PROJECT SYNDICATE – 26 MAGGIO 2026

How to Build a Sovereign AI Stack

di Jayant Sinha

Paradoxically, no economy can build the architecture of AI sovereignty alone. Coalitions of countries must work together to regulate the sector, design and operate the rails, open the model layer, and universalize the agent interface, thereby embedding contestability and interoperability into every layer.

NEW DELHI—Discussions about "AI sovereignty" generally focus on two things: investment and capabilities. Countries are encouraged to fund national champions that can develop frontier AI models, build compute clusters, and assemble domestic data pipelines. But while these steps do matter, they cannot, by themselves, deliver true AI sovereignty. For that, a fully contestable and interoperable AI stack is essential.

An AI stack contains multiple interdependent layers. Energy infrastructure feeds compute clusters. Compute (processing power) enables the training and operation of foundation AI models. Foundation models are wrapped in orchestration layers and deployed through cloud platforms. Applications and agents sit on top. Each layer must be contestable (firms can enter and exit the market easily), which depends partly on interoperability (different tools, models, and systems can exchange information).

A contestable stack expands the market for the leading model providers, lowers input costs for the downstream firms that build on top of them, and prevents the most advanced AI companies from becoming economic chokepoints. If a country secures contestability at only one layer, it might still find itself a price taker at the others, because powerful firms use their dominance in one layer to capture adjacent layers through bundled pricing, interlocking investment, privileged access, and exclusive partnerships.

Vendor lock-in and extraction at every layer of the stack are not in any country's long-term interest. An AI economy in which a few firms can extract rents from every other firm operating on the stack will eventually constrict the very ecosystem that produced those firms.

The first pillar of a sovereign AI stack is thus ex-ante competition regulation. Some jurisdictions are already delivering this. The European Union's Digital Markets Act establishes criteria to identify "gatekeepers," interoperability mandates, and prohibitions on self-preferencing. Similarly, the United Kingdom's Digital Markets, Competition and Consumers Act empowers the Competition and Markets Authority's Digital Markets Unit to assign major tech firms Strategic Market Status, thereby subjecting them to enforceable "conduct requirements," including interoperability and self-preferencing rules.

India is headed in this direction as well. In 2022, the Standing Committee on Finance, which I then chaired, released a report calling for an ex-ante framework for digital markets to prevent anti-competitive practices by Big Tech. This led to the creation of the Committee on Digital Competition Law, which in 2024 produced a draft Digital Competition Bill. While that bill was withdrawn last year, revisions are underway.

In the United States, competition law has long sought to ensure contestability. This was true of the Sherman Antitrust Act of 1890, the AT&T Consent Decree of 1982, and the prosecution of Microsoft in 1998. Most recently, the Federal Trade Commission and Department of Justice under then-President Joe Biden's administration revived antitrust enforcement, with a focus on data network effects, multi-sided platform dynamics, and algorithmic foreclosure in digital markets. But the country has yet to commit to contestability in the AI stack.

In any case, given the speed at which AI markets concentrate, regulation alone cannot preserve contestability. The second pillar of a sovereign AI stack is architectural: contestability must be built into the rails that underpin digital systems.

India's Data Empowerment and Protection Architecture shows how this can work in practice. It establishes consent-based "data rails," which enable individuals to share their information across financial institutions, health-care systems, and telecommunications services in a manner that is secure, transparent, and traceable,

while allowing regulators to gain insight into system behavior without needing vast supervisory agencies.

The same architectural logic can be applied to compute, models, and agents. While these rails are not bound by national borders, each country would use them in determining its preferred balance between risk and innovation, and it would retain sovereignty over the AI systems operating within its jurisdiction, no matter where the underlying compute resides.

The third pillar of a sovereign AI stack also transcends borders: a shared open-weight AI-model pathway available to all countries. Open weights—when the trained parameters that define a model's behavior are publicly released—make models genuinely interoperable in ways that closed application programming interfaces do not. Open-weight models can be inspected, adapted, and deployed within national jurisdictions on terms each country defines.

What the world needs is a privately funded open-weight architecture for AI on the scale of Red Hat's cloud infrastructure. A consortium of middle powers should encourage the creation of such an architecture, and a foundation with a golden share (such as the Robert Bosch Stiftung) should be tasked with governing it, both to preserve its openness and to ensure safe model deployment. A tiered approach to model release, backed by rigorous evaluation and guided by shared safety standards, is essential.

The fourth pillar of sovereign AI is the consumer interface. Agents are becoming the primary gateway through which citizens access information and services. Systemically significant digital enterprises should be required to provide a free agent that meets clearly defined safety standards, preserves privacy, and is interoperable with other agents. India's Unified Payments Interface, for example, is a real-time digital-payments system that enables people to send money, pay bills, and manage multiple bank accounts in a single mobile application. While premium services may sit above the baseline, the most important interface of the AI age cannot become a means of extracting money or data from people.

Paradoxically, no country can build the architecture of "AI sovereignty" alone. Coalitions of countries must work together to regulate the sector, design and operate the rails, open the model layer, and universalize the agent interface, thereby ensuring

contestability, interoperability, and accountability at every layer of the stack. The challenge is formidable, not least because these coalitions will have to manage tense trade-offs—between openness and safety, sovereignty and interoperability, and regulation and innovation. But the imperative is clear.