Reading AI Readiness Backwards

What country experience reveals once AI adoption is already underway

June 2026

Copyright © UNDP 2026. All rights reserved.

One United Nations Plaza, NEW YORK, NY10017, USA

UNDP is the leading United Nations organization fighting to end the injustice of poverty, inequality, and climate change. Working with our broad network of experts and partners in 170 countries, we help nations to build integrated, lasting solutions for people and planet.

Learn more at undp.org or follow @UNDP @UNDPDigital

The views expressed in this publication are those of the author(s) and do not necessarily represent those of the United Nations, including UNDP, or the UN Member States. UNDP does not endorse any entity, brand, product, or service.

Acknowledgements

This report was prepared by the United Nations Development Programme (UNDP), under the strategic direction of Robert Opp, Chief Digital Officer, and Keyzom Ngodup Massally, Director, AI Hub for Sustainable Development and Head of Digital and AI Programmes.

The report was authored by Alexander Hradecky, Maria Paula Mujica Ramirez, McDonald Nyoni, and Helin Su Aslan. It benefited from substantive inputs and review from Jennifer Louie and Navya Alam. Editorial review and direction were provided by Dwayne Carruthers, Romilly Golding and Xiuzhen Li.

The assessments underpinning this report were made possible through the support of UNDP's AI Sprint, funded through contributions from UNDP's core donors, and through the AI Hub for Sustainable Development. These workstreams helped enable the delivery of AI Landscape Assessments at scale and supported the translation of country-level engagements into cross-country learning.

Above all, the report reflects the trust and partnership of the governments, institutions, civil society actors, private-sector representatives, researchers, and development partners who contributed to the 26 AI Landscape Assessments completed to date. Their willingness to share experiences, priorities, constraints, and lessons is what gives this report its grounding in the realities of AI adoption and implementation.

The insights gathered through these engagements are shared here as a contribution to broader learning on AI adoption, implementation, and governance in support of sustainable human development.

About this report

When UNDP began designing the Artificial Intelligence Landscape Assessment (AILA) in 2023, several AI readiness frameworks already existed. Several more have emerged since. The aim was not to add another diagnostic to a crowded field, but to fill a specific gap in how countries and partners understand AI adoption in practice. UNDP AILAs offer an assessment approach anchored in development outcomes, grounded in country context, and connected to the programmes, partnerships, and institutional processes already operating on the ground.

Many AI readiness frameworks assess countries from the outside in, measuring national conditions against global indicators and benchmarks. AILA was designed to complement that perspective by working from the inside out. It brings together quantitative indicators, stakeholder perspectives, desk research, and contextual analysis through a whole-of-society process, producing a practical evidence base for strategies, policies, investments, institutional capacity, and governance priorities.

What makes AILA distinct is also the institutional context in which it is delivered. Because UNDP works through long-term country programmes and partnerships with governments, civil society, the private sector, and development actors, assessment findings can be connected to policy processes, capacity-building efforts, investment planning, and reforms already underway.

This report draws on 26 AILAs completed to identify what these country engagements reveal when read together. Its purpose is to make those lessons visible, transferable, and actionable for governments, development partners, and ecosystem actors working to steer AI adoption towards sustainable human development.

Table of contents

Acknowledgements ii

About this report iii

Executive summary 1

1 Introduction 3

2 What becomes binding once AI adoption begins 5

2.1 Political attention is not the primary constraint 7

2.2 AI adoption often enters through systems, not strategies 9

2.3 Data determines whether adoption becomes implementation 12

2.4 Foundations determine agency, not just capacity 15

2.5 Ecosystem capacity determines whether foundations become local value 19

2.6 Governance becomes consequential through implementation 22

2.7 The implementation test 25

3 Decision points once adoption is underway 27

4 The work ahead 33

Annex I: The AILA evidence base and methodology 34

Further Information 40

0 Executive summary

Artificial Intelligence (AI) adoption is now outpacing AI readiness. Across many developing and emerging economy contexts, AI is entering public institutions, markets, digital systems, and everyday public services while the systems needed to coordinate, govern, evaluate, and sustain adoption are still being built. AI's development potential depends on whether countries can connect adoption to public priorities, institutional capability, local ecosystems, and accountable implementation.

Drawing on 26 UNDP Artificial Intelligence Landscape Assessments (AILA) completed between 2024 and 2026, with 10 more underway, this report examines what AI readiness means once adoption is already underway.

The countries analysed reflect the set of national AILA engagements completed at the time of writing, undertaken in response to government demand and UNDP country-level programming priorities across diverse regional, institutional, and digital maturity contexts. They are not presented as a statistically representative sample, and the report does not rank countries, aggregate readiness scores, or provide a country-by-country review (a full list of completed assessments is provided in Annex I).

Instead, the report uses this evidence base to identify recurring patterns in how AI adoption is moving through public systems, markets, institutions, and wider ecosystems.

This report argues that AI readiness needs to be understood differently once adoption is already underway. Once AI enters public systems through procurement processes, vendor platforms, infrastructure choices, sector programmes, digital public systems, and broader ecosystems, readiness is no longer only about preparation. It becomes a diagnostic tool: a way to identify which conditions become binding, where risks and dependencies are emerging, and what capabilities countries require to steer adoption towards public value and sustainable development.

Across the AILA evidence base, six connected patterns structure the report:

1. **Political attention is not the primary constraint**. In many contexts, AI is already visible as a national priority, but the harder task is building the mandates, financing, coordination, delivery capacity, and accountability needed to implement it responsibly.

2. **AI adoption often enters through systems, not strategies**. Consequential AI-related choices are being made through procurement, software, digital public infrastructure, vendor relationships, sector programmes, and private-sector activity, often before they are visible as formal AI policy decisions.
3. **Data determines whether adoption becomes implementation**. Administrative data, public-sector data, sectoral data, language data, and shared data infrastructure shape whether AI systems can be reliable, reusable, inclusive, monitored, and safely improved over time.
4. **Foundations determine agency**. The availability and affordability of infrastructure, compute, and connectivity, alongside talent, institutional capability, financing, and appropriate governance arrangements, shape whether countries can adapt AI to local needs, negotiate with providers, govern risks, and sustain adoption on terms aligned with national priorities.
5. **Ecosystems determine whether foundations become local value**. Public demand, enterprises, startups, universities, finance, civil society, standards, and procurement pathways need to connect if AI adoption is to build domestic capability, economic opportunity, stronger public systems, and local value creation rather than deepen reliance on externally developed technologies.
6. **Governance becomes consequential through implementation**. Responsible AI matters most when it reaches the operational settings where AI is procured, deployed, monitored, contested, corrected, and held accountable.

Together, these patterns provide a practical lens for assessing whether AI adoption is strengthening the conditions for public value, institutional capability, and sustainable development, or adding another layer of fragmented activity.

Across the contexts examined in this report, **the risk is not that countries are failing to adopt AI**. It is that AI enters public and economic life on terms shaped less by deliberate national choice than by vendor relationships, platform dependencies, imported systems, and short-term project incentives; governed in principle more than in practice; concentrated among those best positioned to benefit; and treated as valuable before there is evidence of who is served, what costs are incurred, and what risks emerge.

The window to influence these systems is open but not indefinitely. Procurement choices, infrastructure investments, decisions around data, vendor relationships, and governance standards being set now will shape the terms of AI adoption for years to come.

Reading AI Readiness Backwards shifts the question away from whether countries are prepared for AI, and towards whether they are building the conditions to steer adoption before early choices become embedded in enduring systems, contracts, institutions, and dependencies. That is where AI's development impact will be decided.

1 Introduction

Much of the global attention on AI has focused on advanced model development, access to compute, and the speed at which AI systems are improving. For many of the countries examined in this report, the AILA evidence points to a more immediate development challenge: how AI is already being adopted, integrated, and governed as it diffuses into public systems, markets, institutions, and everyday services. AI does not wait for national strategies, dedicated legislation, or neatly organized initiatives to be in place. Increasingly, it enters through digital transformation initiatives, vendor platforms, cloud services, informal use of publicly available AI tools, procurement decisions, enterprise software, sector-specific tools, and existing institutional workflows.

This makes AI adoption and diffusion a development challenge, not only a technology challenge. AI has significant potential to support sustainable development, improve people's lives, strengthen public systems, and expand economic opportunity. For countries working to reduce poverty and widen opportunity, the terms of AI adoption matter because they influence how services are delivered, how labour markets evolve, how the private sector competes, and how institutions manage risk. The choices being made now will influence whether AI expands access to services, strengthens local capability, and supports inclusive growth, or whether it concentrates opportunities, deepens dependency, and creates new forms of exclusion or harm.

The aim is not to define a single model of AI readiness or prescribe one pathway for responsible adoption. Countries differ in institutional capacity, digital maturity, infrastructure, market structure, data systems, governance arrangements, and development priorities. Rather, the question is whether AI adoption is becoming visible, deliberate, accountable, locally useful, and open to correction as it moves through real systems. In practice, responsible and sustainable AI adoption depends less on whether every condition is already in place, and more on whether countries are building the capabilities to shape adoption as it unfolds.

To better understand these choices, UNDP completed 26 AILAs in partnership with national counterparts across diverse country contexts between 2024 and 2026, with more than 10 additional assessments underway at the time of writing. AILA is a whole-

of-society diagnostic that examines how AI is entering public systems, markets, and institutions, and the institutional, technical, ecosystem, and governance conditions shaping whether adoption can advance national priorities and sustainable development. Through stakeholder consultations, workshops, surveys, desk research, external indicators, and comparative analysis, it brings together perspectives from government, private sector, academia, civil society, and development partners.

AILA was designed to assess readiness, but the accumulated evidence shows why readiness cannot be understood only as a set of conditions that exist before AI adoption begins. This evolution is reflected in the methodology itself. What began as an AI Readiness Assessment has developed into the Artificial Intelligence Landscape Assessment, recognizing that readiness cannot be understood only as a pre-adoption score, but must be read through the actors, systems, constraints, and choices already shaping AI adoption.

Across country engagements, AILA helps surface where consequential decisions about AI are being made, where trade-offs are emerging, where risks, harms, accountability gaps, and inclusion concerns are being debated or diagnosed, and which conditions are most likely to shape outcomes. This report draws on the primary data, stakeholder insights, and context-specific analysis generated through completed AILA engagements to identify cross-country patterns in how AI adoption is taking shape. It draws on a body of completed AILA engagements, while using selected country examples to illustrate specific patterns. The examples are not intended to represent every completed assessment, nor should the absence of a country example be read as a judgement on that country's relevance or importance to the analysis. A list of completed AILA engagements and a summary of the methodology are provided in Annex I.

This report argues that AI readiness should be read 'backwards'. Once AI is entering public systems through procurement processes, vendor platforms, infrastructure choices, sector programmes, digital public systems, and broader ecosystems, readiness is no longer only a measure of what exists before adoption begins. It becomes a way to identify which conditions become binding as adoption moves through real systems, where risks and dependencies are emerging, and what integrated capabilities are needed to steer AI towards public value and sustainable development.

