The Infrastructure Sovereignty Gap: A Framework for Measuring National AI Dependency through the National AI Sovereignty Index (NASI)

Author: Deusdedit Ruhangariyo,
Independent Researcher, AI Governance & Technology Policy
United States

Abstract

As artificial intelligence systems become embedded within public services, critical infrastructure, and sovereign decision-making processes, nations increasingly rely on AI systems whose underlying infrastructure they neither own nor fully control. While "digital sovereignty" has emerged as a policy concern, existing frameworks fail to measure a state's practical authority over AI execution itself. This paper introduces the **Infrastructure Sovereignty Gap**: the structural disparity between a nation's reliance on AI-enabled systems and its capacity to govern, audit, and sustain the infrastructure on which those systems depend.

To operationalize this risk, the paper proposes the **National AI Sovereignty Index (NASI)**, a composite diagnostic framework designed to assess national exposure across six dimensions: Compute Ownership, Data Center Autonomy, Model Governance Control, Regulatory Leverage, Public Sector Dependency, and Foreign Exposure Risk. Rather than treating sovereignty as a binary attribute, NASI conceptualizes sovereignty as a continuous, multidimensional condition subject to threshold effects and non-compensatory constraints.

The paper further develops a causal mechanism model explaining how infrastructure dependency propagates into governance vulnerability, strategic asymmetry, and systemic public-service risk. Continental exposure patterns are outlined to demonstrate how infrastructure sovereignty gaps manifest differently across geopolitical contexts. Finally, the paper introduces governance pathways—including domestic infrastructure strategies and locally grounded oversight mechanisms—through which states may reduce exposure without rejecting AI adoption itself.

This work is presented as a foundational framework intended for iterative empirical refinement. By shifting the focus from AI adoption to AI infrastructural control, the Infrastructure Sovereignty Gap provides policymakers, regulators, and public institutions with a measurable lens for evaluating sovereignty risk in the age of machine-mediated governance.

Keywords

AI governance; infrastructure sovereignty; national AI strategy; compute dependency; digital public infrastructure; geopolitical AI risk; NASI; public sector AI

---

1. Introduction

Artificial intelligence has moved rapidly from experimental deployment to structural integration within public services. Governments now rely on AI systems to assist in healthcare delivery, education administration, welfare allocation, border control, taxation, elections, and national security. In many cases, these systems operate at scale, in real time, and with direct consequences for citizens' rights and access to essential services.

Yet the infrastructure enabling these systems—compute, data centers, foundational models, and execution environments—is frequently owned, governed, or controlled by foreign entities. As a result, states may exercise policy authority over AI outcomes while lacking material authority over AI execution. This paper argues that this mismatch constitutes a distinct and under-theorized form of systemic risk: the **Infrastructure Sovereignty Gap**.

Current policy discourse often equates AI progress with AI adoption. However, adoption alone does not confer sovereignty. A nation may deploy advanced AI systems extensively while remaining structurally dependent on external infrastructure for computation, model access, updates, and continuity of service. In such cases, governance authority becomes contingent rather than sovereign—subject to foreign policy shifts, commercial incentives, regulatory asymmetries, and infrastructural shocks.

The Infrastructure Sovereignty Gap reframes AI governance as a question of **control without rejection**. The issue is not whether states should adopt AI, but whether they retain sufficient authority to govern AI systems that have become indispensable to public life. This distinction is increasingly urgent as AI transitions from advisory tools to operational systems embedded within core state functions.

This paper advances three central claims. First, infrastructure sovereignty constitutes a distinct dimension of AI governance that cannot be inferred from legal frameworks or ethical principles alone. Second, sovereignty exposure can be measured through a structured, multidimensional index rather than treated as an abstract concept. Third, failure to assess and manage infrastructure sovereignty risks may lead to cascading governance failures even in well-intentioned AI deployments.

To support these claims, the paper introduces the National AI Sovereignty Index (NASI) and a causal model of sovereignty erosion. Together, these tools aim to equip policymakers with a practical framework for diagnosing dependency risks before they materialize as public-sector failures.

