

GLOBAL SOUTH LEGAL ARCHITECTURE FOR ARTIFICIAL INTELLIGENCE GOVERNANCE AND ACCOUNTABILITY

*A Common Minimum Denominators Approach
for the Global South Group of Nations*

Adopted at the
**Second Global South Artificial Intelligence
Law and Governance Dialogue 2026**

Version 6

Foreword

Artificial intelligence will not wait for the Global South to organise its response. Agentic AI is already moving across our jurisdictions, shaping discourse, mediating opportunity, ingesting our data, training upon our languages and cultures, and operating, too often, in a state of legal weightlessness. The question before our group of nations is no longer whether we will be governed by AI; it is whether we will govern AI on our own terms.

This Architecture is offered to the Global South Group of Nations as a working instrument, building upon the instruments adopted by the Global South, including the New Delhi Declaration, the AI Accountability Framework 2026, and the Global Agentic AI Liability Framework, and designed from the outset for collective ownership, sovereign adoption, and operational use. From the moment of its adoption by the Second Dialogue, it belongs to the Global South nations and institutions that take it forward.

What follows is therefore an Architecture in the most operational sense of the word: a structured arrangement of layers, pillars, mechanisms, and working bodies through which the principles already articulated by the Global South can be translated into national statute, sectoral policy, and lived experience. It is offered without proprietary claim, and with the conviction that the twenty-first century's legitimate AI order will be one that the Global South helped to architect.

Dr. Pavan Duggal

Convenor, Second Global South Artificial Intelligence Law and Governance Dialogue,

Architect of the Duggal Doctrine

Pioneer, AI Law and Cyberlaw Jurisprudence

I. Preamble and Context

The twenty-first century has placed the Global South at a defining jurisprudential crossroads. Artificial intelligence, and increasingly agentic artificial intelligence, is reconfiguring economies, rewriting the architecture of public life, and reshaping the relationship between the citizen, the State, and the algorithm. The Global South cannot afford to remain a passive recipient of governance frameworks designed elsewhere; it must become a co-architect of the legal order that will determine how AI is built, deployed, audited, and held to account within and across its sovereign jurisdictions.

This Global South Legal Architecture for Artificial Intelligence Governance and Accountability (hereinafter, the “Architecture”) is conceived not as a declaration, and not as a re-statement of academic principle, but as an operational legal architecture - multi-layered, capable of national absorption, amenable to sector-specific calibration, and built around enforceable common minimum denominators. It builds upon, and does not reinvent, the edifice already laid down by the New Delhi Accord on Artificial Intelligence Law and Governance(2025), the New Delhi Declaration on Artificial Intelligence (February 2026), the New Delhi Compact on Artificial Intelligence Governance and Accountability (May 2026), the AI Accountability Framework 2026, the Global Agentic AI Liability Framework, and the Universal Declaration of AI Accountability Rights adopted at the International AI Accountability Forum 2026. The Architecture functions as the operational synthesis of these instruments , the layer at which their principles meet the statute books of Global South nations.

II. Foundational Premises - Building on, Not Reinventing

The Global South does not need another principles document. What it needs is structural scaffolding - a defined arrangement of layers, instruments, and institutional vehicles through which the principles already articulated can be made operative within national jurisdictions. The Architecture therefore consciously departs from the model of frameworks anchored in foundational principles alone, and offers instead a legal architecture composed of layers, mechanisms, working groups, and ratification pathways.

Comparable architectural precedents are instructive but not determinative. The European Union's tiered, risk-stratified approach under the EU AI Act; the OECD AI Principles read with the OECD Recommendation on AI; the UNESCO Recommendation on the Ethics of Artificial Intelligence; the African Union Continental AI Strategy; the ASEAN Guide on AI Governance and Ethics; the G20 AI Principles; and the Bletchley and Seoul Declarations on Frontier AI - each provides a partial template. None, however, has been designed with the civilizational, linguistic, cultural, and developmental specificities of the Global South at its centre. The Architecture closes that gap.

III. The Three-Layer Legal Architecture

The Architecture is composed of three concentric legal layers. Each layer is independently operative, mutually reinforcing, and designed for incremental adoption.