The analysis follows six connected patterns: political attention, systems of adoption, data, foundations, ecosystems, and governance. Together, these patterns show that readiness depends not on whether each condition exists in isolation, but on whether countries can connect institutions, data, infrastructure, talent, ecosystems, and governance quickly enough to shape adoption while there is still room to influence its terms.

The risk is not that countries fail to adopt AI. In many of the contexts examined here, adoption is already underway. The deeper risk is that AI enters public and economic life on terms shaped less by deliberate national choice than by vendor relationships, platform dependencies, imported systems, and short-term project incentives. *Reading AI Readiness Backwards* shifts the question away from whether countries are prepared for AI in the abstract, and towards whether they are building the conditions to steer adoption before early choices become embedded in systems, contracts, institutions, and dependencies.

2 What becomes binding once AI adoption begins

The AILA evidence suggests that the concept of 'AI readiness' changes once AI adoption is already underway. As AI becomes embedded in the systems through which decisions are made, services are delivered, economies operate, and public institutions manage risk, conditions that may appear abstract in a readiness framework become operational questions. Data, infrastructure, compute, talent, ecosystems, and governance are no longer only indicators of preparedness. They become the foundations that determine whether adoption can be adapted to local needs, governed responsibly, trusted by users, and sustained over time.

Across the AILA evidence base, this shift appears in highly varied contexts. In some countries, AI adoption is visible through national strategies, public-sector use cases, data centre investments, startup activity, or governance discussions. In others, it advances through less visible but consequential channels such as vendor platforms, enterprise software, donor-supported programmes, broader digital transformation efforts, and informal use of publicly available AI tools by civil servants and institutional staff. Across these pathways, readiness depends on whether countries can identify where adoption is taking hold, understand who is shaping it, and build the authority, safeguards, and capabilities needed to influence its direction before decisions become difficult to reverse.

Across the AILA evidence base, six recurring patterns help explain what becomes binding once adoption begins. In this report, 'binding' refers to the conditions that begin to determine what countries can realistically steer, govern, adapt, and sustain as AI moves through real systems. These conditions are not new. But once adoption is underway, they become operationally consequential: they shape whether AI can be implemented responsibly, corrected when risks emerge, and translated into public value.

Pattern	What becomes binding once AI adoption begins
1. **Political attention is not the primary constraint**	Operating authority: mandates, financing, delivery capacity, coordination, and accountability determine whether political attention translates into implementation.
2. **AI enters through systems, not strategies**	Adoption channels: procurement, software, vendors, donor programmes, digital public infrastructure, institutional workflows, and informal use shape adoption before it is always visible as formal AI policy or governance.
3. **Data determines whether adoption becomes implementation**	Data conditions: quality, interoperability, governance, discoverability, and responsible reuse determine whether AI can become reliable, inclusive, monitored, reusable, and sustainable in real systems.
4. **Foundations determine agency**	Foundational capacity: compute access, infrastructure, connectivity, talent, financing, and institutional capability determine whether countries can adapt AI, negotiate with providers, govern risks, and sustain adoption on terms aligned with national priorities.
5. **Ecosystems determine local value**	Ecosystem linkages: public demand, enterprises, startups, universities, finance, civil society, standards, and procurement pathways determine whether adoption builds domestic capability and stronger public systems, or increases reliance on externally developed technologies.
6. **Governance becomes consequential through implementation**	Governance mechanisms: procurement rules, vendor management, monitoring, measurement, redress, oversight, and accountability determine whether responsible AI principles meaningfully shape deployment.

Together, these patterns show that readiness is tested through implementation. Once AI adoption begins, the question is how quickly countries can connect political authority, delivery systems, data, foundations, ecosystems, and governance to shape the choices already being made.

2.1

Political attention is not the primary constraint

Across the assessments reviewed, the primary constraint is rarely a lack of attention to AI among senior decision-makers. Countries included in the AILA evidence base are already asking questions about AI readiness, adoption, and governance.

However, even where AI has political visibility, strategic interest, or national momentum, the challenge is to direct that attention into the institutions, financing, delivery capacity, and accountability needed to implement AI responsibly. This means identifying useful applications, improving public services, strengthening local capability, and managing the risks that emerge as AI enters public and economic systems.

Country evidence from the 26 assessments shows that AI is already being linked to modernization, competitiveness, public-sector reform, digital transformation, and long-term development planning.



Viet Nam is incorporating AI into national strategies on research, development, application, digital transformation, socioeconomic development, and national competitiveness, while also identifying constraints around regulation, funding, data, infrastructure, and workforce capacity.



Malawi shows high-level political recognition of digital transformation and AI, while stressing that political will needs to be channelled into strategy, delivery capacity, data governance, procurement reform, and institutional coordination.



The same dynamic is visible in **Ethiopia**, where strong political commitment and state leadership are helping advance AI as a national priority, supported by continued investments to enhance compute capacity, interoperable data systems, applied skills, and trust frameworks needed for long-term implementation.



Trinidad and Tobago illustrates a similar dynamic, with strong commitment to AI accompanied by priorities around shared infrastructure, trust and safety, public-sector capability, talent retention, private sector co-investment, and public awareness.

The central insight is not that political attention is superficial. It is that attention does not automatically create the conditions for responsible and sustainable implementation. Strategies may be developed, consultations convened, investments considered, use cases tested, and institutions asked to define their roles while the underlying practical delivery questions remain unresolved. Who leads AI deployment across government ministries? How are priorities financed? How are systems procured? Which institutions supervise deployment? Where does accountability sit when systems fail? How are benefits, costs, and risks distributed? Whether countries can answer these questions, and organize institutions, mandates, and resources around them determines whether attention translates into delivery capacity.

Political attention can also create new forms of pressure. When AI becomes a high-priority agenda, multiple actors can begin moving at once, including ministries, regulators, digital agencies, vendors, startups, universities, enterprises integrating AI into existing operations, and development partners supporting country-level initiatives. This activity is often productive, helping surface practical use cases, mobilize partners, test new approaches, and build confidence around AI adoption. But it typically moves faster than the structures and mechanisms needed to keep pace with it. A government may have strong AI narratives, emerging use cases, active international partnerships, or policy momentum while still lacking a lead coordination mechanism, operational roadmap, implementation financing, AI-aware procurement guidance, or clear oversight structures. The result is dispersed action without the requisite institutional arrangements to turn activity into implementation aligned with public priorities, safeguards, and long-term capability.

Fragmentation is not solely a result of weak capacity. It is also due to competing priorities and incentives that often emerge with a fast-moving AI agenda. Ministries want to push through their own mandates. Vendors promote their own proprietary solutions. Businesses adopt tools in response to competitive pressure or operational needs. Development partners support initiatives based on their own funding cycles and thematic priorities. Political pressure often rewards visible progress before coordination mechanisms are fully in place. The task is not to regulate every initiative. It is to identify AI activity early enough, and with enough information about actors, funding, vendors, data use, safeguards, and accountability, to coordinate, finance, govern, and learn from it.

The issue is not whether AI has political visibility. It is whether that visibility is translated into the public interest functions needed to shape adoption responsibly. For governments, this means clarifying who has the mandate to coordinate AI deployment across institutions, how priorities are financed, how procurement of AI systems and tools is guided, how risks are supervised, and how lessons from implementation are fed back into policy. This does not require every AI initiative to be managed by a single ministry or coordination body. It means ensuring that distributed activity is connected to clear priorities, safeguards, accountability, and learning. For development partners and ecosystem actors, this means strengthening the government capabilities needed to coordinate, procure, finance, supervise, and learn from AI adoption, rather than funding parallel pilots or fragmented initiatives. Without this, political attention can generate more AI activity, but that activity may remain scattered, difficult to govern, and weakly connected to public priorities.

2.2

AI adoption often enters through systems, not strategies

The implementation challenge described above is compounded by the way AI adoption takes place across the AILA evidence base. Across the 26 assessments, political attention to AI is increasingly visible, but adoption rarely moves through a single national pathway. While strategies and dedicated AI institutions remain important, many of the decisions that shape adoption are made through the systems that countries already use to deliver services, modernize government, support innovation, and organize economic activity, including procurement processes, digital public infrastructure, service portals, data exchange systems, business registries, payment systems, sector platforms, and enterprise software.

The country evidence from the AILA findings shows that strategies and operational adoption often move in tandem.



In **Uzbekistan**, AI is anchored in Strategy for the Development of Artificial Intelligence Technologies 2030 and Digital Uzbekistan 2030, while implementation is also advancing through sectoral priorities, e-government services, pilots in diagnostics and virtual assistants, and targets to expand AI-based services through the Unified Interactive Public Services Portal.



Bhutan illustrates the same point through the digital public infrastructure (DPI) on which AI adoption is expected to build, including government to citizen (G2C) services in Bhutan as well as payment gateways, government data infrastructure, national data exchange, and digital identity.



Mongolia shows that AI technologies have already entered some public services even as institutional alignment, inter-agency working arrangements, strategic direction, and oversight remain works in progress.



Montenegro similarly links AI adoption to public service integration, government digital infrastructure, open data, public–private partnerships, and potential use cases in health, education, environmental management, tourism, and administrative processes.

Across the AILA evidence base, many consequential decisions shaping AI adoption are made through delivery, procurement, infrastructure, software, and partnership channels before they are recognized as AI-related decisions.

A ministry may procure a decision-support tool as part of service reform. A public agency may introduce analytics or automation through a digital government programme. A sector ministry may adopt AI-enabled tools through a donor-supported initiative. A company may integrate predictive capabilities through enterprise software or cloud services. A startup or university partnership may generate use cases before national priorities, standards, or oversight arrangements are fully defined. These pathways are often most visible in key sectors such as health, agriculture, education, climate, social protection, justice, and public administration, where AI use cases first become concrete but wider implementation depends on the systems around them.

A further complication is that AI is not always classified institutionally, or governed operationally, as artificial intelligence. In some contexts, AI-enabled functionality appears through analytics tools, automation, cloud services, enterprise software, or digitally enhanced workflows that may be understood as part of broader digital transformation rather than as AI adoption. This matters because governance and accountability questions can emerge even where institutions do not yet perceive themselves as making deliberate AI deployment decisions.

This is particularly important for DPI and the wider public digital systems through which services are delivered. **AI is increasingly layered onto systems that shape how people are identified, records are maintained, payments are made, services are accessed and how institutions exchange data**. Digital identity, payment systems, civil, business, land and social registries, data exchange platforms, service portals, and administrative systems for processing applications, claims, or service requests all shape the conditions under which AI can be used. **Where these systems are interoperable, trusted, and well governed, they can provide a stronger basis for responsible AI adoption. Where they are poorly connected, exclusionary, or weakly governed, AI can reproduce or amplify existing weaknesses in service delivery, data sharing, and accountability**, including by making services harder to access, errors harder to detect, or decisions harder to contest.

These channels of adoption can be productive. They can accelerate learning, support locally relevant applications, and bring AI into practical use. The risk emerges when adoption remains disconnected from mechanisms that make it visible, accountable, reusable, and open to correction. AI adoption is rarely one decision; it is the accumulation of decisions about procurement, infrastructure, software, data use, DPI, partnerships, safeguards, service design, and oversight.

