September 2026

Global Nation®

Middle Powers in an Age of AI Superpower Dominance

Consultation Paper

Authors: Anna Hope and Hassan Damluji

Summary

AI is concentrating power fast. Grounded in today's reality and historical experience, this paper sets out a stark warning about where this is likely to lead. The Industrial Revolution shows how rapid technological development, which concentrates wealth and power in a small number of hands, can lead to an age of empire where the standard of living for most of humanity falls. Despite the promise of these new technologies, today's evidence shows that **we are on track for this history to be repeated**.

We also propose concrete solutions. We argue that the only way to prevent the worst geopolitical, societal and economic effects of AI-driven power concentration is through a bold international effort that, rather than seeking "AI sovereignty", can achieve "**AI non-dependence**". This means that no country should be so dependent on a single state or cluster of companies that it is vulnerable to their control. We propose **five critical tests for "AI non-dependence"** that can be used as a yardstick for progress in avoiding these risks.

We argue that meeting these tests requires more than governance and regulation alone. While these are important, countries need to actively **invest to ensure that no element of the "AI stack"**, the set of inputs — hardware, compute, energy and software — that is required to effectively deploy AI at scale, is dependent on just one country or company. Only when countries have real options to choose from can governance and regulation be successful. **Middle Powers are well placed to lead this effort**. They combine a need for effective international cooperation with the collective ability to act as rule-makers.

Contents

The purpose of this paper

There is real momentum behind efforts to curb AI harms and ensure its benefits are widely shared. In July 2026, the AI for Good Global Summit in Geneva ran back-to-back with the inaugural Global Dialogue on AI Governance, convening policymakers, industry, researchers and other stakeholders from around the world. Both the Dialogue

and the Independent International Scientific Panel on AI were established by the same UN General Assembly resolution in 2025, marking a significant step towards giving AI governance a dedicated multilateral, UN-anchored home. The Panel, whose 40 members were appointed in February 2026, presented its first report at the Dialogue, offering an early independent scientific assessment of AI's opportunities, risks and impacts. Earlier in 2026, India's AI Impact Summit in New Delhi concluded with the New Delhi Declaration on AI Impact, endorsed by 89 countries and international organisations within days of the summit and by over 90 since, building on the momentum of the Bletchley Park, Seoul and Paris AI summits. Alongside these efforts, initiatives such as the Rome Call for AI Ethics have sought to build wider consensus around principles including transparency, inclusion, responsibility and impartiality, demonstrating growing convergence on the need for AI to be developed and deployed in ways that serve the common good. In June, UNIDIR's Global Conference on AI, Security and Ethics gathered diplomats, policymakers, military officials, industry, researchers and civil society in Geneva to work through the security and ethical dimensions of the technology. These are all nonetheless voluntary commitments and non-binding frameworks, rather than enforceable rules. And serious questions are being raised about what effect these efforts will ultimately have, particularly given how fast the technology and industry continue to evolve[1].

This paper adds a complementary dimension to this fast-growing field of activity. Where much current work focuses on safety, ethics and governance frameworks, this paper addresses a more structural question: the distribution of power created by the AI stack itself, and what Middle Powers can do together to avoid dependency.

It sets out early thinking on how to avert the negative geopolitical, economic and social consequences of growing power concentration caused by the uneven distribution of Artificial Intelligence (AI) capacities, known collectively as the "AI stack". It describes the risks faced by most countries, sets out what would be needed to mitigate them, and offers suggestions on the approaches countries can take. It is intended to provoke discussion, build consensus on the realities of the risks we face, and move this thinking towards a practical strategy for policymakers navigating security in an AI age.

The problem: an impending age of AI imperialism

Artificial Intelligence is deepening existing fault lines in the international order: economic competition, security dynamics and the distribution of power. Yet debate on managing AI remains overwhelmingly technical, often siloed from the broader political question of how countries can cooperate in an era of fracturing multilateralism to restore order, create stability and ensure all countries can make progress. By "AI imperialism", we do not necessarily mean a plan of conquest. We mean the structural outcome in which a small number of states and firms control the inputs on which

1 Independent International Scientific Panel on AI, Preliminary Report (United Nations, 1 July 2026), available at https://www.un.org/independent-international-scientific-panel-ai/en/preliminary-report

everyone else's economy, administration and defence depend — and can therefore set the terms on which others use them.

