority, the allocation of public funds and the rights of citizens, and oversight of all three is a parliamentary function in the constitutional order of the Republic. An annual report and debate is a modest requirement in resource terms. It also substantially improves the durability of the Strategy across political cycles, which the draft identifies as an explicit aim of the National AI Taskforce. A strategy that Parliament has debated annually is considerably harder for a successor government to abandon without explanation.

Proposal D5. Strengthen the international position

Sections affected: 2.3.4, 2.5.1, 3.12 (all subsections), Annex A.

Current position. The Strategy positions Cyprus as a bridge between the European Union and neighbouring regions, refers to readiness to participate in the Global Partnership on AI, and aligns with the OECD principles. It does not refer to the Council of Europe Framework Convention on Artificial Intelligence and Human Rights, Democracy and the Rule of Law, and it does not describe what the bridge role means in practice.²¹

Proposal. Two additions. First, address the Council of Europe Framework Convention explicitly, stating Cyprus's position and how the Convention relates to obligations arising under Union law. This matters because the Convention is the principal instrument through which non-Union states in the region engage with rights-based AI governance, and it is therefore directly relevant to a bridge role. Second, convert the bridge positioning into concrete commitments, for example a regional forum on AI governance hosted in Cyprus, cooperation programmes with Eastern Mediterranean and Middle Eastern partners on assurance and standards, and participation in international capacity building. A bridge with no traffic on it is a metaphor rather than a strategy.

Why this matters. This is the area where Cyprus has the clearest comparative advantage and the least developed plan. Cyprus is a Member State of the European Union, physically located at the meeting point of three continents, with established professional and commercial links across the Eastern Mediterranean and the Gulf. Regulatory credibility is transferable in a way that compute capacity is not. If Cyprus becomes the place where regional actors learn how European AI governance works in practice, the bridge role generates both influence and revenue. This is also the area where AI³ can contribute directly, and we set out below what we are able to offer.

Proposal D6. Frame international cooperation on AI literacy as open and non-exhaustive

Sections affected: 3.6 (Education and Human Capital Development, International alignment), 2.5.8 (Objective 8), 3.12, Annex C (Education and Human Capital).

Current position. Section 3.12 adopts a broad and outward-looking approach to international engagement, envisaging Cypriot participation in European and global initiatives and cooperation with partners beyond the Union on AI standards, governance and best practice. The treatment of education and human capital is narrower. Under International alignment, the Strategy provides that the UK National AI Skills Framework will inform occupational pathways and competency structures, complemented by collaboration with Finland, France and Estonia on teacher training, AI literacy and ethics in education. Annex C refers additionally to Finland, France, Estonia, the United Kingdom and the United Arab Emirates. The criteria by which these jurisdictions were selected are not stated, and it is not clear whether the list is intended to be illustrative or exhaustive. We note that the Executive Summary describes a wider literature review, covering the national strategies of Cyprus, Greece, Malta, France, Estonia, the United Arab Emirates, Finland and the United Kingdom. The evidence base behind the Strategy is therefore broader than the education provisions suggest.

Proposal. Two changes. First, state the criteria on which international reference points for AI literacy have been selected, and state expressly whether the list is illustrative. Second, reframe the provision as open

²¹ Council of Europe, [Framework Convention on Artificial Intelligence and Human Rights, Democracy and the Rule of Law](#) (CETS No. 225), opened for signature on 5 September 2024.

and non-exhaustive, so that Cyprus is able to draw on relevant expertise, initiatives and good practice from a wider range of European and international jurisdictions as the field develops. Cooperation should be capable of connection to the broader bilateral and international relationships of the Republic in technology, innovation and AI development, including emerging and established technology ecosystems beyond Europe such as India, where these offer complementary expertise, alternative delivery models or opportunities for mutual capacity building. Alignment with Union law, values and standards should remain the governing constraint in all cases.

Why this matters. AI literacy is unusually broad and unusually fast-moving. It spans technical competency, ethical awareness, responsible use, critical thinking, digital inclusion, workforce preparedness and general societal understanding of AI. No single group of four or five jurisdictions is likely to hold the best available practice across all of these dimensions, and the distribution of that practice will change over the life of the Strategy. A closed list therefore risks fixing the Republic to a set of reference points that were sound at the date of drafting and will not remain so.