The practical implication is that strategy, governance, and support need to extend to the places where AI is entering real-world use. For governments, this means looking beyond national AI policy to the procurement processes, service delivery systems, public digital infrastructure, data exchanges, and vendor relationships through which adoption is already being shaped. For development partners and ecosystem actors, it means helping institutions work with those channels directly: making AI use visible, building safeguards into adoption pathways, carrying lessons from one sector or country to another, and avoiding or redesigning applications that lack clear ownership, reliable data, monitoring arrangements, maintenance capacity, or pathways for people to contest and correct errors. Without this, countries may have a strategy for AI while the real-world places where AI is being adopted remain difficult to monitor, evaluate, improve, or translate into public benefit.



2.3

Data determines whether adoption becomes implementation

The first two insights bring data into focus as implementation infrastructure. As political attention grows and AI adoption moves through multiple pathways, the quality, accessibility, and governance of data increasingly determine whether adoption can become usable in real systems. Strategies, procurement decisions, partnerships, and use cases can create momentum. But AI systems depend on the data environment around them. Where data is fragmented, inaccessible, poor quality, non-interoperable, or weakly governed, adoption is less likely to become reliable, inclusive, safe, reusable, or sustainable.

AI systems depend on data at multiple stages of adoption. Data is required to identify priority use cases, train or adapt models, evaluate performance, monitor outcomes, detect errors and continuously improve systems over time. It shapes whether systems can be audited and monitored once deployed. Weak data foundations constrain model development, deployment, oversight and long-term sustainability of governance mechanisms.

However, the challenge is not exclusively about the availability of data. AI adoption depends on whether institutional conditions allow data to be created, accessed, shared, protected, and used across organizations and over time. **Questions of stewardship, interoperability, coordination, and institutional capability often become just as important as the underlying datasets themselves**.

AILA evidence shows this across countries with different levels of digital and AI maturity.



Uzbekistan illustrates the challenge in a context with important data and digital government assets already in place. The AILA for Uzbekistan notes strong scores on interoperability and data quality, the role of state-linked IT company UZINFOCOM, as a central integrator of state IT systems, together with efforts to consolidate state databases and develop Uzbek-language data resources. The next challenge for Uzbekistan is consistency across institutions. The AILA identifies fragmented ministerial and agency data, heterogeneous formats, differing levels of awareness of data-quality standards, gaps in dataset discovery, and the lack of national data lakes or unified quality protocols. The implementation concern is therefore not the absence of data foundations, but whether they are sufficiently connected, discoverable, and governed to support AI use across institutions.



Burundi shows how data can become an underused asset when institutional arrangements for access, quality, sharing, and protection are still developing. Its AILA finds that institutions are producing growing volumes of data across key sectors, but these data remain largely underused. Public data are described as difficult to access or not yet meeting the standards required for advanced analytical use. The assessment also points to the maturity of data governance, sharing and interoperability mechanisms, and concerns around data quality, security, and protection as areas that need strengthening. Framed this way, the data challenge is also an opportunity. Better stewardship, clearer sharing mechanisms, and stronger protection can help turn existing public data into a basis for experimentation, innovation, and more reliable AI adoption.



Côte d'Ivoire and Guinea extend this to public sector capability and governance. Côte d'Ivoire identifies data governance, interoperability, data management capacity, shared platforms, and stronger inter-institutional coordination as priorities for AI adoption in government. Guinea points to the need to strengthen technological maturity and analytical capacity in public institutions, alongside stronger data sharing, institutional coordination, and AI governance covering data protection, cybersecurity, and algorithmic responsibility. In both cases, the data agenda is inseparable from institutional capability. AI adoption depends not only on the availability of data, but on whether public institutions can manage, share, protect, and use it responsibly.

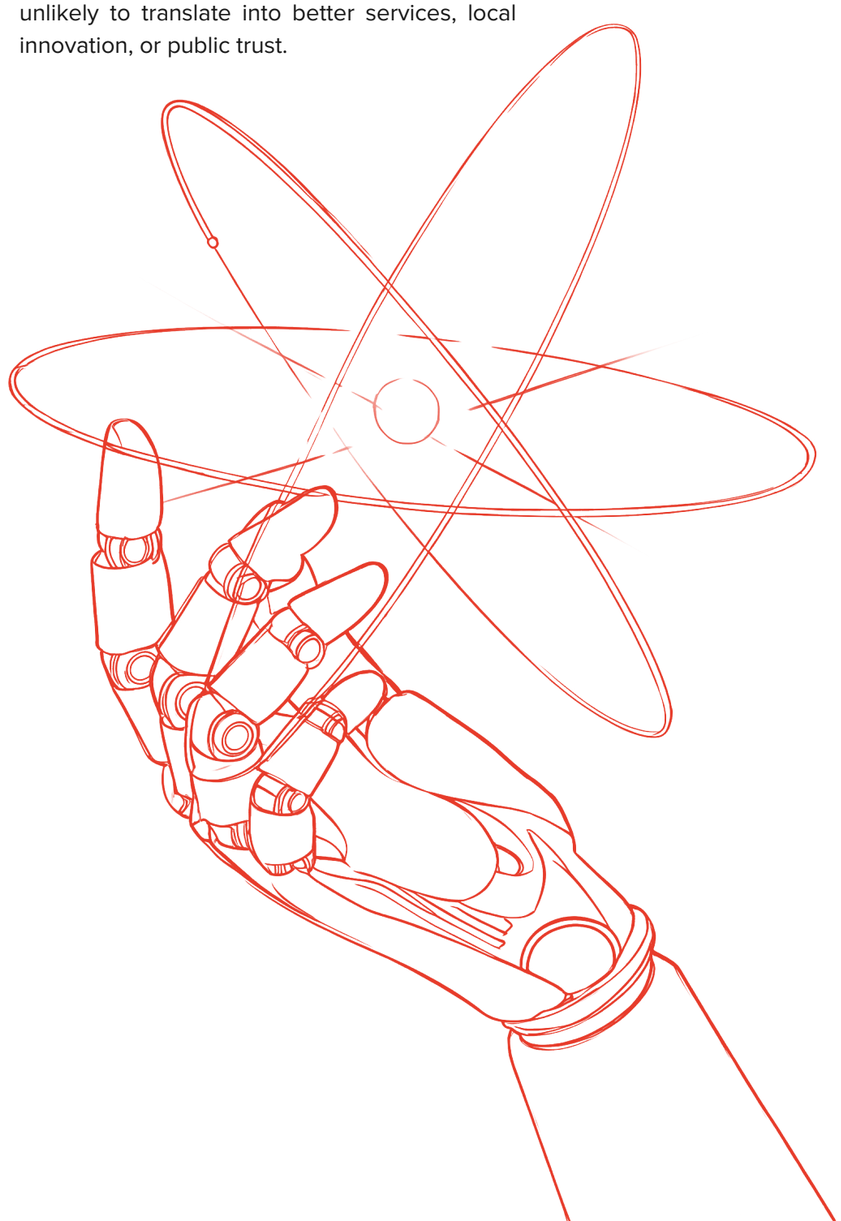


Guatemala shows how DPI can become the bridge between data readiness and responsible AI adoption at scale. Its AILA states that AI readiness depends on the strength of digital public infrastructure and digital maturity, while identifying data interoperability, digital identity, and digital payments as important conditions for responsible AI adoption at scale. It notes that data are often stored in institutional silos, quality is inconsistent, and centralized catalogues and exchange standards are lacking. In the public sector, the assessment similarly points to fragmented and inefficient data exchange between entities, difficulty identifying available datasets, and the absence of interoperability and data-quality standards. These are often described as technical reforms; in practice, they shape whether AI systems can be trained, adapted, monitored, corrected, and reused responsibly across institutions and sectors.

Data is where the opportunities and consequences of AI adoption become concrete. Data affects whether AI systems reflect the people, languages, places, and realities they are meant to support. It determines whether institutions can share information responsibly, coordinate services across agencies, enable local enterprises and researchers to build on public data, and monitor and correct AI-enabled tools over time. Weak data foundations can make systems unreliable, exclusionary, or unsafe. Weak governance can undermine privacy, trust, and accountability.

For governments and development partners, the implication is broader than improving technical data systems. Data readiness should be treated as part of the institutional infrastructure of implementation. This could mean investing in data quality, interoperability, metadata, catalogues, access protocols, privacy safeguards, stewardship models, and institutional data culture. However, it also means recognizing that governance objectives and institutional capability do not evolve at the same pace.

Countries may establish well-defined AI strategies and data governance laws while continuing to face foundational challenges around data quality, stewardship, coordination, or data reuse. Data reform cannot be left to technical teams, because decisions about access, standards, stewardship, privacy, and reuse shape how services are delivered and who benefits from them. Without this layer, political attention, AI strategies, infrastructure investments, and technology partnerships are unlikely to translate into better services, local innovation, or public trust.



2.4

Foundations determine agency, not just capacity

Data is not the only foundation that becomes consequential once AI adoption begins. In this report, foundations refer to the assets, capabilities, and institutional arrangements that shape whether countries can adopt, adapt, govern, and sustain AI on terms aligned with national priorities. Across the AILA evidence, these include infrastructure, compute, connectivity, financing, talent, institutional capability, and shared governance tools that can be reused across sectors and use cases. This is not an exhaustive list, but it captures a recurring pattern across assessments. These factors are often treated as technical prerequisites or general capacity gaps. In practice, they shape whether countries can adapt AI systems to local needs, negotiate effectively with providers, build and retain local capability, manage risks, and sustain implementation over time.

Foundations do not sit before implementation; they shape the terms on which implementation becomes possible. Once AI begins entering public systems, markets, and service delivery, infrastructure choices shape the range of options available to governments, local innovators, public institutions, and users. Connectivity, compute, cloud infrastructure, data centres, energy reliability, and secure digital systems influence whether AI can be tested locally, deployed reliably, adapted to national priorities, and scaled beyond isolated institutional use. **Where foundational choices are constrained by limited compute access, unreliable connectivity, fragmented digital systems, or dependence on a small number of providers, adoption may still occur, but countries have less leverage over how AI is tested, procured, adapted, and governed.** As compute and data infrastructure expand, these choices also need to account for energy demand, environmental sustainability, and alignment with climate commitments.

AILA evidence shows that room to act depends on both the assets countries can build on and the foundational systems that shape how far AI can be tested, adapted, and scaled.



Bhutan illustrates how strong connectivity, government digital infrastructure, national data exchange, and digital identity can provide a basis for public-sector AI adoption, while further investment in computing capacity and network resilience can expand the country's ability to test and apply more compute-intensive AI systems.



Malawi shows why foundational systems such as connectivity, reliable electricity, data availability and quality, interoperable systems, and trusted digital identity remain central if AI adoption is to move beyond isolated or small-scale use.



Viet Nam points to the strategic importance of high-performance computing, advanced AI technologies, digital infrastructure for large-scale deployment, data quality, workforce capacity, and access to advanced servers and AI chips.



Côte d'Ivoire extends the point to public-sector capability, identifying limited technical skills, gaps in structured resources, and the need for better access to shared platforms, data infrastructure, and high-performance computing. Across these examples, foundations matter because they expand or constrain the range of choices available to governments, institutions, and local innovators.



