

For Most Countries, AI Sovereignty Is an Illusion. Resilience Is Real.

By [Nikolaus Lang](#), [Matt Langione](#), [Sesh Iyer](#), [Amartya Das](#), and [David Zuluaga Martínez](#)

ARTICLE MARCH 25, 2026 12 MIN READ

In recent years, many governments have pursued AI sovereignty by focusing on specific layers of the technology stack. But various initiatives have demonstrated the challenges of this strategy: consider Australia’s private sector push to create a “national” large language model (LLM), Germany’s incentives to expand chip production, and India’s efforts to assemble a national GPU cluster, among others. As our previous research has shown, the resource intensity and rapid pace

of AI development means that only a select few superpowers and middle powers have the scale and breadth of capabilities to sustain a sovereignty strategy over time. Yet, even for most of these countries, AI sovereignty conceived as full-stack autarky remains an illusion.

Even if stack-based sovereignty were attainable, it could prove risky: the strategy implicitly relies on the assumption that today's compute-intensive, LLM-based AI paradigm will continue to dominate. Yet, future advances in AI architectures and infrastructure pose the risk of massive sunk costs into stack-based moonshots (see Sidebar).

— What Single-Layer “Sovereignty” Buys, and What It Doesn’t

While concentrating national effort on a single layer of the stack can create useful assets, sustaining “sovereign” control over entire layers is exceptionally resource intensive over time and has proven difficult even in relatively affluent, capable geographies, as the following examples illustrate.

Software Development. Matilda, developed by the startup Maincode, is one of a small number of privately-led Australian initiatives that set out to create a “national” LLM, built and run on Australian infrastructure and tuned on Australian-relevant data and use cases to reflect local language and norms. It has created valuable domestic assets—from local engineering know-how and onshore compute capacity to sector-specific applications such as a pilot with the Australian Football League—but mostly within a proprietary ecosystem rather than as open public goods. Even its backers now acknowledge that sustaining competitive performance requires partnering with global hardware, connectivity, and software providers and drawing on non-Australian training corpora. You can nationalize the artifact, but not the capability, unless you also solve for the complements that keep models improving and make them usable across the economy.

Compute Availability. The IndiaAI program, initially funded at ₹10,300 (~\$1.1 billion) with 38,000 GPUs currently deployed and selective allocations such as 4,096 H100s for a 70-billion-parameter Indian-language model, meaningfully improves domestic bargaining power and lowers barriers to train and fine-tune. As recently as early March 2026, the Indian government added another 24,000 GPUs, bringing total capacity to 62,000. Yet, the scale gap with private fleets remains large: reports indicate that Microsoft alone purchased about 485,000 Hopper-generation GPUs in 2024, while other labs also operate typical clusters exceeding 100,000 GPUs. Public GPU pools therefore work best when they are explicitly designed to *complement* hyperscalers and support enterprise use under

domestic rules, including where inference runs and how firms reach capacity at predictable cost.

Hardware Manufacturing. Germany's marquee effort to anchor an advanced Intel fab in Magdeburg, backed by nearly €10 billion in government subsidies toward a roughly €30 billion project, became entangled in global cost cycles, supplier dependencies, and corporate reprioritization. After prolonged delays and political debate over subsidies, Intel announced that, as part of a broader cost-cutting program, it would not move forward with the Magdeburg project. Even a successful build would have relied on foreign tooling and upstream suppliers, with a long and uncertain path from local fabrication to broad-based AI advantage.

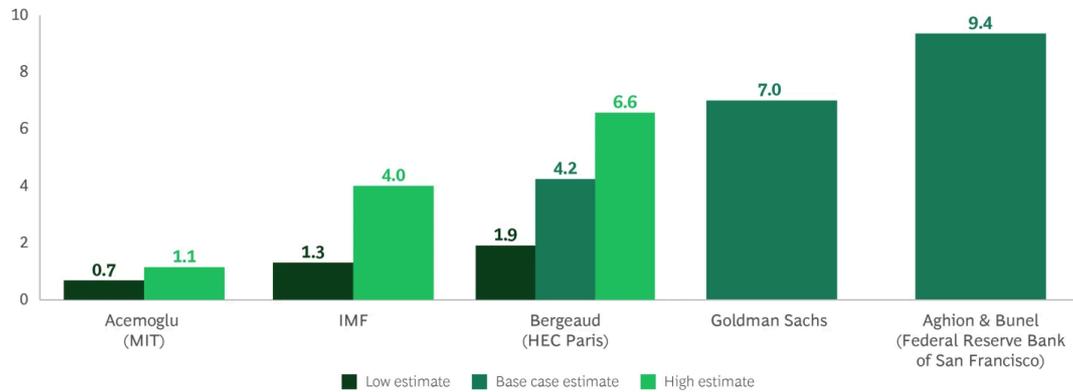
These cases are not arguments against partial plays, but rather are evidence of the resource intensity and difficulty of getting them right. Ultimately, such efforts only work when they are built for adoption and paired with the complements that sustain performance and real economy-wide use.

