

RIHLA
RESEARCH& ADVISORY

Pathways to AI Sovereignty

A Method for Mapping Middle Power AI Options for Sovereign Ecosystems

Jesse Marks
Info@rihlaadvisory.com
MiddlePowerAI.com

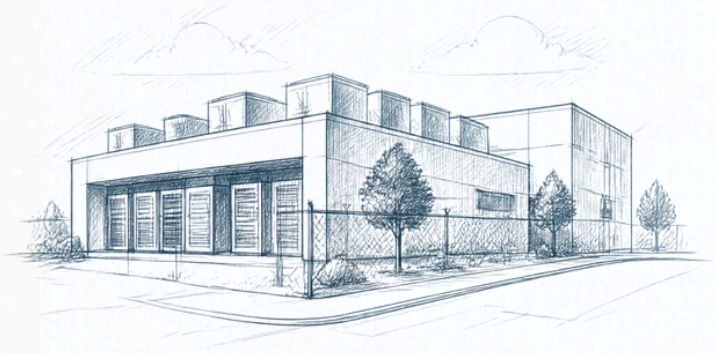
October 2026

Introduction

A Method for Mapping Middle Power AI Options for Sovereign Ecosystems

Governments seeking greater control over artificial intelligence must decide which capabilities to develop domestically and which to obtain through international partnerships. These choices extend across the AI stack, from chips and compute infrastructure to cloud services, AI models, data and applications. A country may host infrastructure operated by foreign hyperscalers while having limited authority over its use. It may also operate AI models developed by foreign frontier labs when domestic teams have sufficient compute, technical expertise, and the necessary permissions.

Rihla Research & Advisory developed *Pathways to AI Sovereignty* to help researchers and policymakers assess these arrangements. The planner maps suppliers and operating control across seven layers of the stack, identifies the capability profile that best fits the selected scenario and tests whether services could continue if access to a supplier changes. It helps governments identify where domestic capability is established, where dependency remains and what they need to secure through investment or negotiation.



"The pathway to AI sovereignty for middle powers depends on a series of sequential decisions about inputs and capabilities across the AI stack."

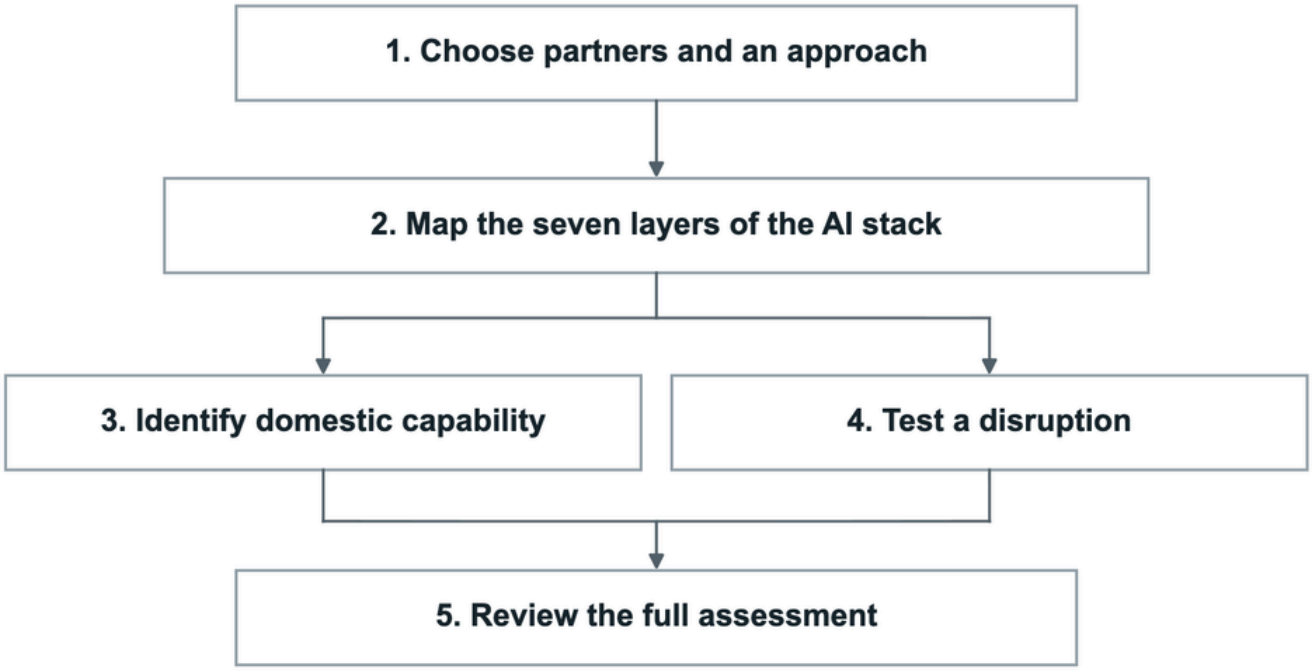
The pathway to AI sovereignty for middle powers depends on a series of sequential decisions about inputs and capabilities across the AI stack. Full autarky is currently beyond individual countries' capabilities, but a growing body of research suggests that governments can gain greater agency and autonomy within interdependent AI value chains. These opportunities depend partly on how they position themselves in relation to AI superpowers such as the United States and China, and how they structure partnerships with hyperscalers and frontier labs.