A broader formulation is also more consistent with the Strategy's own approach elsewhere. Objective 8 envisages partnerships with trusted public, private, academic and international actors, and places emphasis on co-development, knowledge transfer and long-term capability building. The education provisions should reflect the same posture. We would add that AI literacy is one of the few areas where Cyprus can realistically both learn and contribute, since a small and coherent education system can implement and evaluate a national literacy framework faster than a large one.

Proposal D7. Correct a small number of internal inconsistencies in the text

Sections affected: 3.3.1 (list of bodies and Figure 2), 3.3.1.4, 3.4.1, 3.9.6, 3.11.2, Annex A, Annex C, Annex F.

Current position. The following points are drafting matters rather than questions of policy. We raise them because they are inexpensive to correct and because a document of this standing will be read closely by suppliers, by other Member States and by the institutions of the Union.

- The body established under Section 3.3.1.4 is titled the National AI Infrastructure Committee in its own heading and the National AI Infrastructure Council in the Section 3.3.1 list, in Figure 2 and again in Annex A.
- Figure 2 presents the Government AI Accelerator and the Innovation Hub as a single body. The text treats the Government AI Accelerator as a component of the National AI Authority and the Government Innovation Hub as a separate innovation body.
- In Figure 2 the description attached to the AI-Industrial Centre of Excellence is the tracking of talent, skills gaps and workforce readiness, which is the function of the AI Skills Observatory and is stated for that body two positions earlier in the same figure.
- The names given to the centres of excellence in Figure 2 do not match those used in the Section 3.3.1 list.
- The Cyprus AI Security and Certification Authority is referred to in Section 3.9.6 and nowhere else, and is not defined. This is addressed substantively at Proposal B2.
- Industry adoption targets appear at different values and against different dates in Section 5 and in Annex C, and government adoption targets appear at different values in Section 3.4 and in Annex C. This is addressed substantively at Proposal D4.

- The implementation timeline in Section 1.5 combines relative time bands with a calendar band beginning in 2026, which overlaps them. This is addressed substantively at Proposal D4.

Proposal. Adopt a single name for each body and use it consistently in the headings, in the Section 3.3.1 list, in Figure 2 and in the annexes. Redraw Figure 2 so that it matches the text it illustrates and so that each body carries its own description. Reconcile the numerical targets as set out at Proposal D4.

Why this matters. Inconsistent naming of institutions is a small matter in a strategy document and a significant one in the instruments that follow it. Legislation, ministerial decisions, terms of reference and procurement documents will be drafted by reference to this text. Where two names exist for one body, both tend to survive into the implementing instruments, and the question of whether they denote the same body becomes a live one at precisely the moment when it is least convenient. The same applies to a figure that does not match the text, since readers reasonably assume that a diagram carries the authoritative version.

8. Conclusion and offer of contribution

The Cyprus National AI Strategy 2032 is a credible document with a defensible strategic logic. Our proposals do not seek to change its direction. They seek to give the citizen a defined place in it, to convert its ethical commitments into working machinery, to ensure that the obligations it places on private parties are lawful and defined, to make sovereignty something that can be planned and measured rather than asserted, and to ensure that the delivery model is proportionate to the country that has to deliver it.

We would highlight six proposals as the highest value relative to their cost. The Citizen AI Charter at A1 and the public register at A2 are the two most visible commitments the Strategy could make to the public. The clarification of authorities at B2 and the dated sandbox commitment at B3 are necessary for the trusted jurisdiction claim to survive scrutiny. The correction of the private-sector licensing provisions at B7 is, in our assessment, the most urgent single change in this submission, because those provisions carry legal consequences for third parties in their current form. The five-layer definition of sovereignty at C1 gives the rest of the document a coherent spine and would, in our assessment, be a genuine contribution to European practice rather than a restatement of it.

AI³ offers to contribute to implementation in the following ways, at no cost to the Republic.

- Comparative research on the design of the citizen register, the impact assessment methodology and the assurance capability, drawing on the European and international instruments referenced in this submission.
- Convening and facilitation of multi-stakeholder dialogue on implementation, including the participation mechanisms proposed at A6, through our network of researchers, policymakers and practitioners.
- Support for the international positioning proposed at D5 and D6, including regional engagement on AI governance and capacity building. AI³ is already developing cooperation and joint initiatives with organisations and experts in India and other jurisdictions, and would be available to contribute its network in facilitating exchanges or partnerships in support of the objectives of the Strategy.
- Independent review of draft implementation instruments, including the impact assessment template, the procurement clauses and the KPI framework.