For most countries, a more practical AI sovereignty strategy is built around AI resilience—using, adapting, and governing AI domestically at scale, while minimizing strategic dependencies. Despite divergence among forecasters, even relatively conservative estimates view the potential payoff as significant. The International Monetary Fund (IMF), for example, expects effective AI adoption to boost global GDP by 4% over the next decade—roughly \$4.7 trillion (see Exhibit 1). For comparison, that's nearly the size of Germany's economy. Countries that lag behind will risk not only slower economic growth but also diminished competitiveness across industries poised to be disrupted by AI.

EXHIBIT 1

AI is expected to drive global GDP uplift of up to 9% in the next decade

Estimated global GDP uplift in 10 years (%)



Sources: Acemoglu, *The Simple Macroeconomics of AI* (2024); IMF, *The Global Impact of AI: Mind the Gap* (2025); Bergeaud, *The Past, Present, and Future of European Productivity* (2024); Goldman Sachs, *Generative AI could raise global GDP by 7%* (2023); Aghion & Bunel, *AI and Growth: Where Do We Stand?* (2024); BCG Henderson Institute analysis.

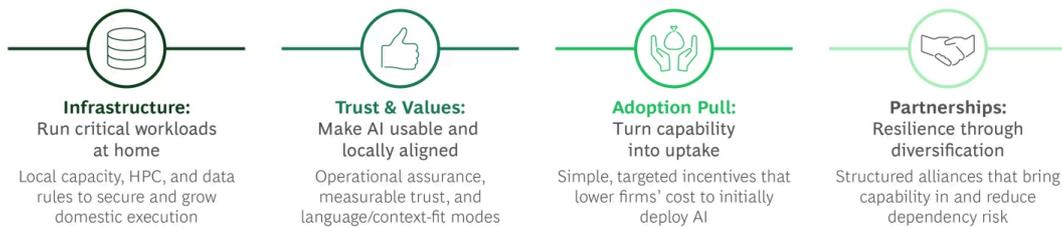
Note: IMF and Goldman Sachs GDP uplift values reported directly; Acemoglu, Bergeaud, and Aghion & Bunel global GDP uplift values calculated based on published regional total factor productivity (TFP) 10-year uplift estimates using TFP to GDP and regional-to-global scaling ratios corresponding to those in the IMF model.

The Four Components of AI Resilience

To assess how nations can support the domestic use, adaptation, and governance of AI, BCG Henderson Institute (BHI) conducted a study of recent AI policy actions in more than 30 countries, representing advanced, emerging, and small-state contexts. We identified pragmatic pathways nations can pursue across four key areas: infrastructure, trust and values, adoption pull, and partnerships (see Exhibit 2). By deliberately shaping how their countries are embedded in the global AI value chain, policymakers can manage their exposure to external control points and maintain reliable local AI deployment in the face of political, economic, and/or other shocks.

EXHIBIT 2

Four levers for building AI resilience



Source: BCG Henderson Institute analysis.

Infrastructure: Securing Domestic Capacity

If a country's strategic goal is AI resilience rather than end-to-end control, the most immediate lever is both physical infrastructure and cloud infrastructure that allow firms and public bodies to use AI at home. The objective is not to catch up with hyperscalers, but rather to ensure that sensitive workloads can run domestically, compliance is feasible under national rules, and enterprises can count on predictable, affordable capacity.

When it comes to infrastructure, Europe's experience captures the shift from symbolism to use. Early "sovereign cloud" schemes such as GAIA-X sought to achieve full continental autonomy but became mired in complex governance committees and branding exercises, producing merely symbolic frameworks rather than usable capacity.