---

2. Related Work and Conceptual Distinctions

The Infrastructure Sovereignty Gap intersects with, but is not reducible to, existing literatures on digital sovereignty, data localization, algorithmic governance, and platform power. Digital sovereignty frameworks typically emphasize control over data flows, regulatory jurisdiction, or content moderation. While important, these approaches often overlook the infrastructural layer at which AI systems are executed.

Similarly, discussions of platform capitalism and cloud dependency highlight market concentration and economic asymmetry but rarely translate these concerns into governance metrics applicable to state decision-making. AI ethics frameworks, meanwhile, focus predominantly on fairness, accountability, and transparency at the model or application level, assuming stable and controllable infrastructure beneath.

What distinguishes AI infrastructure sovereignty from broader technological sovereignty is the **executability constraint**. AI systems do not merely store or transmit information; they require continuous access to high-performance compute, model updates, and tightly integrated execution environments. Loss of control at this layer does not merely degrade service quality—it can disable governance altogether.

The Infrastructure Sovereignty Gap therefore captures a structural condition in which formal authority exists without material control. NASI is introduced to measure this condition explicitly, filling a gap between abstract sovereignty discourse and operational governance risk assessment.

---

3. Defining the National AI Sovereignty Index (NASI)

The **National AI Sovereignty Index (NASI)** is a composite diagnostic framework designed to assess the degree of sovereign control a nation retains over the AI infrastructure supporting its public and strategic systems. NASI does not measure AI capability or innovation capacity; it measures **governance exposure arising from infrastructural dependence**.

NASI is composed of six dimensions:

1. **Compute Ownership (CO):**
   The extent to which a nation owns or controls high-performance compute resources required for AI model training and execution, including sovereign access to GPUs, accelerators, and scheduling authority.

   Compute ownership does not imply full hardware autonomy. Even sovereign data centers remain dependent on global semiconductor supply chains, export controls, and architectural chokepoints. NASI therefore distinguishes between execution control and fabrication independence, treating the latter as an exogenous constraint rather than a prerequisite for governance authority.

2. **Data Center Autonomy (DCA):**
   The degree of domestic control over data center infrastructure supporting AI workloads, including physical location, operational continuity, and insulation from extraterritorial legal or political interventions.

3. **Model Governance Control (MGC):**
   The extent to which a nation can inspect, audit, modify, or constrain the behavior of AI models deployed within its jurisdiction, including access to weights, update pathways, and refusal mechanisms.

4. **Regulatory Leverage (RL):**
   The practical enforceability of domestic AI regulations against infrastructure providers, including the ability to impose conditions, penalties, or suspensions that materially affect system operation.

5. **Public Sector Dependency (PSD):**
   The degree to which essential public services rely on AI systems whose continued operation is contingent on external infrastructure access.

6. **Foreign Exposure Risk (FER):**
   The susceptibility of AI infrastructure to foreign policy actions, sanctions, commercial disputes, or unilateral service withdrawals beyond national control.

NASI is intentionally non-compensatory in critical domains: high performance in one dimension cannot fully offset zero capacity in another where public-service continuity is at stake. The index is designed as a diagnostic tool rather than a ranking instrument, emphasizing threshold effects and failure pathways over aggregate scores.

At this stage, NASI is presented as a structured framework for assessment rather than a finalized quantitative metric. Future work will explore weighting schemes, sector-specific adaptations, and empirical calibration.

---

4. Mechanism Model: From Dependency to Sovereignty Loss

The Infrastructure Sovereignty Gap emerges through a predictable sequence of causal mechanisms:

1. **Infrastructure Dependency Formation:**
   AI systems are adopted through foreign cloud platforms or proprietary model APIs due to cost efficiency, speed, or technical superiority.
2. **Policy Exposure Emergence:**
   Governance authority becomes contingent on infrastructure providers whose incentives and jurisdictions do not align with domestic public obligations.
3. **Strategic Asymmetry Consolidation:**
   Knowledge, bargaining power, and operational control accumulate externally, limiting national ability to negotiate, audit, or adapt systems under stress.
4. **Risk Propagation:**
   Infrastructure disruptions—whether technical, commercial, or geopolitical—cascade into public-service failures with limited domestic recourse.

This mechanism explains why sovereignty loss often remains invisible during pilot deployments and only manifests once AI systems become infrastructurally indispensable.

