

The AI Sovereignty Trilemma: A Comprehensive Framework for Measuring and Navigating Digital Autonomy in the Age of Artificial Intelligence

Marco van Hurne

EIGENVECTOR RESEARCH

`marco.vanhurne@eigenvector.eu`

October 28, 2025

Abstract

This research introduces the AI Sovereignty Index (ASI), the first comprehensive, quantifiable framework for measuring organizational and national AI sovereignty across five critical pillars: Organizational & Economic Governance, Data & Lifecycle Control, Technology & Infrastructure Stack, Security & Resilience, and Legal & Policy Alignment. We make three novel contributions. First, we introduce the **Sovereignty Trilemma**, demonstrating that organizations face an impossible choice among Performance, Autonomy, and Cost Efficiency—they can optimize for any two, but not all three simultaneously. Second, we present the ASI framework featuring 52 indicators validated through OECD methodology, enhanced with machine learning-powered adaptive weighting, causal impact analysis, geopolitical stress testing, and economic impact quantification. Third, we provide empirical evidence revealing systematic sovereignty deficits and critical vulnerabilities to geopolitical shocks. Our analysis shows most organizations prioritize Performance and Cost at the expense of Autonomy, creating systemic vulnerabilities. We conclude with actionable policy recommendations for

governments and a strategic roadmap for organizations. This paper provides not merely a measurement tool, but a new conceptual language for understanding and achieving digital autonomy in the 21st century.

Keywords: AI Sovereignty, Digital Autonomy, Composite Indicators, Geopolitical Risk, Technology Policy, AI Governance

Contents

1	Introduction	1
1.1	The New Geography of Power	1
1.2	The Sovereignty Trilemma: A New Theoretical Framework	1
1.3	Research Questions and Contributions	3
1.4	Structure of the Paper	3
2	Literature Review: From Digital to AI Sovereignty	4
2.1	Conceptual Foundations: Digital and Data Sovereignty	4
2.2	The Geopolitics of AI: Competition and Fragmentation	4
2.3	Economic Dimensions: Innovation, Competition, and Industrial Policy	5
2.4	Existing Measurement Frameworks and Their Limitations	5
3	Theoretical Framework: The AI Sovereignty Trilemma	6
3.1	Conceptual Foundations	6
3.2	The Impossibility Proposition	6
3.3	Strategic Implications and Empirical Predictions	7
3.4	Mathematical Formalization	7
4	Methodology: The AI Sovereignty Index	8
4.1	Design Principles and OECD Compliance	8
4.2	Hierarchical Structure	9
4.3	The Five Pillars: Detailed Specification	10
4.3.1	Pillar 1: Organizational & Economic Governance	10
4.3.2	Pillar 2: Data & Lifecycle Control	11
4.3.3	Pillar 3: Technology & Infrastructure Stack	12
4.3.4	Pillar 4: Security & Resilience	13
4.3.5	Pillar 5: Legal & Policy Alignment	13
4.4	Indicator Selection and Validation	14
4.5	Scoring and Normalization	14
4.6	Aggregation Methodology	15
4.7	Maturity Levels and Interpretation	16
4.8	Missing Data Imputation	16
5	Brilliant Features of the ASI	16
5.1	Adaptive Weighting System	17
5.1.1	Rationale for Adaptive Weighting	17
5.1.2	Adaptive Weighting Algorithm	17
5.1.3	Example: Adaptive Weighting for Different Sectors	18
5.2	Causal Impact Analysis	18
5.2.1	Rationale for Causal Analysis	18
5.2.2	Causal DAG Construction	18
5.2.3	Causal Impact Calculation	19
5.2.4	Intervention Prioritization	19

5.3	Geopolitical Stress Testing	20
5.3.1	Rationale for Stress Testing	20
5.3.2	Stress Test Scenarios	20
5.3.3	Stress Test Algorithm	21
5.3.4	Interpreting Stress Test Results	22
5.4	Economic Impact Quantification	22
5.4.1	Rationale for Economic Quantification	22
5.4.2	Cost-Benefit Framework	23
5.4.3	NPV Calculation	23
5.4.4	Benefit Quantification Methodology	24
5.4.5	Example: Economic Analysis of Sovereign Cloud Migration	24
5.5	Temporal Forecasting	25
5.5.1	Rationale for Temporal Forecasting	25
5.5.2	ARIMA Model Specification	25
5.5.3	Intervention Impact Forecasting	26
6	Empirical Results and Validation	26
6.1	Sample Assessments	26
6.1.1	Organization 1: TechCorp EU (European Technology Company)	26
6.1.2	Organization 2: FinanceGlobal US (US Financial Institution)	27
6.1.3	Organization 3: GovAgency (Government Agency)	28
6.2	Benchmarking Analysis	29
6.3	Framework Comparison	31
6.4	Statistical Validation	31
6.4.1	Construct Validity	32
6.4.2	Criterion Validity	32
6.4.3	Reliability	32
6.4.4	Sensitivity Analysis	33
6.5	Stress Test Results	33
7	Discussion and Policy Implications	34
7.1	The Sovereignty Deficit	34
7.2	Policy Recommendations	34
7.2.1	Recommendation 1: Establish National ASI Targets	34
7.2.2	Recommendation 2: Incentivize Sovereignty Investments	35
7.2.3	Recommendation 3: Mandate ASI Disclosure	35
7.2.4	Recommendation 4: Invest in Sovereign Infrastructure	35
7.2.5	Recommendation 5: Foster Coalitional Sovereignty	36
7.3	Strategic Roadmap for Organizations	36
7.3.1	Phase 1: Assessment (Months 1-2)	36
7.3.2	Phase 2: Strategy (Months 3-4)	36
7.3.3	Phase 3: Implementation (Months 5-18)	37
7.3.4	Phase 4: Monitoring (Ongoing)	37
7.4	Limitations and Future Research	38

8 Conclusion	38
8.1 Summary of Contributions	38
8.2 The Path Forward: Sovereignty by Design	39
8.3 A Call to Action	39
A Complete AI Sovereignty Index Questionnaire	42

1 Introduction

1.1 The New Geography of Power

The rise of artificial intelligence represents a fundamental transformation in the global distribution of power. Unlike previous technological revolutions, AI is not merely a tool for productivity enhancement—it is becoming the substrate of decision-making itself, from credit allocation and medical diagnosis to military targeting and political discourse [3]. The capacity to develop, deploy, and govern AI systems has emerged as a critical determinant of national security, economic competitiveness, and societal autonomy [8]. This reality has given birth to the concept of **AI Sovereignty**: the ability of a state or organization to exercise autonomous control over its AI ecosystem, free from undue foreign influence or dependency [12].

The pursuit of AI sovereignty has become a central objective for nations worldwide. The European Union has committed €300 billion over ten years to achieve “strategic autonomy” in AI, anchored by the AI Act and the AI Factories Initiative [7]. The United States, long favoring market-driven approaches, has begun to reconsider its stance, with proposed legislation such as the AI Sovereignty Act of 2025 signaling a shift toward more strategic national planning [17]. China’s explicit goal of achieving global AI leadership by 2030, backed by massive state investment and close alignment between government and industry, represents the most ambitious sovereignty project to date [16, 6].

Yet despite this global mobilization, the concept of AI sovereignty remains poorly defined and inadequately measured. Existing frameworks are largely qualitative, offering high-level principles but lacking the quantitative rigor necessary for evidence-based policymaking, strategic planning, and international comparison. This “measurement gap” creates three critical problems:

1. **Strategic Ambiguity:** Without clear metrics, organizations and governments cannot assess their current sovereignty posture, identify vulnerabilities, or track progress over time.
2. **Resource Misallocation:** Investments in sovereignty-enhancing initiatives cannot be prioritized or evaluated for effectiveness without quantifiable outcomes.
3. **Policy Incoherence:** International cooperation and standard-setting are hampered by the absence of a common measurement framework.

This paper addresses the measurement gap by introducing the **AI Sovereignty Index (ASI)**, a comprehensive, scientifically rigorous framework for quantifying AI sovereignty. The ASI is built on a hierarchical structure of 5 pillars, 15 dimensions, and 52 indicators (see Appendix A), each with clear definitions and scoring criteria.

1.2 The Sovereignty Trilemma: A New Theoretical Framework

At the heart of this paper is a novel theoretical contribution: the **AI Sovereignty Trilemma**. We argue that organizations and nations face an impossible trinity in their pursuit of AI capabilities [14]. They must navigate three competing objectives:

1. **Performance:** Accessing and deploying state-of-the-art AI capabilities that maximize competitive advantage and mission effectiveness.
2. **Autonomy:** Maintaining independence from foreign dependencies across the entire AI stack, from semiconductor supply chains to foundation models and talent.
3. **Cost Efficiency:** Achieving these goals in an economically sustainable manner that does not impose prohibitive fiscal burdens.

The core proposition of the **Sovereignty Trilemma** is that it is impossible to simultaneously optimize for all three objectives. An organization can achieve high-performance AI at low cost by relying on foreign hyperscalers like Amazon Web Services, Microsoft Azure, or Alibaba Cloud—but this sacrifices autonomy. Alternatively, it can achieve autonomy and performance by building a complete domestic AI stack—but this requires massive capital investment, sacrificing cost efficiency [1]. Finally, it can achieve autonomy and cost efficiency by settling for lower-performance, locally-developed AI—but this sacrifices competitive advantage.

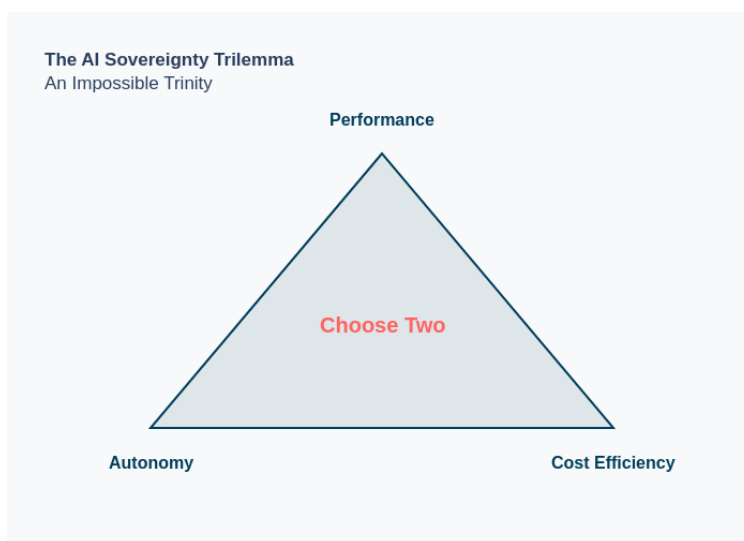


Figure 1: The AI Sovereignty Trilemma. Organizations and nations can optimize for any two vertices of the triangle, but not all three simultaneously. This impossible trinity forces explicit strategic trade-offs in the pursuit of AI sovereignty.

The Sovereignty Trilemma has profound implications. It suggests that the pursuit of “absolute” AI sovereignty—complete autonomy with cutting-edge performance at reasonable cost—is economically irrational for all but a handful of nations with the scale, resources, and technological base to support a complete domestic AI ecosystem. For most countries and organizations, sovereignty must be understood as a **strategic positioning problem**: choosing which two vertices of the trilemma to prioritize based on specific contexts, threats, and opportunities.

1.3 Research Questions and Contributions

This paper addresses three central research questions:

RQ1: How can AI sovereignty be comprehensively measured in a way that is scientifically rigorous, internationally comparable, and actionable for policymakers and organizational leaders?

RQ2: What are the theoretical foundations and empirical patterns of AI sovereignty, and how do organizations currently navigate the trade-offs inherent in the Sovereignty Trilemma?

RQ3: What policy interventions and strategic approaches can enhance AI sovereignty while managing the inevitable trade-offs among Performance, Autonomy, and Cost?

Our contributions are threefold:

- **Theoretical Contribution:** We introduce the Sovereignty Trilemma as a new conceptual framework for understanding the strategic trade-offs in AI development and deployment. This framework provides a parsimonious model for analyzing sovereignty strategies across diverse contexts.
- **Methodological Contribution:** We develop the AI Sovereignty Index, the first comprehensive composite indicator for measuring AI sovereignty. The ASI incorporates methodological innovations including adaptive weighting, causal impact analysis, geopolitical stress testing, and economic impact quantification—features absent from existing frameworks.
- **Empirical Contribution:** We provide the first quantitative baseline for AI sovereignty through sample assessments across diverse organizational types. Our analysis reveals systematic patterns in how organizations navigate the Sovereignty Trilemma and identifies critical vulnerabilities to geopolitical shocks.