On our current trajectory, AI's benefits will accrue primarily to superpowers and large technology firms, entrenching inequality within and between countries. This is often described as the "power concentration risk". We believe its likelihood and severity have both been consistently underestimated. The first industrial revolution offers a warning: a technological lead opened by a small number of European powers translated into military invasion, mass killing, mass unemployment, starvation, and the loss of political and cultural rights for large parts of humanity. There is no reason to assume a benign outcome this time — though power would concentrate in different regions than it did two centuries ago.

Concretely, if AI comes to play as central a role in government and business as expected, and if those functions in most countries depend on critical elements of the "AI stack" based in the US or China, we should expect:

- Jobs lost to AI everywhere, with the new jobs that replace them concentrated in the US and China rather than in the countries where jobs were lost – meaning fewer well-paid opportunities and rising unemployment elsewhere.
- An increasing share of economic value flowing to organisations based in the US and China, driving long-term economic divergence and falling living standards elsewhere (for comparison, over the period in which Britain pulled decisively ahead of the rest of the world, estimated income per capita fell in both China and India[23]).
- Governments or businesses seen as hostile or unhelpful by the US or China (or by companies based there) risk having critical functions disabled remotely, causing catastrophic failures.
- Any military conflict being so imbalanced that the US and China could pursue harmful policies towards other countries with effective impunity.

Why governance alone is not enough

There are a range of initiatives and ideas to advance AI governance at the global level to help protect countries from its harms and help ensure its benefits are shared. This is necessary, but not sufficient.

Regulation shapes behaviour within a system of power. It does not, by itself, alter the distribution of that power. The businesses that build frontier models — as well as the firms that design, fabricate and license the chips, energy and compute on which everything else depends — remain subject to the jurisdiction, and ultimately the

2 Broadberry, S. N., Custodis, Johann and Gupta, Bishnupriya (2015) India and the great divergence : an Anglo-Indian comparison of GDP per capita, 1600–1871. Explorations in Economic History, Volume 55 . pp. 58-75. doi:10.1016/j.eeh.2014.04.003 <https://doi.org/10.1016/j.eeh.2014.04.003> ISSN 0014-4983.

3 Broadberry S, Guan H, Li DD. China, Europe, and the Great Divergence: A Study in Historical National Accounting, 980–1850. The Journal of Economic History. 2018;78(4):955-1000. doi:10.1017/S0022050718000529

direction, of the states in which they are based, whether they are private or state-owned. Regulation alone cannot confine that power unless it too is backed by power.

The Qing Dynasty sought to prohibit the opium trade that had grown around British-controlled Indian opium. When it moved to enforce the ban, Britain responded with military force in the Opium Wars, compelling China to accept a new trading order that ultimately legalised the opium trade. We are already seeing a similar approach emerging in the AI age, with the United States showing that it is willing to threaten countries that try to regulate US technology firms in order to bring them to heel. The more dependent other countries become on the US or China for technology, the harder it will be to create governance frameworks that are enforceable.

What does security in the AI age require?

A central priority for those seeking to realise AI's benefits while mitigating its harms must be to prevent high concentration of power amongst a few states and companies. As an absolute minimum, this means no outside actor should be able to access or weaponise a nation's data, sabotage its public or private sectors, or switch off the systems its society depends on.

We propose five tests for non-dependence in the AI age. We deliberately use the phrase 'non-dependence' to mean the absence of external control over the systems a country depends on – not self-sufficiency across the whole AI stack, which, as we set out below, no Middle Power can realistically achieve:

1. Data ownership: a country's data cannot be accessed, manipulated, or used against it by an external power.

2. Operational independence: the functioning of critical public and private services cannot be disrupted or sabotaged by a foreign actor.