Costa Rica adds a further nuance. Its AILA points to strong general connectivity and favourable conditions for sustainable digital infrastructure, including extensive 4G coverage, high household internet access, and a largely renewable electricity matrix. At the same time, the assessment finds that high-performance computing, AI-ready data centre services, and geographically accessible compute resources remain limited. This shows that although a country's general digital strength can create a platform for AI, constraints can remain around the specialized infrastructure needed to test, deploy, and scale more compute-intensive AI systems.

The capability agenda extends beyond the supply of technical talent. Engineers, data scientists, and AI specialists remain essential, but harnessing their true value depends on the quality of the AI-ready institutional frameworks and environment around them. **Adoption requires teams that can define meaningful problems, understand service contexts, commission or build appropriate systems, manage vendors, evaluate performance, interpret outputs, protect rights, and integrate AI into workflows**. Some of these capabilities are technical; others sit in procurement, policy, law, service design, management, and frontline delivery. The implementation question is whether countries can bring these capabilities together in the institutions where AI systems will be used.

Several assessments point to this broader capability agenda:



Rwanda shows how investment in technical capability, AI training, university programmes, and specialized talent can strengthen participation in AI adoption, while also highlighting that non-technical competencies such as project management and ethical governance need to be scaled if pilots are to translate into repeatable institutional capability.



Malawi links skills, institutions, market dynamics, digital public infrastructure, public sector capability, startup development, access to compute, public sector data, and safe testing environments as interconnected parts of the AI adoption challenge. Guinea similarly links human capital, sustainable digital skills, local innovation, quality jobs, and reduced technological dependency.

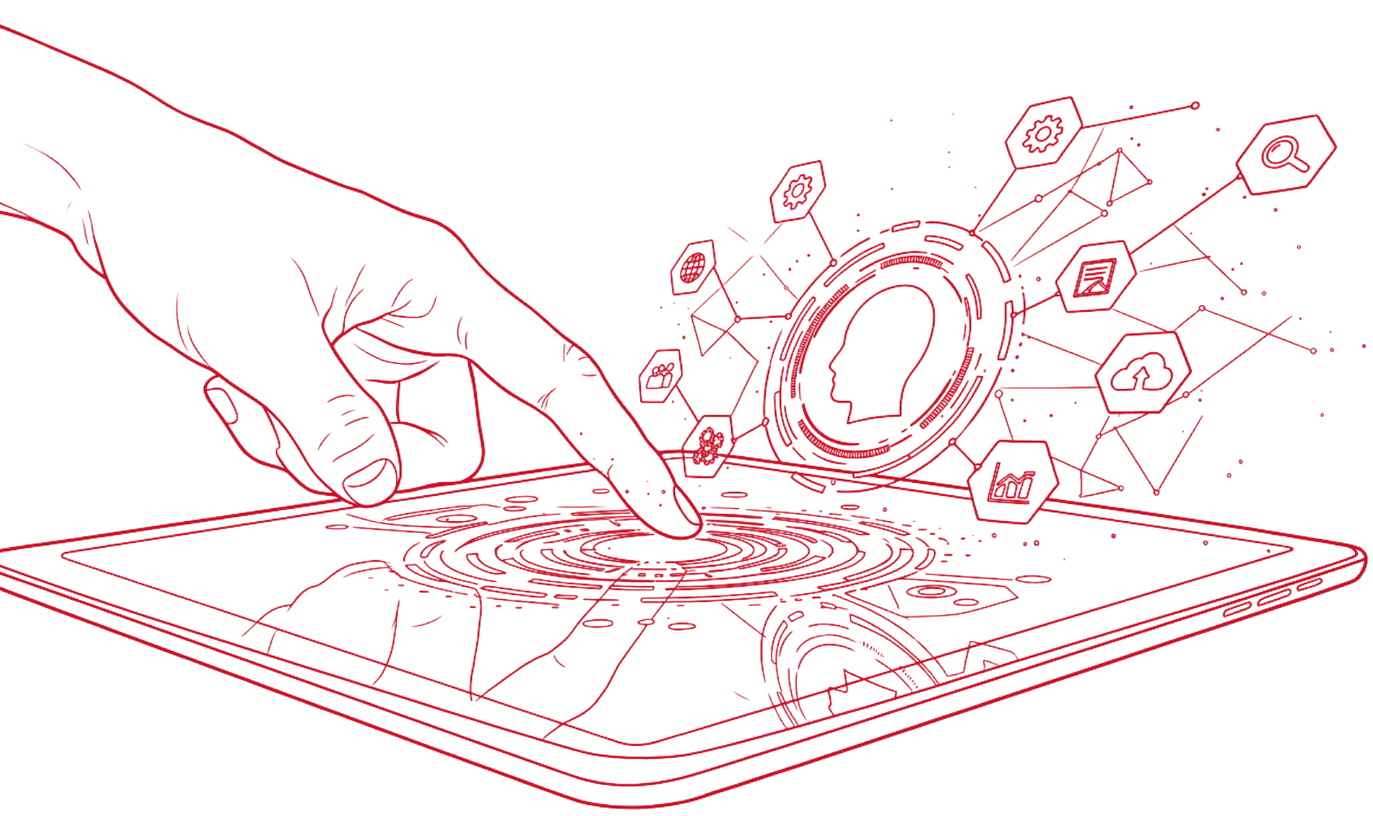
Foundations also include cross-cutting governance capabilities that can be reused across sectors and use cases. As AI adoption expands, countries cannot design procurement rules, evaluation approaches, monitoring systems, escalation pathways, incident response procedures, and redress mechanisms separately for every application. These shared capabilities sit alongside DPI and help determine whether AI systems can be assessed, corrected, contested, and governed. They are especially important where institutional capacity is constrained, because they reduce duplication, create more consistent expectations across vendors and sectors, and help institutions respond when systems fail.

This is why foundations should be judged by the agency they create. Foundations determine whether governments and national ecosystems can choose which systems to adopt, how they adapt them to local needs, whether they can procure and supervise them effectively, build local capability around them, and ensure they serve their national development priorities. AI can still be adopted where these foundations are limited, but the terms of adoption are more likely to be shaped by external providers, short-term projects, or disconnected institutional decisions.

Where foundations are stronger, more interoperable, and connected to public demand and local capability, countries have more agency to build markets, retain talent, support innovation, govern risks, and align AI adoption with national priorities. Agency is the difference between having access to AI and having the capacity to choose how it is used.

Agency should not be confused with full technological self-sufficiency. For many countries, the goal is to build enough domestic capability, trusted partnerships, interoperability, and shared governance tools to make deliberate choices within an interdependent global AI ecosystem. That may mean building strengths in data, applications, governance, talent, procurement standards, or evaluation capacity, while securing access to compute, cloud services, and frontier capabilities on terms that preserve flexibility and public benefit. It may also require shared or regional approaches to compute, assurance, vendor intelligence, procurement standards, language resources, and regulatory learning, especially where countries face similar constraints and can benefit from pooled assets and South–South exchange.

For governments, development partners, and ecosystem actors, the implication is that foundational investment should not be treated as progress by default. **A data centre, cloud partnership, training programme, procurement reform, or AI sandbox should be assessed on whether it expands the ability to choose, adapt, supervise, contest, and sustain AI systems over time**. The most useful foundations are those that can be reused across sectors, reduce avoidable dependency, support local capability, and make adoption safer and more governable as it scales.



2.5

Ecosystem capacity determines whether foundations become local value

Foundational investments create value only when they connect to demand, financing, research, enterprise development, safeguards, and delivery pathways. Infrastructure, compute, data systems, talent, and institutional capability matter most when there are actors and incentives able to turn these assets into useful applications, local capability, better services, and economic opportunity. In this report, the AI ecosystem refers less to a map of stakeholders than to the relationships that determine whether public institutions, enterprises, startups, universities, finance, civil society, technology providers, and development partners can organize around concrete adoption goals rather than isolated projects.

Governments remain central. They set priorities, regulate markets, procure systems, convene actors, deliver services, and often create the first significant demand for AI-enabled solutions. In many contexts, governments are also major economic actors, data holders, and service providers. This means public demand can shape the market for AI, but only after institutions have defined priority problems, opened appropriate procurement pathways, created safe testing environments, and set standards that local actors can respond to. At the same time, enterprises, researchers, communities, and civil society can surface use cases, constraints, and risks that governments may not see or be privy to. Ecosystem capacity depends on whether these signals meet in ways that produce coherent implementation strategies rather than scattered, piecemeal activity.



The **Montenegro** AILA connects AI readiness to innovation policy, regulatory modernization, digital business models, local competitiveness, public procurement, funding for innovation, and public–private collaboration. The recommendations include public–private partnerships, a national open data platform, support for AI-focused startups and small- to medium-sized enterprises (SMEs), expanded connectivity, and AI research and innovation centres, alongside measures to help domestic enterprises participate in public procurement and large-scale projects.



AILA results from **Malawi** make the market–system challenge especially explicit. The assessment describes skills, institutions, and market dynamics as forming the bedrock of a functioning AI ecosystem. It identifies positives, such as emerging technical talent, university programmes, and innovation nodes, while also highlighting gaps in faculty capacity, compute access, research commercialization, cross-sector co-development, financing, public sector data access, procurement pathways for startups, and safe testing environments. Rather than focusing on individual skills or isolated use cases, the focus is on strengthening the system that connects research, training, finance, public demand, and adoption.





