

THE AMPLIFIED SOVEREIGNTY PARADOX: AI GOVERNANCE LESSONS FROM SMALL ISLAND DEVELOPING STATES

Mohamed Shareef

Nexia Maldives

ABSTRACT

Global AI governance frameworks designed for wealthy nations fail in contexts where technology trade-offs are existential rather than regulatory. The question is not whether these frameworks are well-intentioned but whether they assume state capacities that simply do not exist. Our contribution targets this structural mismatch. First, we introduce the "Amplified Sovereignty Paradox" (ASP), a framework showing how AI transforms binary sovereignty tensions into a three-dimensional trilemma involving individual rights, national autonomy, and technical controllability. Second, we demonstrate why Small Island Developing States (SIDS) are not edge cases but structural stress tests for global governance. Their constraints make invisible trade-offs explicit. Third, we translate these insights into five standards-grade requirements (R1-R5). These include tamper-evident decision logging and procurement transparency, complete with conformance classes and KPIs. These are mapped to specific ITU-T Study Groups (SG13, SG17, SG16, SG20) for 2027-2029 timelines. The practical consequence is clear. Solutions developed for SIDS offer governance models for the global majority because they solve for constraints that exist everywhere, even if they remain hidden in the world's wealthiest capitals.

Keywords – AI governance, digital public infrastructure, digital sovereignty, ITU-T standards, small island developing states

1. INTRODUCTION

The ITU's "AI for good" vision confronts a fundamental challenge: global AI governance frameworks designed for wealthy nations fail to address contexts in which sovereignty trade-offs are most acute [1-3]. While the EU, US, and China advance competing regulatory models [4-7], a critical research gap remains. The 39 SIDS nations, despite UNESCO findings that 71.2% lack AI data [8] and UNDP recognition as "digital trailblazers" [10, 11], remain under-theorized in AI governance literature.

This gap matters. The constraints in SIDS turn AI governance from a regulatory exercise into an existential trilemma. With 1-5 technical staff in data protection

authorities and mandatory digital ID systems, SIDS cannot implement frameworks assuming hundreds of regulators and optional digital services. The 2024 Antigua Agenda highlights digital transformation as essential for SIDS [9], but no framework explains how AI alters their sovereignty calculus.

This paper addresses three research questions: (RQ1) How does AI transform digital sovereignty challenges into qualitatively different problems for resource-constrained states? (RQ2) What makes SIDS critical cases for understanding AI governance trade-offs? (RQ3) How can these challenges be operationalized for ITU-T standardization?

Our contribution moves the debate from normative principles to structural realities. First, we introduce the Amplified Sovereignty Paradox (ASP), which redefines sovereignty not as a binary tension but as a trilemma of competing constraints, individual rights, national autonomy, and technical control. Second, we ground this theory in the operational experience of SIDS (2016–2025). We argue these contexts are not anomalies to be ignored but early warning systems for the global governance architecture, revealing dependencies that remain invisible in larger economies. Third, we advance from critique to codification. We operationalize these insights into five standards-grade requirements (R1-R5) with specific KPIs, offering the ITU-T a concrete roadmap for 2027–2029. We provide not just a diagnosis of the trap, but the specifications for the exit.

Section 2 analyzes the governance blind spots in existing frameworks. Section 3 presents the ASP framework with its three dimensions and two amplification mechanisms: Section 4 establishes SIDS as critical cases through evidence-grounded vignettes. Section 5 operationalizes the requirements for the ITU-T. Finally, Section 6 concludes with actionable recommendations.

2. THE GOVERNANCE BLIND SPOT

The contemporary AI governance debate is defined by three geopolitical actors with a shared limitation: their models are products of, and designed for, large, wealthy, and powerful states. For small states, the EU model may be normatively

appealing but practically impossible due to a lack of regulatory capacity. The EU AI Act envisions conformity assessments and monitoring systems, which many SIDS cannot afford. A typical SIDS data protection authority operates with 1-5 technical staff members, compared to EU GDPR supervisory authorities. The question is not whether these frameworks are well-intentioned but whether they assume a state capacity that simply does not exist in SIDS.

The digital sovereignty challenge is particularly acute for SIDS. Recent analysis from ISO/IEC JTC 1/SC 27 (cybersecurity) and SC 37 (biometrics) shows that many low and middle-income countries depend on concentrated biometric algorithm vendors, with billions of digital IDs issued globally [12, 13]. European initiatives, such as EuroStack, highlight concerns about dependence on US technology for identity management, payments, and cloud infrastructure [14]. These concerns amplify manifold for SIDS, with even less bargaining power and fewer alternative vendors.