3. Access to competitive technology: businesses and governments can access AI tools that let them compete economically and militarily.

4. Capacity: countries have access to sufficient sustainable energy to power AI and a sufficiently skilled workforce to use it.

5. Hardware: no single external actor controls a country's access to the hardware necessary to deploy AI effectively.

Where any of these tests are not met, countries cannot be truly independent and secure. Instead, they are in a meaningful way *dependent* on other states and foreign companies to carry out the basic functions of a modern economy. That dependence can create unacceptable vulnerability.

Meeting these tests does not guarantee good governance, but without them, effective AI governance is not possible. Only once these conditions are in place can AI's benefits be shared by the majority, and its worst harms – including entrenched economic

inequality, an escalation of interstate conflict, and environmental disaster – be mitigated.

The real choice facing Middle Powers is not sovereignty versus dependency, but cooperation versus dependency. When pooled with others, the five conditions above become theoretically achievable. Alone, they are not. And crucially, shared interest will not automatically produce coordinated action.

Middle Powers: the path to independence

The retreat of the United States from multilateralism has underlined the important role of Middle Powers. These are countries with significant influence but without hegemonic capability that can play a role in driving international cooperation. Middle Powers combine a strong need for effective international cooperation with the collective ability to act as rule-makers. We believe they are uniquely placed to offer a credible path to curbing AI's harms and ensuring its benefits are genuinely shared.

None have any realistic chance of achieving full AI sovereignty across energy, chips, compute, and frontier models. For every country other than the US and China, going it alone is not a viable path to security: it guarantees falling further behind while still bearing the cost [4]of trying. Nor does buying access to US frontier models under a "sovereign AI" label solve this, since that access can be withdrawn by the US government at any time.

As recently as June 2026, the US Commerce Department directed Anthropic to suspend global access to two of its most advanced models (Claude Fable 5 and Mythos 5) for all foreign nationals, citing national security concerns; Anthropic complied, disabling access for all customers worldwide[5]. While this episode did not last long, it established the fact that the most capable models on the market can be switched off globally, at a few hours' notice, by one government. The same dynamic runs into hardware, which is arguably the most concerning. Chip design, fabrication and advanced lithography sit with a handful of firms in just a few jurisdictions (Nvidia in the US, TSMC in Taiwan, ASML in the Netherlands). When the UAE's G42 and Saudi Arabia's Humain were cleared in November 2025 to buy Nvidia Blackwell chips, the licences were granted subject to security and reporting conditions monitored by Washington[6]. The US government is directly controlling what US-headquartered private companies are allowed to sell, to whom, and on what terms. A country dependent on US-hosted models or production has no guarantee of the access it might want or need.

4 https://www.anthropic.com/news/fable-mythos-access

5 Anthropic, "Statement on the US government directive to suspend access to Fable 5 and Mythos 5", 12 June 2026, https://www.anthropic.com/news/fable-mythos-access.

6 https://www.commerce.gov/news/press-releases/2025/11/statement-uae-and-saudi-chip-exports

The challenges to getting this done are familiar to solving any collective action problem: why move first, why pay in if results aren't guaranteed, and why not wait for someone else to carry the cost? But given the risks we face in an AI age, we must find a way through.

What might a Middle Power coalition response look like? The following sets out a starting point for discussion, structured around the tests:

- **On data ownership:** countries could agree that all national public data and sensitive private data be hosted within the collective jurisdiction of the coalition, rather than with any single external provider. A "data sovereignty bloc" of this kind would mean no member's data sits under the legal or physical control of a country outside the alliance.
- **On operational independence and access to competitive technology:** rather than each country competing to build frontier capability alone, members could establish a cross-border research and compute consortium dedicated to open-source, high-performance models, reducing reliance on any single country or firm.
- **On capacity:** members, many of which run fossil-heavy grids, could coordinate siting and sourcing low-carbon energy to power AI systems (see Brazil's hydro-powered grid[7]), and coordinate investment in the skills needed to use AI effectively.
- **On hardware:** through pooled purchasing power and diversified sourcing across allied suppliers, members could ensure that no single external actor controls all of the hardware and compute on which a country's ability to deploy AI effectively depends. Because competing hardware cannot realistically be developed at the speed required, this is resilience rather than independence – but combined with the other four tests, it materially reduces the leverage any one supplier or government holds.