---

5. Continental Exposure Patterns

Infrastructure sovereignty gaps manifest unevenly across regions due to historical investment patterns, capital access, regulatory capacity, and geopolitical positioning. While high-income states may exhibit partial compute ownership but model dependency, lower-income regions often face compounded exposure across all NASI dimensions.

These patterns are not indicators of governance failure but of structural asymmetry within the global AI ecosystem. NASI is designed to surface these asymmetries without presuming uniform development pathways or policy objectives.

---

6. Governance Responses to the Infrastructure Sovereignty Gap

Reducing infrastructure sovereignty risk does not require rejecting AI adoption. Instead, it requires intentional governance strategies aligned with national capacity and public-

service priorities. These include domestic compute strategies, sovereign cloud initiatives, procurement standards tied to auditability, and international cooperation frameworks that preserve execution authority.

---

7. Tradeoffs, Constraints, and Limitations

Infrastructure sovereignty entails economic, technical, and temporal tradeoffs. Full autonomy may be infeasible or undesirable for smaller states, and strategic dependency may be rational in non-critical sectors. NASI does not prescribe uniform sovereignty targets but provides a structured method for identifying unacceptable risk thresholds.

NASI is intentionally not optimized for inter-state ranking. Its primary function is diagnostic: to identify sovereignty-breaking thresholds within national AI deployments. Quantitative weighting is deferred to sector-specific applications where failure costs can be meaningfully specified.

Measurement challenges, data availability, and dynamic infrastructure markets present limitations that future empirical work must address.

---

8. Conclusion

The Infrastructure Sovereignty Gap represents a structural governance risk inherent to AI-driven public systems. By introducing NASI, this paper provides a practical framework for diagnosing that risk before it materializes as systemic failure. The goal is not to halt AI progress, but to ensure that progress does not erode the conditions under which sovereign governance remains possible.

Appendix A: Illustrative NASI Assessment (Worked Example)

This appendix provides an illustrative application of the National AI Sovereignty Index (NASI) to demonstrate operational feasibility. The example is intentionally transparent and non-exhaustive, serving as a methodological template rather than a definitive national assessment.

Illustrative Case: Kenya (Hypothetical Assessment)

Kenya represents a mid-sized economy with advanced digital adoption, growing AI usage in public services, and significant reliance on foreign infrastructure providers.

A1. Compute Ownership (CO): Low

Kenya does not possess sovereign control over high-performance AI compute at scale. GPU-intensive workloads are predominantly executed through foreign hyperscalers, with limited domestic scheduling authority or guaranteed access during global demand spikes.

Implication: National AI execution capacity is externally contingent, particularly in crisis scenarios.

A2. Data Center Autonomy (DCA): Medium

Domestic data centers exist, but advanced AI workloads frequently rely on foreign-owned facilities or hybrid architectures subject to extraterritorial legal regimes.

Implication: Partial operational continuity exists, but sovereignty over AI-critical workloads remains constrained.

A3. Model Governance Control (MGC): Low

Public-sector AI deployments rely primarily on proprietary foundation models accessed via APIs. Model weights, update mechanisms, and refusal logic are externally governed.

Implication: Auditability and behavioral constraint are limited to contractual terms rather than sovereign enforcement.

A4. Regulatory Leverage (RL): Medium

Kenya possesses formal regulatory authority over AI applications within its jurisdiction, but enforcement against foreign infrastructure providers is limited by market concentration and cross-border dependencies.

Implication: Regulatory authority exists in principle but is weakly coupled to execution control.

A5. Public Sector Dependency (PSD): High

AI systems are increasingly embedded in health diagnostics, agricultural advisory platforms, identity systems, and public-service delivery.

Implication: Infrastructure disruptions would have immediate public welfare consequences.

A6. Foreign Exposure Risk (FER): High

AI infrastructure access is vulnerable to pricing changes, policy shifts, service withdrawal, or geopolitical realignment beyond national influence.

Implication: Systemic exposure to external shocks.

A7. Diagnostic Outcome

Under NASI's non-compensatory logic, high Public Sector Dependency combined with low Compute Ownership and Model Governance Control indicates a **high Infrastructure Sovereignty Gap**, despite moderate regulatory capacity.