We thank the National AI Taskforce, the Chief Scientist for Research, Innovation and Technology, and the Deputy Ministry of Research, Innovation and Digital Policy for a substantive and candid draft, and for a consultation process that has attracted serious engagement. We would welcome the opportunity to discuss any of these proposals in more detail.

Annex 1. Mapping of proposals to affected sections

The table below maps each proposal to the sections of the draft Strategy it affects and indicates the type of change required. Proposals marked as low cost require no new institution and no significant new expenditure, and could be adopted in the final text without further analysis.

Ref	Proposal	Sections of the draft Strategy affected	Type of change	Cost
A1	Citizen AI Charter with service commitments	1.1, 1.4.8, 2.5.7, 3.1.2, 3.4.3, 3.9.3 to 3.9.6, Annex A, Annex F	New annex to Strategy	Low cost
A2	Public register of AI systems used by public bodies	3.3.2, 3.4, 3.9.3, 3.9.6, Section 5 item 4, Annex B	New commitment	
A3	Human route, non-digital route and accessible design	2.5.4, 3.1.1, 3.4.1, 3.4.3, Annex A, Annex C	Amend text	Low cost
A4	Greek language as citizen right and sovereign capability	2.5.6, 3.2, 3.5, 3.7, 3.8.6, Annex C	New programme	
A5	Local government as a delivery partner	2.5.4, 3.3.1, 3.4, 3.4.4, Annex C	Amend text	Low cost
A6	Standing role for citizens and civil society	3.3.1, 3.3.1.3, 3.3.1.5, 3.3.1.6, 3.11.3, 4.1.6	New body, low cost	Low cost
B1	Reflect Regulation (EU) 2026/1744 and restate calendar	1.4.10, 2.3.3, 3.11.2, 3.12.1, Section 5, Annex A	Factual update	Low cost
B2	Clarify National AI Authority and designated authorities	3.3.1, 3.3.1.1, 3.9.6, Annex A	Clarify text	Low cost
B3	Dated, operational AI regulatory sandbox	3.8.5, Section 5 item 2, Annex A	Add commitment	
B4	One national impact assessment methodology	3.9, 3.3.2, 3.4.2, Annex B, Annex F	New instrument	
B5	Incident reporting, redress and public disclosure	3.3.2, 3.9.5, 3.11.3, 3.13, Annex B, Annex F	New mechanism	
B6	Procurement clauses, problem-first framing and access for Cypriot suppliers.	2.5.8, 3.1.2, 3.3.2, 3.4.1, 3.6.3, Annex B, Annex F	Amend and add	
B7	Private disclosure and operating licence conditionality	3.4.2, 3.5, Annex A, Annex C	Remove and define	Low cost
B8	Judicial AI terminology and AIREG scope	3.5, Annex C	Amend text	Low cost
C1	Five-layer definition of AI sovereignty	1.1, 2.4, 2.5.6, 2.5.8, 3.7.3, Annex G	Definition	Low cost
C2	Testable infrastructure sovereignty and concentration risk	3.7.1 to 3.7.5, 3.10.2, Annex F	Amend and measure	

Ref	Proposal	Sections of the draft Strategy affected	Type of change	Cost
C3	Data sequencing, access framework and EHDS alignment	2.5.6, 3.2, Annex C, Annex D	Amend and add	
C4	Add the missing model layer	2.5.8, 3.7, 3.8, 3.8.6, Annex D	New subsection	
C5	National AI assurance and conformity capability	2.5.1, 3.8, 4.3, Section 5 item 4, Annex C	New programme	
D1	Costed delivery plan with lifecycle costs	1.5, 3.10.3, Section 5, Annex F	Companion document	
D2	Right-size the institutional architecture	3.3.1 and subsections, 4.2, Section 5 items 1 and 3	Consolidate	Low cost
D3	Operational environmental sustainability requirements	3.4.1, 3.7, 3.10.1, Annex D, Annex F	Amend and measure	
D4	Dated milestones, target arithmetic, annual report to Parliament	1.5, 2.5.2, 3.4, 3.5, 3.11, Section 5, Annex C, Annex F	Correct and add	Low cost
D5	Strengthen the international position	2.3.4, 2.5.1, 3.12, Annex A	Amend and add	Low cost
D6	Open, non-exhaustive framing of AI literacy cooperation	2.5.8, 3.6, 3.12, Annex C	Amend text	Low cost
D7	Correct internal naming and figure inconsistencies	1.5, 3.3.1, 3.3.1.4, 3.9.6, Annex A, Annex C, Annex F	Drafting	Low cost

Annex 2. International reference points

The table below summarises the principal comparators referred to in this submission. They are included because they are operating instruments rather than proposals, and because each is proportionate to what a state of Cyprus's size can implement.