This tool is the first version of a developing research project that aims to help governments assess how their current AI position could evolve as they build their stack. The current version, while simple in its construction, provides users a taxonomy of five middle-power capability profiles based on the specific capabilities they select in the planning tool. Through it, users can examine how changes in infrastructure, access to AI models, and legal operating authority affect domestic capability and dependency on foreign providers.

Methodology

This tool uses a multi-level decision framework to examine how a country's choice of AI partners shapes subsequent decisions across the AI stack (see figure 1). Macro-level choices, such as whether to build on infrastructure supplied by U.S. or Chinese hyperscalers or cooperate with other middle powers, influence which technologies a country can access and under what conditions. Users can identify their country's position across seven layers of the AI stack to assess how ownership, operating control and imported capabilities shape dependency and agency. They can then run a simple disruption test to assess whether essential AI services could continue if a hyperscaler or AI service provider restricted access to models, cloud services or compute hardware.

Figure 1: Overview of the Multi-level Decision Framework



User Assessment Questions

	Assessment Question	User Response
1	**Which partnership path does the country choose?**	Users select a U.S., Chinese or middle-power path as the starting point. They can then adjust suppliers and operating arrangements across the stack. The path alone does not determine the capability profile.
2	**What does the country seek from its AI partners?**	Users choose an approach to working with hyperscalers, model developers or partner states. Options include using integrated AI services, obtaining selected technologies, building domestic capability or sharing compute and development. Suggested settings can be customized.
3	**Who supplies each layer of the stack, and who controls its operation?**	Users identify suppliers and domestic, joint or foreign operating control across chips, compute, cloud, AI models, data, applications and talent. Additional inputs cover compute location, domestic development, sufficient compute capacity and locally stored model weights licensed for continued use.
4	**What restrictions or permissions affect access?**	The tool identifies general conditions concerning advanced chips, domestic operating permissions and restrictions on combining suppliers. These highlight rules and contractual provisions that users need to verify for the partnership being assessed.
5	**What happens if a supplier withdraws or access ends?**	Users can explore how a disruption will shape their access. This includes a halt to chip deliveries, a cloud provider withdrawal, loss of model API access, or a cooperation partner withdrawal over a time window of 30 days, six months, or twelve months. It also considers replacement providers, data and application portability, and hardware maintenance to identify which functions can continue, face limits, or are at risk.
6	**Which capability profile best fits the scenario?**	The tool identifies the closest match among AI Importer, Hybrid Adopter, Compute Host, Domestic Operator and Technology Developer. It reports remaining dependencies and disruption results alongside the profile. It assesses supplier diversification and potential bargaining power separately.

Taxonomy of Capability Profiles

These profiles are descriptive, not prescriptive. In practice, each state must make its own national decisions about its sovereign AI stack based on its domestic capabilities and international dependencies. While some cases may be similar in terms of their challenge, the pathway a state takes ultimately reflects its domestic conditions. This tool supports that decision-making process.