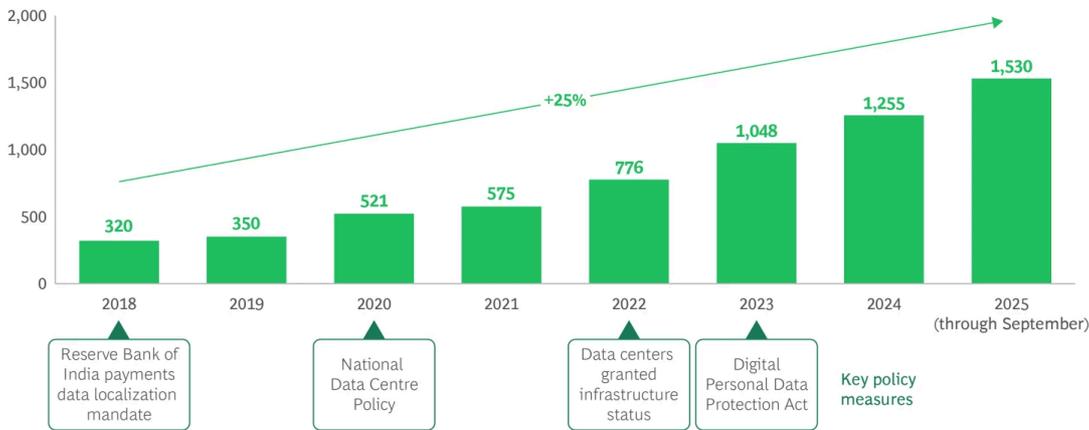
European leaders have instead seen greater success more recently in creating shared, high-performance capacity that researchers, small to medium enterprises (SMEs), and agencies can book and access today. The LUMI system in Finland and Leonardo in Italy, developed by the EuroHPC program, provide exactly that: large national and cross-border machines that put compute in the hands of real users. They do not rival hyperscalers for commercial elasticity, but they do create a baseline of domestic capability, skills formation, and experience with AI workloads that can be stepped up over time. In this way, the EuroHPC program is a pragmatic move toward resilience.

India's experience shows a complementary route that relies less on continental coordination and more on targeted rules that pull capacity onshore. The Reserve Bank of India's 2018 directive required payments data to be stored locally, which effectively created an "anchor tenant" for domestic processing and compliance. As global providers have expanded to meet that requirement, the country has been able to layer AI-relevant workloads behind existing data boundaries and create fiscal/regulatory incentives to encourage continued capacity growth (see Exhibit 3).

EXHIBIT 3

India's data center capacity has grown at an average rate of 25% per year since 2018

Data Center Capacity (megawatts)



Sources: CBRE; Colliers; JLL; Jefferies; CareEdge; BCG Henderson Institute analysis.

Note: 2020–2025 values reported by CBRE directly; 2018–2019 values triangulated across JLL, Colliers, and CareEdge reports.

Google's recent commitment to invest roughly \$15 billion to build an AI infrastructure hub and 1 GW of data center capacity in Visakhapatnam—alongside similar investments announced by Microsoft, OpenAI with Tata Consultancy Services, and India's Adani Group and Reliance Industries—underscores how domestic data-localization requirements and overall policy pull mechanisms create clear demand signals that can catalyze large-scale private infrastructure investment. Between 2018 and 2025, India's data center capacity grew roughly 66% faster than the global average, and it is poised to accelerate even further, reaching more than 8 GW by 2030.

The outcome is practical sovereignty: regulated inference runs at home, compliance risk is lower, and firms can adopt with fewer frictions, even if the underlying accelerators are globally sourced. This play is aided by India's market size, although smaller economies can adapt it by aggregating demand in regulated sectors, requiring in-country processing of priority workloads or pooling access through regional blocs.

The policy logic tying these examples together is straightforward: make AI work for the domestic economy by shaping where inference runs rather than attempting to replicate hyperscalers. Domestic inference enables governments to set clear standards to encourage AI adoption by firms. This infrastructure foundation sets up the next two levers.