Layer One - National Law on Artificial Intelligence and Accountability

The foundational layer envisages the enactment, by each Global South nation, of a dedicated national statute on artificial intelligence and accountability. Such legislation shall define artificial intelligence systems, agentic AI systems, foundation models, and frontier AI; designate competent national regulators; set out duties of AI developers, deployers, and integrators; codify the rights of affected persons and communities; and create dedicated mechanisms for redress, audit, and enforcement. The national statute is the load-bearing wall of the Architecture.

Layer Two - National Sector-Specific Policies on AI

The second layer comprises sector-specific national policy instruments addressing the role of AI in health, education, finance, agriculture, public administration, law enforcement, judicial processes, social welfare, electoral integrity, defence and national security, and digital public infrastructure. These policies operationalise the national statute within the unique risk profile, capacity reality, and developmental priorities of each sector.

Layer Three - Evolution of Best Practices and Living Standards

The third layer is an evolutive layer - a continuously updated body of best practices, technical standards, model contractual clauses, audit methodologies, model risk assessments, and benchmark compliance protocols. This layer recognises that AI is a living technology and that its governance must be Living Law: responsive, iterative, and capable of doctrinal refresh without requiring statutory amendment at every turn.

IV. Pillars of the Architecture - The Common Minimum Denominators

The Architecture rests on seven non-negotiable pillars. These are the common minimum denominators that must inform every national instrument, every sectoral policy, and every best practice adopted within the Global South Group of Nations.

Pillar 1 - Algorithmic Sovereignty

The sovereign rights of every Global South nation include the right to algorithmic sovereignty: the right to govern, regulate, supervise, audit, restrict, and where necessary prohibit, the artificial intelligence algorithms operating within its national jurisdiction. Algorithmic sovereignty is not discretionary; it is a constitutive element of national sovereignty in the age

of agentic AI. Every AI algorithm operating within a Global South jurisdiction shall be amenable to, and compliant with, the norms, rules, and regulations of that jurisdiction.

Pillar 2 - Prevention of Cognitive Colonialism

The Architecture stands against cognitive colonialism. being the silent transfer of epistemic, narrative, linguistic, and value-laden authority to AI systems trained, hosted, and governed beyond Global South frontiers. Global South nations shall adopt legal and policy measures to ensure that AI systems deployed in education, public discourse, journalism, judicial reasoning, and cultural production do not function as instruments of cognitive subordination.

Pillar 3 - Prohibition of Unconsented Data Distillation and Extraction

The unconsented distillation, scraping, ingestion, fine-tuning, and commercial extraction of personal data, cultural data, indigenous knowledge, and behavioural data from Global South populations, particularly populations that are not yet positioned to comprehend the nuances of AI training and inference, by Big Tech and frontier AI providers, is prohibited. Consent must be informed, granular, withdrawable, and meaningful within the cultural and linguistic context of the data subject.

Pillar 4 - Mandatory Compliance of Big Tech AI Algorithms with National Law

Big Tech AI algorithms accessible within or operating upon Global South jurisdictions shall comply with the applicable norms, statutes, regulations, judicial pronouncements, and constitutional standards of those jurisdictions. Extraterritorial corporate practice cannot displace constitutional duty. Non-compliance shall attract proportionate civil, regulatory, and where warranted, criminal consequences.

Pillar 5 - Protection of Linguistic, Cultural, Tribal, and Indigenous Data Heritage

The linguistic, cultural, tribal, and indigenous data heritage of the Global South , including oral traditions, indigenous knowledge systems, sacred and ceremonial content, traditional medicine, customary legal practices, and community narratives , shall not be ingested into foreign training corpora, fine-tuned upon, transformed, derived from, or commercialised by AI systems without the free, prior, and informed consent of the relevant communities and the explicit authorisation of the relevant national governments. This Pillar governs what AI systems may take from Global South communities.

Pillar 6 - Agentic AI Liability Anchoring

The Architecture squarely addresses the contentious question of liability for agentic artificial intelligence. The Global Agentic AI Liability Framework is incorporated by reference. Agentic AI systems, capable of autonomous goal-directed action across digital and physical environments, shall not operate within Global South jurisdictions in a state of legal weightlessness. Liability shall attach across the chain of design, deployment, integration, fine-

tuning, and operational use, with principal liability anchored upon those who derive the principal economic benefit and exercise the principal control.