A summary of constraints

A key constraint in driving this agenda forward is the current lack of consensus. First, there are two camps whose approaches are at odds with this paper.

1. Those focused on maximising AI capability as fast as possible, for whom dependency on the fastest available provider is an acceptable trade-off rather than a risk worth solving.
2. Those pursuing full independent capability rather than pooled resilience, for whom "cooperation" would be seen as settling for less.

This paper has argued that both proposals are flawed. Dependency on the US or China is already putting countries – even their traditional allies – in a position of huge

7 International Energy Agency (IEA), "Brazil," Countries & Regions, https://www.iea.org/countries/brazil

vulnerability. And AI sovereignty is an impossible dream. But unless more consensus is built around this view, the political will to implement it will be scarce.

Second, even if consensus were to be achieved, implementing an effective Middle Power AI coalition to achieve the tests laid out in this paper is fiendishly difficult. For example:

Data ownership requires the political will and cost of building or trusting sovereign hosting infrastructure, at a moment when relying on incumbent providers is far cheaper and easier.

Operational independence requires countries to absorb short-term disruption ahead of the crisis this independence is designed to avoid, meaning that many may feel that these costs are not urgent or justified (this could be compared to the difficulty of encouraging investments in pandemic preparedness in advance of a pandemic).

Access to competitive technology depends on pooling research and compute at a scale no single Middle Power can reach alone – which means ceding some control to a shared, cross-border effort.

Capacity requires the slower, harder work of building sustainable energy infrastructure and a skilled workforce, but on a timeline dictated by the pace of AI development and at a time when public budgets are tight.

Hardware requires reducing single-point dependency by diversifying beyond a small number of suppliers that are concentrated at multiple points of the value chain in a handful of countries including China and the US[8]. This requires sustained coalition investment in expanding fabrication capacity among members who already hold a foothold, such as South Korea and Japan[9], alongside diversifying sources of critical inputs, so that the pool of independent options grows rather than staying fixed. It should be acknowledged that the most capable potential partners here are US treaty allies.

The underlying constraint running through all these challenges is the limits to international solidarity. Building a coalition means persuading countries and their publics to see international cooperation as a route to security and prosperity, not a threat to sovereignty. That belief is contested in many countries, but recent polling[10] suggests support for cooperation is starting to rise in a diverse range of countries across the world – something to harness and build on as a core foundation for success.

8 OECD (2025), "Mapping the semiconductor value chain: Working towards identifying dependencies and vulnerabilities", OECD Science, Technology and Industry Policy Papers, No. 182, OECD Publishing, Paris, https://doi.org/10.1787/4154cdbf-en.

9 CSIS, "Mapping the Semiconductor Supply Chain: The Critical Role of the Indo-Pacific Region": https://www.csis.org/analysis/mapping-semiconductor-supply-chain-critical-role-indo-pacific-region

10 Hope, A and Damluji, H "At its darkest hour, is internationalism back?," Global Nation, May 2026,https://globalnation.world/publications/at-its-darkest-hour-is-internationalism-back-2/

The theory of change

The following theory of change visualises the proposition set out in this paper to assist with interrogating the steps to success.