Trust and Values: Enabling Adoption by Shaping Norms

Resilience depends as much on confidence as it does on capacity. Firms adopt faster when they can test systems against clear, operational standards and when models reflect local languages and social norms. The goal is not to legislate everything in advance. It is to reduce perceived risk for adopters and encode national priorities into the AI systems that people will actually use.

Singapore has become a reference point for "assurance that travels." Rather than stopping at high-level principles, its Infocomm Media Development Authority (IMDA) issued a Model AI

Governance Framework tailored to generative AI and paired it with AI Verify, an open toolkit that organizations can run on real systems. IMDA has more recently expanded this approach to agentic AI, issuing updated governance guidance on accountability, oversight, and risk management for autonomous systems. Beyond framework setting, Singapore has also launched a Global AI Assurance Pilot with 34 organizations to test production applications and generate practical guidance on “what and how to test,” including for LLM use cases.

By packaging test suites, convening industry to trial them, and documenting results, Singapore has made trust operational for adopters across sectors. For small and mid-sized countries without deep regulatory capacity, this is a repeatable play: stand up light-weight assurance, invite vendors and local firms to participate, and spread the practices through regional forums.

This outward-facing strategy is reinforced by Singapore’s coalition work with peers. The Small States AI Playbook, developed with Rwanda, translates governance concepts into implementable steps for governments with limited resources and is designed explicitly to be “picked up and used” by administrations that cannot build everything themselves.

South Africa has taken a similarly practical approach to operationalizing trust. Its National AI Policy Framework emphasizes inclusive, context-appropriate AI, an objective that is supported in practice through national digital language initiatives such as the South African Centre for Digital Language Resources, which develops corpora across the country’s 11 official languages. Together, these efforts help ensure that AI systems will reflect local linguistic norms and eventually be safely deployed in public-facing services.

The broader lesson for policymakers is that, by making trust testable and collaborating on standards that reflect local priorities, governments can tilt norms and adoption regardless of whether they train their own frontier models.

Adoption Pull: Turning Capacity and Confidence Into Use

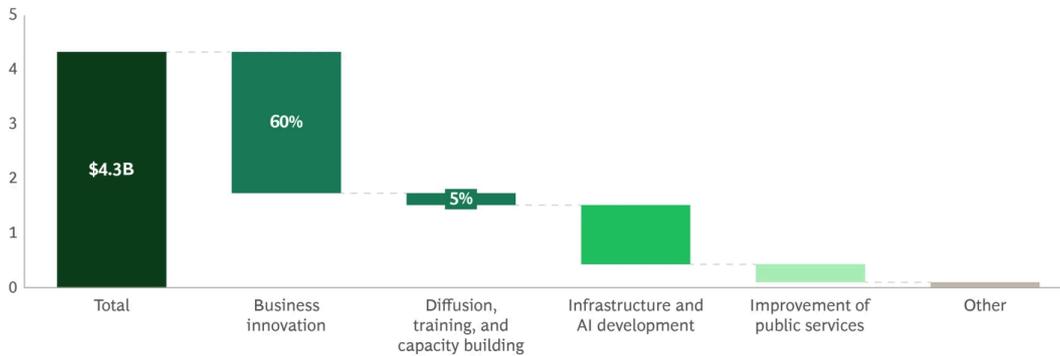
Capacity and assurance only matter if organizations actually deploy AI. In many lower-wage or slack labor markets, the business case for tech deployment may be weaker, and adoption stalls even when tools exist. Demand-side policy needs to do real work. The goal is boosting the adoption rate by firms, not just modernizing the state.

Brazil is an important bellwether because it is pairing values, infrastructure, and adoption pull at scale to drive firm-level uptake—not chase a sovereign stack. Its national AI program, running through 2028, commits roughly \$4.3 billion—equivalent to about 0.5% of the government’s annual budget, a significant allocation by emerging-market standards. Crucially, it directs roughly 65% to business innovation and upskilling projects via grants, co-funding, and subsidized credit so companies in priority sectors deploy AI in production. An additional 25% is invested in the enabling infrastructure that underpins resilient access—including upgraded regional HPC centers—with the balance for public service and additional areas (see Exhibit 4). By tying most of its funds to deployment and use rather than pure research, Brazil’s strategy is designed to generate early productivity gains in the real economy and seed a domestic supplier base that can serve regulated workloads.