These governance models assume that institutional capacities are unavailable to SIDS. The NIST AI Risk Management Framework requires "AI system impact assessments" and "governance structures" that presuppose dedicated teams [38]. UNESCO's Recommendation on AI Ethics emphasizes "algorithmic transparency" which becomes meaningless without source code access [40]. The OECD AI Principles promote 'human-centered AI' while SIDS face a binary choice: accept vendor terms or exclude citizens from digital services [41].

3. CONCEPTUAL FRAMEWORK: THE AMPLIFIED SOVEREIGNTY PARADOX

3.1 The three sovereignty dimensions




Figure 1 – The Amplified Sovereignty Paradox (ASP) trilemma

We propose that AI-driven Digital Public Infrastructure (DPI) creates trade-offs across three distinct sovereignty dimensions (see Figure 1). These dimensions do not compete in theory alone. They collide in every major DPI deployment across SIDS.

National sovereignty (state autonomy): The capacity of states to determine digital policies, maintain independence from foreign technology providers, and exercise meaningful control over critical digital infrastructure.

Technical sovereignty (system controllability): The capacity to audit, verify, modify, and understand algorithmic systems deployed within a jurisdiction, including access to model parameters and decision logic.

3.2 From paradox to trilemma: The amplification mechanism

The original sovereignty paradox identified how state-provided digital infrastructure, necessary for citizen empowerment, simultaneously creates an unprecedented surveillance architecture [15, 16]. AI transforms this two-way tension into a three-dimensional trilemma through two amplifiers:

Amplifier 1: Surveillance automation. AI systems transform surveillance from a latent capacity into an active, preemptive governance layer. GPU-accelerated facial recognition operates at 30fps, evaluating citizens in real time rather than retrospectively [17]. Transformer models identify behavioral anomalies across multimodal streams that exceed human analytical capacity. Edge AI inference enables 24/7 monitoring without human supervision. Together, these mechanisms automate what Zuboff [3] terms "extraction asymmetry," creating a governance environment in SIDS where the velocity of algorithmic decision-making outstrips the state's capacity for human review. In practice, review is no longer a safeguard when the system moves faster than the regulator.

Amplifier 2: Structural vendor dependence. In the SIDS context, reliance on foreign tech providers creates a "jurisdictional hollow" in which technical reality overrides national law across three dimensions. The compute sovereignty gap, defined by foundation model training requirements of 10^{23} - 10^{24} FLOPs, renders domestic AI development economically impossible, forcing the outsourcing of critical state functions. Model immobility arises from proprietary APIs that prevent on-premise deployment; this forces data export to foreign clouds, compelling regulators to tacitly ignore their data residency laws to maintain service continuity. Algorithmic lock-in occurs when fine-tuned models create deep path dependencies. This "data gravity" anchors sovereignty to the vendor, meaning the state technically retains the right to govern but loses the material capacity to switch providers. Sovereignty on paper does not survive when the technical substrate is elsewhere.

3.3 The ASP trilemma: Structural trade-offs

The ASP describes a structural trilemma. AI-driven DPI adoption to enhance public services simultaneously creates unprecedented automated surveillance capabilities, diminishes individual sovereignty, and engenders critical dependencies on foreign technology vendors, diminishing national and technical sovereignty.

Consider the impossible procurement specification for a SIDS: "Provide a biometric facial recognition system that enables seamless citizen access, does not create surveillance capacity, uses algorithms we can fully audit, and costs less than \$500K." No vendor can fulfill this requirement, and SIDS know this before the tender is even drafted.

The trilemma manifests differently depending on which dimension is prioritized. Prioritize individual sovereignty and reject biometric AI systems; citizens are excluded from digital services. Prioritize national sovereignty and demand data residency with algorithm escrow; vendors refuse to bid or charge prohibitive premiums. Prioritize technical sovereignty and require open-source algorithms with model transparency; no commercial vendors meet requirements. The trade-off is structural, not temporary. It does not improve with scale or better documentation.

4. SIDS AS GOVERNANCE LABORATORIES

4.1 Why SIDS are critical cases for AI governance

SIDS are critical cases for AI governance because unique constraints magnify the ASP, making trade-offs clear and the stakes existential.