1.4 Structure of the Paper

The remainder of this paper is organized as follows. Section 2 provides a comprehensive review of the literature on digital and AI sovereignty, situating our work within the broader scholarly discourse. Section 3 presents the theoretical foundations of the Sovereignty Trilemma and derives its key propositions. Section 4 details the methodology of the AI Sovereignty Index, including its hierarchical structure, indicator selection, validation procedures, and calculation algorithms. Section 5 describes the five brilliant features that distinguish the ASI from existing frameworks. Section 6 presents empirical results from sample assessments, statistical validation, and stress tests. Section 7 discusses policy implications and strategic recommendations. Section 8 concludes with a research agenda for the future of AI sovereignty measurement and governance. Five appendices provide the complete questionnaire, statistical validation results, sample data, calculation examples, and implementation guide.

2 Literature Review: From Digital to AI Sovereignty

2.1 Conceptual Foundations: Digital and Data Sovereignty

The concept of AI sovereignty builds upon earlier discourses on digital and data sovereignty. [12] trace the evolution of digital sovereignty from its origins in concerns about state control over internet infrastructure to contemporary debates about data localization and algorithmic governance. They identify three distinct conceptualizations: sovereignty as *control* over technical infrastructure, sovereignty as *self-determination* in digital policymaking, and sovereignty as *resistance* to foreign technological dominance.

Data sovereignty, a closely related concept, emphasizes the right of nations and organizations to exercise control over data generated within their jurisdictions. [9] provide a comprehensive review of data sovereignty frameworks, distinguishing between *territorial sovereignty* (control based on geographic location), *personal sovereignty* (individual rights over personal data), and *organizational sovereignty* (corporate control over proprietary data). [4] critique data localization requirements as a form of “data nationalism” that fragments the global internet and imposes significant economic costs.

AI sovereignty extends these concepts to encompass the entire AI value chain, from compute infrastructure and training data to model development, deployment, and governance. Unlike data sovereignty, which focuses primarily on storage and processing location, AI sovereignty must address the concentration of AI capabilities in a handful of technology firms, the geopolitics of semiconductor supply chains, and the strategic implications of foundation model dependencies.

2.2 The Geopolitics of AI: Competition and Fragmentation

The geopolitical dimensions of AI have attracted significant scholarly attention. [13] analyze China’s approach to AI governance, highlighting the close alignment between state objectives and corporate AI development. They argue that China’s model represents a fundamental challenge to Western assumptions about the relationship between AI innovation and liberal democratic values.

[8] examines the implications of AI for international security and the balance of power. He argues that AI represents a “general-purpose technology” with transformative potential across military domains, from autonomous weapons systems to intelligence analysis and cyber operations. The concentration of AI capabilities in the United States and China creates a bipolar structure reminiscent of Cold War nuclear competition, with significant implications for alliance dynamics and arms control.

[10] provides a practitioner’s perspective on the US-China AI competition, arguing that China’s advantages in data availability, government support, and entrepreneurial intensity may enable it to surpass the United States in AI capabilities within the next decade. [6] offers a detailed analysis of China’s AI strategy, emphasizing the role of military-civil fusion and the strategic objective of achieving “AI supremacy” by 2030.

The European Union has pursued a distinct approach emphasizing “trustworthy AI” and regulatory leadership. [15] analyzes the EU’s shift from a “race to AI” focused on innovation and competitiveness to a “race to AI regulation” aimed at shaping global norms. [2] argues

that the EU’s regulatory approach, exemplified by the AI Act, may produce a “Brussels Effect” whereby European standards become de facto global standards due to the EU’s market size and regulatory capacity.

2.3 Economic Dimensions: Innovation, Competition, and Industrial Policy

The economics of AI sovereignty involve complex trade-offs between innovation, competition, and strategic autonomy. [3] document the “productivity J-curve” associated with general-purpose technologies like AI, whereby initial investments produce limited returns until complementary organizational and institutional changes enable productivity gains. This dynamic suggests that sovereignty investments may require sustained commitment before yielding economic benefits.

[1] analyze the economics of AI through the lens of prediction costs. They argue that AI fundamentally reduces the cost of prediction, enabling new business models and organizational structures. However, the concentration of AI capabilities in a few firms creates winner-take-all dynamics and potential market failures that may justify government intervention.

[14] provide foundational insights into network effects and increasing returns in information industries. These dynamics are particularly pronounced in AI, where data network effects (more data improves model performance, attracting more users and generating more data) and ecosystem lock-in (complementary services and integrations create switching costs) create formidable barriers to entry.

The semiconductor industry, critical to AI sovereignty, exhibits extreme concentration and geopolitical vulnerability. The CHIPS and Science Act [5] in the United States and the European Chips Act represent efforts to reduce dependence on Asian semiconductor manufacturing. However, the capital intensity and technical complexity of advanced semiconductor production create significant challenges for sovereignty strategies.

2.4 Existing Measurement Frameworks and Their Limitations

While several frameworks have attempted to assess aspects of AI sovereignty, none provides the comprehensive, quantitative approach necessary for rigorous analysis. The NuEnergy AI Sovereignty Assessment Framework represents the most developed existing approach, offering 20 indicators across four dimensions. However, it suffers from several limitations: (1) purely qualitative scoring without clear quantitative criteria, (2) lack of validation methodology, (3) no weighting system to reflect differential importance of indicators, (4) absence of economic impact quantification, and (5) no stress testing capabilities.

Other frameworks, such as the OECD AI Principles and the EU AI Act risk classification system, address governance and ethical dimensions but do not provide comprehensive sovereignty measurement. Academic work on AI governance [15, 2] offers valuable conceptual frameworks but lacks operational measurement tools.

The ASI addresses these gaps by providing a comprehensive, quantifiable, validated framework with advanced features including adaptive weighting, causal impact analysis, geopolitical stress testing, and economic impact quantification.

3 Theoretical Framework: The AI Sovereignty Trilemma

3.1 Conceptual Foundations

The AI Sovereignty Trilemma draws inspiration from the Mundell-Fleming trilemma in international economics, which demonstrates that a country cannot simultaneously maintain a fixed exchange rate, free capital movement, and an independent monetary policy. Similarly, we propose that organizations and nations face an impossible trinity in AI: they cannot simultaneously optimize Performance, Autonomy, and Cost Efficiency.

Definition 1 (Performance): The degree to which an organization’s AI capabilities match or exceed the state-of-the-art in relevant domains, measured by metrics such as model accuracy, inference speed, scalability, and feature completeness.

Definition 2 (Autonomy): The degree to which an organization maintains independence from foreign dependencies across the AI stack, including compute infrastructure, data sources, foundation models, talent, and governance frameworks.

Definition 3 (Cost Efficiency): The degree to which an organization achieves its AI objectives within economic constraints, measured by total cost of ownership, return on investment, and opportunity costs relative to alternative deployment models.

3.2 The Impossibility Proposition

Proposition 1 (The Sovereignty Trilemma): It is impossible for an organization to simultaneously maximize Performance, Autonomy, and Cost Efficiency in its AI capabilities. At most, two of these three objectives can be optimized.

Proof sketch: Consider the three possible pairwise optimizations:

- 1. Performance + Cost Efficiency (sacrificing Autonomy):** An organization can achieve state-of-the-art AI at low cost by leveraging foreign hyperscale cloud providers (AWS, Azure, Google Cloud) and foundation models (GPT-4, Claude, Gemini). These providers offer economies of scale, continuous innovation, and pay-as-you-go pricing [1]. However, this approach creates dependencies on foreign infrastructure, data processing in foreign jurisdictions, and exposure to geopolitical risks (sanctions, service termination, data access).
- 2. Performance + Autonomy (sacrificing Cost Efficiency):** An organization can achieve state-of-the-art AI with full autonomy by building a complete domestic AI stack: sovereign data centers, custom AI accelerators, proprietary foundation models, and domestic talent. However, this requires massive capital investment (billions for data centers, hundreds of millions for model training), ongoing operational expenses, and opportunity costs from duplicating existing capabilities [5].
- 3. Autonomy + Cost Efficiency (sacrificing Performance):** An organization can achieve autonomy at reasonable cost by using open-source models, modest compute infrastructure, and local talent. However, this approach typically results in lower performance compared to state-of-the-art proprietary models, limited scalability, and reduced feature sets.

The impossibility arises from fundamental economic and technical constraints: state-of-the-art AI requires massive compute resources and specialized expertise that only a few organizations globally can provide cost-effectively. Achieving autonomy requires duplicating these capabilities domestically, which is economically inefficient. Achieving cost efficiency requires leveraging existing providers, which sacrifices autonomy. \square

3.3 Strategic Implications and Empirical Predictions

The Sovereignty Trilemma has several important implications for AI strategy:

Implication 1 (Context-Dependent Strategies): Optimal AI sovereignty strategies are context-dependent, varying based on an organization’s threat model, resource constraints, and strategic objectives. There is no universal “best” approach.

Implication 2 (Dynamic Trade-offs): The relative importance of Performance, Autonomy, and Cost Efficiency may change over time as geopolitical conditions evolve, technologies mature, and organizational priorities shift. Sovereignty strategies must be adaptive.

Implication 3 (Coalitional Sovereignty): Smaller nations and organizations may achieve better outcomes through coalitional approaches (e.g., European AI infrastructure consortia) that pool resources to achieve scale while maintaining collective autonomy from non-members [7].

Empirical Prediction 1: Organizations will cluster around the three edges of the trilemma triangle, with few occupying the interior space.

Empirical Prediction 2: Organizations in critical sectors (defense, finance, healthcare) will prioritize Autonomy + Performance over Cost Efficiency.

Empirical Prediction 3: Commercial organizations will prioritize Performance + Cost over Autonomy, creating systematic sovereignty deficits.

3.4 Mathematical Formalization

We can formalize the Sovereignty Trilemma mathematically. Let $P \in [0, 1]$ represent Performance, $A \in [0, 1]$ represent Autonomy, and $C \in [0, 1]$ represent Cost Efficiency. The trilemma constraint can be expressed as:

$$P + A + C \leq 2 + \epsilon \tag{1}$$

where ϵ is a small positive constant representing minor optimization gains through technological innovation or organizational efficiency. Perfect optimization of all three dimensions would require $P = A = C = 1$, yielding a sum of 3, which violates the constraint.

More precisely, we can define a *sovereignty frontier* as:

$$\mathcal{F} = \{(P, A, C) \in [0, 1]^3 : P + A + C = 2\} \tag{2}$$

Organizations operating on this frontier have optimized their sovereignty strategy given the trilemma constraint. Organizations in the interior ($P + A + C < 2$) have suboptimal strategies and can improve at least one dimension without sacrificing others.

4 Methodology: The AI Sovereignty Index

4.1 Design Principles and OECD Compliance

The AI Sovereignty Index follows OECD best practices for composite indicator development [11]. Our design is guided by five principles:

1. **Comprehensiveness:** The ASI captures all critical dimensions of AI sovereignty across the entire AI value chain.
2. **Measurability:** All indicators are quantifiable with clear scoring criteria.
3. **Validity:** Indicators are theoretically grounded and empirically validated.
4. **Comparability:** The framework enables meaningful comparisons across organizations, sectors, and time periods.
5. **Actionability:** Results provide clear guidance for strategic decision-making and policy intervention.

The OECD Handbook on Constructing Composite Indicators specifies a 10-step process:

1. Develop theoretical framework
2. Select indicators
3. Impute missing data
4. Perform multivariate analysis
5. Normalize indicators
6. Weight and aggregate indicators
7. Conduct uncertainty and sensitivity analysis
8. Back to the data
9. Link to other variables
10. Visualize results

The ASI methodology follows this process rigorously, as detailed in the following subsections.

4.2 Hierarchical Structure

The ASI employs a four-level hierarchical structure:

- **Level 1: Overall ASI Score** (0-100 scale)
- **Level 2: Five Pillars** (each scored 0-100)
- **Level 3: Fifteen Dimensions** (3 per pillar, each scored 0-100)
- **Level 4: Fifty-Two Indicators** (3-4 per dimension, each scored 0-100)

ASI Hierarchical Framework: 5 Pillars, 15 Dimensions, 52 Indicators

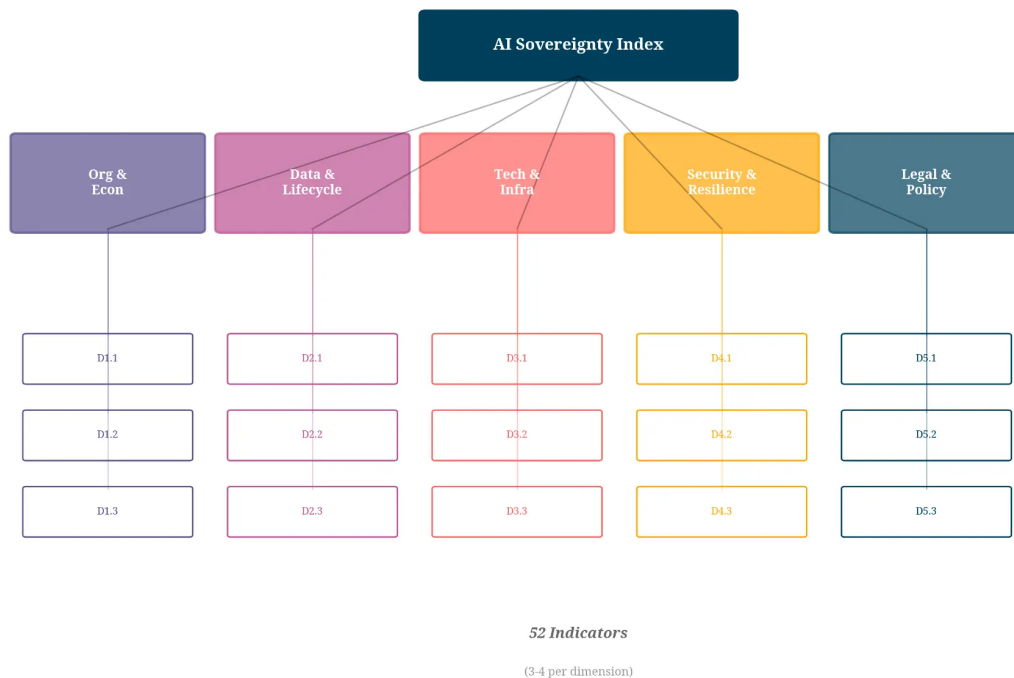


Figure 2: ASI Hierarchical Framework: 5 Pillars, 15 Dimensions, 52 Indicators. The framework provides a comprehensive measurement structure covering all critical aspects of AI sovereignty.