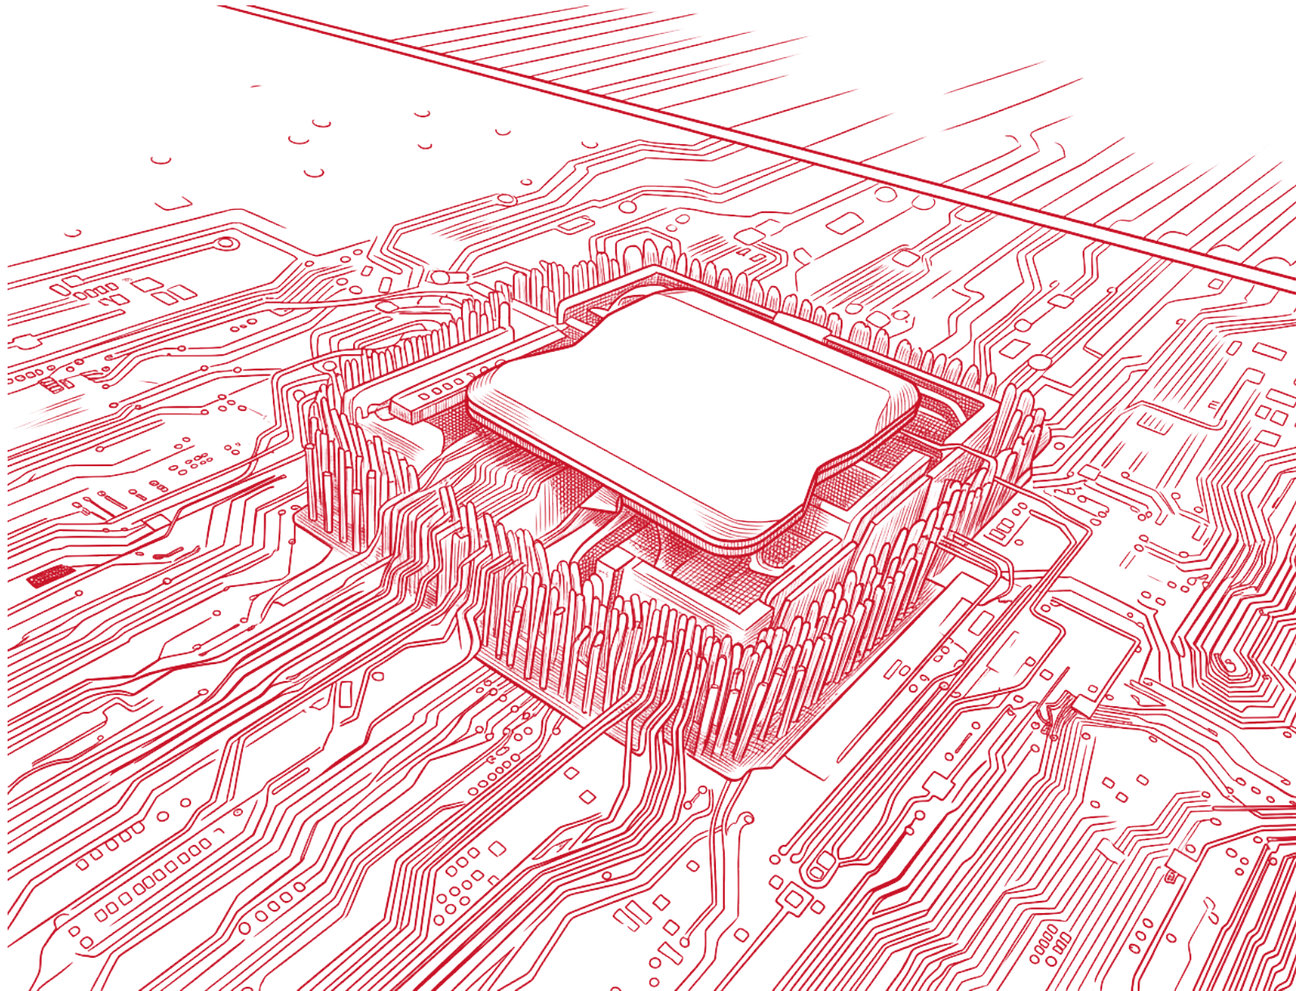
AILA evidence from **Colombia** and **Paraguay** adds a broader regional perspective. Colombia identifies the need to strengthen private sector capacity for AI innovation, improve access to financing, and better articulate policies that can mature the innovation ecosystem and support AI-enabled economic and social development. Paraguay adds a participation and capability perspective, highlighting training on AI adoption, rights, and ethical use across the public sector, private sector, academia, and wider society, with roles for universities, industry associations, software and scientific communities, inclusion-focused organizations, and public institutions.



The **Dominican Republic** adds the implementation link between ecosystem activity and local value creation. Its AILA identifies opportunities to strengthen incentives, expand access to financing, enhance coordination among ecosystem actors, and deepen the links between training, research, entrepreneurship, and productive sector needs to accelerate applied AI innovation. The proposed pathway for the Dominican Republic focuses on connecting talent, financing, entrepreneurship, research, and public and private demand, so that AI adoption can generate economic value, new ventures, and more sustainable productive capabilities rather than a set of isolated initiatives.

Ecosystem capacity is whether each wave of AI adoption builds more local capability over time, or drains it. A country may have a strategy, improving infrastructure, emerging talent, or visible use cases, but the value of these investments is limited if those efforts do not create demand for local problem-solving, routes for enterprises to grow, stronger university and research links, more capable public institutions, and feedback from users and affected communities. The question is whether each wave of adoption leaves the ecosystem better able to solve the next problem, rather than dependent on a new project, platform, or vendor each time.

In many contexts, strengthening ecosystem capacity will create more space for enterprises, startups, universities, and researchers to respond to public and market needs. This does not mean the private sector should lead by default, or that government capacity matters less. Government capacity and ecosystem capacity are mutually dependent. Governments need the capability to set priorities, procure intelligently, regulate markets, and protect the public interest. Enterprises, universities, researchers, and civil society need real problems to work on, fair access to opportunities, usable data where appropriate, and standards that make participation possible. Development partners can strengthen this cycle by helping public institutions define demand clearly, opening space for local actors to respond, supporting finance and procurement models that do not exclude smaller players, and ensuring that safeguards are built into adoption from the start.



2.6

Governance becomes consequential through implementation

The previous sections show that AI adoption is shaped by political attention, delivery systems, data foundations, infrastructure, talent, and ecosystems. Governance is what gives those conditions direction, accountability, and limits. High-level principles and commitments do not govern systems on their own. They become consequential when they shape how AI systems are selected, procured, deployed, made operational, monitored, contested, corrected, and improved over time.

As AI becomes embedded in public services, enterprise software, digital public infrastructure, and sector systems, governance depends on more than regulation or strategy. It depends on the operating environment around the technology: data practices, procurement arrangements, institutional responsibilities, oversight mechanisms, human review, escalation pathways, and the ability of affected people and institutions to detect, contest, correct failures, and improve the overall performance over time.

Across the AILA evidence base, references to responsible AI are increasingly present in strategies, consultations, stakeholder concerns, and policy discussions. This language often does more than signal alignment with global norms. It is a way countries articulate concerns about power, agency, dignity, sovereignty, accountability, and public trust in a context where many AI systems are externally developed, vendor-mediated, or shaped by standards set elsewhere. The implementation challenge is translating those concerns into shared obligations that governments, vendors, developers, and partners can act on. Governance therefore needs to reach the operational settings where AI adoption is shaped, including procurement requirements, vendor expectations, data governance, oversight, redress, and lifecycle monitoring. Without these channels, governance can remain concentrated in high-level frameworks while adoption advances through procurement processes, software deployments, vendor relationships, digital transformation initiatives, and sector-specific implementation. As a result, consequential decisions shaping AI outcomes may be made before they are recognized as governance decisions.

AILA findings show this gap across country contexts:



In **Bhutan**, the assessment points to the need to move from responsible AI principles towards practical governance mechanisms, including stronger oversight arrangements, public engagement, safety protocols, legal frameworks, risk categorization, algorithmic impact assessments, monitoring, and explainability.



AILA results from **Mongolia** similarly highlight gaps in transparency, safety, and accountability, including limited mechanisms for algorithmic transparency, explainability, impact assessment, risk classification, and dedicated oversight.



The assessment from **Viet Nam** points to the need for accountability, transparency, safety, and inclusivity mechanisms in public sector AI adoption, including transparency around algorithms, data use and decision-making, bias mitigation, safety protocols, ongoing evaluation, and engagement with diverse stakeholders.



Trinidad and Tobago adds a trust and safety perspective, with the assessment identifying gaps in AI-specific regulation, oversight, accountability, public trust, cybersecurity, incident reporting, rights and recourse, and mechanisms for escalation, opt-out, and explainability. It also shows how trust and safety gaps can move the governance conversation from diagnosis towards practical questions of procurement, vendor expectations, public sector guidance, and institutional readiness for responsible AI use.



The **Dominican Republic** makes this implementation gap especially explicit. Its AILA finds that references to responsible AI and strategic intent are present, but that ethics and trust and safety remain weakly institutionalised. Expert inputs point to limited mechanisms for accountability, transparency, human oversight, incident reporting, post-deployment risk monitoring, impact assessment, and procedures for individuals to challenge algorithmic decisions. Workshop proposals focused less on additional high-level governance principles than on practical instruments, including updated data protection rules, audits and certifications, information classification, observatories, data-quality standards, and clearer institutional responsibilities.

This operational view of governance is also reflected in UNDP's AI Trust & Safety Re-Imagination Programme,[1] which treats responsible AI as part of the full system around adoption, rather than as a layer above it. For most countries, the immediate governance challenge is less about building their own AI models than about the adoption, procurement, integration, and oversight of AI systems built, supplied, or updated by external providers. Procurement, vendor management, and lifecycle oversight therefore become some of the most immediate governance levers available. Contractual requirements, disclosure obligations, audit access, performance expectations, data-handling commitments, escalation protocols, exit conditions, and redress mechanisms can shape outcomes before systems are deployed at scale. This is especially important where governments remain accountable for effects on people while many design choices, updates, and performance trade-offs sit with external providers.

Governance is therefore operational as much as regulatory. Institutions need to ask the right questions to inform responsible adoption, monitor systems after deployment, respond when harms occur, and ensure that affected groups have meaningful avenues for explanation, contestation, and redress. This is not only about constraining risk. Clear rules, trusted processes, and predictable safeguards can also give governments, domestic enterprises, startups, researchers, and people more confidence to experiment, invest, and adopt AI in ways that serve public purposes.

Measurement is part of that governance machinery. If AI systems are introduced to improve services, reduce administrative burdens, expand access, better identify needs, or support decision-making, countries need ways to assess whether those outcomes are realistically being achieved. This means looking beyond deployment and short-term efficiency gains to service quality, error rates, equity of access, user experience, staff capability, costs, trade-offs, and whether benefits persist over time. Without these feedback loops, governments may be aware that AI has been introduced but not know whether it is improving services, shifting burdens onto workers or people, or creating risks that only become visible after deployment.

Accountability also depends on institutional design. Where policy, implementation, and oversight responsibilities are unclear or concentrated in the same institutions, accountability becomes harder to sustain. This is especially important when AI systems are procured or integrated by one part of government, used by another, and affect end users who may have limited visibility into how decisions are made.

Governance cannot be treated as a separate workstream alongside implementation. It ought to be built into procurement, data systems, DPI, service design, vendor management, monitoring, oversight, and redress. Private sector actors need clearer expectations around transparency, performance, auditability, escalation, and exit. Civil society, academia, and affected communities need channels to surface risks and contest failures. Development partners should support governance capabilities that can be reused across sectors, rather than frameworks that sit above adoption. Governance becomes effective when it gives institutions the capacity to adopt, learn, correct, and remain accountable over time.

[1] https://www.undp.org/digital/ai-trust-and-safety

2.7 The implementation test

The six patterns covered in this report are not separate findings; they describe a single implementation problem.

Once AI adoption is underway, readiness depends on whether political attention, delivery systems, data, infrastructure, talent, ecosystems, and governance can be connected quickly enough to influence the choices being made around AI adoption. Political attention can create mandates and momentum, but it does not organize financing, procurement, oversight, or learning. Adoption through platforms, procurement, public digital systems, and sector-specific tools can create practical experience, but it can also create dependencies before they are recognised as policy choices.

The implementation test is therefore not whether each condition exists in isolation. It is whether the elements of the adoption system work together. A promising use case can fail if the data are poor, inaccessible, or not trusted. Improved infrastructure may have limited development value if institutions cannot procure, supervise, or adapt systems effectively. Technical talent may not become local capability if there is no demand, financing, or pathway into public and market use. Governance principles may remain symbolic if they do not reach procurement, vendor management, monitoring, and redress. Weakness in one part of the system changes what the others can achieve.