Pillar 7 - Protection of Indigenous Communities and Cultural Practices from Algorithmic Harm

Artificial intelligence systems shall not be deployed, instrumentalised, weaponised, or operated in a manner that destroys, dilutes, distorts, or prejudicially impacts indigenous communities, tribal populations, or traditional cultural practices of the Global South. AI shall not function as an engine of cultural erasure, linguistic displacement, or epistemic subordination upon communities whose worldviews predate the algorithmic age. This Pillar governs how AI systems may operate upon and around Global South communities.

V. Novel and Innovative Elements - The New Building Blocks

The Architecture introduces, for the first time, the following innovative elements as new building blocks for the Global South Group of Nations. These elements are original contributions intended to extend the doctrinal frontier and to differentiate the Architecture from instruments already in circulation.

- **Global South Algorithmic Sovereignty Index.** An annual, comparative, publicly published index measuring the degree to which each participating Global South jurisdiction has operationalised the seven pillars in domestic law, regulation, and enforcement practice.
- **Inter-Global South AI Mutual Legal Assistance Compact.** A standing mechanism for cross-border cooperation in AI-related investigations and enforcement, elaborated in full at Section VI below as a designated early deliverable of the Architecture.
- **Civilizational AI Audit.** A dedicated audit standard examining whether an AI system, by design, training data composition, or output behaviour, has been calibrated to respect the civilizational, linguistic, and cultural traditions of the deploying jurisdiction.
- **Doctrine of Algorithmic Non-Alignment.** A foreign-policy posture for Global South nations: the right not to be conscripted, by virtue of platform dependence or model dependence, into the AI-strategic competitions of dominant powers.
- **Right to Algorithmic Refusal.** A collective right vested in indigenous and tribal communities to refuse the deployment of AI systems upon their data, lands, languages, or cultural production, with statutory effect.
- **Reverse Burden of Proof for Foreign AI Providers.** In disputes concerning harm caused by AI systems originating outside the deploying Global South jurisdiction, the burden of demonstrating safety, lawful data provenance, and absence of cognitive or cultural harm shall rest with the foreign provider.

- **Sovereign Data Repository Principle.** Each Global South nation shall be entitled to designate categories of national, cultural, or indigenous data that may not be exported, ingested into foreign training corpora, or used for foreign fine-tuning without sovereign authorisation.
- **Global South AI Harm Notification Protocol.** A standing inter-jurisdictional protocol for the rapid notification, classification, and coordinated response to material AI-induced harms - whether informational, economic, electoral, cultural, or physical.

VI. The Inter-Global South AI Mutual Legal Assistance Compact

The enforcement teeth of the Architecture, particularly against extraterritorial AI providers, depend upon collective action. No Global South nation acting alone possesses the market-access leverage necessary to compel compliance from frontier AI developers operating from foreign jurisdictions. Collective leverage, exercised through harmonised legal cooperation, can. The Inter-Global South AI Mutual Legal Assistance Compact (the “MLA Compact”) is therefore designated as an early deliverable of the Architecture, to be developed and opened for signature within twelve months of adoption.

The MLA Compact shall establish, at minimum: (i) reciprocal recognition by participating Global South States of one another's AI regulatory determinations, audit findings, and enforcement orders; (ii) standardised cross-border information-sharing protocols for AI incident investigations, including agentic AI incidents involving cross-jurisdictional harm; (iii) a harmonised audit methodology, drawing upon the Civilizational AI Audit standard, that permits a single audit conducted in one participating State to be recognised as evidence in another; (iv) coordinated enforcement protocols against extraterritorial AI providers, including the option of joint regulatory action and synchronised market-access conditions; (v) mutual technical assistance, including capacity-sharing arrangements for smaller Global South jurisdictions that lack independent technical audit capability; and (vi) a standing inter-State liaison mechanism, hosted by the Architecture's permanent secretariat, to facilitate ongoing cooperation.