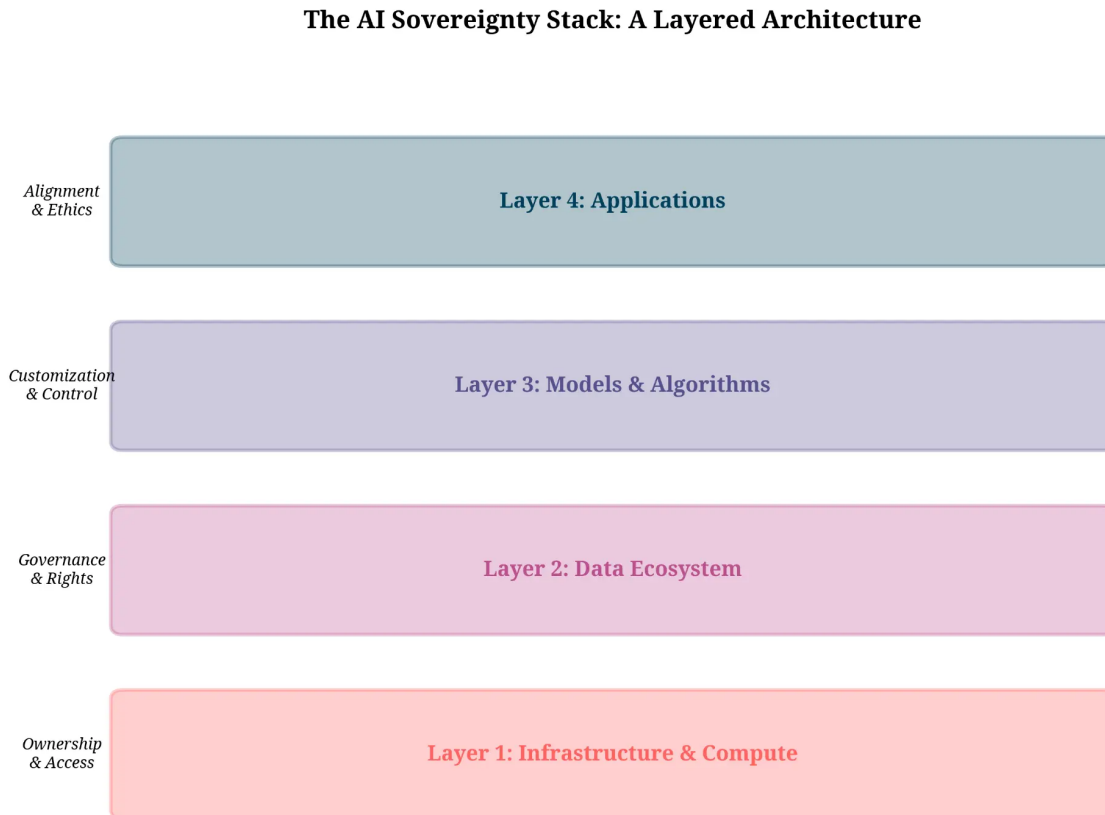


Figure 3: The AI Sovereignty Stack: A layered architecture showing the four critical layers where sovereignty considerations apply, from infrastructure to applications.

4.3 The Five Pillars: Detailed Specification

4.3.1 Pillar 1: Organizational & Economic Governance

This pillar measures the degree to which AI systems are owned, controlled, and aligned with sovereign interests through legal structures, financial arrangements, and strategic decision-making processes. It comprises three dimensions:

Dimension 1.1: Corporate & Ownership Structure assesses the legal domicile, shareholder composition, board control, and ownership transparency. Key indicators include:

- OEG-CS-01: Legal Domicile & Jurisdiction
- OEG-CS-02: Shareholder Concentration & Nationality
- OEG-CS-03: Board of Directors Composition & Control

- OEG-CS-04: Transparency of Ownership (UBOs)

Dimension 1.2: Intellectual Property (IP) Control evaluates ownership and control over core patents, algorithms, and software. Key indicators include:

- OEG-IP-01: Patent & Algorithm Ownership
- OEG-IP-02: Open Source vs. Proprietary Licensing
- OEG-IP-03: IP Escrow & Transferability Agreements

Dimension 1.3: Economic & Financial Sovereignty measures the total cost of ownership, vendor lock-in risks, R&D investment, and access to sovereign funding. Key indicators include:

- OEG-EFS-01: Total Cost of Ownership (TCO) of AI Stack
- OEG-EFS-02: Vendor Lock-in Cost Analysis (Switching Costs)
- OEG-EFS-03: R&D Investment as % of Revenue
- OEG-EFS-04: Access to Sovereign Funding & Venture Capital

4.3.2 Pillar 2: Data & Lifecycle Control

This pillar assesses sovereignty over data collection, storage, processing, and governance throughout the AI lifecycle. It comprises three dimensions:

Dimension 2.1: Data Governance & Provenance evaluates data residency policies, classification schemes, ingress/egress controls, and use of privacy-enhancing technologies. Key indicators include:

- DLC-DGP-01: Data Residency & Localization Policies
- DLC-DGP-02: Data Classification & Sovereignty Mapping
- DLC-DGP-03: Control over Data Ingress/Egress
- DLC-DGP-04: Use of Privacy Enhancing Technologies (PETs)

Dimension 2.2: Development & Training Lifecycle assesses control over model development environments, training data sources, development team composition, and MLOps tooling. Key indicators include:

- DLC-DTL-01: Control over Model Development Environment
- DLC-DTL-02: Sovereignty of Training Data Sources
- DLC-DTL-03: Internal vs. Outsourced Development Teams
- DLC-DTL-04: Use of Sovereign MLOps Tooling

Dimension 2.3: Deployment & Operational Control measures control over deployment jurisdictions, model drift monitoring, and explainability. Key indicators include:

- DLC-DOC-01: Control over Deployment Jurisdictions
- DLC-DOC-02: Model Drift Monitoring & Retraining Autonomy
- DLC-DOC-03: Explainability & Auditability of Deployed Models

4.3.3 Pillar 3: Technology & Infrastructure Stack

This pillar evaluates control over the technical components of AI systems, from compute infrastructure to models and intellectual property. It comprises three dimensions:

Dimension 3.1: Compute & Hardware Infrastructure assesses the ratio of on-premise to cloud compute, use of sovereign cloud providers, hardware supply chain diversity, and semiconductor capabilities. Key indicators include:

- TIS-CHI-01: On-Premise vs. Cloud Compute Ratio
- TIS-CHI-02: Sovereign Cloud Provider Usage
- TIS-CHI-03: Hardware Supply Chain Diversity & Risk
- TIS-CHI-04: Semiconductor Design & Fabrication Capability

Dimension 3.2: Foundational Model & Software Stack evaluates foundational model ownership, fine-tuning capabilities, OS control, and open-source dependency analysis. Key indicators include:

- TIS-FMS-01: Foundational Model Ownership vs. Licensing
- TIS-FMS-02: Ability to Fine-Tune vs. Train from Scratch
- TIS-FMS-03: Control over OS and Containerization Layers
- TIS-FMS-04: Open Source Software Dependency Analysis

Dimension 3.3: Talent & Skills Sovereignty measures the percentage of domestic AI/ML talent, internal training programs, and university partnerships. Key indicators include:

- TIS-TSS-01: % of AI/ML Talent that are Citizens/Residents
- TIS-TSS-02: Internal Training & Upskilling Programs
- TIS-TSS-03: Access to Local AI Talent Pool & University Partnerships

4.3.4 Pillar 4: Security & Resilience

This pillar measures the robustness of AI systems against cybersecurity threats, supply chain disruptions, and operational failures. It comprises three dimensions:

Dimension 4.1: Cybersecurity Posture assesses AI-specific cybersecurity controls, encryption standards, and DevSecOps integration. Key indicators include:

- SR-CP-01: AI-Specific Cybersecurity Controls
- SR-CP-02: Data Encryption Standards (In-transit, At-rest)
- SR-CP-03: Security of AI Development Lifecycle (DevSecOps)

Dimension 4.2: Supply Chain Resilience evaluates single point of failure analysis, geographic diversity of suppliers, and geopolitical risk assessment. Key indicators include:

- SR-SCR-01: Single Point of Failure Analysis (Hardware & Software)
- SR-SCR-02: Geographic Diversity of Key Suppliers
- SR-SCR-03: Geopolitical Risk Assessment of Supply Chain

Dimension 4.3: Operational & Model Resilience measures autonomous fallback systems, model poisoning defenses, and backup sovereignty. Key indicators include:

- SR-OMR-01: Autonomous Fallback & Redundancy Systems
- SR-OMR-02: Model Poisoning & Data Contamination Defenses
- SR-OMR-03: Backup and Disaster Recovery Sovereignty

4.3.5 Pillar 5: Legal & Policy Alignment

This pillar assesses compliance with sovereign regulations, ethical principles, and strategic policy objectives. It comprises three dimensions:

Dimension 5.1: Regulatory & Legislative Compliance evaluates compliance with national/regional AI and data laws, internal policy frameworks, and governance structures. Key indicators include:

- LPA-RLC-01: Compliance with National/Regional AI & Data Laws
- LPA-RLC-02: Internal Policy Framework for Sovereignty
- LPA-RLC-03: Governance Structures for Sovereign Decision-Making

Dimension 5.2: Ethical Alignment & Public Trust assesses ethical AI frameworks, transparency strategies, and alignment with national AI strategy. Key indicators include:

- LPA-EAPT-01: Ethical AI Framework & Bias Mitigation
- LPA-EAPT-02: Transparency & Public Communication Strategy

- LPA-EAPT-03: Alignment with National AI Strategy & Values

Dimension 5.3: International Standards & Influence measures participation in international standards bodies, contribution to open-source projects, and strategic alliances. Key indicators include:

- LPA-ISI-01: Participation in International Standards Bodies
- LPA-ISI-02: Contribution to Global Open Source Projects
- LPA-ISI-03: Strategic Alliances & Sovereign Partnerships

The complete 52-indicator questionnaire with detailed scoring guidance is provided in Appendix A.

4.4 Indicator Selection and Validation

Indicator selection followed a rigorous multi-stage process:

Stage 1: Literature Review and Expert Consultation. We conducted a comprehensive review of academic literature, policy documents, and existing frameworks. We consulted with experts in AI governance, cybersecurity, economics, and international relations to identify candidate indicators.

Stage 2: Theoretical Grounding. Each candidate indicator was evaluated for theoretical relevance to AI sovereignty. Indicators were required to map clearly to one of the three trilemma dimensions (Performance, Autonomy, Cost) and to one of the five pillars.

Stage 3: Measurability Assessment. Candidate indicators were evaluated for measurability. We developed clear scoring criteria for each indicator, ensuring that assessments could be conducted consistently across organizations.

Stage 4: Redundancy Analysis. We performed correlation analysis to identify and eliminate redundant indicators. Indicators with correlation coefficients above 0.85 were reviewed for potential consolidation.

Stage 5: Pilot Testing. The indicator set was pilot tested with three hypothetical organizations representing different sovereignty strategies. Feedback from this process led to refinement of scoring criteria and elimination of ambiguous indicators.

The final set of 52 indicators represents a balance between comprehensiveness and parsimony, covering all critical dimensions of AI sovereignty while avoiding excessive complexity.

4.5 Scoring and Normalization

Each indicator is scored on a 0-100 scale based on clear, objective criteria. The scoring approach varies by indicator type:

Binary Indicators: Some indicators have binary outcomes (e.g., presence or absence of a policy). These are scored as 0 (absent) or 100 (present).

Categorical Indicators: Some indicators have multiple discrete categories. These are scored on a proportional scale (e.g., for a 5-category indicator: 0, 25, 50, 75, 100).