Table 3. AI capability profiles and partnership strategies

Type	Defining characteristics	Principal dependencies	Policy priorities
 **AI Importer**	Domestic institutions use externally operated AI services and may develop applications, but lack sufficient capability or permissions to operate the underlying systems.	Foreign cloud operation, model access and provider support.	Secure dependable access, establish migration arrangements and develop operating expertise for essential services.
 **Hybrid Adopter**	Domestic teams operate or adapt imported AI technology under local or joint control, combining domestic capability with continued reliance on foreign technology.	Foreign hardware, model licenses, and decisions or support shared with partners.	Strengthen domestic operating capability, clarify shared decision rights and test service continuity.
 **Compute Host**	The country hosts AI infrastructure, while foreign providers or joint partners retain control over compute or cloud operations.	External operating authority, capacity allocation and access conditions.	Secure domestic access to compute and connect hosting agreements to operating expertise and continuity rights.
 **Domestic Operator**	Domestic institutions have the expertise, capacity and permissions to operate essential AI services, including systems built with imported hardware and foreign-developed models.	Replacement hardware, maintenance inputs and external software support.	Sustain technical staffing, preserve operating permissions and prepare for supply interruptions.
 **Technology Developer**	Domestic institutions sustain AI technology development independently or through joint development, with continuing technical contributions and rights to modify and maintain the technology.	Research compute, imported hardware and any development resources or permissions provided by partners.	Maintain research capacity, retain technical expertise and support development linked to defined public or economic needs.

AI Importer

An AI Importer relies on foreign hyperscalers and model providers to operate essential parts of its AI stack. Domestic institutions access AI models through APIs or managed cloud services and build applications around them. They may develop sophisticated public services without possessing the model weights, compute infrastructure or operating permissions needed to run the underlying systems themselves.

This creates a dependency between domestic applications and decisions made by external providers. If a frontier lab withdraws an API, applications using that model may lose their ability to generate outputs. If a hyperscaler suspends a cloud account, the government may lose access to the environment hosting its applications and data. Developing an application domestically does not remove these dependencies.

The scope for greater agency depends on whether domestic institutions can change providers or take over selected operations. Switching models may require rewriting application code and checking the replacement's accuracy. Moving cloud providers may require exporting databases and rebuilding software that uses the original hyperscaler's proprietary tools. An alternative becomes credible when domestic teams have completed this work and demonstrated that the replacement can support the essential service.

Hybrid Adopter

A Hybrid Adopter uses foreign AI technology while domestic teams operate or adapt parts of the stack. A university might fine-tune an imported model on domestic datasets, while a public institution deploys its weights on locally operated servers. A joint venture might combine a foreign partner's cloud technology with domestic staff responsible for deployment and maintenance. These arrangements establish domestic capability while retaining dependencies on foreign hardware, software or development support.

The boundary of domestic control depends on the permissions and technical responsibilities attached to each layer. A model license may allow local inference but restrict modification or redistribution. Domestic operators may manage applications while the foreign cloud partner controls administrator access to the underlying platform. Shared ownership can therefore coexist with unequal authority over the system.

This pathway provides greater agency when domestic teams can maintain essential functions without repeatedly seeking external approval. Relevant evidence includes whether local staff can deploy model updates, restore a failed service and transfer workloads to another environment. Where a partner's approval remains necessary, the assessment should identify the decision involved and what would happen if that approval were withheld.

Compute Host

A Compute Host provides the domestic infrastructure on which AI workloads run, while a foreign hyperscaler or joint partner controls compute or cloud operations. The country may supply electricity and land, finance construction or own the data-center buildings. The operator may still decide which customers receive accelerator capacity and whether domestic researchers can use it.

Hosting infrastructure creates different capabilities from controlling the workloads it supports. A data center can operate within national borders while serving primarily foreign customers. A domestic institution may have to purchase access on commercial terms despite public support for the facility. The hosting arrangement must establish whether the government can obtain compute for its own needs and who can approve or deny that access.

Reserved capacity makes this distinction concrete. An agreement could give domestic universities, research institutions, or local labs access to a defined pool of accelerators for model training or guarantee capacity for public-service inference. It would need clear parameters to govern when and how much compute these institutions can access. If the hyperscaler leaves or closes down a model, continued operation would also depend on access to the cloud software and technical staff needed to use the installed hardware.