EXHIBIT 4

65% of Brazil's national AI program funds are focused on raising firm-level uptake of AI

National AI Program Committed Funds, 2024–2028 (USD billions)



Sources: Brazilian Artificial Intelligence Plan (PBIA); BCG Henderson Institute analysis.

South Korea's "AI Voucher" program tackles the same problem at the SME level. The scheme offers SMEs vouchers worth up to roughly 200 million won (approximately \$140,000) to acquire AI solutions from an approved supplier roster, complemented by sector-specific initiatives (for example, data vouchers that subsidize the processing and analysis of clinical data for medical AI development). Evaluations and industry reporting indicate that such vouchers accelerate adoption and improve performance in firms that would otherwise postpone investment. For countries whose enterprise capex is tight, micro-subsidies linked to real deployments can move the adoption needle faster than top-down moonshots can.

The common thread is that adoption does not emerge organically simply because tools exist. Governments that want the productivity dividend can lower first-mile costs, reduce uncertainty, and target sectors in which the public benefits are largest. When adoption programs are paired with domestic execution environments and practical assurance, the result is a higher national diffusion rate—exactly the outcome "sovereign AI" seeks, but is achieved through use rather than ownership.

Partnerships and Diversification: Designing Interdependence, Not Denying It

No country can eliminate interdependence in AI. The strategic question is whether that interdependence is a vulnerability or a design choice. For most economies, the fastest path to "minimum viable sovereignty" is networked capacity: attract foreign capital and know-how to operate in-country under domestic rules while placing selective bets abroad to secure access during shocks.

Japan's emerging approach is instructive: operationalizing sovereignty through alliance rather than autarky. Instead of attempting to onshore every layer, Tokyo is moving to underwrite outward foreign direct investment (FDI) so that Japanese firms build critical nodes—semiconductors, rare earth processing, data infrastructure, 5G—in trusted partner countries. The strategy is to design interdependence that reduces exposure to individual chokepoints while locking in reciprocal access during shocks.

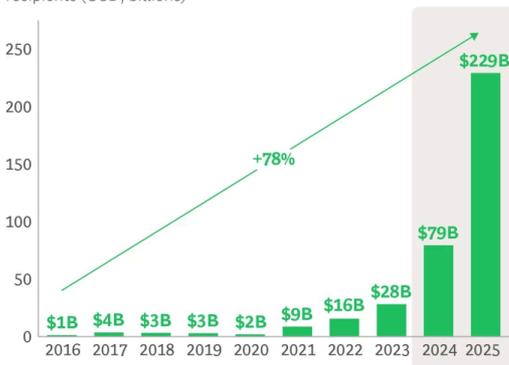
Japan is not alone in this approach: AI-related FDI has surged even to recipient countries other than the leading AI “superpowers” US and China. Between 2016 and 2025, transnational AI investment grew by a factor of roughly 200%, with approximately 20 times as many projects per year. Today, major blocs are already channeling flows toward partners such as India, South Korea, and Malaysia (see Exhibit 5).

EXHIBIT 5

A growing number of countries are benefiting from rising foreign direct investment in AI

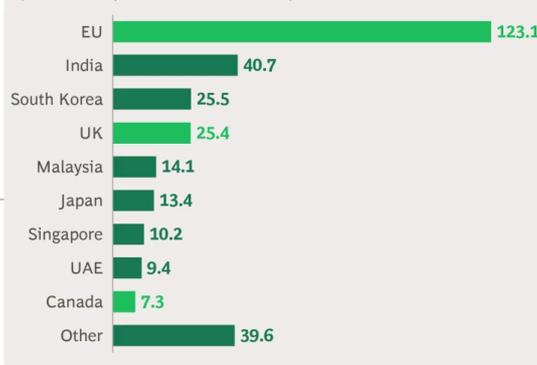
Global foreign direct investment in AI has accelerated dramatically ...

Total foreign direct investment in AI, excluding US and China as recipients (USD, billions)



... benefitting many beyond the US and China

Total FDI in AI by recipient country, 2024–2025 (USD billions; excludes US and China)



Sources: fDi markets; BCG Henderson Institute analysis.

Note: Figures refer to announced greenfield FDI projects that fDi Markets classified with their 'AI' tag; other forms of FDI (e.g., Brownfield, M&A) are not included.