The central development risk is not that countries fail to adopt AI. In many contexts, adoption is already entering public institutions, markets, services, and digital systems. The deeper risk is unmanaged adoption, where systems are introduced through vendor relationships, platform dependencies, short-term projects, or institutional gaps, with limited visibility into who benefits, what costs are incurred, what risks emerge, and whether improvements endure.

The task is not to wait for perfect conditions. It is to make adoption visible enough to influence while there is still room to do so. Governments and national ecosystems need to know where AI is entering, who is shaping decisions, what dependencies are being created, what capabilities are being built, who may benefit or be harmed, and where accountability sits when systems fail.



This is the practical meaning of ‘reading readiness backwards’. **Readiness is no longer a threshold before action. It is the capacity to connect institutions, data, infrastructure, talent, ecosystems, and governance as adoption unfolds, so that AI strengthens public systems, expands economic opportunity, supports local innovation, and advances human development**.

The next section turns this into decision points for governments, partners, and ecosystem actors.

3 Decision points once adoption is underway

If a country’s AI readiness only becomes truly visible through adoption, the next question is: What should governments, development partners, and ecosystem actors pay attention to as this process unfolds? The aim is not to define a one-size-fits-all template for countries to follow. Countries differ in institutional capacity, digital maturity, market structure, infrastructure, governance arrangements, and development priorities. A more useful approach is to examine whether AI-related decisions are strengthening the conditions for responsible and sustainable adoption, or adding layers of disconnected activity.

These are not universal recommendations or a prescriptive checklist. They are decision points for governments, development partners, and ecosystem actors. Each is framed as a question to help identify what needs to be made visible, organized, financed, governed, measured, or reused as adoption moves forward. Together, they test whether AI activity is visible enough for institutions to influence, governed enough to build trust among people, reusable enough to build capability, and measured enough to learn from. The purpose is to identify where attention, resources, and accountability need to concentrate as adoption moves forward.

Key questions for AI adoption decisions:

Guiding question	What to examine
1. Where is AI already entering systems?	Map the pathways through which AI-enabled tools are already being used or considered. This includes public institutions, service-delivery systems, digital public infrastructure, procurement processes, vendor platforms, enterprise software, research partnerships, and development projects, and informal use of publicly available AI tools by civil servants or institutional staff.
	AI adoption is often underway before it is visible as an AI policy issue. Understanding where adoption is happening helps identify which systems are affected, what dependencies are being created, and where risks or opportunities are emerging while there is still room to shape decisions.
2. Who is shaping AI adoption, and around whose priorities?	Identify the actors influencing AI adoption across government, markets, and the wider national ecosystem, including public institutions, regulators, digital agencies, technology providers, startups, universities, investors, civil society, telecommunications and cloud providers, businesses integrating AI into existing operations, and development partners.
	Adoption can move quickly without necessarily building local capability or serving public priorities. The question is whether public demand, local innovation, research capability, financing, safeguards, and accountability are connected around problems that matter, or whether local actors remain peripheral to the value being created.
3. Is there operating authority behind the strategy?	Strategies can create direction, but they do not create implementation capacity on their own. Countries need arrangements that clarify which entity has the mandate to lead, how institutions work together, how priorities are financed, how systems are procured, who supervises deployment, how risks are escalated, and where accountability sits when systems fail.
	This matters most where multiple institutions and partners are already moving at the same time. Operating authority gives governments a way to convert political attention into delivery capacity, and to ensure that lessons from implementation inform future decisions rather than remain trapped in individual projects or institutions.
4. Do foundational investments expand agency to act?	Foundational investments in data, compute, connectivity, cloud infrastructure, digital public infrastructure, talent, and institutional capability should be assessed on the basis of whether they expand a country's ability to make deliberate choices about AI adoption.
	The issue is not whether every country builds every layer of the AI stack. It is whether countries have enough access, leverage, interoperability, institutional capability, and trusted partnerships to use AI on terms aligned with national priorities. A data centre, cloud partnership, training programme, procurement reform, or AI sandbox should not be treated as a standalone intervention or a signal of progress on its own. Each should strengthen the ability to choose, adapt, supervise, contest, and sustain AI systems over time.

Guiding question	What to examine
5. Do early initiatives generate reusable capacity?	AI initiatives in health, agriculture, education, climate, social protection, justice, or public administration should not be treated as standalone applications or sector pilots. A useful initiative should generate stronger data assets, improved institutional workflows, procurement learning, local partnerships, safeguards, evaluation practices, and delivery capacity. This is the difference between a demonstration and an implementation pathway. A demonstration may prove that a tool can work in one setting. An implementation pathway builds capabilities that can be used again across institutions, sectors, and future adoption decisions. It should also clarify who owns, maintains, evaluates, and adapts the system after a pilot, vendor contract, donor project, or research partnership ends.
6. What can be shared, pooled, or reused?	Because not every country can or should build every layer of the AI stack alone, some capabilities may be more effective when shared, pooled, or reused across countries and institutions. This is especially relevant where countries face similar constraints around compute access, assurance, evaluation, language resources, procurement intelligence, talent, or vendor oversight. Shared approaches can include procurement standards, pooled assurance and evaluation capacity, regional talent initiatives, vendor intelligence, common benchmarks, reusable risk assessment tools, language resources, shared compute arrangements, and South–South learning. These assets can reduce duplication, strengthen leverage with providers, and create more consistent expectations across vendors, sectors, and countries. They can also help smaller or more resource-constrained countries retain agency without pursuing full technological self-sufficiency.
7. How will value, cost, and risk guide decisions?	AI adoption should be assessed against the outcomes it is meant to improve. Governments and development partners should define early what success means in concrete terms, such as better service quality, reduced errors, improved accessibility, faster response times, lower administrative burden, stronger institutional capability, or more equitable access. Costs, trade-offs, and burdens should be tracked alongside benefits. Who benefits? Who bears new burdens? Are workers, users, and affected communities involved in identifying what should be measured? What evidence would justify scaling, adapting, pausing, or stopping a system? Without this feedback loop, countries risk scaling AI systems based on assumed value rather than demonstrated public benefit.
8. Are safeguards embedded where adoption decisions are made?	Responsible AI needs to appear in the operational settings where AI is selected and used. This includes procurement clauses, vendor obligations, data governance arrangements, service design, model documentation, auditability requirements, impact assessments, monitoring systems, escalation pathways, oversight mandates, and redress mechanisms. These safeguards make adoption contestable, correctable, and trusted. They also allow governance to support responsible innovation by giving institutions clearer ways to detect failure, respond to harm, learn from deployment, and maintain accountability over time. When safeguards are operational, institutions can adopt AI with more confidence because they are not relying on principles alone.

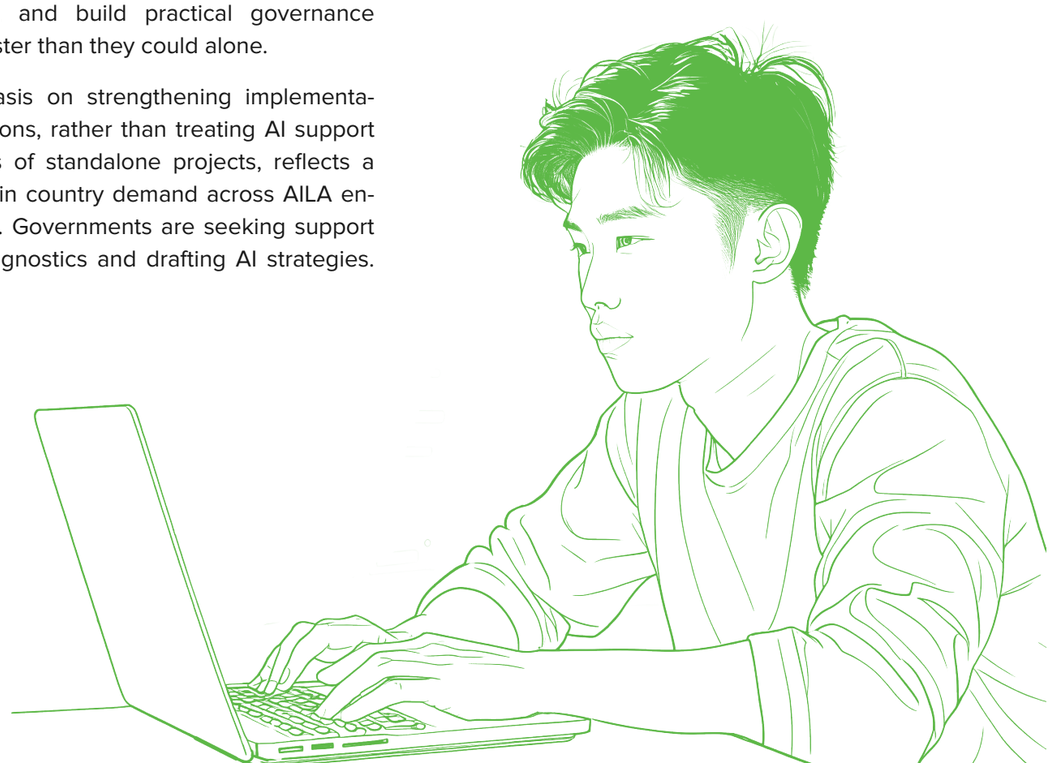
What this means for support

For governments, these guiding questions and decision points help move towards a clearer view of what needs to be organized, financed, governed, measured, and reused for sustainable and responsible AI adoption. For ecosystem actors, they clarify where participation matters, from local innovation and research to accountability, inclusion, and trust. For development partners, they provide a way to assess whether the support required lies in strengthening the conditions around AI adoption or funding standalone AI projects.

Support should be judged by what it generates. A compute investment should support more demand, talent, sustainability, and access models. A governance framework should encourage better procurement and deployment. A sector initiative should build reusable data, institutional learning, safeguards, and delivery capacity, not just a single application. Ecosystem support should help local actors participate in adoption pathways. Shared assets should help countries learn from one another, reduce duplication, and build practical governance capacity faster than they could alone.

The emphasis on strengthening implementation conditions, rather than treating AI support as a series of standalone projects, reflects a wider shift in country demand across AILA engagements. Governments are seeking support beyond diagnostics and drafting AI strategies. They are increasingly focused on the implementation conditions around adoption, including data reform, institutional coordination, governance instruments, procurement capability, public sector literacy, ecosystem partnerships, and sector-specific pathways for responsible deployment. What is valued most is support that connects AI strategy to the systems through which AI is adopted.

For UNDP, AILA provides a structured way to support countries as they work through these decision points. Its primary value lies in helping governments and partners move beyond a static assessment of readiness to identify where AI is already entering systems, where constraints are emerging, and which institutional, technical, ecosystem, and governance conditions need to be strengthened. In that sense, AILA helps shift support from readiness diagnosis alone towards the practical work of making AI adoption visible, accountable, locally useful, and aligned with sustainable development.



4 The work ahead

AI adoption is no longer a distant policy question for many countries. It is already entering public institutions, markets, digital systems, and everyday services. The development significance of this moment lies in the terms on which adoption takes place, and whether countries can shape those terms as AI moves through public and economic systems.