Technology Developer

A Technology Developer has domestic institutions that make continuing technical contributions to AI models or the systems supporting their development. These institutions can conduct training or substantial adaptation, evaluate the resulting models and maintain their work. International partnerships can support this capability when domestic teams hold development responsibilities and rights to modify the technology.

The distinction concerns what domestic institutions can develop themselves. A locally developed model based on imported weights may involve substantial domestic training and evaluation, or it may involve little more than repackaging. The assessment must establish what the domestic team contributed, which parts of the model it can change and whether it can repeat the development process. Delivering one model does not establish that the team can sustain further development.

Dependence on foreign inputs can be extensive. Domestic researchers may need a hyperscaler's accelerator clusters to train models or rely on a hyperscaler's base model. If access ends, further development may stop even though the existing model remains accessible. This profile therefore requires examining the resources available for subsequent development, rather than judging capability solely from a completed release. In the planner, it concerns model and service development rather than independent production across the entire stack.

Domestic Operator

A Domestic Operator can run essential AI services without relying on a foreign provider to operate them day to day. Domestic teams manage the compute and cloud environment, control the data and applications, and have permission to use the AI model. The chips and model may still come from abroad. What distinguishes this profile is that domestic institutions can operate the service themselves.

This requires more than owning servers or storing model weights locally. Domestic teams need administrator access, enough compute to meet demand, and the skills to diagnose problems, restore service after an outage and update the system. They also need to identify any externally managed components the service depends on. For example, a foreign authentication service could interrupt an application even when its model and compute are operated domestically.

The main challenge is sustaining that control over time. Imported chips eventually need replacement, model licenses must remain valid, and repairs may require foreign software or specialist support. During a disruption, essential services can continue only if domestic teams have the equipment, technical expertise and operating rights needed for the period being tested.

Conclusion

This tool will continue to be developed and adjusted as research on middle-power AI grows and as we conduct more case studies on how governments implement different partnership agreements. The first version offers a starting point for thinking through a middle power's path to AI sovereignty and assessing the utility of a multilevel decision framework approach. This tool has limitations as a "plug-and-play" guide because each state's capacities and domestic composition must be considered when making practical decisions about trade-offs, dependencies, and the adoption of AI systems.

Future versions of this tool will aim to use country cases and observed disruptions to refine the decision rules and test whether the profiles capture the choices governments actually face in practice. Rihla Research & Advisory is also developing a larger, more targeted toolkit for governments exploring their best path to AI sovereignty, with clear questions at each stage, from assessing needs to negotiating contracts, to support better local ownership over their AI ecosystem.



Further Reading

Davies, Matt, and Jai Vipra. “Computing Commons: Designing Public Compute for People and Society". Ada Lovelace Institute, February 7, 2025. https://www.adalovelaceinstitute.org/report/computing-commons/

Lehdonvirta, Vili, Boxi Wu, Zoe Jay Hawkins, Celine Caira, and Lucia Russo. Measuring Domestic Public Cloud Compute Availability for Artificial Intelligence. OECD Artificial Intelligence Papers, no. 49. Paris: OECD Publishing, October 29, 2025. https://www.oecd.org/en/publications/measuring-domestic-public-cloud-compute-availability-for-artificial-intelligence_8602a322-en.html

Leicht, Anton, and Dean Ball. The Race Worth Winning: Middle Powers in the Age of Machine Intelligence. Foundation for American Innovation, February 13, 2026. https://www.thefai.org/posts/the-race-worth-winning-middle-powers-in-the-age-of-machine-intelligence

Leicht, Anton, and Sam Winter-Levy. AI Breakout Capacity: A Middle-Power Path Between Racing and Dependence. Carnegie Endowment for International Peace, September 28, 2026. https://carnegieendowment.org/research/2026/09/artificial-intelligence-ai-development-breakout-germany-india-japan-uk-middle-power

Marks, Jesse. “America’s Middle Power AI Export Challenge.” Stimson Center, July 27, 2026. https://www.stimson.org/2026/americas-middle-power-ai-export-challenge/

Marks, Jesse. The Gulf AI Monitor Database. Rihla Research & Advisory, 2026. https://www.gulfaimonitor.com/

