

16TH ITU ACADEMIC CONFERENCE


ITUKALEIDOSCOPE

GENEVA2026

*AI and frontier
technologies for good*

7-9 July 2026
Geneva, Switzerland

Hosted by  **AI for Good**
Global Summit




THE AMPLIFIED SOVEREIGNTY PARADOX

AI Governance Lessons from Small Island Developing States

7 July 2026

The core claim.

SIDS are not marginal cases in AI governance. They are **stress tests** – exposing the trade-offs that scale, money, and institutions let wealthy states hide.





Mohamed Shareef

Nexia Maldives

Session #1.1



A governance blind spot shaped by high-capacity assumptions

Three actors set the global agenda and share one limitation.

The EU's rights-based model, the US market model, and China's state model differ in philosophy – but their implementation assumptions reflect large, high-capacity states [4–7]. SIDS face the same AI risks with thinner institutions, smaller markets, and deeper vendor dependence.

It is not a policy gap. It is a sovereignty gap.

Algorithmic transparency stays weak when states lack audit access, technical capacity, or vendor cooperation — so the practical choice often narrows to accepting vendor terms or limiting citizens' access.

THE CAPACITY GAP, IN NUMBERS

1-5

specialist staff in many SIDS data-protection or digital regulators – versus much larger teams at major regulatory authorities


39

Small Island Developing States, still under-theorized in AI governance

71.2%

of SIDS face AI data-quality constraints (UNESCO, 2024) [8]

The Amplified Sovereignty Paradox



Strengthen one dimension and another tends to weaken
— unless compensating safeguards are built in.

AI turns the old two-way sovereignty paradox into a three-way trilemma — through two amplifiers:

AMPLIFIER 1 · Surveillance Automation

AI shifts surveillance from a latent capacity to an always-on governance layer — real-time facial recognition and edge inference 24/7. *Review stops being a safeguard once the system moves faster than the regulator.*

AMPLIFIER 2 · Structural Vendor Dependence

Foundation-model-scale AI (10^{23} – 10^{24} FLOPs) sits beyond domestic compute, forcing dependence on foreign platforms for advanced services; model immobility pushes data offshore; lock-in creates “data gravity.” *Sovereignty on paper does not survive when the technical substrate sits elsewhere.*

Why SIDS are structural stress tests

Geographic fragmentation

Maldives: 1,190 islands over 90,000 km². Fiji: 330 islands. Sovereign hosting is possible in theory, but costly and fragile to sustain across dispersed islands – pushing governments toward cloud, hybrid, or vendor-managed infrastructure.

Climate vulnerability

AI is both survival tool and sovereignty risk. The minister's dilemma: a foreign black-box cyclone model, or no early-warning system at all. Governance principles are tested by survival imperatives.

Thin institutional capacity

DPA's with very small technical teams cannot run EU-style bureaucracies – making SIDS the ideal testbed for lightweight, agile governance that scales up, not down.

Why this matters for theory. In larger states, spare capacity hides the paradox. In SIDS, the same trade-offs become operational and visible –which is exactly what makes them early-warning systems for everyone else.

METHOD: 14 SIDS strategies reviewed · practitioner consultations (Maldives, Tonga, PNG) · three evidence-grounded vignettes [23–33] → 5 RECURRING PATTERNS

1 Surveillance amplification

2 Structural vendor dependence

3 Geopolitical vulnerability

4 Procurement impossibility

5 Emergency bypass

Five standards-grade requirements

INDIVIDUAL SOVEREIGNTY

R1 Decision Transparency

Know when and how an automated decision was made.

Tamper-evident logs + citizen portal ≤24h

R4 Human Redress

Preserve the right to human review and override.

Review median ≤72h; officer override

NATIONAL SOVEREIGNTY

R2 Data Residency & Escrow

Reduce hollow legal sovereignty over citizen data.

In-jurisdiction storage + model escrow

R5 Procurement Transparency

Reveal vendor and supply-chain dependencies.

Machine-readable SBOM (SPDX/CycloneDX)

TECHNICAL SOVEREIGNTY

R3 Model Change Notice

Stop silent model updates from becoming invisible policy.

MCN ≥30 days; fairness ΔTPR ≤5pp

Why R3 matters:

Full algorithm access is rarely attainable – R3 preserves accountability through disclosure of change, not source code.

One adoption ladder, every requirement.

CC-A baseline compliance → CC-B citizen-facing + SLAs → CC-C automated verification & public registries. Meet A today; build to C as capacity grows.

Honest about maturity. R1 + R4 are near-term, deployable as lightweight overlays. R5 needs vendor integration. R2 + R3 are structural – achievable as regional cloud options, market power, and procurement leverage improve.

What ITU-T can do next

1

Create a cross-SG work item

“AI Governance for Resource-Constrained Environments” — coordinating SG13, SG16, SG17, SG20.

2

Pilot the lightweight requirements

Field-test R1 (Decision Transparency) and R4 (Human Redress) — the most immediately deployable requirements.

3

Back regional sovereignty stacks

Pool audit capacity, procurement leverage, and technical assurance across CARICOM, PIF, and the Indian Ocean Commission.

Standards alone will not solve this. A three-person regulator may receive the standard but cannot meaningfully audit a live neural network without pooled support. The answer is not only national capacity — it is **pooled regional capacity**.

Pathway: near-term profiles (2027) · interop testing + Apache-2.0 reference code (2027–28) · SIDS Governance Profile Pack + ISO/IEC JTC 1 liaison (2028–29). Full Study-Group mapping in backup.




Three contributions, one frontier

- 1 Names the problem – the Amplified Sovereignty Paradox.
- 2 Shows why SIDS are not edge cases – but early-warning systems for everyone.
- 3 Turns critique into standards – R1–R5 the ITU-T can act on.

Thank you!

*The aim is not to lower the ambition of AI governance
— it is to make it work where capacity is scarce and dependence is real.*

*“The unseen frontier of AI governance lies not in Silicon Valley
or Brussels – but in Bridgetown, Malé, and Kingston.”*



ITU-T Study Group mapping (backup)

Req	Study Group	Proposed work item	Deliverable	
R1	SG17/Q6	Tamper-evident logging for AI-driven identity systems	Recommendation (X)	Q2–Q4 2027
R2	SG13/Q4	Data residency & sovereignty for gov cloud — SIDS profile	Tech Report + Conformance Profile	Q2–Q4 2027
R3	SG16/Q26-27 · SG17/Q6	AI model-change management: fairness testing & disclosure	Recommendation (F/X)	Q3 2027–Q1 2028
R4	SG20/Q2 · SG16/Q26	Explainability & human review for public-service decisions	Tech Report + Guidelines	Q2–Q3 2027
R5	SG17/Q6 · SG13/Q4	Software Bill of Materials (SBOM) for AI procurement	Procurement Standard + Schema	Q1–Q3 2027

Coordination: cross-SG group with quarterly meetings; liaison to W3C CCG and ISO/IEC JTC 1 SC 27/37/42. Open-source reference code (Apache 2.0): Decision-Log Module, Redress Workflow, SBOM Validator, Test Harness Suite.