



HOW TO WIN FRIENDS AND EXPORT AI

BY SATVIK PENDYALA AND RYAN FEDASIUK

Satvik Pendyala (Satvik.Pendyala@AEI.org) is a research associate at the American Enterprise Institute. Ryan Fedasiuk (Ryan.Fedasiuk@AEI.org) is a fellow for China and Technology at the American Enterprise Institute and an adjunct professor at Georgetown University's Security Studies Program. He previously served as an advisor for US-Chinese bilateral affairs at the U.S. Department of State.

The world is running [short](#) on computing power needed to sustain the AI boom. Overbooked logic nodes, constrained packaging, and a massive memory crunch are starting to plague software developers and strain markets. Faced with immense local demand, Washington might be tempted to curtail [exports](#) of advanced chips and preserve high-end compute for data centers located in the United States.

That would be a colossal mistake. As countries around the world race to build and install data centers capable of serving more demanding AI workloads, the world is on track to run short on energy far sooner than silicon. In a new [report](#) for the American Enterprise Institute, we find that leading US AI chip companies may face a shortage of more than 4 million high-end chips by 2028, while new datacenters also face a roughly 2- to 4-year grid connection delay.

This is why the Trump administration works to [export AI infrastructure](#) abroad: Chips earmarked for exclusive US use risk sitting idle for years before being plugged into the grid and running AI workloads. Keeping chips at home without the necessary energy infrastructure risks leaving them to waste.

So, who should be eligible to purchase scarce American compute, on what timescale, and under which conditions?

The case for AI exports

Picture a data center groundbreaking outside Jakarta in 2030, but the silicon, the AI models running through it, and underlying energy infrastructure are all Chinese. Washington must prepare to offer an alternative if it wants to shape the tangible infrastructure of AI buildout. Otherwise, the risk is a repeat of the American experience

in 5G, where Chinese infrastructure became embedded in critical digital supply chains [throughout](#) the developing world.

America's best strategy over the next decade will be in shaping the AI standards defining the next generation of technology development. China possesses a major advantage in open-source AI, physical AI, and industrial systems—the United States will likely run into domestic political hurdles building large numbers of data centers as more states [null](#) restrictions. American statecraft should therefore focus on building computing hubs abroad to process American AI workloads and spread its AI stack to foreign markets. The best way to ensure that hyperscalers in foreign markets run AI models on American chips—and contract US tech companies to serve local customers—is creating a distributed network of computing nodes based on American hardware.

Markets not adequately served by American AI infrastructure invite Chinese competitors to stake out entry points. Many countries may accept a performance cut and deal with China's smaller supply for at least some amount of compute. With Chinese chip companies facing an eager domestic market, countries will compete to be first in line for a limited amount of new Chinese AI chips. AI infrastructure gets procured in multi-year contracts, and every contract signed with the Chinese AI ecosystem encourages partnering with China's extensive green energy industry, with transformer [exports](#) alone rising 36% in 2025. Without an explicit export strategy, American firms risk being sidestepped by Beijing's ability to package sectoral industrial champions as part of its technodiplomacy.

Being selective about compute recipients

If AI chip exports are essential to America's tech strategy, how should Washington best deploy them? As Washington considers markets to serve first, partners in Asia should expect Washington to ask about local interconnection timelines and infrastructure assessments to evaluate who best to triage limited export capacity. A second, related question: how to balance the upside of sharing frontier chips with new partners against the risk of [remote access](#) by Chinese AI labs, [smuggling](#), or diversion to third parties? American compute should be the easiest to sell and hardest to steal.

The United States will not get every country on board with extensive end-user monitoring and restricted localization. For Washington, it's worth being clear about which ones won't, why, and whether it might still be worth selling them high-end compute anyway.

India, South Korea, and France have AI developer ecosystems, tightly enforced data localization requirements, and deeply rooted preferences for maintaining some autonomy from Washington and Beijing. Although these countries will balk at [American demands](#) for regular customer audits and similar “know-your-customer” safeguards, the United States still has an interest in furnishing them with high-end compute. Bangalore, Seoul, and Paris all have strong domestic tech industries and AI markets poised to grow in ways complementary with San Francisco. These countries will continue insisting that US frontier models and user data be hosted locally—acceptable conditions for premier markets of importance. The United States should allow for chip exports here, provided implemented measures against transshipment and preventing entity-list access to data centers. The same countries should also be first in line for licensed/waived-license export of the best American chips.