Martinet, Charles, and Yohann Ralle. “Can Mid-Sized Economies Come Together to Build Frontier AI?” The AI Wonk, OECD.AI Policy Observatory, December 16, 2025. https://oecd.ai/en/wonk/can-mid-sized-economies-come-together-to-build-frontier-ai

Phillips-Robins, Alasdair, Teddy Tawil, and Sam Winter-Levy. The Compute Coalition: How to Build the Future of AI in the Free World. Carnegie Endowment for International Peace, June 8, 2026. https://carnegieendowment.org/research/2026/06/the-compute-coalition-how-to-build-the-future-of-ai-in-the-free-world

Richardson, Aris, Haley Yi, Michelle Nie, Simon Wisdom, Casey Price, Ruben Weijers, Steven Veld, and Mauricio Baker. “How Sovereign Is Sovereign Compute? A Review of 775 Non-U.S. Data Centers.” arXiv preprint arXiv:2508.00932, July 30, 2025. Revised May 3, 2026. https://arxiv.org/abs/2508.00932

Seger, Elizabeth, and Guy Ward-Jackson. “Build the Roads, Not the Engine: An Open-Source Strategy for Middle-Power AI Sovereignty.” Tony Blair Institute for Global Change, September 2026. https://institute.global/insights/tech-and-digitalisation/build-the-roads-not-the-engine-an-open-source-strategy-for-middle-power-ai-sovereignty