Geographic fragmentation: Archipelagic nations exemplify the constraints of the ASP. The Maldives spans 1190 islands across 90000 km². Fiji comprises 330 islands. Samoa includes 10 islands across 2831 km². Caribbean states, such as Barbados and Saint Lucia, face similar challenges. Mauritius's digital ID covers 1.3 million citizens across multiple islands, requiring cloud-based biometrics. Fiji's Walesi platform serves 900000 people across scattered islands through a satellite-dependent infrastructure. These nations cannot maintain on-premise data centers on each island, forcing them to rely on foreign cloud providers. This creates a sovereignty trilemma, where digital inclusion requires surveillance infrastructure and foreign dependencies.

Climate vulnerability: For SIDS facing existential climate threats, AI is simultaneously a survival tool and a sovereignty risk. Climate modeling, cyclone early warning systems, and sea-level monitoring require sophisticated AI systems [11, 19]. This dramatically raises the stakes. What degree of sovereignty is worth trading for climate-adaptation capacity? When a government minister must choose between a foreign vendor's proprietary hurricane prediction model and no early warning system at all, abstract principles about algorithmic transparency give way to pragmatic calculations about citizen survival.

Institutional capacity constraints: SIDS cannot afford the sprawling regulatory agencies envisioned by frameworks such as the EU AI Act [4]. Most SIDS data protection authorities have 1-5 technical staff members. This capacity constraint makes them crucial testbeds for lightweight and agile governance models that function without extensive bureaucracy. UNESCO's Caribbean AI Policy Roadmap (2024) offers a framework that specifically addresses these constraints through regional cooperation, graduated implementation, and vendor accountability mechanisms [20].

These realities are not edge cases. They are the operating conditions under which digital governance must function.

4.2 Methodological approach

This analysis employed a mixed-methods approach. We systematically reviewed 14 SIDS digital transformation strategies (2020-2024) from Caribbean, Pacific, and AIMS regions, coding them for sovereignty assertions and vendor relationships [21]. We synthesized the ASP framework by building on surveillance theory [16] and digital colonialism literature [21]. Drawing on extensive senior-level practitioner experience in SIDS contexts, validated through consultations with digital leaders from Maldives, Tonga, and Papua New Guinea, we developed three evidence-grounded vignettes synthesized from documented patterns in SIDS digital identity initiatives [23-33]. These vignettes serve as analytical tools to illustrate governance challenges while protecting specific national contexts.

To be clear about what this analysis is and is not: the review of 14 national digital strategies constitutes our empirical baseline; practitioner consultations with digital leaders ground the operational insights; and the amplified sovereignty paradox is the theoretical abstraction that results. Together, these form a conceptual proposal intended to guide future ITU-T standardization work.

4.3 Evidence: Three vignettes demonstrating the ASP

To illustrate the ASP, we present three composite vignettes synthesized from documented patterns in SIDS implementations [23-33]. The purpose here is not to name countries. It is to show patterns that are visible to anyone working in this space.

VIGNETTE A: Geopolitically-sensitive procurement

A SIDS nation procured a World Bank-funded (\$10M) national digital ID mandatory for 300+ services. The system, involving facial recognition and liveness detection, was initially awarded to a UK-based vendor with geopolitically-sensitive ownership. Contract terms were not disclosed, and political pressure later forced a vendor switch. With a DPA staff of only three and limited audit capacity, the state faces surveillance amplification via always-on biometrics and a "black box" contract that limits oversight [26-33].

VIGNETTE B: Emergency governance bypass

During the COVID-19 pandemic, a SIDS government accelerated a biometric ID rollout to achieve 90% coverage, integrating contact-tracing analytics. The system was hosted by a European cloud provider with offshore processing. Legislation was passed during lockdown with minimal consultation, bypassing a DPA that lacked enforcement power. The result was a permanent surveillance infrastructure built under an emergency pretext, with data residency laws tacitly ignored to maintain service [23-25, 30].

VIGNETTE C: Constitutional and procurement challenges

A national biometric ID pilot launched in 2022 faced "procurement impossibility." An initial 2016 proposal with a foreign vendor was struck down as unconstitutional in 2019. The government revised the model in 2021 to a hybrid approach – keeping biographic data local while hosting facial algorithms foreign-hosted due to technical constraints. This compromise maintains foreign algorithm dependency, proving that sovereignty on paper often yields to technical reality [22, 23, 34-37].