For mid-sized economies, plugging into these friend-shoring corridors—with terms such as local cloud regions, data residency, skills transfer, and multi-vendor rules—can build resiliency without requiring them to finance the stack themselves.

The Spain-IBM memorandum sits at the other end of the same spectrum. Madrid is not trying to own every layer; it is bringing models in Spanish and co-official languages to life through a partnership, alongside supercomputing and tools that operate under national rules. That is diversification by design: a global partner for capability with local control for relevance and compliance.

In practice, these partnership models can also complement both the infrastructure and values/trust strategies described earlier. When countries specify that sensitive workloads must run domestically, FDI in local regions becomes a resilience asset rather than a dependency. When they publish assurance toolkits and language resources, partners can meet national expectations without bespoke regulation. The result is optionality: multiple routes to capacity, multiple vendors aligned to domestic rules, and clearer fallbacks if geopolitical shocks interrupt one channel.

Putting Use Before Ownership

The AI sovereignty debate will continue to be framed in the language of control. But for the majority of countries, the most meaningful form of control is ensuring that their businesses and public bodies can use AI reliably under domestic rules. That is what resiliency delivers. It is not a retreat from sovereignty, but rather a more effective way to achieve it when end-to-end control is out of reach.

While countries' individual approaches will necessarily differ based on resourcing and priorities, the overall sequence is clear:

1. Start by considering how critical workloads can be executed domestically/regionally on modern capacity.
2. Focus on making AI usable and locally aligned so adopters move from pilots to full-scale production.
3. Pull demand through the economy with targeted, easy-to-use programs that lower first-mile costs.
4. Design partnerships that bring capability in while keeping control over data location, assurance, and continuity.

With those pieces in place, countries can raise the national diffusion rate of AI and bank the productivity dividends that long-run growth depends on.

For businesses, the message is equally practical. Companies have the most to gain when policymakers get applied AI resilience right. Doing so lowers compliance burdens, reduces latency and risk for sensitive workloads, and creates clear procurement paths for AI systems that are vetted and supportable. The most successful firms will meet governments where they are heading—with modular products, localized partnerships, and deployment models that respect domestic rules while preserving the advantages of global scale.

Ultimately, interdependence is not something to be wished away in modern AI. It is a structural reality that governments can shape to capture more of AI's benefits for both their citizens and their firms.

The BCG Henderson Institute is Boston Consulting Group’s strategy think tank, dedicated to exploring and developing valuable new insights from business, technology, and science by embracing the powerful technology of ideas. The Institute engages leaders in provocative discussion and experimentation to expand the boundaries of business theory and practice and to translate innovative ideas from within and beyond business. For more ideas and inspiration from the Institute, please visit our [website](#) and follow us on [LinkedIn](#) and [X \(formerly Twitter\)](#).

Authors



Nikolaus Lang

Managing Director & Senior Partner; Global Leader, BCG Henderson Institute; Global Vice Chair, Global Advantage Practice
Munich



Matt Langione

Managing Director & Partner
Boston



Sesh Iyer

Managing Director & Senior Partner
Washington, DC



Amartya Das

Principal
San Francisco - Bay Area



David Zuluaga
Martínez

Senior Director, BCG Henderson Institute
New York



ABOUT BOSTON CONSULTING GROUP

Boston Consulting Group partners with leaders in business and society to tackle their most important challenges and capture their greatest opportunities. BCG was the pioneer in business strategy when it was founded in 1963. Today, we work closely with clients to embrace a transformational approach aimed at benefiting all stakeholders—empowering organizations to grow, build sustainable competitive advantage, and drive positive societal impact.

Our diverse, global teams bring deep industry and functional expertise and a range of perspectives that question the status quo and spark change. BCG delivers solutions through leading-edge management consulting, technology and design, and corporate and digital ventures. We work in a uniquely collaborative model across the firm and throughout all levels of the client organization, fueled by the goal of helping our clients thrive and enabling them to make the world a better place.

© Boston Consulting Group 2026. All rights reserved.

For information or permission to reprint, please contact BCG at permissions@bcg.com. To find the latest BCG content and register to receive e-alerts on this topic or others, please visit bcg.com. Follow Boston Consulting Group on [Facebook](#) and [X \(formerly Twitter\)](#).