Those terms are being set now. Procurement choices, infrastructure investments, data governance arrangements, vendor relationships, institutional mandates, and ecosystem partnerships will shape AI adoption for years to come. They will influence whether AI strengthens public systems, expands economic opportunity, supports local innovation, and improves people's lives, or whether it deepens dependency, concentrates benefits, and makes existing inequalities harder to correct.

This is why this report reads readiness backwards. Once adoption is underway, readiness is no longer about what was in place beforehand – it becomes visible in what AI adoption requires as it diffuses through real systems. Institutions need to see where AI is entering systems. Data and infrastructure need to support reliable use, local actors need pathways to participate in value creation. Governance needs to reach operational decisions. And people affected by AI-enabled systems need ways to question, contest, and correct outcomes.

Advances in technological capability will shape what AI systems can do, but they will not determine their development impact. As adoption accelerates, that impact will depend on institutional and operational choices about who benefits, whose risks are surfaced, whose needs are reflected, and who can influence systems once they are deployed. Acceleration is not neutral. Where institutions, data systems, governance arrangements, and ecosystems can absorb and direct adoption, faster diffusion can expand opportunity. Where they cannot, speed can compound dependency, exclusion, and unaccountable risk.

The window to shape these systems is still open, but not indefinitely. The work ahead is not to wait for perfect conditions, nor to treat AI diffusion as inevitable or believe it will be self-correcting. It is to make adoption visible, deliberate, and accountable before early choices become embedded in infrastructure, contracts, standards, institutional habits, and dependencies.

Reading AI Readiness Backwards shifts the question away from whether countries are prepared for AI, and towards whether they are building the conditions to steer adoption before early choices become embedded in systems, contracts, institutions, and dependencies. That is where AI's development impact will be decided.

Annex I:

The AILA evidence base and methodology

About the Artificial Intelligence Landscape Assessment (AILA)

The Artificial Intelligence Landscape Assessment (AILA) is UNDP's whole-of-society diagnostic framework for understanding the conditions shaping AI adoption, governance, and ecosystem development. Developed to support countries in navigating AI in ways that align with national priorities and sustainable development objectives, AILA combines quantitative and qualitative evidence to assess how institutional, technical, ecosystem, and governance factors influence the adoption and use of AI.

AILA is structured around three interconnected pillars:

- ✓ **AI Ecosystem**, which examines the foundational conditions that support AI development and adoption, including data, talent, enabling environments, and sustainable computing infrastructure.
- ✓ **AI for Government**, which focuses on the use of AI in public sector transformation, including government capability, governance arrangements, technology adoption, and data governance.
- ✓ **AI Regulation and Ethics**, which examines regulatory frameworks, ethical principles, trust and safety considerations, and the broader governance arrangements needed for responsible AI adoption.

Together, these pillars provide a structured way to examine AI readiness as a combination of technical, institutional, ecosystem, and governance conditions that shape how AI is adopted and governed in practice.

Methodological approach

AILA uses a mixed-methods methodology designed to combine quantitative indicators with stakeholder perspectives and contextual analysis.

The assessment draws on four primary sources of evidence:

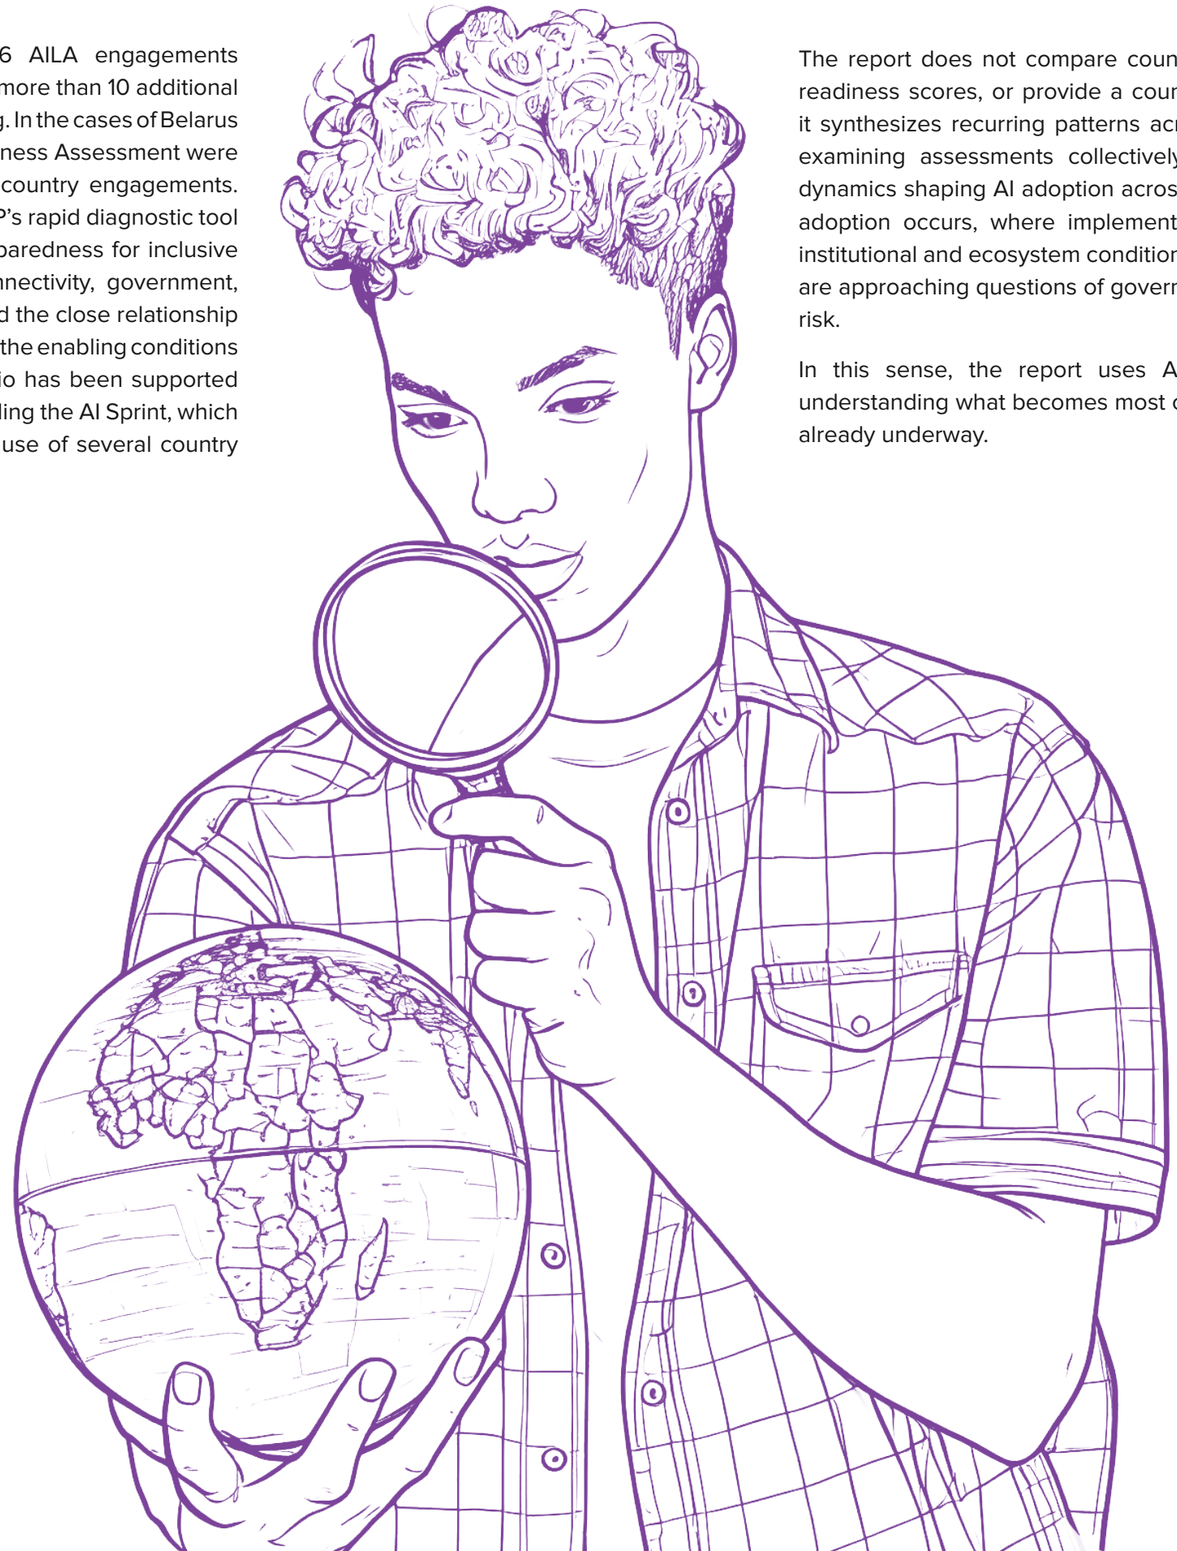
- ✓ Expert consultations with government officials, academics, researchers, and private-sector specialists;
- ✓ Stakeholder surveys spanning government, private sector, academia, and civil society;
- ✓ Desk research examining policies, strategies, legislation, institutional arrangements, and publicly available evidence;
- ✓ External indicators drawn from international datasets, including measures related to digital infrastructure, governance, innovation, connectivity, cybersecurity, statistical capacity, research and development, and talent.

These inputs are complemented by a multi-stakeholder validation workshop that brings together representatives from government, private sector, academia, civil society, and development partners to review findings, validate evidence, identify priorities, and co-create recommendations.

Beyond generating a readiness assessment, the process is intended to surface opportunities, constraints, implementation challenges, governance questions, emerging risks, and areas where coordination is required to support responsible AI adoption.