This assessment demonstrates how NASI surfaces governance risk not visible through adoption metrics or regulatory inventories alone.

Appendix B: Aggregation Logic and Weighting Principles

NASI is designed as a diagnostic framework rather than a competitive ranking system. Accordingly, aggregation follows three guiding principles:

1. **Non-Compensatory Constraints:**
   Certain dimensions—particularly Compute Ownership, Model Governance Control, and Public Sector Dependency—exhibit threshold effects. Zero capacity in these domains cannot be offset by strength elsewhere when critical services are involved.
2. **Sector Sensitivity:**
   Weighting varies by sector. For example, health, elections, and security systems require stricter sovereignty thresholds than non-critical applications.
3. **Failure-Oriented Evaluation:**
   NASI prioritizes identification of failure pathways over aggregate optimization. The objective is early warning, not numerical maximization.

Future work will explore formal aggregation functions, including hybrid threshold–additive models and scenario-based stress testing.

Appendix C: Sovereignty–Efficiency Tradeoff Matrix

Infrastructure sovereignty entails measurable economic and operational tradeoffs. Table 1 outlines a conceptual matrix distinguishing acceptable strategic dependency from unacceptable governance risk.

Dependency Level	Cost Efficiency	Governance Risk	Policy Posture
Low	Low	Low	Sovereign Core
Medium	High	Medium	Managed Hybrid
High	Very High	High	Critical Risk

NASI enables states to make explicit, informed choices about where dependency is tolerable and where it becomes a systemic threat.

Appendix D: Scope, Limitations, and Future Work

This framework is subject to several limitations:

- Data availability varies across jurisdictions.
- Infrastructure control is dynamic and market-sensitive.
- Sovereignty thresholds are context-dependent rather than universal.

Future research directions include empirical NASI calibration across multiple regions, longitudinal tracking of sovereignty erosion, and integration with procurement and resilience planning tools.

Appendix E: The Collaborative Reflection Protocol (CRP) as a Participatory Safeguard Mechanism

CRP (Collaborative Reflection Protocol) refers to a governance framework for embedding locally legitimate decision authority and ethical oversight into AI system deployment, particularly within public and sovereign institutions. Within this framework, NASI functions as a diagnostic instrument, while CRP provides procedural legitimacy for decisions made under infrastructural constraint.

NASI identifies where sovereignty exposure exists; CRP governs how authority is exercised within those conditions.

Final Statement

The Infrastructure Sovereignty Gap reframes AI governance as a question of execution authority rather than abstract control. Through NASI, this paper provides a structured, defensible method for identifying when AI adoption outpaces sovereign capacity. The

framework is intended to evolve through empirical application, policy engagement, and iterative refinement.

Appendix F: NASI Application Protocol (NAP)

F.1 Purpose and Scope

The NASI Application Protocol (NAP) provides a **practical procedure** for governments, regulators, and public-sector institutions to apply the National AI Sovereignty Index as a **diagnostic and early-warning tool**, rather than as a comparative ranking instrument.

NAP is designed for:

- Ministries of ICT, Finance, Health, Interior
- National procurement authorities
- Supreme audit institutions
- Multilateral technical assistance missions

It assumes **no domestic AI manufacturing capability** and does not require advanced technical modeling.

F.2 Application Principle

NASI is applied under **non-compensatory logic**.

Failure in any single dimension may constitute a sovereignty risk **regardless of strength in other dimensions**, particularly for critical public services.

NASI is therefore used to:

- Identify **exposure**, not optimize performance
- Detect **execution-time dependency**, not regulate innovation
- Trigger **policy attention**, not prescribe isolation

F.3 Step-by-Step Application Procedure

Step 1: Define the Assessment Boundary

The assessing authority must specify:

- Sector (e.g., health, welfare, identity, taxation)
- AI function (decision support, eligibility determination, automation)
- Criticality level (advisory vs decision-executing)

NASI should be applied **per sector**, not nationally in aggregate.

Step 2: Dimension-by-Dimension Assessment

For each NASI dimension, the authority collects **documentary and operational evidence**, not self-reported claims.