Instrument	Jurisdiction	Relevance
Algoritmeregister, national algorithm register	Netherlands	Operating since 2022 with approximately one thousand public sector entries. The data protection authority has called for registration to be made mandatory. Directly relevant to Proposal A2.
Algorithmic Transparency Recording Standard	United Kingdom	Mandatory for central government departments and for arm's length bodies delivering public services. Provides a tested schema and scope and exemptions policy. Relevant to Proposal A2.
Impact Assessment for Human Rights and Algorithms	Netherlands	Applied to impactful public sector systems. A worked methodology rather than a principle. Relevant to Proposal B4.
Algorithmic Impact Assessment and Directive on Automated Decision-Making	Canada	Mandatory assessment tool for federal automated decision systems, with published results and proportionate requirements. Relevant to Proposal B4.
ALIA and the Salamandra family of models	Spain	Publicly funded open models with a focus on Spanish and co-official languages, developed at the Barcelona Supercomputing Centre. Relevant to Proposals A4 and C4.
GPT-NL	Netherlands	Public interest language model developed by TNO, SURF and the Netherlands Forensic Institute, with attention to copyright and transparency. Relevant to Proposals A4 and C4.
OpenEuroLLM	European Union	Consortium developing model families covering official Union languages with access to EuroHPC resources. The appropriate participation route for a state of Cyprus's size. Relevant to Proposal C4.
Burokratt and the Kratt action plan	Estonia	A single modular public assistant used across many agencies, supported by an action plan with a stated investment package of approximately eighty-five million euro. Relevant to Proposals A3, D1 and D2.
Federal AI use case inventories, OMB Memorandum M-25-21	United States	Federal agencies compile and publish annual inventories of their AI use cases, with expanded disclosure for uses affecting rights or safety. A non-European precedent for the same transparency commitment. Relevant to Proposal A2.
Blueprint for an AI Bill of Rights	United States	A short public-facing statement of expectations covering notice, explanation, human alternatives and protection from discriminatory outcomes. Relevant to the format of Proposal A1.
AI Risk Management Framework	United States, NIST	A voluntary framework structuring the govern, map, measure and manage functions, widely used and increasingly referenced by European organisations. Relevant to Proposals B4 and B5.
Article 57 AI regulatory sandboxes	European Union	Each Member State must ensure at least one national sandbox is operational, with the option of a joint sandbox and priority access for small enterprises. Relevant to Proposal B3.

Instrument	Jurisdiction	Relevance
Regulation (EU) 2026/1744, Digital Omnibus on AI	European Union	Amends the AI Act and revises the high-risk compliance calendar. In force since 27 July 2026. Relevant to Proposal B1.
Energy Efficiency Directive Article 12 and Delegated Regulation (EU) 2024/1364	European Union	Energy and water reporting obligations for data centres above 500 kilowatts, with a Union rating scheme and labels in development. Relevant to Proposal D3.
Cloud Sovereignty Framework and SEAL assurance levels	European Union	Eight scored sovereignty objectives covering jurisdiction of decisive authority, exposure to extraterritorial law, key control and operational autonomy. Applied by the Commission in April 2026 to a sovereign cloud award of up to one hundred and eighty million euro. Relevant to Proposals C1 and C2.
Digital Operational Resilience Act third party risk regime	European Union	Established discipline for concentration and exit risk in the financial sector, capable of extension to government AI procurement. Relevant to Proposal C2.
Framework Convention on Artificial Intelligence	Council of Europe	The principal rights-based instrument through which non-Union states in the region engage with AI governance. Relevant to Proposal D5.

We would note one general observation from this comparison. The strategies that have delivered measurable results in small and medium states have concentrated on a small number of instruments applied consistently, rather than on comprehensive institutional design. Estonia has one virtual assistant platform used across many agencies, one action plan and one funding package. The Netherlands has one register and one impact assessment methodology. The United Kingdom has one transparency standard made mandatory. The common factor is not ambition. It is repetition of the same instrument across the whole of government.