Note: These vignettes are analytical constructs synthesized from documented patterns across multiple SIDS to illustrate the ASP's operation while protecting specific national contexts.

4.4 Synthesis: What the vignettes reveal

These vignettes reveal five recurring ASP patterns.

- (1) **Surveillance amplification:** AI biometrics create persistent monitoring infrastructure (Vignettes A, B).
- (2) **Structural vendor dependence:** The lack of domestic AI capacity forces reliance on foreign providers (all vignettes).
- (3) **Geopolitical vulnerability:** Technology procurement becomes entangled with foreign policy (Vignette A, [27, 32]).
- (4) **Procurement impossibility:** SIDS must often choose between accepting vendor terms or having no system (Vignette C).
- (5) **Emergency bypass:** Crises (COVID-19, climate events) accelerate AI adoption while bypassing governance safeguards (Vignette B).

The sequence varies by country but the logic does not.

5. OPERATIONALIZING THE ASP: STANDARDS REQUIREMENTS

5.1 Five standards-grade requirements

Drawing from the W3C VC [34], ISO/IEC 5259 [29], NIST AI RMF [38], and ISO/IEC 42001 [30], we propose five core requirements (per RFC 2119 [31]):

To provide a consistent adoption pathway for resource-constrained governments, conformance classes across all five requirements follow a unifying three-tier progression. Class A establishes baseline internal compliance: the minimum viable level of accountability. Class B introduces citizen-facing transparency and defined service-level agreements, representing intermediate maturity. Class C mandates fully automated cryptographic verification, public registries, or structural sovereignty measures, representing advanced maturity. The logic is deliberate: governments can meet Class A requirements with existing infrastructure today, then build toward Class C as institutional capacity grows.

R1 – Individual sovereignty: Decision Transparency

Scope: AI systems that make automated eligibility decisions (benefits, healthcare, and identity).

Normative requirement: Systems MUST produce tamper-evident append-only logs of data access, SHALL provide a citizen-accessible portal within 24h, MUST publish daily cryptographic anchors to the public ledger.

Conformance classes: CC-A (logging only), CC-B (+ citizen portal + weekly anchor), and CC-C (+ daily anchor + audit API).

Test assertions: TA-R1-01: 100 test cases → 100 log entries; TA-R1-02: Verify the Merkle proof for 50 entries.

KPI: Log availability $\geq 99.5\%$; portal response $\leq 2s$ (95th percentile).

The intent is straightforward: citizens should not enter black boxes to access essential services.

R2 – National sovereignty: Data residency & escrow

Scope: Procurement of AI-driven DPI systems.

Normative requirement: Raw citizen personal data SHALL be stored in the jurisdiction. Model artifacts MUST be escrowed with verified release conditions. Quarterly attestation reports are required.

Conformance classes: CC-A (biographic only), CC-B (biographic + biometric; annual), and CC-C (full residency; quarterly; automated verification).

Test assertions: TA-R2-01: Validate JSON attestation schema; TA-R2-02: Verify escrow manifest hashes.

KPI: 100% data residency compliance; escrow verification $\geq 99\%$.

This requirement forces a conversation that most SIDS have postponed for a decade.

R3 – Technical sovereignty: Model change notice

Scope: AI systems in which model updates materially affect eligibility outcomes.

Normative requirement: Deployers MUST publish Model Change Notice (MCN) ≥ 30 days before significant changes, SHALL include the version delta, signed hash, fairness regression results ($\Delta TPR \leq 5pp$), and rollback conditions.

Conformance classes: CC-A (14-day internal notice), CC-B (30-day notice; basic public disclosure), and CC-C (30-day; 5+ strata; public registry; automated third-party verification).

Test assertions: TA-R3-01: Verify ≥ 30 -day median; TA-R3-02: Confirm $\max(|\Delta \text{TPR}|) \leq 0.05$.

KPI: 100% MCN coverage; mean notice ≥ 35 days; fairness pass $\geq 95\%$.

Without this, model updates become invisible policy changes.

R4 – Individual sovereignty: Human redress

Scope: All automated decisions affecting citizens' access to services.

Normative requirement: Automated decisions MUST include an explainability artifact (outcome, top three factors, review officer contact, decision ID) within 24h. Human review median ≤ 72 h, 95th percentile ≤ 120 h. Officers MUST have override authority.

