

Four Layers of AI Sovereignty

Can Japan Establish AI Sovereignty, or Just a New Dependency? — Policy Brief

A two-page summary of the full analysis, "Four Layers of AI Sovereignty"

Published: July 22, 2026 (Japanese original) / English edition · Publisher: Japan Institute for Crisis Management

The Issue

In July 2026, Noetra — backed by SoftBank, Sony, NEC, and Honda — announced Japan's largest national AI compute project (the "FRONTia Project"): roughly 27,500 NVIDIA GPUs and a 140MW data center, with government support reported at nearly \$6 billion over five years (a figure not stated in the ministry's own published text, but reported consistently by multiple independent outlets including Reuters). This brief breaks down how much of this project's claim to "sovereignty" is substantive, across four layers: data/model, weights/operations, compute, and power/institutions.

Figure 1: Four-Layer AI Sovereignty Framework

Four-Layer Framework for Evaluating AI Sovereignty		
Not a binary "domestic AI or not" — sovereignty assessed across four axes		
Axis	Key Question	Noetra Status
① Data & Model Sovereignty	What data trains the model, and where is it developed?	Substantive (domestic)
② Weights & Operational Sovereignty	Who controls trained weights, and release risk?	Unverified
③ Compute Sovereignty	Who supplies the hardware for training and inference?	Foreign-dependent
④ Power & Institutional Sovereignty	Can the state regulate the power and infrastructure supporting it?	Regulatory gap

This framework is designed for general applicability to sovereign AI policies beyond Noetra.

Not a binary "sovereign or not" — sovereignty assessed across four independent axes (full analysis, Section 3)

Where Noetra Stands, By Layer

① Data & Model	What data trains the model, and where is it developed?	Substantive
② Weights & Operations	Who controls released weights, and release risk?	Unverified
③ Compute	Who supplies the hardware for training/inference?	Foreign-dependent
④ Power & Institutions	Can the state regulate the power/infrastructure behind it?	Regulatory gap

This framework is designed for general applicability beyond Noetra, to other national sovereign-AI programs (full analysis, Sections 1 and 8).

Current Assessment

Noetra's "domestic AI" claim is substantive at the data and model layer, but its compute infrastructure (GPUs and CPUs) is entirely US-made. Japan's Economic Security Promotion Act does not name data centers as a critical-infrastructure sector, and that gap was not closed by the Diet session that closed on July 17, 2026 — data-security provisions were explicitly deferred as a "future issue," per an official House of Councillors research report. The EU (NIS2) and US (CISA) both explicitly designate data centers as critical infrastructure. As a preliminary two-country test of whether any state can

satisfy all four layers, this paper examined China and Russia: even authoritarian control does not guarantee compute sovereignty (semiconductor self-sufficiency). For democracies, the realistic policy priority is not simultaneous control of all four layers, but incremental expansion of existing critical-infrastructure law.

Policy Recommendations

- **1. Physical-AI safety standards:** Establish cross-industry safety-verification and red-teaming frameworks, incorporating OT/ICS domains, for autonomous robot-control AI.
- **2. Model weight release guidelines:** Develop technical guidelines — covering inversion-resistance testing and signature verification — before releasing models trained on sensitive data.
- **3. Critical-infrastructure clarification:** Resolve, without waiting for the next legislative cycle, whether data centers/AI compute facilities will be addressed through infrastructure designation or a data-protection regime.
- **4. Supply-chain diversification:** Treat diversification away from single-vendor compute dependence, and development of domestic alternatives, as a medium- to long-term policy priority.

About This Brief

This brief summarizes "Four Layers of AI Sovereignty," a full analysis covering international comparison (France, Saudi Arabia, India, the UK, Germany, Canada), physical-AI security, model weight-release risk, semiconductor supply chains, a Japan-EU-US critical-infrastructure comparison, and a preliminary test of the framework against China and Russia. The full paper documents primary sources and verification for each recommendation above.