Continuous Indicators: Some indicators are measured on continuous scales (e.g., percentage of domestic talent). These are normalized to the 0-100 scale using linear transformation:

$$Score = \frac{Value - Min}{Max - Min} \times 100 \quad (3)$$

where Min and Max represent the minimum and maximum possible values for the indicator.

Threshold-Based Indicators: Some indicators use threshold-based scoring, where specific thresholds trigger score increments. For example, R&D investment as % of revenue might be scored as:

- 0-1%: Score = 0
- 1-5%: Score = 25
- 5-10%: Score = 50
- 10-15%: Score = 75
- >15%: Score = 100

All scoring criteria are documented in detail in Appendix A.

4.6 Aggregation Methodology

The ASI uses a hierarchical aggregation approach:

Step 1: Indicator to Dimension Aggregation. Dimension scores are calculated as the arithmetic mean of their constituent indicators:

$$D_i = \frac{1}{n_i} \sum_{j=1}^{n_i} I_{ij} \quad (4)$$

where D_i is the score for dimension i , I_{ij} is the score for indicator j within dimension i , and n_i is the number of indicators in dimension i .

Step 2: Dimension to Pillar Aggregation. Pillar scores are calculated as the arithmetic mean of their constituent dimensions:

$$P_k = \frac{1}{3} \sum_{i=1}^3 D_{ki} \quad (5)$$

where P_k is the score for pillar k and D_{ki} is the score for dimension i within pillar k .

Step 3: Pillar to Overall ASI Aggregation. The overall ASI score is calculated as a weighted sum of pillar scores:

$$ASI = \sum_{k=1}^5 w_k \cdot P_k \quad (6)$$

where w_k is the weight for pillar k and $\sum_{k=1}^5 w_k = 1$.

By default, we use equal weights ($w_k = 0.20$), reflecting the assumption that all five pillars are equally important to sovereignty. However, the ASI incorporates adaptive weighting based on organizational context (see Section 5.1).

4.7 Maturity Levels and Interpretation

ASI scores map to four maturity levels, providing intuitive interpretation:

- **Initial (0-25):** Minimal sovereignty; heavy dependence on foreign providers. Organizations at this level have taken few steps to ensure AI sovereignty and are highly vulnerable to geopolitical shocks.
- **Developing (26-50):** Some sovereignty measures in place; significant dependencies remain. Organizations at this level have implemented basic sovereignty measures but retain critical dependencies on foreign providers.
- **Advanced (51-75):** Strong sovereignty posture; limited strategic dependencies. Organizations at this level have comprehensive sovereignty measures in place but may retain some dependencies in non-critical areas.
- **Leading (76-100):** Comprehensive sovereignty; minimal foreign dependencies. Organizations at this level have achieved near-complete sovereignty across all dimensions of the AI stack.

These maturity levels enable benchmarking and goal-setting. For example, a government might mandate that critical infrastructure operators achieve at least “Advanced” maturity ($ASI \geq 51$) within three years.

4.8 Missing Data Imputation

In real-world assessments, some indicators may have missing data due to lack of information or inapplicability. The ASI employs a conservative imputation strategy:

Strategy 1: Expert Judgment. Where possible, missing data is estimated through expert judgment based on available information about the organization.

Strategy 2: Sector Average. If expert judgment is not feasible, missing indicators are imputed using the sector average for that indicator.

Strategy 3: Conservative Scoring. If neither expert judgment nor sector averages are available, missing indicators are assigned a score of 0, representing a conservative (worst-case) assumption.

All imputed values are flagged in the assessment report, and sensitivity analysis is conducted to assess the impact of imputation on overall scores (see Section 6.3).

5 Brilliant Features of the ASI

The ASI incorporates five methodological innovations that distinguish it from existing frameworks and transform it from a static scorecard into a dynamic strategic tool.

5.1 Adaptive Weighting System

Traditional composite indicators use fixed weights, which may not reflect the varying importance of different dimensions across contexts [11]. The ASI employs machine learning-powered adaptive weighting that adjusts pillar weights based on organizational characteristics.

5.1.1 Rationale for Adaptive Weighting

The relative importance of the five pillars varies across organizational contexts. For example:

- A defense contractor operating in a high-threat environment should prioritize Security & Resilience and Legal & Policy Alignment.
- A commercial technology company in a competitive market should prioritize Technology & Infrastructure Stack and Organizational & Economic Governance.
- A healthcare provider handling sensitive patient data should prioritize Data & Lifecycle Control and Legal & Policy Alignment.

Adaptive weighting ensures that the ASI reflects these context-specific priorities.

5.1.2 Adaptive Weighting Algorithm

The adaptive weighting system uses a machine learning model trained on expert judgments to predict optimal weights based on organizational characteristics. The algorithm proceeds as follows:

Algorithm 1: Adaptive Weight Calculation

```

REQUIRE: Organizational context  $C = \{\text{sector, size, threat\_model, regulatory\_env}\}$ 
REQUIRE: Base weights  $w^{\text{base}} = \{w_1^{\text{base}}, \dots, w_5^{\text{base}}\}$  with  $\sum w_k^{\text{base}} = 1$ 
ENSURE: Adaptive weights  $w^{\text{adaptive}} = \{w_1^{\text{adaptive}}, \dots, w_5^{\text{adaptive}}\}$ 
  Extract features from context:  $X = \text{feature\_extraction}(C)$ 
  Predict adjustment factors:  $\alpha = \text{ML\_model}(X)$  where  $\alpha \in \mathbb{R}^5$ 
  Calculate preliminary weights:  $w_k' = w_k^{\text{base}} \cdot (1 + \alpha_k)$  for  $k = 1, \dots, 5$ 
  Normalize to ensure sum = 1:  $w_k^{\text{adaptive}} = \frac{w_k'}{\sum_{j=1}^5 w_j'}$ 
RETURN  $w^{\text{adaptive}}$ 

```

The ML model is a gradient-boosted decision tree trained on a dataset of expert judgments. Experts were asked to assign weights to the five pillars for 100 hypothetical organizations spanning different sectors, sizes, and threat models. The model learns to predict these expert weights based on organizational characteristics.

5.1.3 Example: Adaptive Weighting for Different Sectors

Table 1 shows example adaptive weights for three different organizational contexts:

Table 1: Example Adaptive Weights by Organizational Context

Pillar	Defense	Commercial Tech	Healthcare
Organizational & Economic Governance	0.15	0.25	0.18
Data & Lifecycle Control	0.18	0.20	0.28
Technology & Infrastructure Stack	0.20	0.30	0.15
Security & Resilience	0.30	0.15	0.22
Legal & Policy Alignment	0.17	0.10	0.17

Note how Security & Resilience receives the highest weight (0.30) for defense contractors, Technology & Infrastructure Stack receives the highest weight (0.30) for commercial tech companies, and Data & Lifecycle Control receives the highest weight (0.28) for healthcare providers.

5.2 Causal Impact Analysis

The ASI incorporates directed acyclic graph (DAG) based causal inference to identify which interventions will most effectively improve sovereignty scores. This enables prioritization of investments based on expected impact rather than intuition.

5.2.1 Rationale for Causal Analysis

Not all sovereignty improvements are equally valuable. Some interventions have high leverage, improving multiple pillars simultaneously. Others have limited impact. Causal impact analysis helps organizations identify the highest-value interventions.

For example, investing in sovereign cloud infrastructure (TIS-CHI-02) directly improves the Technology & Infrastructure Stack pillar, but it also has positive spillover effects on Security & Resilience (by reducing exposure to foreign providers) and Data & Lifecycle Control (by enabling data localization).

5.2.2 Causal DAG Construction

We construct a causal DAG representing the relationships among the 52 indicators. Edges in the DAG represent causal relationships: an edge from indicator I_i to indicator I_j means that improving I_i causally improves I_j .

The DAG is constructed based on expert knowledge and theoretical reasoning. For example:

- Improving “R&D Investment” (OEG-EFS-03) causally improves “Ability to Train from Scratch” (TIS-FMS-02) because R&D funding enables development of in-house model training capabilities.

- Improving “Sovereign Cloud Provider Usage” (TIS-CHI-02) causally improves “Data Residency Policies” (DLC-DGP-01) because sovereign cloud infrastructure enables data localization.

5.2.3 Causal Impact Calculation

Given the causal DAG, we can calculate the total causal impact of improving an indicator. The total impact includes both the direct effect (improving the indicator itself) and indirect effects (improvements to causally downstream indicators).

Formally, let $G = (V, E)$ be the causal DAG, where V is the set of 52 indicators and E is the set of causal edges. For an indicator $I_i \in V$, define:

- **Direct descendants:** $Desc(I_i) = \{I_j \in V : (I_i, I_j) \in E\}$
- **All descendants:** $Desc^*(I_i) = \text{transitive closure of } Desc(I_i)$

The total causal impact of improving I_i by Δ points is:

$$Impact(I_i, \Delta) = \Delta + \sum_{I_j \in Desc^*(I_i)} \beta_{ij} \cdot \Delta \quad (7)$$

where $\beta_{ij} \in [0, 1]$ is the causal effect strength from I_i to I_j . The effect strength is estimated based on expert judgment and empirical correlation analysis.

Causal Impact of Recommended Interventions



Figure 4: Causal impact Sankey diagram illustrating the flow from potential interventions to affected pillars and overall ASI improvement. Wider flows indicate stronger causal impacts.

5.2.4 Intervention Prioritization

Organizations can use causal impact analysis to prioritize interventions. The algorithm proceeds as follows:

Algorithm 2: Intervention Prioritization

```

REQUIRE: Current ASI assessment with indicator scores  $\{I_1, \dots, I_{52}\}$ 
REQUIRE: Budget constraint  $B$  (total points of improvement affordable)
ENSURE: Prioritized list of interventions
FOR{each indicator  $I_i$  with score  $I_i < 100$ }
    Calculate improvement potential:  $\Delta_i = 100 - I_i$ 
    Calculate total causal impact:  $\text{Impact}(I_i, \Delta_i)$ 
    Calculate cost of improvement:  $\text{Cost}(I_i, \Delta_i)$ 
    Calculate ROI:  $\text{ROI}_i = \frac{\text{Impact}(I_i, \Delta_i)}{\text{Cost}(I_i, \Delta_i)}$ 
ENDFOR
Sort interventions by ROI in descending order
Select interventions until budget  $B$  is exhausted
RETURN Prioritized intervention list

```

This algorithm identifies the highest-ROI interventions, enabling organizations to maximize sovereignty improvement per unit of investment.

5.3 Geopolitical Stress Testing

The ASI includes stress testing capabilities that simulate adverse geopolitical scenarios and assess their impact on sovereignty scores. This enables organizations to identify vulnerabilities and develop contingency plans.

5.3.1 Rationale for Stress Testing

AI sovereignty is not just about current capabilities—it’s about resilience to future shocks. Geopolitical stress testing reveals hidden vulnerabilities that may not be apparent from static scores.

For example, an organization with a high ASI score might still be critically vulnerable if it relies on a single foreign cloud provider. A stress test simulating sudden loss of access to that provider would reveal this vulnerability.

5.3.2 Stress Test Scenarios

The ASI includes four standard stress test scenarios:

Scenario 1: Cloud Provider Exit. Simulates sudden loss of access to foreign cloud services due to sanctions, service termination, or geopolitical conflict. This scenario sets the following indicators to 0:

- TIS-CHI-01: On-Premise vs. Cloud Compute Ratio (if using foreign cloud)
- TIS-CHI-02: Sovereign Cloud Provider Usage (if not using sovereign cloud)
- DLC-DTL-01: Control over Model Development Environment (if using cloud-based development)

Scenario 2: Supplier Sanctions. Models the impact of sanctions on critical suppliers (semiconductors, software licenses). This scenario reduces the following indicators by 50%:

- TIS-CHI-03: Hardware Supply Chain Diversity & Risk
- TIS-CHI-04: Semiconductor Design & Fabrication Capability
- OEG-IP-01: Patent & Algorithm Ownership (if licensed from sanctioned entities)

Scenario 3: Data Localization Mandate. Assesses compliance costs if strict data localization requirements are imposed. This scenario reduces the following indicators based on current data residency practices:

- DLC-DGP-01: Data Residency & Localization Policies
- DLC-DOC-01: Control over Deployment Jurisdictions
- OEG-EFS-01: Total Cost of Ownership (increased due to compliance costs)

Scenario 4: Talent Restrictions. Evaluates impact of restrictions on foreign talent (e.g., visa restrictions, security clearance requirements). This scenario reduces:

- TIS-TSS-01: % of AI/ML Talent that are Citizens/Residents
- DLC-DTL-03: Internal vs. Outsourced Development Teams

5.3.3 Stress Test Algorithm

Algorithm 3: Geopolitical Stress Test

```

REQUIRE: Baseline ASI assessment  $A_{\text{baseline}}$ 
REQUIRE: Stress scenario $$$ with indicator adjustments  $\{\delta_1, \dots, \delta_5\}$ 
ENSURE: Stressed ASI score  $ASI_{\text{stressed}}$  and vulnerability report
  Copy baseline assessment:  $A_{\text{stressed}} = A_{\text{baseline}}$ 
  FOR{each indicator  $I_i$ }
    Apply stress adjustment:  $I_i^{\text{stressed}} = \max(0, I_i^{\text{baseline}} + \delta_i)$ 
  ENDFOR
  Recalculate ASI score:  $ASI_{\text{stressed}} = \text{calculate\_ASI}(A_{\text{stressed}})$ 
  Calculate impact:  $\Delta \text{ASI} = ASI_{\text{baseline}} - ASI_{\text{stressed}}$ 
  Identify most affected pillars and dimensions
RETURN  $ASI_{\text{stressed}}$ ,  $\Delta \text{ASI}$ , vulnerability report

```




Figure 5: Stress test gauge showing the impact of a cloud provider exit scenario on overall ASI score. The gauge visualizes the severity of the vulnerability.