Conformance classes: CC-A (binary outcome; no SLA), CC-B (top three factors; ≤ 72 h median), and CC-C (+ interactive explainability; ≤ 48 h guarantee).

Test assertions: TA-R4-01: Validate 100 artifacts for three factors; TA-R4-02: Compute the median from 200 cases.

KPI: 100% artifact coverage; median ≤ 72 h; 95th percentile ≤ 120 h; Satisfaction $\geq 75\%$.

The intent is simple: automation cannot be allowed to remove a citizen's right to be heard by a human.

R5 – National sovereignty: Procurement transparency

Scope: AI procurement for the DPI.

Normative requirement: Vendors SHALL provide a machine-readable software bill of materials (SPDX/CycloneDX) with training data sources, subprocessors, and jurisdictions. Quarterly updates.

Conformance classes: CC-A (annual; non-public), CC-B (quarterly; SBOM; notifications), and CC-C (+ public summaries; vulnerability scanning).

Test assertions: TA-R5-01: Validate the SBOM schema; TA-R5-02: confirm notification $\leq$ seven days.

KPI: 100% SBOM coverage, quarterly updates, disclosure ≤ 90 days, and vulnerability ID $\geq 95\%$.

The intent is clear: governments cannot govern what they cannot see.

5.2 Mapping requirements to sovereignty dimensions

The five requirements systematically address the three dimensions of ASP:

- **Individual sovereignty** is protected by R1 (decision transparency) and R4 (human redress), ensuring that citizens can understand and contest automated decisions.
- **National sovereignty** is protected by R2 (data residency), which ensures that data remain in the jurisdiction, and R5 (SBOM), which provides visibility of vendor dependencies and supply chain risks.
- **Technical sovereignty** is protected by R3 (model change notice), which maintains some algorithmic accountability even when full algorithm access is unattainable.

To reflect operational readiness accurately, we categorize these requirements by maturity tier. Requirements R1 (decision transparency) and R4 (human redress) are designed to be near-term deployable as lightweight software overlays. Requirement R5 (SBOM) requires initial vendor integration testing before wider adoption. Requirements R2 (data residency) and R3 (model change notice), by contrast, are aspirational targets: their implementation depends on structural shifts in regional cloud infrastructure and further market evaluation.

5.3 ITU-T Study Group mapping

The ITU-T Work Items with deliverables and timelines are summarized in Table 1.

Standardization roadmap

Near-term (2027): Conformance profiles and implementation guidelines for R1, R4, and R5 in SG17, SG20, and SG13.

Mid-term (2027-2028): Interoperability testing, reference implementation, and evaluation through a cross-SG coordination group.

Long-term (2028-2029): SIDS governance profile pack combining all five requirements; liaison with ISO/IEC JTC 1 for harmonization with AI management system standards (ISO/IEC 42001) and data quality standards (ISO/IEC 5259).

Coordination mechanism: Cross-SG group "AI Governance for Resource-Constrained Environments" with quarterly meetings and liaison to W3C CCG, ISO/IEC JTC 1 SC 27/37/42.

5.4 Implementation pathway

We propose a phased implementation approach that recognizes the capacity constraints of SIDS while enabling progressive adoption.

Table 1 – ITU-T Work Items with deliverables and timelines

Req.	ITU-T SG/Question	Proposed Work Item Title	Deliverable Type	Timeline
R1	SG17/Q6 (Identity management, tele-bio-metrics)	Tamper-evident logging for AI-driven identity systems: Requirements and test procedures	ITU-T Recommendation (X-series)	Q2-Q4 2027
R2	SG13/Q4 (Cloud computing)	Data residency and sovereignty for government cloud services: SIDS profile	Technical Report + Conformance Profile	Q2-Q4 2027
R3	SG16/Q26-27 (Multimedia AI); SG17/Q6	AI model change management for public sector: Fairness testing and disclosure requirements	ITU-T Recommendation (F/X-series)	Q3 2027-Q1 2028
R4	SG20/Q2 (Smart sustainable cities); SG16/Q26	Explainability and human review mechanisms for automated public service decisions	Technical Report + Implementation Guidelines	Q2-Q3 2027
R5	SG17/Q6 (Identity); SG13/Q4 (Cloud)	Software Bill of Materials (SBOM) for AI systems in government procurement	Procurement Standard + Schema Specification	Q1-Q3 2027

