

Institutional Readiness as a Layer of Defence in Frontier AI Governance


A practical argument for strengthening the institutions expected to govern increasingly capable AI systems.

ANTHONY ONDICHO, PMP®

Founder & CEO, DFIP Blue Growth Partners
Kenya | August 2026

CORE THESIS

Frontier AI safety is not only a technical challenge. It is also an institutional challenge. If the capability of AI systems advances faster than the capacity of governments and organisations to understand, govern, supervise and respond to them, weak institutional readiness can become a pathway through which technical risks escalate into systemic harm.

A decorative graphic consisting of several concentric circles in a light blue color, located in the bottom right corner of the page. The circles are centered around a point and expand outwards, creating a ripple effect.

THE ARGUMENT

Why institutional readiness belongs in frontier AI governance

My interest in institutional AI readiness comes from more than a concern about whether organisations adopt new technology successfully. After more than eighteen years working in international development, infrastructure, humanitarian response and organisational leadership, I have repeatedly seen that institutions can possess resources, policies and technologies yet still fail to achieve intended outcomes when leadership, accountability, capability and implementation systems are weak.

AI intensifies that concern. Increasingly capable systems are moving from experimentation into operational decision-making, public services, financial systems, knowledge work and critical organisational processes. The speed of adoption can be far greater than the speed at which institutions build governance capability. This creates a dangerous asymmetry: organisations can acquire powerful AI faster than they can develop the judgement, structures and controls needed to govern it.

For frontier AI, that gap matters because the consequences may no longer remain local. A weakly governed system deployed within a highly connected institution can affect data, finance, information, infrastructure, public decision-making or cyber environments far beyond the original use case. Institutional weakness can therefore become an amplifier of technical risk.

“The speed of adoption can be far greater than the speed at which institutions build governance capability.”

INSTITUTIONAL CAPACITY

The institutional gap

The core risk is simple: frontier capability may advance faster than institutional capability. Governments, regulators and organisations may be expected to supervise systems whose capabilities, limitations and failure modes they do not fully understand. In such conditions, formal governance can exist on paper while practical control remains weak.

KNOWLEDGE ASYMMETRY

Frontier AI developers and vendors may understand their systems far better than the institutions procuring, regulating or overseeing them. Institutions that cannot independently interrogate evidence, evaluations and limitations risk becoming dependent on the assurances of the actors they are expected to govern.

LEADERSHIP AND ACCOUNTABILITY

Where executive ownership is unclear, responsibility can fragment across technology, legal, data protection, procurement, risk and business teams. High-impact AI requires a clearly accountable human authority with the competence and mandate to challenge, pause or terminate deployment.

GOVERNANCE CAPABILITY

Policies alone are insufficient. Institutions need AI-system inventories, risk classification, procurement controls, human-oversight arrangements, incident escalation, monitoring and audit mechanisms that work in practice.

WORKFORCE CAPABILITY

Boards, executives, regulators, procurement teams, programme managers and frontline users require different levels of AI literacy. A generic awareness course cannot substitute for role-specific competence.

RESPONSE CAPACITY

When an AI system behaves unexpectedly, institutions need the ability to detect the problem, understand its consequences, escalate it and act quickly. Weak response capacity can allow local failures to spread.

COUNTERARGUMENT

Why technical safeguards alone may not be enough

The strongest argument against my position is that catastrophic frontier-AI risk may arise primarily from technical loss of control or other model-level failures. If that is correct, technical alignment, capability evaluations and hard deployment safeguards may be more important than institutional readiness.

I take that counterargument seriously. Institutional readiness is not a substitute for technical safety. It is a complementary layer of defence. Even strong technical controls still have to be required, interpreted, audited, enforced and updated by institutions. A regulator must know what evidence to demand. A board must know what risk it is accepting. A procuring institution must know what controls belong in a contract. A human reviewer must have both competence and authority. An incident-response mechanism must exist before the incident occurs.

I would change my view if strong technical safeguards were shown to reliably contain catastrophic risks regardless of institutional competence. Until then, I believe technically safer systems still require institutions capable of governing their deployment and responding when assumptions fail.

Institutional readiness is not a substitute for technical safety. It is a complementary layer of defence.