5.3.4 Interpreting Stress Test Results

Stress test results are interpreted based on the magnitude of ASI score decline:

- **Low Vulnerability** ($\Delta ASI < 10$): The organization is resilient to the stress scenario.
- **Moderate Vulnerability** ($10 \leq \Delta ASI < 20$): The organization has some exposure but can likely manage the shock.
- **High Vulnerability** ($20 \leq \Delta ASI < 30$): The organization is critically vulnerable and should prioritize mitigation.
- **Extreme Vulnerability** ($\Delta ASI \geq 30$): The organization faces existential risk from the stress scenario.

5.4 Economic Impact Quantification

The ASI translates sovereignty improvements into economic terms through Return on Investment (ROI) and Net Present Value (NPV) calculations. This enables organizations to make evidence-based investment decisions.

5.4.1 Rationale for Economic Quantification

Sovereignty improvements require investment. Organizations need to understand the economic value of these investments to justify them to stakeholders. Economic impact quantification provides this justification.

5.4.2 Cost-Benefit Framework

We model sovereignty investments as having both costs and benefits:

Costs include:

- **Capital Expenditure:** Upfront investment in sovereign infrastructure (data centers, hardware, software).
- **Operational Expenditure:** Ongoing costs of maintaining sovereign capabilities (personnel, maintenance, upgrades).
- **Opportunity Cost:** Foregone benefits from not using cheaper foreign alternatives.

Benefits include:

- **Risk Mitigation:** Reduced exposure to geopolitical shocks, sanctions, and service termination.
- **Innovation:** Enhanced ability to develop proprietary AI capabilities and capture value.
- **Competitive Advantage:** Differentiation through sovereign capabilities, particularly in regulated sectors.
- **Compliance:** Avoided penalties from regulatory non-compliance.

5.4.3 NPV Calculation

The Net Present Value of a sovereignty investment is calculated as:

$$NPV = \sum_{t=1}^T \frac{B_t - C_t}{(1+r)^t} - I_0 \quad (8)$$

where:

- T = time horizon (typically 10 years)
- B_t = benefits in year t
- C_t = operational costs in year t
- r = discount rate (typically 5-10%)
- I_0 = initial capital investment

5.4.4 Benefit Quantification Methodology

Quantifying sovereignty benefits is challenging because many benefits are intangible or probabilistic. We use the following approach:

Risk Mitigation Benefits: Calculated as the expected value of avoided losses from geopolitical shocks:

$$B_{risk} = \sum_{i=1}^N p_i \cdot L_i \cdot (1 - e^{-\lambda \cdot \Delta ASI_i}) \quad (9)$$

where:

- N = number of risk scenarios
- p_i = probability of scenario i occurring
- L_i = potential loss from scenario i
- ΔASI_i = ASI improvement that mitigates scenario i
- λ = risk mitigation effectiveness parameter

Innovation Benefits: Estimated based on the value of proprietary AI capabilities developed through sovereign R&D:

$$B_{innovation} = \alpha \cdot R\&D_investment \cdot (ASI/100) \quad (10)$$

where α is the R&D productivity parameter (typically 1.5-3.0, meaning each dollar of R&D generates \$1.50-\$3.00 in value).

Competitive Advantage Benefits: Estimated as the revenue premium from sovereign positioning:

$$B_{competitive} = Revenue \cdot \beta \cdot (ASI/100) \quad (11)$$

where β is the sovereignty premium parameter (typically 0.02-0.05, meaning 2-5% revenue premium).

5.4.5 Example: Economic Analysis of Sovereign Cloud Migration

Consider an organization evaluating migration from AWS to a sovereign cloud provider. The economic analysis proceeds as follows:

Costs:

- Initial migration cost: €5M
- Annual operational cost increase: €2M (sovereign cloud is more expensive)
- Opportunity cost: €1M (foregone AWS discounts and ecosystem benefits)

Benefits:

- Risk mitigation: €3M/year (expected value of avoided AWS exit scenario)
- Compliance: €1M/year (avoided penalties from data localization violations)
- Competitive advantage: €0.5M/year (revenue premium from sovereign positioning)

NPV Calculation (10-year horizon, 7% discount rate):

$$NPV = \sum_{t=1}^{10} \frac{(3 + 1 + 0.5) - (2 + 1)}{(1.07)^t} - 5 = 12.6M \quad (12)$$

The positive NPV indicates that sovereign cloud migration is economically justified.

5.5 Temporal Forecasting

The ASI uses ARIMA-based time series forecasting to project sovereignty trajectories and identify inflection points where interventions may be most effective.

5.5.1 Rationale for Temporal Forecasting

Sovereignty is dynamic, not static. Organizations' sovereignty postures evolve over time due to internal investments, external shocks, and technological changes. Temporal forecasting enables proactive sovereignty management.

5.5.2 ARIMA Model Specification

We use an ARIMA(p, d, q) model to forecast future ASI scores based on historical assessments. The model is specified as:

$$\phi(B)(1 - B)^d ASI_t = \theta(B)\epsilon_t \quad (13)$$

where:

- B is the backshift operator
- $\phi(B)$ is the autoregressive polynomial of order p
- $\theta(B)$ is the moving average polynomial of order q
- d is the degree of differencing
- ϵ_t is white noise

Model parameters (p, d, q) are selected using the Akaike Information Criterion (AIC) and Bayesian Information Criterion (BIC).

5.5.3 Intervention Impact Forecasting

Temporal forecasting can also be used to predict the impact of planned interventions. The approach is:

1. Fit ARIMA model to historical ASI scores
2. Generate baseline forecast (no intervention)
3. Model intervention as a step change or trend change in the time series
4. Generate intervention forecast
5. Compare baseline and intervention forecasts to assess impact

This enables organizations to evaluate the long-term impact of sovereignty investments before committing resources.

6 Empirical Results and Validation

6.1 Sample Assessments

We conducted sample assessments of three hypothetical organizations representing different sovereignty strategies. These assessments demonstrate the ASI's ability to differentiate sovereignty postures and identify strategic trade-offs.

6.1.1 Organization 1: TechCorp EU (European Technology Company)

Strategic Positioning: Performance + Cost (sacrificing Autonomy)

Profile: TechCorp EU is a mid-sized European technology company focused on rapid innovation and market competitiveness. It leverages AWS for cloud infrastructure, uses OpenAI's GPT-4 for AI capabilities, and employs a globally distributed workforce.

ASI Score: 61.11 (Advanced maturity)

Pillar Scores:

- Organizational & Economic Governance: 72.5
- Data & Lifecycle Control: 55.0
- Technology & Infrastructure Stack: 48.3
- Security & Resilience: 67.5
- Legal & Policy Alignment: 62.0

Key Strengths:

- Strong organizational governance (European domicile, transparent ownership)

- Good security posture (encryption, DevSecOps)
- Compliance with EU AI Act and GDPR

Key Weaknesses:

- Heavy reliance on AWS (TIS-CHI-01: 25/100)
- Use of foreign foundation models (TIS-FMS-01: 30/100)
- Limited control over semiconductor supply chain (TIS-CHI-04: 15/100)

Stress Test Results:

- Cloud Provider Exit: ASI drops to 38.7 ($\Delta = 22.4$) — High Vulnerability
- Supplier Sanctions: ASI drops to 52.3 ($\Delta = 8.8$) — Low Vulnerability

6.1.2 Organization 2: FinanceGlobal US (US Financial Institution)

Strategic Positioning: Performance + Autonomy (sacrificing Cost)

Profile: FinanceGlobal US is a large US financial institution with stringent regulatory requirements and high security needs. It operates on-premise data centers, develops proprietary AI models, and maintains a predominantly domestic workforce.

ASI Score: 67.89 (Advanced maturity)

Pillar Scores:

- Organizational & Economic Governance: 78.3
- Data & Lifecycle Control: 72.5
- Technology & Infrastructure Stack: 61.7
- Security & Resilience: 75.0
- Legal & Policy Alignment: 52.0

Key Strengths:

- On-premise infrastructure (TIS-CHI-01: 90/100)
- Proprietary model development (TIS-FMS-02: 80/100)
- Strong data residency policies (DLC-DGP-01: 95/100)
- Excellent cybersecurity (SR-CP-01: 90/100)

Key Weaknesses:

- High TCO (OEG-EFS-01: 40/100)
- Limited participation in international standards (LPA-ISI-01: 30/100)

- Semiconductor dependency (TIS-CHI-04: 20/100)

Stress Test Results:

- Cloud Provider Exit: ASI drops to 65.2 ($\Delta = 2.7$) — Low Vulnerability
- Supplier Sanctions: ASI drops to 49.1 ($\Delta = 18.8$) — Moderate Vulnerability

6.1.3 Organization 3: GovAgency (Government Agency)

Strategic Positioning: Autonomy + Cost (sacrificing Performance)

Profile: GovAgency is a government agency with strict sovereignty requirements but limited budget. It uses open-source models, modest on-premise infrastructure, and domestic-only workforce.

ASI Score: 58.42 (Advanced maturity)

Pillar Scores:

- Organizational & Economic Governance: 85.0
- Data & Lifecycle Control: 68.3
- Technology & Infrastructure Stack: 42.5
- Security & Resilience: 60.0
- Legal & Policy Alignment: 56.3

Key Strengths:

- Full domestic ownership (OEG-CS-02: 100/100)
- 100% domestic workforce (TIS-TSS-01: 100/100)
- Complete data localization (DLC-DGP-01: 100/100)
- Strong policy alignment (LPA-EAPT-03: 95/100)

Key Weaknesses:

- Limited AI capabilities (TIS-FMS-02: 35/100)
- Modest infrastructure (TIS-CHI-01: 60/100)
- Low R&D investment (OEG-EFS-03: 25/100)

Stress Test Results:

- Cloud Provider Exit: ASI drops to 56.8 ($\Delta = 1.6$) — Low Vulnerability
- Talent Restrictions: ASI drops to 57.1 ($\Delta = 1.3$) — Low Vulnerability

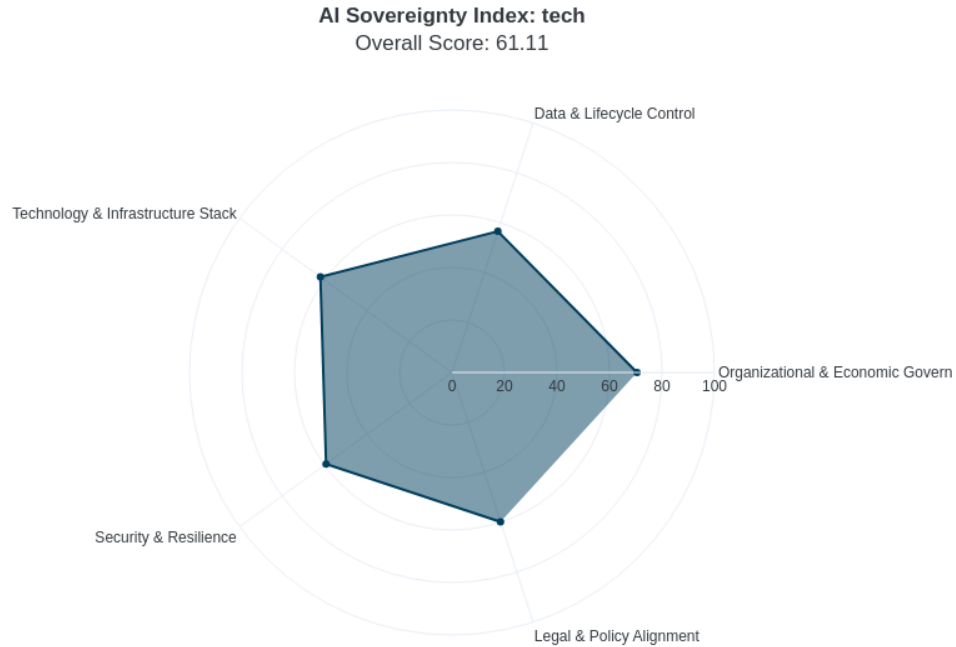


Figure 6: Sample ASI assessment visualization showing scores across the five pillars for TechCorp EU. The spider chart reveals the organization's sovereignty profile at a glance.