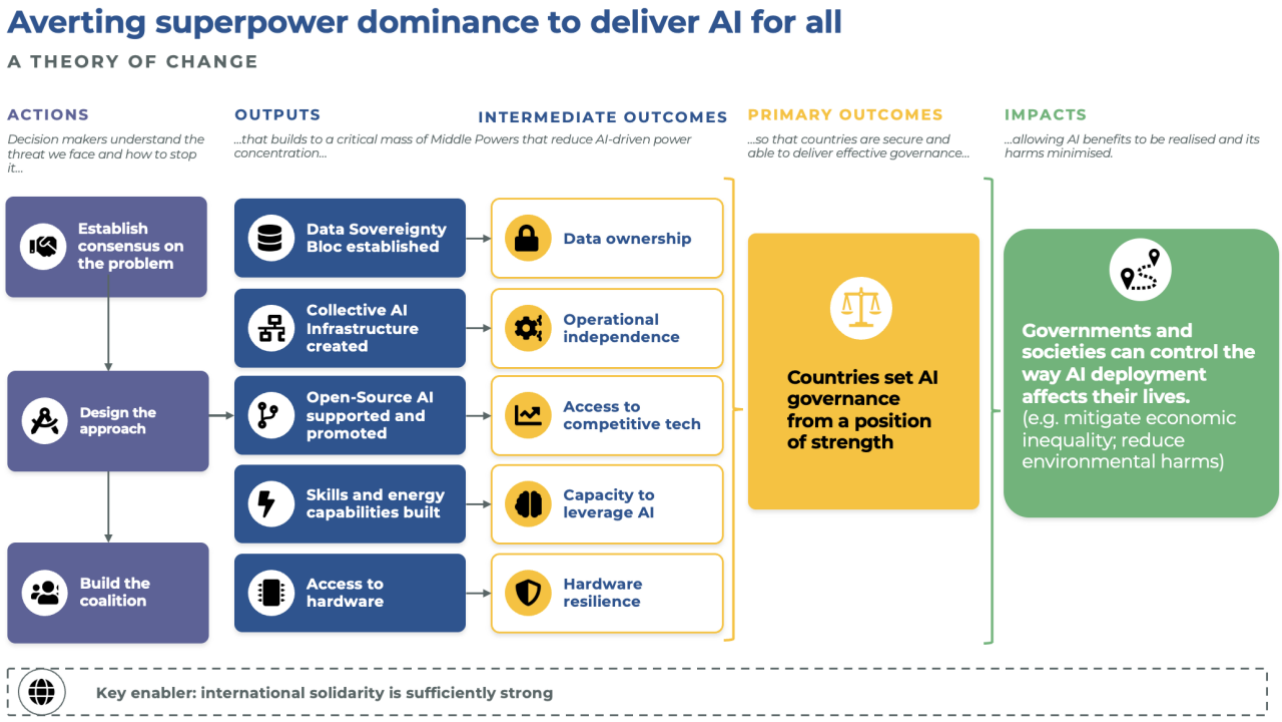

We propose that there are three key initial stages that begin sequentially but overlap in execution, with each building on the one before.

First, **building consensus on the problem**: ensuring policymakers, business and the public alike sufficiently recognise power concentration as a serious and immediate risk, not a speculative one.

Second, **designing the approach**: working out the specific mechanics of a coalition that avoids the collective-action problems that have undermined past cooperation efforts. This would include identifying which countries could act as credible anchors, what incentives would bring early movers, how access could realistically be restricted to members to mobilise more states over time, and how contribution could be tied to benefit so the burden doesn't fall unfairly on those who commit first.

Third, **building membership**: converting design into an active coalition, starting with a small number of countries whose early commitment gives others confidence to follow.

This paper is aimed at the first stage. We see the second and third as the crucial next piece of work.

Acknowledgements

While the paper reflects the views of Global Nation alone, we are hugely grateful for the input and advice of Christian Schlaepfer, Policy Advisor, Intergovernmental and UN Affairs at Starling Institute; Gordon LaForge, Co-Director at New America and Senior Fellow for AI Policy at Starling Institute; Andrea Venzon, Founder of Avantgarde; Peter Reeves, Data and AI Lead at Seismic; and Olivia Jennings, Principal Consultant at Seismic.

Feedback and collaboration

We are keen to hear from those who have feedback on the ideas set out in this paper or would like to collaborate on this work.

Please email Anna Hope, Policy Lead at Global Nation, at Anna.Hope@globalnation.world.