The MLA Compact is the instrument through which the seven pillars become enforceable beyond the borders of any single State. Without it, sovereignty over algorithms remains aspirational; with it, the Global South acquires, for the first time, a collective enforcement architecture proportionate to the scale of the actors it must regulate.

VII. Capacity Building, Awareness, and Implementation

Law without capacity is aspiration without execution. The Architecture commits the Global South Group of Nations to a sustained, structured programme of capacity building across judiciaries, regulators, parliamentarians, civil servants, journalists, civil society actors, and

educational institutions. Public awareness - in the languages of the people, not merely the language of the law - is the indispensable bedrock upon which algorithmic sovereignty becomes lived experience.

Implementation of the New Delhi Accord on Artificial Intelligence Law and Governance, the New Delhi Compact on Artificial Intelligence Accountability, the AI Accountability Framework 2026, and the Global Agentic AI Liability Framework shall be actuated through the three-layer Architecture set out herein, supported by model legislative templates, model regulatory protocols, and shared training curricula developed within the Working Groups.

VIII. The Working Group Architecture

The Architecture establishes two complementary Working Groups. They are distinct in composition, complementary in mandate, and jointly constitute the institutional engine of the Architecture.

Working Group on Global South Algorithmic Sovereignty

Composition: the initial members of this Working Group shall be drawn from the speakers, domain experts, scholars, and practitioners participating in the Second Global South Artificial Intelligence Law and Governance Dialogue 2026. This Working Group will operate as the doctrinal, scholarly, and technical engine of the Architecture, charged with the elaboration of model statutory text, model audit protocols, the Global South Algorithmic Sovereignty Index methodology, the Civilizational AI Audit standard, and the technical annexes of the MLA Compact.

Working Group on Global South AI Governance

Composition: this Working Group shall be a State-anchored body. Formal invitations shall be extended to government ministries, ministers, and designated nodal officers holding the AI, digital, technology, law, or data protection portfolio across all nations of the Global South Group. Its mandate is to translate the doctrinal output of the first Working Group into governmental recommendations, harmonised national positions, the opening for signature of the MLA Compact, and multilateral submissions including submissions to the United Nations preparatory processes.

A standing sub-track on South Asia within the broader Global South Group of Nations shall be constituted within the second Working Group. Regional sub-tracks may further be constituted for Africa, Latin America and the Caribbean, Southeast Asia, and Small Island Developing States, with seat distribution to be finalised at the inaugural session of the second Working Group.

IX. Founding Endorsement and Signatory States

Recognising that the durable power of any Global South instrument rests upon the visibility of State endorsement, this Architecture is opened for adoption for Global South nations.

The Architecture remains open for endorsement, signature, accession, and national adoption by all States of the Global South Group of Nations, and by relevant regional and multilateral institutions of the Global South. Subsequent State endorsements shall be recorded by the permanent secretariat established under Section X, and shall be published periodically as the Architecture's expanding coalition of Signatory States.

X. Ratification and Adoption -The Online Dialogue Pathway

The Architecture has been formally presented and adopted at the Second Global South Artificial Intelligence Law and Governance Dialogue 2026, conducted as an online convening. To preserve the legitimacy of the outcome and to ensure that the Architecture is referable as a genuine instrument of a genuine event, the following ratification mechanism was followed:

- Plenary deliberation at the Second Dialogue, with on-record interventions by participating delegates, scholars, regulators, and civil society representatives.
- Formal adoption by acclamation at the plenary.
- There would be publication of the adopted Architecture as the operative reference instrument of the Second Global South AI Law and Governance Dialogue 2026, with an open invitation for subsequent endorsement by Global South States, regulators, ministries, and institutions.
- It is aimed to work towards the establishment of a permanent digital secretariat to maintain the Architecture as a living instrument, to publish periodic updates, to administer the work of the two Working Groups, and to host the MLA Compact liaison function.

XI. Pathway to the United Nations AI Scientific Model Preparatory Process

The deliberations and outputs of the Architecture, including the doctrinal output of the Algorithmic Sovereignty Working Group, the governmental output of the AI Governance Working Group, and the MLA Compact text, shall be channelled into the United Nations preparatory process for the global AI scientific model and related multilateral AI governance instruments. The ambition, however, is wider: the Architecture is to be operative for the Global South Group of Nations whether or not it is absorbed by or referenced to by other multilateral processes. Its primary loyalty is to the constituencies it serves.