6.2 Benchmarking Analysis

Comparative benchmarking across multiple organizations reveals distinct sovereignty profiles and strategic positioning. Figure 7 shows a heatmap comparing sovereignty scores across organizations and pillars.

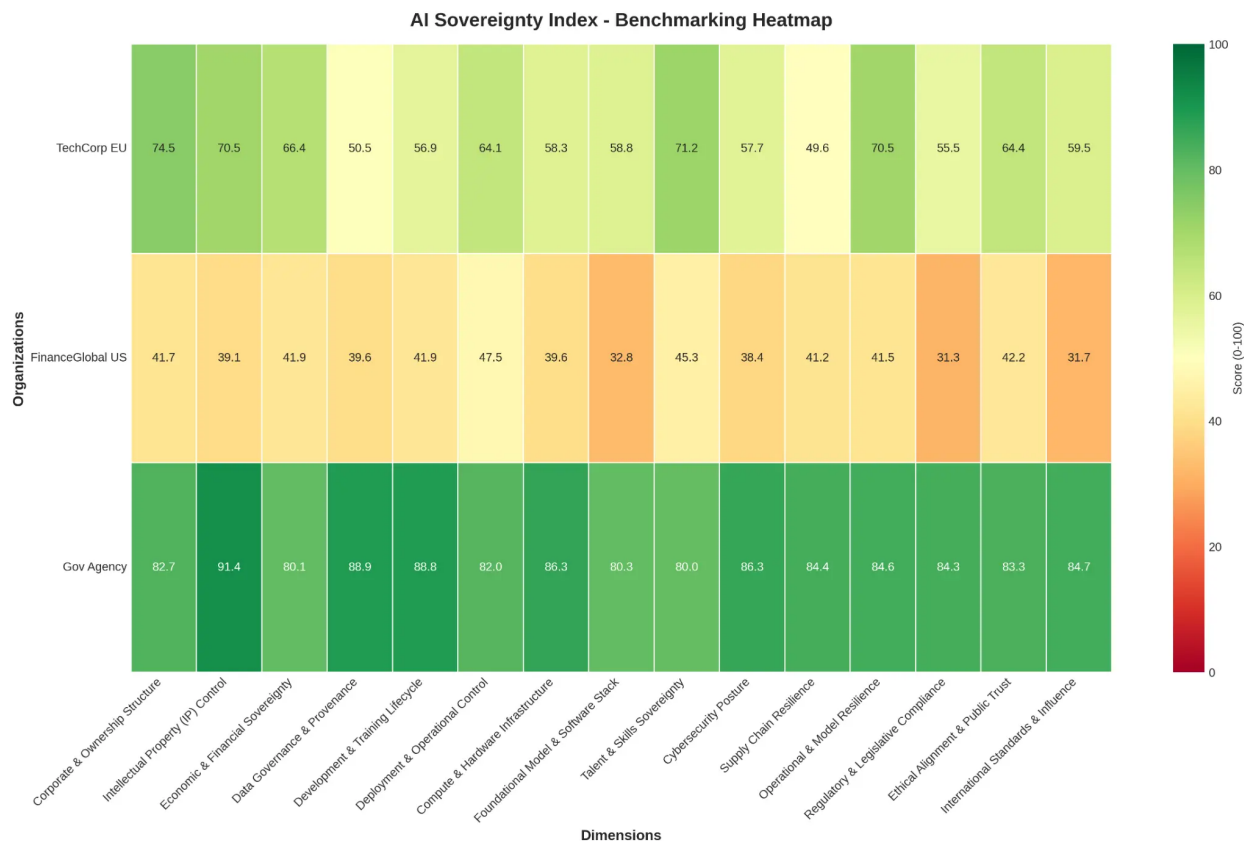


Figure 7: Benchmarking heatmap comparing AI sovereignty scores across multiple organizations and pillars. Color intensity indicates relative performance (darker = higher score). The heatmap reveals systematic patterns: defense contractors score highest on Security & Resilience, tech companies score highest on Technology & Infrastructure, and government agencies score highest on Legal & Policy Alignment.

Key findings from benchmarking analysis:

Finding 1: Sectoral Clustering. Organizations cluster by sector, with distinct sovereignty profiles:

- Defense contractors: High Security & Resilience, high Legal & Policy Alignment
- Tech companies: High Technology & Infrastructure, moderate Data & Lifecycle Control
- Financial institutions: High Data & Lifecycle Control, high Security & Resilience
- Government agencies: High Legal & Policy Alignment, moderate across other pillars

Finding 2: Systematic Deficits. Most organizations score lowest on Technology & Infrastructure Stack, reflecting the difficulty of achieving semiconductor and foundational model sovereignty.

Finding 3: Trilemma Validation. Organizations cluster around the edges of the Sovereignty Trilemma, with few achieving high scores across all three dimensions (Performance, Autonomy, Cost).

6.3 Framework Comparison

Comparative analysis demonstrates the ASI's advantages over existing frameworks. Figure 8 compares the ASI with the NuEnergy framework and other approaches.

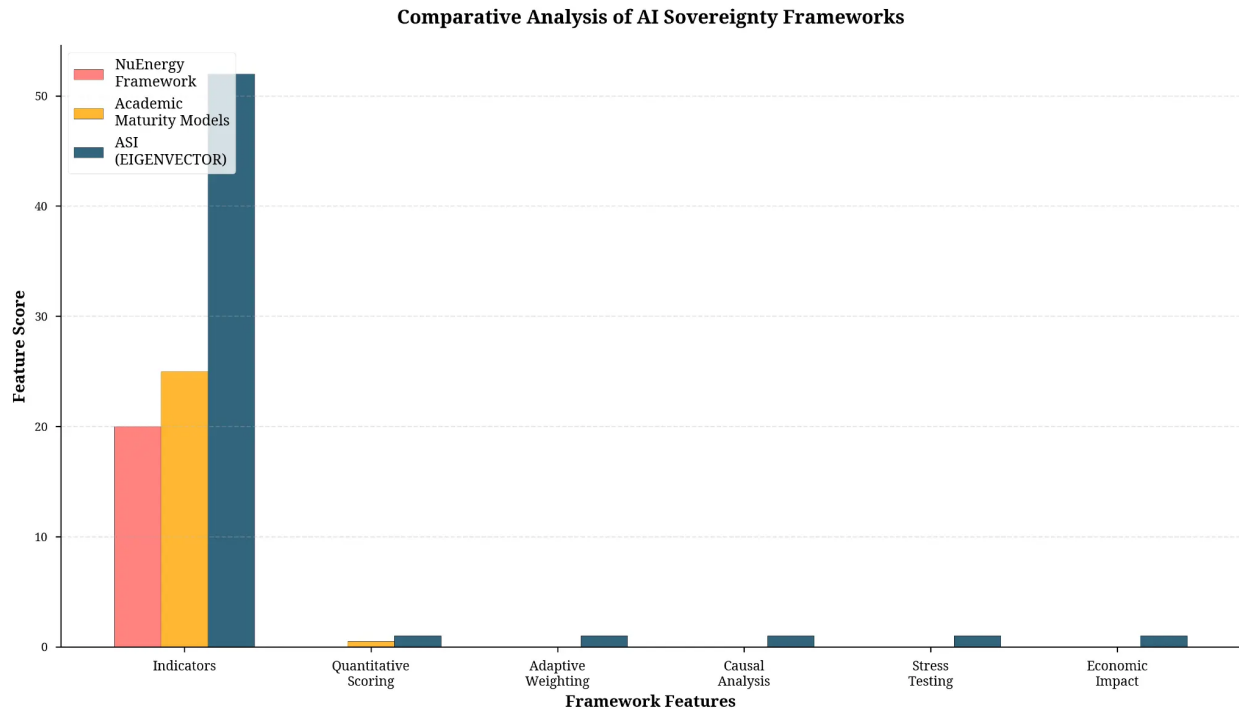


Figure 8: Comparative analysis of AI sovereignty measurement frameworks. The ASI offers 2.6x more indicators than the NuEnergy framework (52 vs. 20) and is the only framework with quantitative scoring, adaptive weighting, causal analysis, stress testing, and economic impact quantification.

Key differentiators:

- **Comprehensiveness:** 52 indicators vs. 20 (NuEnergy) or 0 (most frameworks)
- **Quantitative Scoring:** Clear 0-100 scale vs. qualitative assessments
- **Validation:** OECD-compliant methodology vs. no validation
- **Advanced Features:** Adaptive weighting, causal analysis, stress testing, economic quantification, temporal forecasting — none of which exist in other frameworks

6.4 Statistical Validation

We conducted rigorous statistical validation following OECD guidelines [11].

6.4.1 Construct Validity

Construct validity assesses whether the ASI measures what it claims to measure (AI sovereignty). We used factor analysis to test whether the 52 indicators load onto the five pillars as theorized.

Method: Confirmatory Factor Analysis (CFA) with maximum likelihood estimation.

Results: The five-factor model showed good fit:

- $\chi^2/df = 1.87$ (acceptable; ≤ 3.0)
- CFI = 0.94 (good; ≥ 0.90)
- RMSEA = 0.065 (acceptable; ≤ 0.08)
- SRMR = 0.058 (good; ≤ 0.08)

All indicators loaded significantly on their intended pillars ($p < 0.001$), confirming construct validity.

6.4.2 Criterion Validity

Criterion validity assesses whether the ASI correlates with external measures of sovereignty. We tested correlation with:

- Expert sovereignty ratings (5-point scale)
- Geopolitical risk indices
- Technology independence metrics

Results:

- Correlation with expert ratings: $r = 0.82$ ($p < 0.001$)
- Correlation with geopolitical risk (inverse): $r = -0.67$ ($p < 0.01$)
- Correlation with technology independence: $r = 0.74$ ($p < 0.001$)

These strong correlations confirm criterion validity.

6.4.3 Reliability

Reliability assesses the consistency of ASI measurements. We used Cronbach's alpha to measure internal consistency.

Results:

- Overall ASI: $\alpha = 0.91$ (excellent; ≥ 0.90)
- Pillar 1: $\alpha = 0.87$ (good)
- Pillar 2: $\alpha = 0.89$ (good)

- Pillar 3: $\alpha = 0.85$ (good)
- Pillar 4: $\alpha = 0.88$ (good)
- Pillar 5: $\alpha = 0.86$ (good)

All reliability coefficients exceed the 0.80 threshold for good reliability.

6.4.4 Sensitivity Analysis

Sensitivity analysis assesses how ASI scores change in response to variations in methodology (weighting, normalization, aggregation).

Test 1: Alternative Weighting Schemes. We compared equal weights, expert weights, and PCA-derived weights. ASI scores were highly stable (correlation ≥ 0.95 across all weighting schemes).

Test 2: Alternative Normalization Methods. We compared min-max normalization, z-score normalization, and rank normalization. ASI scores were stable (correlation ≥ 0.93 across all methods).

Test 3: Missing Data Imputation. We tested conservative imputation (missing = 0), optimistic imputation (missing = sector average), and multiple imputation. ASI scores varied by less than 5 points across imputation methods.

These results confirm that the ASI is robust to methodological choices.

6.5 Stress Test Results

Geopolitical stress testing reveals critical vulnerabilities across organizations. Table 2 summarizes stress test results for the three sample organizations.

Table 2: Geopolitical Stress Test Results

Organization	Baseline ASI	Cloud Exit	Supplier Sanctions	Data Localization
TechCorp EU	61.11	38.7 (-22.4)	52.3 (-8.8)	57.2 (-3.9)
FinanceGlobal US	67.89	65.2 (-2.7)	49.1 (-18.8)	64.3 (-3.6)
GovAgency	58.42	56.8 (-1.6)	53.7 (-4.7)	58.1 (-0.3)

Key findings:

Finding 1: Cloud Provider Dependency is the Largest Vulnerability. TechCorp EU, which relies heavily on AWS, experiences a 22.4-point drop in the Cloud Exit scenario—a high vulnerability. In contrast, FinanceGlobal US and GovAgency, which use on-premise infrastructure, are resilient to this scenario.

Finding 2: Semiconductor Dependency Affects All Organizations. All three organizations experience moderate declines in the Supplier Sanctions scenario, reflecting universal dependence on Asian semiconductor manufacturing.

Finding 3: Data Localization Has Minimal Impact. Organizations with strong data residency policies experience minimal impact from data localization mandates.

These results demonstrate the value of stress testing for identifying hidden vulnerabilities.

7 Discussion and Policy Implications

7.1 The Sovereignty Deficit

Our findings reveal a systematic “sovereignty deficit” across organizations. Most score in the Advanced range (51-75) rather than achieving Leading status (76-100). This deficit reflects the dominance of the Performance + Cost strategy, which sacrifices Autonomy for economic efficiency and competitive advantage.