Phase 1 (2027): Focus on foundational requirements R1 (decision transparency) and R4 (human redress), as they provide immediate citizen-facing accountability without requiring significant infrastructure changes. Development of open-source reference implementations including

- Decision log module (Python/Flask with PostgreSQL + Merkle anchoring + W3C VC integration)
- Redress workflow tool (React frontend + REST API for case management + SLA tracking)
- SBOM validator (Python with SPDX/CycloneDX parsers + NVD vulnerability checking)
- Test harness suite (pytest-based automated conformance testing for TA-R1-01 through TA-R5-02)

All implementations were released under the Apache 2.0 license and managed via GitHub by ITU-T SG coordinators, with documentation in English, French, and Spanish. These are not theoretical artefacts, they are the minimum viable tools needed for regulators to function in low-capacity environments.

Phase 2 (2027-2028): Progressive adoption of R3 (model change notice) and R5 (SBOM) as procurement cycles allow,

with R2 (data residency) integrated where new system deployments or contract renewals occur.

Phase 3 (2028-2029): Regional coordination through CARICOM, the Pacific Islands Forum, and the Indian Ocean Commission to share implementation experiences, develop shared infrastructure, and build collective bargaining power with vendors.

5.5 Limitations and implementation challenges

Although the proposed requirements (R1-R5) are designed to be more lightweight than comprehensive frameworks, such as the EU AI Act, they are not without cost and do not resolve the ASP trilemma. Their implementation poses significant challenges for SIDS.

Human capital cliff: The most immediate barrier is not merely a headcount but a **specialization gap**. While our analysis confirms that SIDS DPAs operate with 1-5 technical staff, the *type* of expertise required for ASP compliance is often absent. Training programs alone cannot close this gap. The issue is depth, not headcount.

The Challenge: Validating a model change notice (R3) or auditing an SBOM (R5) requires forensic data science and cybersecurity expertise. In Vignette A, the reliance on a foreign vendor was driven by a lack of domestic capability to audit the "liveness detection" algorithm. *Implication:* Without a "regional resource pool" (proposed in Section 5.4), mandating R3 may result in "compliance theater," where regulators receive technical notices they cannot cognitively process, fulfilling the letter of the standard but not its intent.

Financial diseconomies of scale: SIDS face a unique "scale penalty" where fixed costs for sovereignty infrastructure must be amortized over a tiny tax base.

The Challenge: Requirement R2 (data residency) implies local hosting. However, the energy cost in many SIDS (often diesel-dependent) can exceed \$0.30/kWh, compared to <\$0.06/kWh for hyperscale cloud providers. *Implication:* Hosting data locally to satisfy R2 significantly increases the Total Cost of Ownership (TCO). As seen in Vignette C, the shift to a "hybrid model" (local biographic, foreign biometric) was an economic necessity, not just a technical choice. Strict enforcement of R2 could bankrupt digital modernization projects, forcing a reversion to paper-based systems.

Energy determines sovereignty more than policy in these *environments*. This is rarely acknowledged but always felt.

Asymmetric market power: The ASP framework assumes that vendors will comply with new standards (e.g., R5 SBOMs) to win contracts. In SIDS markets, this assumption often fails because of low contract values.

The Challenge: A SIDS digital ID contract (often <\$5M) is a rounding error for major global vendors. If a SIDS regulator rigidly enforces "R3: Model Change Notice" with

penalties, large vendors may view the compliance cost as exceeding the contract value. *Implication:* This creates a "sovereignty paradox": strict standards might drive reputable global vendors out of the market, leaving SIDS to rely on lower-tier, less scrupulous vendors willing to sign contracts but unable to deliver secure systems, exacerbating the very surveillance risks (Amplifier 1) the standards aim to prevent.

Therefore, these standards are not a complete solution but a necessary first step in managing ASP. Their viability depends on parallel efforts in regional capacity-building, as discussed in the conclusion section.

6. CONCLUSION

Global AI governance risks building a paradigm irrelevant to the global majority. The question is not whether these frameworks are well-intentioned, but whether they assume state capacities that simply do not exist. This study redirects attention to these overlooked contexts.

We introduced the "Amplified Sovereignty Paradox" (ASP) to describe the structural trilemma governments face when adopting AI-driven digital public infrastructure. By analyzing documented implementations across SIDS (2016–2025), we demonstrated that these nations are not edge cases. They are structural stress tests. Geographic fragmentation and climate vulnerability reveal trade-offs that remain latent in resource-rich contexts but become existential in island environments.