Example data sources (illustrative):

- Procurement contracts
- Cloud service agreements
- Regulatory enforcement powers
- Data residency audits
- Model update logs
- Incident response authority

Each dimension is classified qualitatively as:

- **High Control**
- **Partial Control**
- **Minimal or No Control**

No numerical aggregation is required at this stage.

Step 3: Threshold Screening

If **any** of the following conditions hold, the sector is provisionally classified as **High Sovereignty Risk**:

- Compute execution cannot be guaranteed during geopolitical disruption
- Model behavior cannot be frozen, audited, or overridden domestically

- Regulatory authority cannot enforce remedies at execution time
- Public services would halt if external providers withdraw access

This screening step operationalizes the **executability constraint**.

Step 4: Risk Classification

Based on threshold screening:

- **Green**: No single-point execution dependency for critical functions
- **Yellow**: Dependency exists but is bounded, contractually mitigated, or non-critical
- **Red**: One or more dimensions represent an unmitigated execution-time dependency

This classification is **contextual**, not comparative.

Step 5: Policy Response Mapping

For each Red or Yellow classification, the authority identifies **targeted interventions**, such as:

- Dual-vendor execution pathways
- Domestic failover requirements
- Model escrow or audit rights
- Execution locality guarantees
- Sector-specific sovereignty thresholds

NASI does **not** mandate domestic ownership—only **governable execution**.

F.4 Use Constraints and Limitations

NAP explicitly does **not**:

- Rank countries
- Mandate protectionism
- Assume hardware self-sufficiency

- Replace cost–benefit analysis

It is a **risk visibility tool**, not a development index.

F.5 Iterative Use and Versioning

Governments are encouraged to:

- Reapply NASI annually or upon major procurement changes
- Publish partial findings where appropriate
- Use NASI profiles as inputs into national AI strategies

Empirical weighting and quantitative scoring are reserved for **future versions** once sufficient comparative data exists.

References

1. Floridi, L., Cowls, J., Beltrametti, M., et al. (2018). *AI4People—An Ethical Framework for a Good AI Society: Opportunities, Risks, Principles, and Recommendations*. Minds and Machines, 28(4), 689–707.
2. United Nations Secretary-General. (2023). *Governing AI for Humanity: Final Report of the High-Level Advisory Body on Artificial Intelligence*. United Nations.
3. OECD. (2023). *AI Compute and Climate: Policy Considerations for Governments*. Organisation for Economic Co-operation and Development.
4. ITU. (2023). *AI and Data Governance: Cross-Border Challenges and Policy Approaches*. International Telecommunication Union.
5. European Commission. (2024). *Regulation (EU) 2024/1689 Laying Down Harmonised Rules on Artificial Intelligence (AI Act)*. Official Journal of the European Union.
6. World Bank. (2022). *Digital Public Infrastructure for Inclusive Development*. World Bank Group.
7. Kwet, M. (2019). *Digital Colonialism: US Empire and the New Imperialism in the Global South*. Race & Class, 60(4), 3–26.

8. UNDP. (2023). *Human Development Report 2023/2024: Breaking the Gridlock—Reimagining Cooperation in a Polarized World*. United Nations Development Programme.
9. Varian, H., Shapiro, C., & Shoven, J. (2021). *The Economics of Platforms and AI Infrastructure*. Journal of Economic Perspectives, 35(4), 3–26.
10. Rahman, K. S., & Thelen, K. (2019). *The Rise of the Platform Business Model and the Transformation of Twenty-First-Century Capitalism*. Politics & Society, 47(2), 177–204.
11. DeNardis, L. (2020). *The Internet in Everything: Freedom and Security in a World with No Off Switch*. Yale University Press.
12. G7 Digital and Technology Ministers. (2023). *Hiroshima AI Process: Guiding Principles for Generative AI*. Government of Japan.
13. Brundage, M., Avin, S., Clark, J., et al. (2018). *The Malicious Use of Artificial Intelligence: Forecasting, Prevention, and Mitigation*. arXiv:1802.07228.
14. Crawford, K. (2021). *Atlas of AI: Power, Politics, and the Planetary Costs of Artificial Intelligence*. Yale University Press.
15. Acemoglu, D., & Johnson, S. (2023). *Power and Progress: Our Thousand-Year Struggle Over Technology and Prosperity*. PublicAffairs.