AILA evidence base and scope

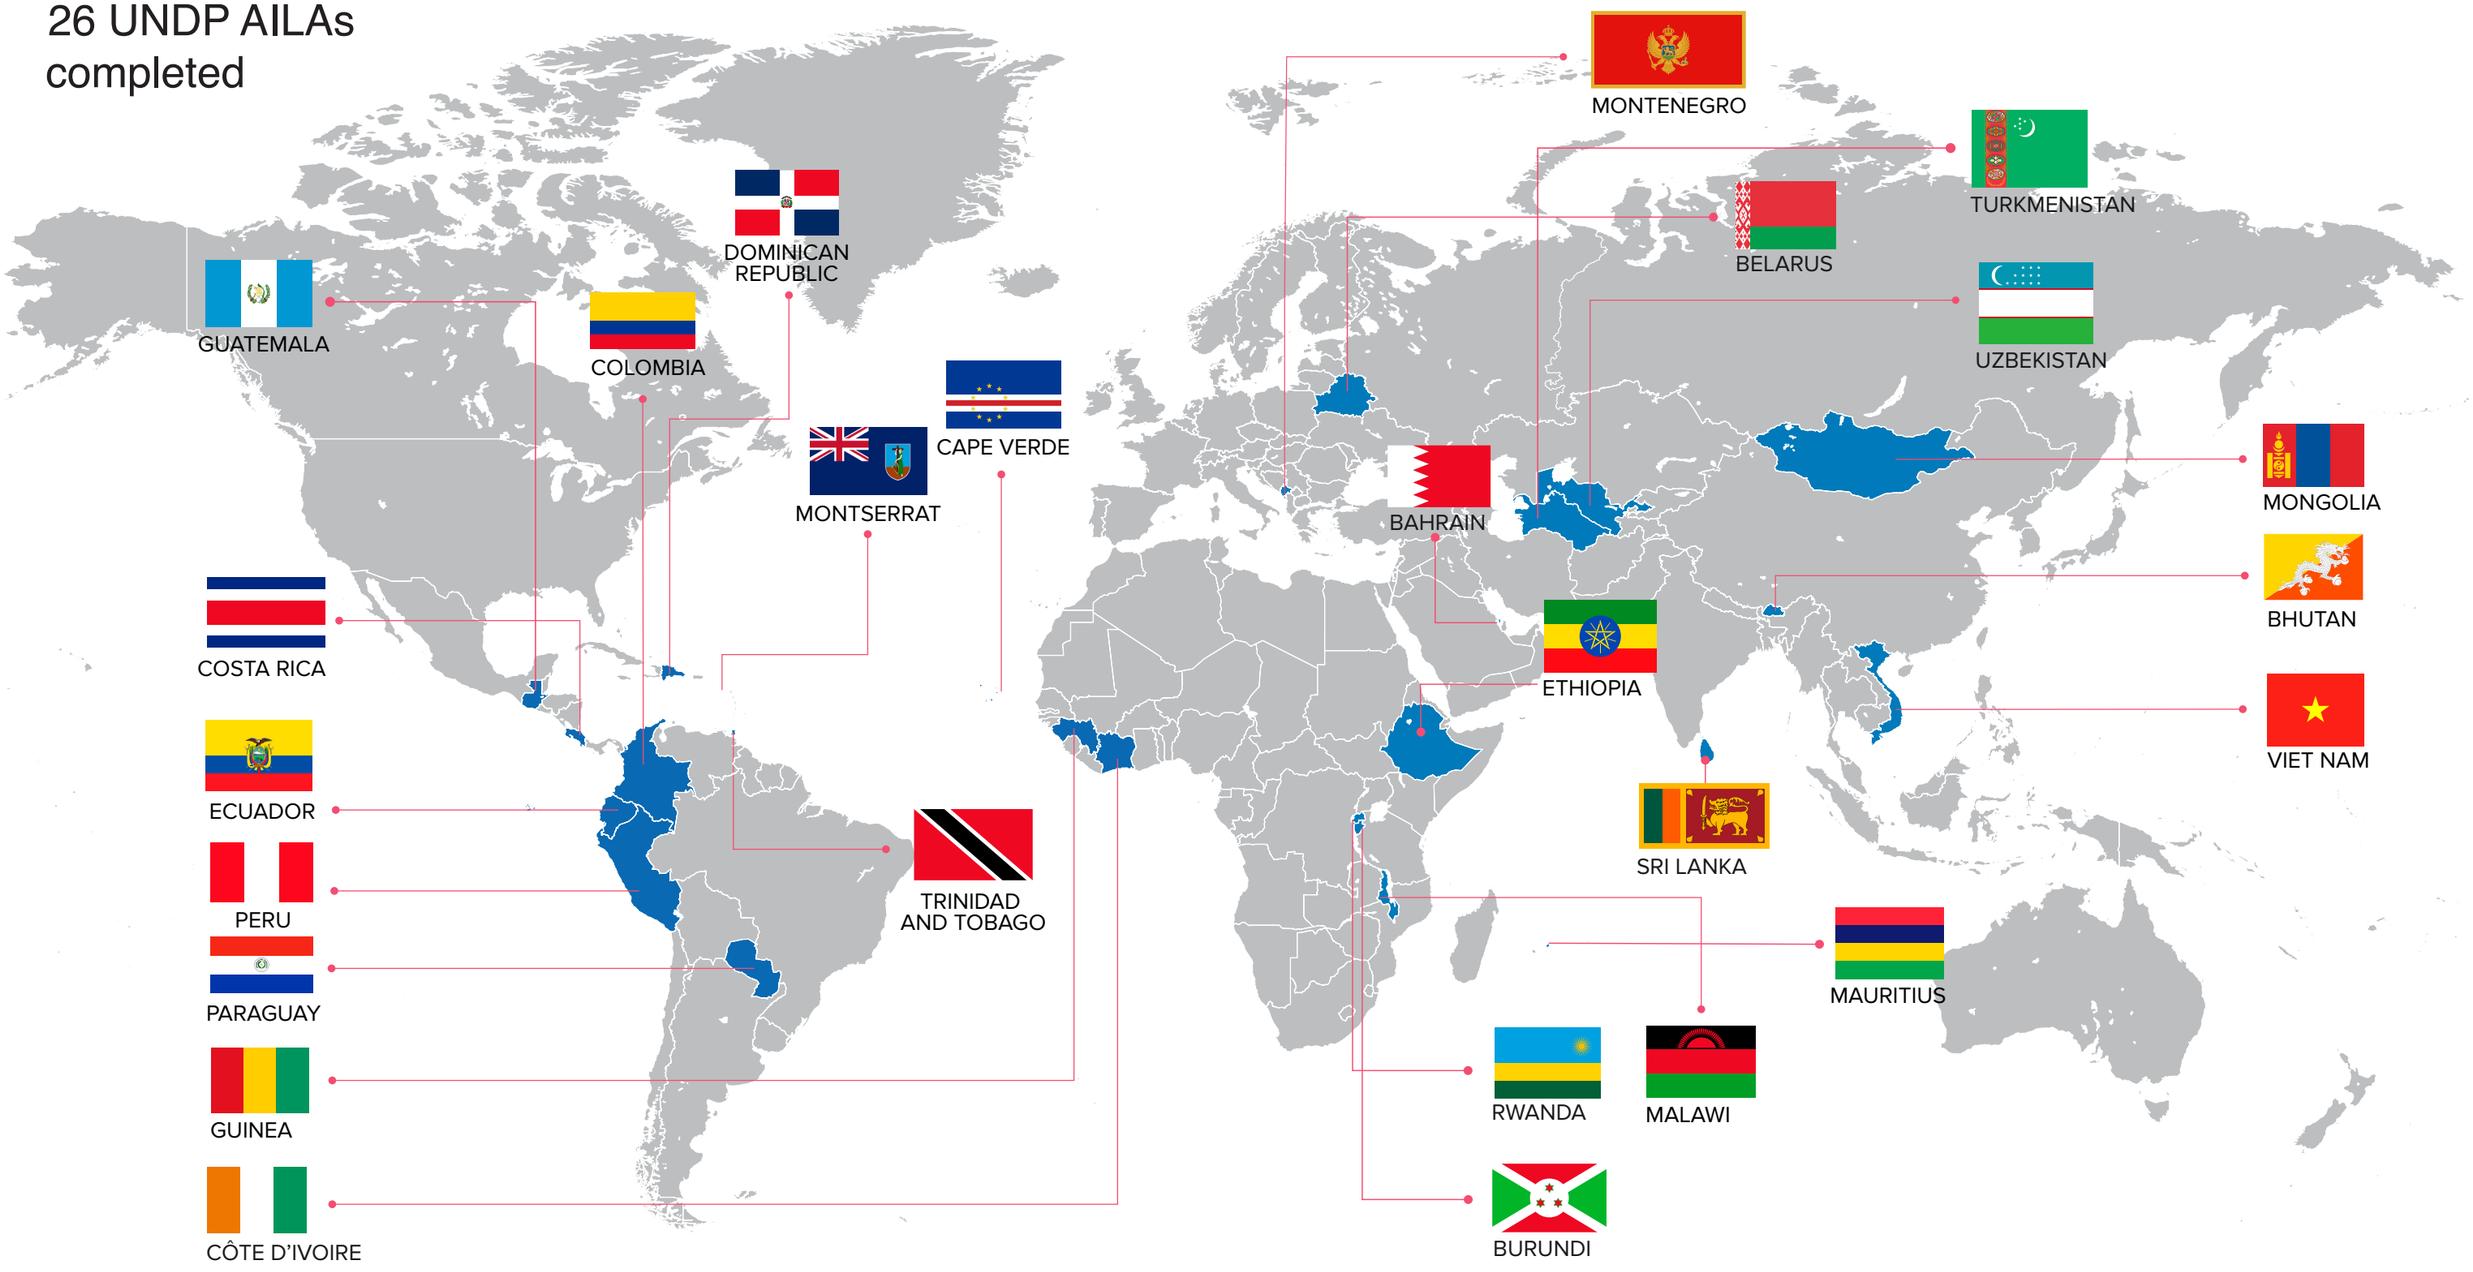
This report draws on findings from 26 AILA engagements completed between 2024 and 2026, with more than 10 additional assessments underway at the time of writing. In the cases of Belarus and Montserrat, AILA and the Digital Readiness Assessment were delivered together as part of integrated country engagements. The Digital Readiness Assessment is UNDP's rapid diagnostic tool for helping governments assess their preparedness for inclusive digital transformation across people, connectivity, government, regulation, and the economy. This reflected the close relationship between digital readiness, AI adoption, and the enabling conditions examined through AILA. The AILA portfolio has been supported through several UNDP workstreams, including the AI Sprint, which contributed to the delivery and follow-on use of several country assessments.



The report does not compare countries, assign rankings, aggregate readiness scores, or provide a country-by-country summary. Instead, it synthesizes recurring patterns across the AILA evidence base. By examining assessments collectively, the report identifies common dynamics shaping AI adoption across diverse contexts, including how adoption occurs, where implementation constraints emerge, which institutional and ecosystem conditions matter most, and how countries are approaching questions of governance, public value, inclusion, and risk.

In this sense, the report uses AILAs as an evidence base for understanding what becomes most consequential once AI adoption is already underway.

26 UNDP AILAs completed



Regional Bureau for Africa (RBA)
Burundi
Cabo Verde
Côte d'Ivoire
Ethiopia
Guinea
Malawi
Mauritius
Rwanda

Regional Bureau for Europe and the Commonwealth of Independent States (RBEC)
Belarus
Montenegro
Turkmenistan
Uzbekistan

Regional Bureau for Arab States (RBAS)
Bahrain

Regional Bureau for Asia and the Pacific (RBAP)
Bhutan
Mongolia
Sri Lanka
Viet Nam

Regional Bureau for Latin America and the Caribbean (RBLAC)
Colombia
Costa Rica
Dominican Republic
Ecuador
Guatemala
Montserrat
Paraguay
Peru
Trinidad and Tobago

Further Information

AILA engagements are conducted in partnership with national stakeholders and may serve different purposes depending on country context. Publication arrangements therefore vary. Publicly available AILA country reports, along with further information on the AILA framework and methodology, are available on: https://www.undp.org/digital/aila

Further information on the UNDP AILA, including the AILA framework, methodology, and publicly available country reports is available at: https://www.undp.org/digital/aila. Further information on UNDP's Digital Readiness Assessment is available at: https://www.undp.org/digital/dra.

AI Hub for Sustainable Development

UNDP