We moved beyond diagnosis by translating the ASP into five standards-grade requirements (R1–R5). These operationalize sovereignty through concrete mechanisms like tamper-evident decision logging and procurement transparency. Each is mapped to specific ITU-T Study Group questions for the 2027–2029 cycle. We provide not just a critique of the trap, but the specifications for the exit.

However, technical standards alone cannot resolve the "Human Capital Cliff" facing small nations. A regulator with three staff members cannot audit a constantly evolving neural network, no matter how perfect the standard. This is not a temporary gap, it is a permanent feature of small state governance. Therefore, our proposal relies on a shift from national isolation to regional interdependence. We call for shared "sovereignty stacks", regional technical bodies capable of performing complex audits that no single island nation can sustain alone. Sovereignty in the AI age requires collective capacity, not just individual autonomy.

The unseen frontier of AI governance lies not in Silicon Valley or Brussels but in Bridgetown, Malé, and Kingston. Solutions developed for these constraints will work everywhere. Solutions developed for the wealthy will work only for the wealthy. The choice facing SIDS is not whether to adopt AI but is whether they retain any meaningful control once they do.

REFERENCES

- [1] ITU, "Artificial Intelligence for Good: Accelerating the SDGs," International Telecommunication Union, Geneva, Switzerland, 2024. [Online]. Available: <https://aiforgood.itu.int>
- [2] V. Eubanks, *Automating Inequality: How High-Tech Tools Profile, Police, and Punish the Poor*. New York, NY, USA: St. Martin's Press, 2018.
- [3] S. Zuboff, *The Age of Surveillance Capitalism: The Fight for a Human Future at the New Frontier of Power*. New York, NY, USA: PublicAffairs, 2019.
- [4] *Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024 laying down harmonised rules on artificial intelligence (Artificial Intelligence Act)*, OJ L, 2024/1689, 12.7.2024.
- [5] A. Bradford, *The Brussels Effect: How the European Union Rules the World*. New York, NY, USA: Oxford University Press, 2020.
- [6] Exec. Order No. 14110, 88 Fed. Reg. 75191 (Oct. 30, 2023), "Safe, Secure, and Trustworthy Development and Use of Artificial Intelligence."
- [7] R. Creemers, G. Webster, and H.-Y. Liu, "Translation: Personal Information Protection Law of the People's Republic of China (PIPL)," *DigiChina*, Stanford University, Aug. 20, 2021.
- [8] S. Neupane, *Needs Assessment for Artificial Intelligence, Digital Transformation and Open Data for Small Island Developing States*. Paris, France: UNESCO, 2024.
- [9] United Nations, "The Antigua and Barbuda Agenda for SIDS (ABAS) – A Renewed Declaration for Resilient Prosperity," Fourth International Conference on Small Island Developing States, A/CONF.223/3, St. John's, Antigua and Barbuda, 2024.
- [10] UNDP, "Small Island Digital States: How Digital Can Catalyse SIDS Development," United Nations Development Programme, New York, NY, USA, Policy Brief, 2024.
- [11] World Economic Forum, "Why Small Island Developing States are Digital Trailblazers," *WEF Agenda*, Geneva, Switzerland, May 2024. [Online]. Available: <https://www.weforum.org/stories/2024/05/sids-small-island-developing-states-digital-climate/>
- [12] World Bank, "Principles on Identification for Sustainable Development: Toward the Digital Age," World Bank Group, Washington, DC, USA, 2021.
- [13] *Information technology — Biometric presentation attack detection — Part 3: Testing and reporting*, ISO/IEC 30107-3:2023, Feb. 2023.