Several other countries will wait out American monopoly pricing, preferring future Chinese compute with fewer strings. Advanced chips already have multi-year lead times, so a foreign government weighing US terms today measures them against the future availability of Chinese hardware. For Washington, aligning major markets on the larger US ecosystem might make it worth tolerating partner hedging on some specific asks. The [US-UAE AI Acceleration Partnership exemplifies](#) this dynamic: While Abu Dhabi generally [failed](#) to offload Chinese telecommunication networks, Washington could live with some degree of uncertainty or opacity in Emirati compliance. Some degree of influence over international developer ecosystems was preferable to potentially ceding the market to China. If Washington shows no flexibility, countries well-positioned to adopt AI infrastructure could become Huawei’s or Alibaba’s first major international AI chip customers. If Washington demands guardrails against remote access or transshipment (as it should), it should offer near-frontier capabilities to sweeten the term sheet. For these markets, access to the best US chips should be understood as a function of compliance with those guardrails.

A final tranche of states face such high demand for AI compute that they may accept American demands in exchange for high-end chips access. Australia courts significant American AI investment and is considering a [change to its copyright law](#) to accommodate American demands. Other Asian markets, including Singapore, Malaysia, and Thailand are major finance and digital services hubs needing large compute clusters to offer globally competitive products in the AI era. These hubs also come under suspicion of being major training hubs for Chinese AI labs. Leading Chinese labs maintain [remote access](#) to data centers in countries in Southeast Asia, with Malaysia the biggest [destination](#) for overseas Chinese data center investment. If the United States wants its larger export control regime on China to have teeth, it needs to demonstrate that compliance with remote access restrictions on Chinese firms is the price for access to US chip exports. These markets remain key for the US digital ecosystem and the best way to enable their own growth is to structure chip exports so American AI technology can best help develop their own domestic and regional workloads, reducing dependence on Chinese investment as a driver of AI demand.

Executing export promotion

American AI export policy should focus on keeping America’s leading-edge compute competitive with an emerging Chinese stack. The focus should be on persuading potential partners to accept American asks during the limited period when the United States holds a relative monopoly over global compute production. US chip exports should thus be paired with demonstrated milestones and achievement-based funds.

The immediate goal of export promotion is finding countries most ready to install large quantities of US chips. The vehicle for that is the Export-Import Bank (EXIM), whose [ExportAI](#) initiative

provides export credit, insurance, and loans foreign customers to purchase US AI chips. EXIM’s programs are best-geared for customers with capable infrastructure to jump to the front of the queue. In addition to existing financing initiatives, EXIM should adopt guidelines for choosing who and how to finance. EXIM should provide special consideration to countries with sub-three-year grid connection delays for datacenters, with more lenient repayment conditions for AI initiatives in countries that have a high chance of [contributing back](#) to the US tech ecosystem. Countries that haven’t yet purchased significant Chinese tech infrastructure also should be given preferential payout terms, with extensive post-hoc Chinese AI infrastructure purchases automatically triggering a stricter financing regime.

While there may be a rather limited list of immediate markets ready for EXIM’s financing, there are countries that will require systemic energy and telecom buildout to become sponges for US compute. The Development Finance Corporation (DFC), which supports capital-intensive projects like power infrastructure and grid modernization through flexible loans and guarantees, is an ideal vehicle. Whereas EXIM’s direct AI export financing machinery targets immediately accessible markets, DFC’s mandate emphasizes the longer-term infrastructure buildout needed to bring countries to where they can buy American chips. To amend its guidelines, DFC’s funding can come with preferential repayment rules conditional on countries not accepting large imports of Chinese AI hardware, reduced interest rates for buying energy and telecom hardware from [Pax-Silica](#) members, and direct grants for states to develop talent and infrastructure to build supply chain nodes needed for an independent electric stack.

Conclusion

Return to that groundbreaking outside Jakarta. The concrete has not been poured, the racks have not been ordered, and the engineers who will spend their careers optimizing for one architecture or another are still university students. If the United States decides to curb its chip exports abroad, the global intelligence economy will begin to bend, slowly and imperceptibly, away from the United States. Washington should take steps to lock in capital commitments before Chinese chips hit the global market at scale.

The Pilot commentaries and responses represent the views of the respective authors. Alternative viewpoints are always welcomed and encouraged. Please write to rob@pacforum.org for more information on how to contribute.