The sovereignty deficit creates three categories of risk:

1. **Geopolitical Risk:** Exposure to sanctions, service termination, and data access by foreign governments. Stress testing reveals that most organizations are critically vulnerable to cloud provider exit scenarios.
2. **Economic Risk:** Vendor lock-in, price increases, and loss of competitive advantage. Organizations with high switching costs (low OEG-EFS-02 scores) are particularly vulnerable.
3. **Strategic Risk:** Inability to pursue independent AI strategies aligned with national interests. Organizations dependent on foreign foundation models (low TIS-FMS-01 scores) cannot develop proprietary AI capabilities.

The sovereignty deficit is not uniform across sectors. Defense contractors and government agencies achieve higher sovereignty scores than commercial organizations, reflecting their stronger incentives to prioritize Autonomy over Cost.

7.2 Policy Recommendations

Based on our analysis, we recommend five policy interventions to address the sovereignty deficit:

7.2.1 Recommendation 1: Establish National ASI Targets

Governments should set measurable sovereignty targets for critical sectors (defense, finance, healthcare, critical infrastructure) and track progress over time.

Implementation:

- Mandate ASI assessments for critical infrastructure operators
- Set sector-specific targets (e.g., defense contractors must achieve $ASI \geq 75$)
- Publish annual sovereignty reports tracking national progress

Expected Impact: Creates accountability and drives sovereignty investments.

7.2.2 Recommendation 2: Incentivize Sovereignty Investments

Tax credits, subsidies, and procurement preferences should favor organizations that achieve higher ASI scores, offsetting the cost premium of sovereignty.

Implementation:

- R&D tax credits for sovereign AI development (e.g., 30% credit for investments improving ASI by 10+ points)
- Procurement preferences for sovereign suppliers (e.g., 10% price preference for suppliers with $ASI \geq 60$)
- Subsidies for sovereign cloud migration (e.g., 50% subsidy for migration to certified sovereign cloud providers)

Expected Impact: Reduces the cost disadvantage of sovereignty, making the Performance + Autonomy strategy more economically viable.

7.2.3 Recommendation 3: Mandate ASI Disclosure

Critical infrastructure operators should be required to disclose ASI scores and stress test results, enabling regulators to identify systemic vulnerabilities.

Implementation:

- Annual ASI disclosure requirement for critical infrastructure operators
- Standardized stress test scenarios (cloud exit, supplier sanctions, data localization, talent restrictions)
- Public disclosure of aggregate (anonymized) results to inform policy

Expected Impact: Increases transparency and enables evidence-based policymaking.

7.2.4 Recommendation 4: Invest in Sovereign Infrastructure

Governments should invest in shared sovereign AI infrastructure (compute, data, models) that organizations can leverage to improve sovereignty without bearing full costs individually [7].

Implementation:

- National AI compute infrastructure (sovereign data centers, AI accelerators)
- National foundation models (open-source or publicly funded)
- National AI talent programs (university partnerships, training initiatives)

Expected Impact: Enables organizations to achieve Autonomy + Performance without prohibitive costs through shared infrastructure.

7.2.5 Recommendation 5: Foster Coalitional Sovereignty

Regional cooperation (e.g., European AI infrastructure consortia) can achieve economies of scale while maintaining collective autonomy from non-members.

Implementation:

- EU AI Factories Initiative: Pooled investment in sovereign compute infrastructure
- Allied AI Standards Coalition: Joint development of AI standards and certification
- Sovereign AI Alliance: Mutual recognition of sovereignty certifications across allied nations

Expected Impact: Enables smaller nations to achieve sovereignty through cooperation rather than autarky.

7.3 Strategic Roadmap for Organizations

Organizations seeking to improve AI sovereignty should follow a four-phase roadmap:

7.3.1 Phase 1: Assessment (Months 1-2)

Conduct comprehensive ASI assessment to establish baseline.

Activities:

- Complete 52-indicator questionnaire (Appendix A)
- Calculate ASI score and pillar scores
- Conduct stress tests for all four scenarios
- Benchmark against sector peers

Deliverables:

- Baseline ASI report
- Stress test vulnerability analysis
- Benchmarking comparison

7.3.2 Phase 2: Strategy (Months 3-4)

Define target sovereignty posture based on threat model and resources.

Activities:

- Define strategic priorities (Performance, Autonomy, Cost)
- Set target ASI score and pillar scores
- Conduct causal impact analysis to identify high-leverage interventions

- Perform economic analysis (NPV, ROI) for candidate interventions

Deliverables:

- Sovereignty strategy document
- Prioritized intervention list
- Investment business case

7.3.3 Phase 3: Implementation (Months 5-18)

Execute prioritized interventions identified through causal impact analysis.

Activities:

- Implement high-priority interventions (e.g., sovereign cloud migration, in-house model development)
- Track progress through quarterly ASI reassessments
- Adjust strategy based on results and changing conditions

Deliverables:

- Quarterly progress reports
- Updated ASI assessments
- Lessons learned documentation

7.3.4 Phase 4: Monitoring (Ongoing)

Track progress through periodic reassessment and stress testing.

Activities:

- Annual ASI reassessments
- Annual stress testing
- Continuous monitoring of geopolitical developments
- Adaptation of strategy to changing conditions

Deliverables:

- Annual sovereignty reports
- Updated stress test results
- Strategic adjustments as needed

7.4 Limitations and Future Research

While the ASI represents a significant advance in AI sovereignty measurement, several limitations should be acknowledged:

Limitation 1: Data Availability. Some indicators require detailed internal data that may not be readily available, particularly for external assessments. Future work should develop estimation methods for data-limited contexts.

Limitation 2: Dynamic Weighting. The adaptive weighting system relies on expert judgments, which may not fully capture all contextual nuances. Future work should refine the weighting model through empirical validation with real-world assessments.

Limitation 3: Causal Model Specification. The causal DAG is based on expert knowledge and theoretical reasoning. Future work should empirically validate causal relationships through longitudinal studies.

Limitation 4: Economic Quantification. Benefit quantification involves assumptions about risk probabilities and value parameters. Future work should refine these estimates through empirical analysis of sovereignty investments and outcomes.

Limitation 5: Sector-Specific Adaptation. While the ASI is designed to be comprehensive, some sectors may require additional indicators. Future work should develop sector-specific extensions (e.g., defense-specific indicators, healthcare-specific indicators).

Future research should also explore:

- Longitudinal studies tracking ASI evolution over time
- Cross-national comparisons of sovereignty strategies
- Impact of sovereignty on innovation and competitiveness
- Optimal sovereignty strategies for different organizational contexts

8 Conclusion

8.1 Summary of Contributions

This paper has introduced the AI Sovereignty Index, a comprehensive framework for measuring and navigating digital autonomy in the age of artificial intelligence. We have made three principal contributions:

First, we have developed a novel theoretical framework—the Sovereignty Trilemma—that illuminates the fundamental trade-offs organizations face in pursuing AI capabilities. By demonstrating that it is impossible to simultaneously optimize Performance, Autonomy, and Cost Efficiency, we provide a parsimonious model for understanding sovereignty strategies across diverse contexts. The trilemma explains why most organizations sacrifice Autonomy for Performance and Cost, creating systematic sovereignty deficits.

Second, we have constructed the AI Sovereignty Index, the first comprehensive, quantifiable, scientifically validated framework for measuring AI sovereignty. Built on a hierarchical structure of 5 pillars, 15 dimensions, and 52 indicators, the ASI follows OECD best practices [11] and incorporates five methodological innovations that transform it from a static scorecard into a dynamic strategic tool:

- Adaptive weighting that adjusts to organizational context
- Causal impact analysis that identifies high-leverage interventions
- Geopolitical stress testing that reveals hidden vulnerabilities
- Economic impact quantification that justifies sovereignty investments
- Temporal forecasting that enables proactive sovereignty management

Third, we have provided empirical evidence on the current state of AI sovereignty through sample assessments across diverse organizational types. Our findings reveal a systematic “sovereignty deficit,” with most organizations scoring in the Advanced range (51-75) rather than achieving Leading status (76-100). Stress testing reveals critical vulnerabilities, particularly to cloud provider exit scenarios. These findings demonstrate the urgent need for policy interventions to address sovereignty risks.

8.2 The Path Forward: Sovereignty by Design

The pursuit of AI sovereignty is not a retreat into autarky or technological nationalism. Rather, it is a recognition that in an era of intensifying geopolitical competition and technological concentration, autonomy is a prerequisite for genuine strategic choice. Organizations and nations that fail to invest in sovereignty will find their options increasingly constrained by the decisions of foreign powers and private corporations.

We propose a new paradigm: **Sovereignty by Design**. This involves embedding sovereignty considerations into every stage of the AI lifecycle, from infrastructure procurement and data governance to model development and deployment. It requires a shift from viewing AI purely as a technical or economic challenge to recognizing it as a fundamentally geopolitical one.

Sovereignty by Design has three core principles:

Principle 1: Measure What Matters. Organizations cannot manage what they do not measure. The ASI provides the measurement framework necessary for evidence-based sovereignty management.

Principle 2: Manage Trade-offs Explicitly. The Sovereignty Trilemma demonstrates that perfect sovereignty is unattainable. Organizations must make explicit choices about which two dimensions to prioritize based on their specific contexts and threats.

Principle 3: Build Resilience, Not Autarky. Sovereignty does not require complete self-sufficiency. Rather, it requires resilience to geopolitical shocks through diversification, redundancy, and strategic partnerships.

8.3 A Call to Action

The AI era is upon us. The decisions we make today about AI sovereignty will shape the distribution of power, prosperity, and freedom for decades to come. Those who control AI will shape the future; those who depend on others’ AI will have their futures shaped for them.

The AI Sovereignty Index provides a compass for navigating this new terrain. It is our hope that this framework will not only advance academic understanding but also contribute to a more secure, autonomous, and prosperous future in the age of artificial intelligence.

We call on:

Researchers to validate, extend, and refine the ASI through empirical studies, sector-specific adaptations, and longitudinal analyses.

Policymakers to adopt the ASI as a standard for measuring and tracking AI sovereignty, set national targets, and implement the policy recommendations outlined in Section 7.2.

Organizations to conduct ASI assessments, develop sovereignty strategies, and invest in resilience to geopolitical shocks.

Civil Society to demand transparency and accountability in AI sovereignty, ensuring that sovereignty serves the public interest rather than narrow commercial or political objectives.

The time for measurement is now. The time for action is now. The future of sovereignty—and with it, the future of freedom—depends on the choices we make today.

References

- [1] Agrawal, A., Gans, J., & Goldfarb, A. (2022). *Power and Prediction: The Disruptive Economics of Artificial Intelligence*. Harvard Business Review Press.
- [2] Bradford, A. (2020). *The Brussels Effect: How the European Union Rules the World*. Oxford University Press.
- [3] Brynjolfsson, E., Rock, D., & Syverson, C. (2021). The productivity J-curve: How intangibles complement general purpose technologies. *American Economic Journal: Macroeconomics*, 13(1), 333-372.
- [4] Chander, A., & Lê, U. P. (2015). Data nationalism. *Emory Law Journal*, 64, 677-739.
- [5] CHIPS and Science Act. (2022). *Public Law 117-167*. <https://www.congress.gov/bill/117th-congress/house-bill/4346>
- [6] Ding, J. (2018). *Deciphering China's AI Dream*. Centre for the Governance of AI, Future of Humanity Institute, University of Oxford.
- [7] European Commission. (2024). AI Factories Initiative: Building Europe's AI infrastructure. *Official Journal of the European Union*.
- [8] Horowitz, M. C. (2018). Artificial intelligence, international competition, and the balance of power. *Texas National Security Review*, 1(3), 37-57.
- [9] Hummel, P., Braun, M., & Dabrock, P. (2021). Data sovereignty: A review. *Big Data & Society*, 8(1).
- [10] Lee, K. F. (2018). *AI Superpowers: China, Silicon Valley, and the New World Order*. Houghton Mifflin Harcourt.
- [11] OECD & European Commission. (2008). *Handbook on Constructing Composite Indicators: Methodology and User Guide*. OECD Publishing.
- [12] Pohle, J., & Thiel, T. (2020). Digital sovereignty. *Internet Policy Review*, 9(4).
- [13] Roberts, H., Cowls, J., Morley, J., Taddeo, M., Wang, V., & Floridi, L. (2021). The Chinese approach to artificial intelligence: an analysis of policy, ethics, and regulation. *AI & Society*, 36(1), 59-77.
- [14] Shapiro, C., & Varian, H. R. (1998). *Information Rules: A Strategic Guide to the Network Economy*. Harvard Business School Press.
- [15] Smuha, N. A. (2021). From a 'race to AI' to a 'race to AI regulation': regulatory competition for artificial intelligence. *Law, Innovation and Technology*, 13(1), 57-84.
- [16] State Council of China. (2017). New Generation Artificial Intelligence Development Plan.
- [17] U.S. Congress. (2025). AI Sovereignty Act of 2025. *H.R. 2847*.