XII. Towards a Global South Model Law on AI Accountability

A defined and time-bound deliverable of the two Working Groups, jointly, shall be a Global South Model Law on Artificial Intelligence Accountability (the “Model Law”). The Model Law shall offer a ready, adaptable statutory text - covering definitions, regulatory architecture, duties of AI developers and deployers, rights of affected persons and communities, agentic AI liability, audit, redress, enforcement, and inter-State cooperation, that any Global South legislature may adopt, customise, and enact as its national statute. The Model Law shall consciously reflect the seven pillars and the novel building blocks set out in this Architecture.

Drafting Timeline - Preferable Nine Months from Adoption

The Model Law shall be aimed and encouraged to be delivered on a preferable nine-month timeline from the date of adoption of this Architecture, organised as follows:

- **Months 1 to 3 - Scoping Phase.** Establishment of the Drafting Committee; finalisation of the Model Law outline, table of contents, and section-by-section drafting allocation; consultative round with the Global South States; publication of a public consultation paper.
- **Months 4 to 6 - First-Draft Phase.** Production of the first complete draft of the Model Law; internal review by both Working Groups; expert technical review; release of the first draft for public and State consultation.
- **Months 7 to 9 - Finalisation Phase.** Incorporation of consultation feedback; legal review; finalisation of the consolidated Model Law text; formal publication and release for adoption by Global South legislatures.

Composition of the Model Law Drafting Committee

The Drafting Committee shall be a joint body drawing from both Working Groups, configured as follows:

- **Chair:** to be selected by consensus at the inaugural meeting of the Drafting Committee.
- **Rapporteur:** drawn from the Working Group on Global South Algorithmic Sovereignty and designated by that Working Group, responsible for the consolidation and drafting of the Model Law text.
- **Members:** a balanced membership of not fewer than fifteen and not more than twenty-one persons, with regional distribution across Africa, Latin America and the Caribbean, South Asia, Southeast Asia, the Pacific, and Small Island Developing States, and balanced expertise across constitutional law, AI law, data protection, indigenous rights law, technology audit, and parliamentary drafting.

The Drafting Committee shall be constituted within thirty days of the adoption of this Architecture, and its work shall be supported by the permanent digital secretariat established under Section X.

XIII. The Hopes and Aspirations of the Citizenry of the Global South

The Architecture exists, finally, for the citizen, the farmer whose crop pricing is increasingly shaped by algorithmic intermediaries, the student whose learning is being mediated by foreign models, the patient whose diagnosis is being processed through systems trained on data that does not look like them, the artisan whose craft is being scraped into training corpora, the journalist whose work competes with synthetic content, the voter whose information environment is being algorithmically curated, and the child whose worldview is being calibrated by recommendation systems answerable to no one in their own country.

For each of them, the Architecture affirms a foundational promise: that the algorithms which shape your life shall be subject to the laws of your land, the values of your civilization, the consent of your community, and the dignity of your person. This is the benchmarking standard of the Architecture, and it is the benchmark against which every subsequent instrument, every deployment, and every contestation shall be measured.

XIV. Closing Imperatives

The Architecture is not a final text; it is a foundational scaffold. It will be tested, refined, and extended by the work of the two Working Groups, by the MLA Compact opened for signature, by the Model Law preferably to be delivered within nine months, and by the lived experience of the citizens it serves. It is offered to the Global South Group of Nations as a common minimum denominator compact - practical, ratifiable, evolvable, and as the foundation for a Global South Model Law on AI Accountability.

Algorithmic sovereignty, the prevention of cognitive colonialism, the protection of indigenous and cultural heritage, the prohibition of unconsented data distillation, the assertion of compliance over Big Tech, the anchoring of agentic AI liability, and the dignity of the Global South citizen -these are not aspirations. Under this Architecture, they are commitments.

***Adopted at the Second Global South AI Law and Governance Dialogue 2026,
on this 28th day of May, 2026, at Virtual Convening.***