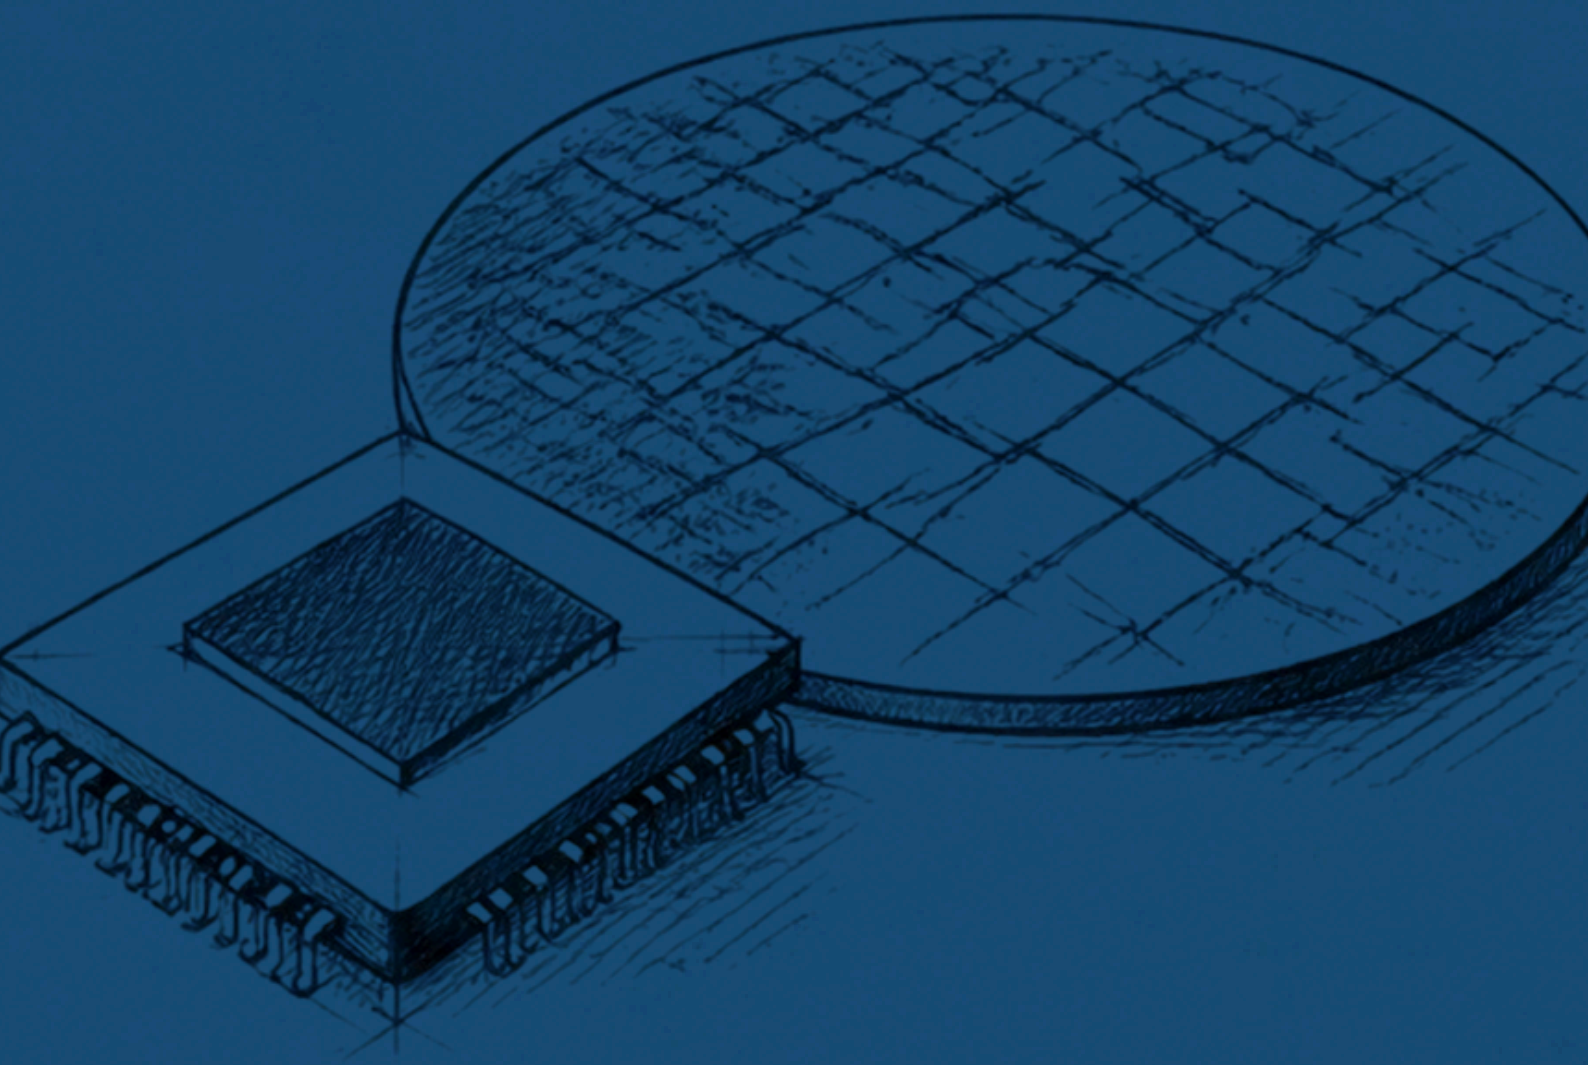
Tanner, Brooke, Cameron F. Kerry, Andrew W. Wyckoff, Nicoleta Kyosovska, Andrea Renda, and Elham Tabassi. Is AI Sovereignty Possible? Balancing Autonomy and Interdependence. Brookings Institution, February 17, 2026. https://www.brookings.edu/articles/is-ai-sovereignty-possible-balancing-autonomy-and-interdependence/

Varela Sandoval, Francisco Javier, Isabella Wilkinson, Alex Krasodomski, and Rowan Wilkinson. How Middle Powers Can Weather US and Chinese AI Dominance: The Case for “Sovereign AI” Strategies. Chatham House, February 16, 2026. Updated March 13, 2026. https://www.chathamhouse.org/2026/02/how-middle-powers-can-weather-us-and-chinese-ai-dominance

Widder, David Gray, Meredith Whittaker, and Sarah Myers West. “Why ‘Open’ AI Systems Are Actually Closed, and Why This Matters.” Nature 635 (2024): 827–33. https://www.nature.com/articles/s41586-024-08141-1

Winter-Levy, Sam, and Anton Leicht. “The AI Divide: How U.S.-Chinese Competition Could Leave Most Countries Behind.” Foreign Affairs, February 10, 2026. https://www.foreignaffairs.com/united-states/ai-divide

RIHLA
RESEARCH&
ADVISORY



For more information on Rihla Research & Advisory or our services, please contact info@rihlaadvisory.com

Rihla Research & Advisory
MiddlePowerAI.com

October 2026