- [14] Bertelsmann Stiftung, "EuroStack – A European Alternative for Digital Sovereignty," Bertelsmann Stiftung and University College London (UCL), Gütersloh, Germany, Rep., Feb. 13, 2025.
- [15] T. O'Reilly, "Government as a Platform," *Innovations: Technology, Governance, Globalization*, vol. 6, no. 1, pp. 13–40, 2011.
- [16] D. Lyon, *Identifying Citizens: ID Cards as Surveillance*. Cambridge, UK: Polity Press, 2009.
- [17] M. Mann and A. Daly, "(Big) Data and the North-in-South: Australia's Informational Imperialism and Digital Colonialism," *Television & New Media*, vol. 20, no. 4, pp. 379–395, 2019.
- [18] R. Khera, "The Dissent on Aadhaar: Big Data Meets Big Brother," *Economic and Political Weekly*, vol. 54, no. 21, May 2019.
- [19] ITU, "Facts and Figures: Focus on Small Island Developing States," International Telecommunication Union, Geneva, Switzerland, 2024.
- [20] UNESCO, "Artificial Intelligence Policy Roadmap for the Caribbean," UNESCO Cluster Office for the Caribbean, Kingston, Jamaica, 2024.
- [21] N. Couldry and U. A. Mejias, *The Costs of Connection: How Data Is Colonizing Human Life and Appropriating It for Capitalism*. Stanford, CA, USA: Stanford University Press, 2019.
- [22] B. Flyvbjerg, "Five Misunderstandings About Case-Study Research," *Qualitative Inquiry*, vol. 12, no. 2, pp. 219–245, 2006.
- [23] Office of the Prime Minister, "National Identification System (NIDS)," Government of Jamaica, Kingston, Jamaica, 2024. [Online]. Available: <https://opm.gov.jm/portfolios/national-identification-system/>
- [24] Access Now, "The Jamaica NIDS Digital Identification Program: A Cautionary Tale," *Access Now*, Oct. 2023. [Online]. Available: <https://www.accessnow.org/jamaica-nids-digital-id/>
- [25] Supreme Court of Judicature of Jamaica, *Julian J. Robinson v. The Attorney General of Jamaica*, [2019] JMFC Full 04, Judgment, Apr. 12, 2019.
- [26] "Maldives Launches eFaas: Your Key to a Digital Maldives," *Corporate Maldives*, Malé, Maldives, Jun. 2023.
- [27] "Regula to Replace SumSub as Face Biometrics Provider for Maldives," *Biometric Update*, Dec. 2, 2024. [Online]. Available: <https://www.biometricupdate.com>
- [28] "NCIT Resumes eFaas Online Face Verification with New Provider," *Corporate Maldives*, Malé, Maldives, Dec. 2, 2024.
- [29] *Artificial intelligence – Data quality for analytics and machine learning (ML) – Part 1: Overview, terminology and examples*, ISO/IEC 5259-1:2024, Feb. 2024.
- [30] *Information technology – Artificial intelligence – Management system*, ISO/IEC 42001:2023, Dec. 2023.
- [31] S. Bradner, "Key words for use in RFCs to Indicate Requirement Levels," IETF, RFC 2119, Mar. 1997.
- [32] "NCIT Introduces New Security Features for eFaas," *PSM News*, Malé, Maldives, Oct. 2024.
- [33] "Maldives ID Transformation: Digital First Approaches," *Biometric Update*, Jul. 2024.
- [34] M. Sporny, D. Longley, and D. Chadwick, "Verifiable Credentials Data Model v1.1," W3C Recommendation, Mar. 03, 2022.
- [35] M. Sporny et al., "Decentralized Identifiers (DIDs) v1.0," W3C Recommendation, Jul. 19, 2022.
- [36] C. Allen, "The Path to Self-Sovereign Identity," *Life With Alacrity*, Apr. 25, 2016. [Online]. Available: <http://www.lifewithalacrity.com>
- [37] A. Mühle, A. Grüner, T. Gayvoronskaya, and C. Meinel, "A Survey on Essential Components of a Self-Sovereign Identity," *Computer Science Review*, vol. 30, pp. 80–86, 2018.
- [38] "Artificial Intelligence Risk Management Framework (AI RMF 1.0)," National Institute of Standards and Technology (NIST), Gaithersburg, MD, USA, NIST AI 100-1, Jan. 2023.
- [39] *Artificial intelligence – Data quality for analytics and machine learning (ML) – Part 2: Data quality measures*, ISO/IEC 5259-2:2024, Nov. 2024.
- [40] UNESCO, "Recommendation on the Ethics of Artificial Intelligence," UNESCO, Paris, France, SHS/BIO/PI/2021/1, 2021.
- [41] OECD, "Recommendation of the Council on Artificial Intelligence," OECD, Paris, France, OECD/LEGAL/0449, 2019.
- [42] Partnership on AI, "AI and Climate Adaptation in Small Island States," Partnership on AI, San Francisco, CA, USA, Case Study, 2023.
- [43] ITU-T Focus Group on Artificial Intelligence for Health (FG-AI4H), "Ethics and Governance of AI for Health," ITU, Geneva, Switzerland, Technical Report, 2024.