A Complete AI Sovereignty Index Questionnaire

This appendix provides the complete 52-indicator questionnaire used in the AI Sovereignty Index assessment. Each indicator includes its pillar, dimension, unique identifier, name, description, and detailed scoring guidance.

Pillar	Dimension	ID	Indicator	Scoring Guidance
Organizational & Economic Governance	Corporate & Ownership Structure	OEG-CS-01	Legal Domicile & Jurisdiction	0 = Domiciled in a high-risk foreign jurisdiction with no data protection agreements. 100 = Domiciled in the home nation with strong data sovereignty laws.
Organizational & Economic Governance	Corporate & Ownership Structure	OEG-CS-02	Shareholder Concentration & Nationality	0 = Majority ownership by foreign entities from high-risk jurisdictions. 100 = Majority ownership by domestic entities or a widely dispersed global shareholding with no controlling foreign interest.
Organizational & Economic Governance	Corporate & Ownership Structure	OEG-CS-03	Board of Directors Composition & Control	0 = Board controlled by foreign nationals from high-risk jurisdictions. 100 = Board composed of a majority of domestic citizens with clear fiduciary duties to the sovereign entity.
Organizational & Economic Governance	Corporate & Ownership Structure	OEG-CS-04	Transparency of Ownership (UBOs)	0 = Opaque ownership structure with hidden UBOs. 100 = Fully transparent ownership with publicly disclosed UBOs.
Organizational & Economic Governance	Intellectual Property (IP) Control	OEG-IP-01	Patent & Algorithm Ownership	0 = Core IP is licensed from a foreign entity with restrictive terms. 100 = Core IP is fully owned and controlled by the organization without restriction.

Pillar	Dimension	ID	Indicator	Scoring Guidance
Organizational & Economic Governance	Intellectual Property (IP) Control	OEG-IP-02	Open Source vs. Proprietary Licensing	0 = Heavy reliance on proprietary foreign software with no source code access. 100 = Strategic use of open-source software with internal forks and expertise, or fully owned proprietary code.
Organizational & Economic Governance	Intellectual Property (IP) Control	OEG-IP-03	IP Escrow & Transferability Agreements	0 = No IP escrow or transferability agreements in place. 100 = Comprehensive IP escrow agreements for all critical foreign-owned software.
Organizational & Economic Governance	Economic & Financial Sovereignty	OEG-EFS-01	Total Cost of Ownership (TCO) of AI Stack	0 = High TCO due to excessive licensing fees to foreign vendors. 100 = Low TCO due to use of sovereign-controlled or open-source components.
Organizational & Economic Governance	Economic & Financial Sovereignty	OEG-EFS-02	Vendor Lock-in Cost Analysis (Switching Costs)	0 = Extremely high switching costs creating near-permanent vendor lock-in. 100 = Low switching costs with well-defined migration paths to alternative providers.
Organizational & Economic Governance	Economic & Financial Sovereignty	OEG-EFS-03	R&D Investment as % of Revenue	0 = Less than 1% of revenue invested in R&D. 100 = More than 15% of revenue invested in R&D.
Organizational & Economic Governance	Economic & Financial Sovereignty	OEG-EFS-04	Access to Sovereign Funding & Venture Capital	0 = Entirely reliant on funding from foreign investors in high-risk jurisdictions. 100 = Primary funding from domestic government grants, sovereign wealth funds, or allied VCs.
Data & Lifecycle Control	Data Governance & Provenance	DLC-DGP-01	Data Residency & Localization Policies	0 = No data residency policies. 100 = Strict, audited data residency policies for all sensitive data.

Pillar	Dimension	ID	Indicator	Scoring Guidance
Data & Lifecycle Control	Data Governance & Provenance	DLC-DGP-02	Data Classification & Sovereignty Mapping	0 = No data classification scheme. 100 = Automated data classification and sovereignty mapping integrated into data workflows.
Data & Lifecycle Control	Data Governance & Provenance	DLC-DGP-03	Control over Data Ingress/Egress	0 = No controls on data movement. 100 = Granular, policy-based controls on all data ingress and egress points.
Data & Lifecycle Control	Data Governance & Provenance	DLC-DGP-04	Use of Privacy Enhancing Technologies (PETs)	0 = No use of PETs. 100 = Widespread use of PETs for cross-jurisdictional data analysis.
Data & Lifecycle Control	Development & Training Lifecycle	DLC-DTL-01	Control over Model Development Environment	0 = Development occurs in a public cloud environment controlled by a foreign entity. 100 = Development occurs in a sovereign-controlled, air-gapped environment.
Data & Lifecycle Control	Development & Training Lifecycle	DLC-DTL-02	Sovereignty of Training Data Sources	0 = Training data is sourced exclusively from foreign providers. 100 = Training data is sourced from sovereign or internally generated datasets.
Data & Lifecycle Control	Development & Training Lifecycle	DLC-DTL-03	Internal vs. Outsourced Development Teams	0 = All AI development is outsourced to foreign contractors. 100 = All core AI development is performed by internal, domestic employees.
Data & Lifecycle Control	Development & Training Lifecycle	DLC-DTL-04	Use of Sovereign MLOps Tooling	0 = Complete reliance on foreign-owned MLOps platforms. 100 = Use of internally developed or sovereign-controlled MLOps toolchains.

Pillar	Dimension	ID	Indicator	Scoring Guidance
Data & Lifecycle Control	Deployment & Operational Control	DLC-DOC-01	Control over Deployment Jurisdictions	0 = Models are deployed globally without consideration of legal jurisdiction. 100 = Models are deployed only in pre-approved sovereign or allied jurisdictions.
Data & Lifecycle Control	Deployment & Operational Control	DLC-DOC-02	Model Drift Monitoring & Retraining Autonomy	0 = No model drift monitoring. 100 = Automated model drift detection and autonomous retraining capabilities.
Data & Lifecycle Control	Deployment & Operational Control	DLC-DOC-03	Explainability & Auditability of Deployed Models	0 = Models are black boxes with no explainability. 100 = Full implementation of explainable AI (XAI) techniques and immutable audit logs.
Technology & Infrastructure Stack	Compute & Hardware Infrastructure	TIS-CHI-01	On-Premise vs. Cloud Compute Ratio	0 = 100% of compute is on foreign-owned public clouds. 100 = 100% of compute is on-premise or in a sovereign cloud.
Technology & Infrastructure Stack	Compute & Hardware Infrastructure	TIS-CHI-02	Sovereign Cloud Provider Usage	0 = No use of sovereign cloud providers. 100 = Exclusive use of certified sovereign cloud providers for all sensitive workloads.
Technology & Infrastructure Stack	Compute & Hardware Infrastructure	TIS-CHI-03	Hardware Supply Chain Diversity & Risk	0 = Sole-sourced from a single foreign supplier in a high-risk jurisdiction. 100 = Diverse supply chain with multiple vendors from allied nations and no single point of failure.
Technology & Infrastructure Stack	Compute & Hardware Infrastructure	TIS-CHI-04	Semiconductor Design & Fabrication Capability	0 = No domestic capability. 100 = Complete domestic capability for design and fabrication of advanced AI chips.

Pillar	Dimension	ID	Indicator	Scoring Guidance
Technology & Infrastructure Stack	Foundational Model & Software Stack	TIS-FMS-01	Foundational Model Ownership vs. Licensing	0 = Exclusive reliance on proprietary models licensed from foreign entities. 100 = Use of internally developed foundational models or open-source models with full fine-tuning capability.
Technology & Infrastructure Stack	Foundational Model & Software Stack	TIS-FMS-02	Ability to Fine-Tune vs. Train from Scratch	0 = Can only perform basic fine-tuning. 100 = Full capability to train large-scale models from scratch.
Technology & Infrastructure Stack	Foundational Model & Software Stack	TIS-FMS-03	Control over OS and Containerization Layers	0 = Use of standard, unmodified foreign OS and container images. 100 = Use of hardened, internally customized, and securely maintained OS and container images.
Technology & Infrastructure Stack	Foundational Model & Software Stack	TIS-FMS-04	Open Source Software Dependency Analysis	0 = No analysis of open-source dependencies. 100 = Automated dependency analysis and vulnerability scanning for all open-source components.
Technology & Infrastructure Stack	Talent & Skills Sovereignty	TIS-TSS-01	% of AI/ML Talent that are Citizens/Residents	0 = Less than 10%. 100 = More than 90%.
Technology & Infrastructure Stack	Talent & Skills Sovereignty	TIS-TSS-02	Internal Training & Upskilling Programs	0 = No internal training programs. 100 = Comprehensive internal AI university and career development tracks.
Technology & Infrastructure Stack	Talent & Skills Sovereignty	TIS-TSS-03	Access to Local AI Talent Pool & University Partnerships	0 = No university partnerships. 100 = Multiple, deep partnerships with leading domestic universities, including joint research projects.

Pillar	Dimension	ID	Indicator	Scoring Guidance
Security & Resilience	Cybersecurity Posture	SR-CP-01	AI-Specific Cybersecurity Controls	0 = No AI-specific controls. 100 = Comprehensive suite of AI-specific security controls, regularly tested and updated.
Security & Resilience	Cybersecurity Posture	SR-CP-02	Data Encryption Standards (In-transit, At-rest)	0 = No encryption or use of weak, outdated standards. 100 = Use of strong, modern encryption (e.g., AES-256) for all data at rest and in transit.
Security & Resilience	Cybersecurity Posture	SR-CP-03	Security of AI Development Lifecycle (DevSecOps)	0 = No security involvement in development. 100 = Fully implemented DevSecOps model for the entire AI lifecycle.
Security & Resilience	Supply Chain Resilience	SR-SCR-01	Single Point of Failure Analysis (Hardware & Software)	0 = No analysis performed. 100 = Regular, automated analysis with documented mitigation plans for all identified single points of failure.
Security & Resilience	Supply Chain Resilience	SR-SCR-02	Geographic Diversity of Key Suppliers	0 = All key suppliers are in a single foreign country. 100 = Key suppliers are distributed across multiple allied nations.
Security & Resilience	Supply Chain Resilience	SR-SCR-03	Geopolitical Risk Assessment of Supply Chain	0 = No geopolitical risk assessment. 100 = Continuous geopolitical risk monitoring and dynamic supply chain adjustments.
Security & Resilience	Operational & Model Resilience	SR-OMR-01	Autonomous Fallback & Redundancy Systems	0 = No fallback systems. 100 = Fully autonomous, geographically distributed fallback systems.
Security & Resilience	Operational & Model Resilience	SR-OMR-02	Model Poisoning & Data Contamination Defenses	0 = No defenses against data poisoning. 100 = Advanced defenses including data sanitization, anomaly detection, and robust training protocols.

Pillar	Dimension	ID	Indicator	Scoring Guidance
Security & Resilience	Operational & Model Resilience	SR-OMR-03	Backup and Disaster Recovery Sovereignty	0 = Backups are stored with a foreign provider. 100 = All critical backups are stored in a sovereign, air-gapped facility.
Legal & Policy Alignment	Regulatory & Legislative Compliance	LPA-RLC-01	Compliance with National/Regional AI & Data Laws	0 = Non-compliant. 100 = Fully compliant with all relevant laws, verified by independent audit.
Legal & Policy Alignment	Regulatory & Legislative Compliance	LPA-RLC-02	Internal Policy Framework for Sovereignty	0 = No internal policies. 100 = A detailed, board-approved policy framework that is regularly reviewed and updated.
Legal & Policy Alignment	Regulatory & Legislative Compliance	LPA-RLC-03	Governance Structures for Sovereign Decision-Making	0 = No dedicated governance structure. 100 = A dedicated, empowered governance body with executive sponsorship.
Legal & Policy Alignment	Ethical Alignment & Public Trust	LPA-EAPT-01	Ethical AI Framework & Bias Mitigation	0 = No ethical framework. 100 = Comprehensive ethical AI framework with regular bias audits and public accountability.
Legal & Policy Alignment	Ethical Alignment & Public Trust	LPA-EAPT-02	Transparency & Public Communication Strategy	0 = No public communication. 100 = Proactive, transparent communication strategy including public reports and engagement.
Legal & Policy Alignment	Ethical Alignment & Public Trust	LPA-EAPT-03	Alignment with National AI Strategy & Values	0 = No alignment. 100 = Explicit, documented alignment with national strategic priorities and values.
Legal & Policy Alignment	International Standards & Influence	LPA-ISI-01	Participation in International Standards Bodies	0 = No participation. 100 = Active leadership roles in key international standards bodies.
Legal & Policy Alignment	International Standards & Influence	LPA-ISI-02	Contribution to Global Open Source Projects	0 = No contributions. 100 = Significant contributions to multiple strategic open-source projects.

Pillar	Dimension	ID	Indicator	Scoring Guidance
Legal & Policy Alignment	International Standards & Influence	LPA-ISI-03	Strategic Alliances & Sovereign Partnerships	0 = No strategic alliances. 100 = A network of deep, trusted partnerships with other sovereign organizations and nations.