FROM PRINCIPLE TO PRACTICE

A practical readiness framework

My policy contributions in Kenya have increasingly focused on making AI governance operational rather than abstract. In comments on Kenya's Draft Artificial Intelligence and Other Emerging Technologies Policy, 2026, I argued for minimum sector-readiness frameworks, clear institutional responsibilities, delivery stage gates, role-based AI fluency, proportionate governance and post-deployment review. In comments on the Office of the Data Protection Commissioner's Draft Guidance Note on Artificial Intelligence, I proposed institutional AI-readiness checkpoints, meaningful human oversight, executive accountability, AI-system inventories and stronger vendor due diligence.

Taken together, these ideas suggest a practical institutional-readiness model for high-impact and frontier AI. Before deployment, an institution should be able to answer six questions:

LEADERSHIP

Who is accountable for the decision to deploy the system, and who can stop it?

GOVERNANCE

What risk classification, oversight, audit, procurement and incident-response controls apply?

PEOPLE

Do the board, executives, technical teams, reviewers and end users have the competence required for their roles?

DATA AND TECHNOLOGY

Are the data, infrastructure, security and technical-assurance mechanisms adequate for the intended use?

HUMAN OVERSIGHT

Is human review meaningful, informed and empowered rather than a procedural rubber stamp?

VALUE AND EXIT

What public or business value is expected, how will it be measured, and under what conditions will the system be modified, suspended or terminated?

This is not intended to become another compliance checklist. The purpose is to create a decision discipline: institutions should demonstrate that they are capable of governing what they intend to deploy. Readiness should be proportionate to risk, but the burden should rise as potential consequences increase.

AFRICAN CONTEXT

Why this matters for Africa

Africa has an opportunity to benefit enormously from AI in agriculture, health, education, financial inclusion, public administration and infrastructure. But the continent also faces uneven institutional capacity, digital infrastructure, regulatory capability and access to specialised expertise. A governance model that assumes every institution has the same resources risks either encouraging ungoverned adoption or creating requirements that only the largest organisations can satisfy.

The response should not be to slow African innovation. It should be to build institutional capability alongside technological capability. Governments, development institutions, universities, private companies and funders should treat AI-readiness capacity as part of digital infrastructure: leadership capability, governance templates, shared assurance services, role-based learning, technical assistance and context-specific readiness assessments.

This is particularly important where AI influences access to finance, public benefits, social protection, health, employment or essential livelihood opportunities. In my own work through DFIP, the development of DFIP One has reinforced this point. A platform intended to improve the finance readiness of underserved smallholder farmers could itself create exclusion if incomplete data or poorly designed algorithms are allowed to become automatic gatekeepers. Responsible deployment therefore requires human review, transparency, correction mechanisms, data protection and conscious attention to vulnerable users.

PRIORITY INTERVENTION

The highest-leverage intervention

The highest-leverage work I would prioritise now is context-specific AI readiness assessment followed by leadership and institutional capacity building before high-impact deployment. The assessment should not ask only whether an organisation possesses technology. It should test whether leadership, governance, people, data, technical controls, risk management and operational processes are capable of supporting responsible use.

The strongest objection is that readiness assessments could become bureaucratic exercises that delay beneficial innovation without reducing real risk. That is a legitimate danger. To avoid it, readiness should be risk-based, time-bound and tied to actual deployment decisions. If evidence showed that such assessments consistently failed to influence decisions or safety outcomes, I would favour stronger mandatory technical assurance and deployment controls instead.

READINESS BEFORE DEPLOYMENT

Assess leadership, governance, people, data, technical controls, risk management and operational processes before high-impact AI deployment.

CONCLUSION

From adoption to stewardship

The central question for the AI era is not simply whether institutions can adopt AI. It is whether they can remain accountable for what they adopt.

Frontier AI governance will require technical research, international coordination, regulation and engineering safeguards. But those mechanisms ultimately operate through institutions. If institutions are unable to understand the systems they oversee, challenge powerful vendors, assign accountability, protect affected people or respond when failures occur, the governance architecture remains fragile.

Institutional readiness should therefore be treated as a layer of defence: not the only layer, and not a replacement for technical AI safety, but a necessary capability for turning governance principles into effective human control. As AI capability grows, strengthening the institutions expected to govern it should be treated as part of the safety agenda itself.

ABOUT THE AUTHOR

ANTHONY ONDICHO, PMP®

Founder & CEO of DFIP Blue Growth Partners. Anthony has more than 18 years of experience in international development, infrastructure, climate resilience, humanitarian response and programme leadership. His current work focuses on institutional AI readiness, responsible AI transformation, development finance and digital solutions for inclusive development in Africa.

anthony@dfipbluegrowth.com | dfipbluegrowth.com