



MEMORANDUM

October 29, 2025

To: U.S. Secretary of Commerce Howard Lutnick; U.S. Secretary of the Treasury Scott Bessent; U.S. Trade Representative Jamieson Greer
From: Americans for Responsible Innovation
Re: Exports of NVIDIA's B30A AI chips to China threaten to erode the United States' decisive compute advantage

EXECUTIVE SUMMARY

As top U.S. and Chinese economic officials negotiate a trade deal that will address issues ranging from tariffs to rare earth exports, a relaxation of U.S. AI chip export controls is likely among China's top negotiation priorities. This memo warns against authorizing exports of NVIDIA's B30A or similarly powerful chips to the People's Republic of China (PRC), a move that would meaningfully narrow, and potentially fully erase, the US lead in AI compute.

While official technical specifications for the B30A have not been released by Nvidia, multiple sources have reported on its likely capabilities and pricing based on information from industry insiders and prospective Chinese purchasers.¹² This reporting indicates the B30A delivers roughly B300-class capabilities at nearly the same price-performance ratio, enabling Chinese labs to assemble training clusters that closely match US state-of-the-art with only modest extra spend.³

New research from the Institute for Progress finds that without B30A or other AI chip exports, the US is on track to build out 11x-30x more compute than China in 2026, which will be a critical advantage in competing for AI capabilities and global market share. However, B30A exports threaten to substantially reduce or entirely eliminate that advantage. In a best-case scenario, with conservative B30A sales to China, the US compute lead would shrink to <4× Chinese capabilities. In a worst-case scenario, the sale of B30As and similar AI chips from other companies threatens to flip the compute advantage to China.

The sale of B30A chips to the PRC would also hand Chinese cloud providers price-competitive US chips to win global market share and diffuse advanced computing across PRC state and defense ecosystems. None of these outcomes align with the Administration's objective to "maintain, obtain, and enhance" our technological edge.

¹ Reuters, "[Exclusive: Nvidia working on new AI chip for China that outperforms the H20, sources say](#)," Aug. 19, 2025.


² Tom's Hardware, "[Nvidia's next-gen AI chip could double the price of H20 if China export is approved – Chinese firms still consider Nvidia's B30A a good deal](#)," Sept. 4, 2025.

³ Institute for Progress (IFP), "[Should the US Sell Blackwell Chips to China? Assessing the impacts of exporting the NVIDIA B30A AI chip](#)," Oct. 25, 2025.

KEY FINDINGS

- B30A ≈ B300 on price-performance.** If B30As with half the B300's total processing power and memory bandwidth are sold at roughly half the B300's price, China can build clusters with similar capabilities to US B300 clusters. Dividing the capabilities of one B300 into two B30As will impose a cost penalty of around only 20 percent to account for the extra networking and server equipment required to build a cluster with the same total compute power—well within reach of state subsidies.
- The B30A is a giant step up from H20 and over current control thresholds.** According to reported technical specifications, the B30A offers >12× the H20's capability and >18× the US export-control performance thresholds that currently gate China-bound AI chips. It is not a modest, compliance-friendly chip—it is a frontier-class accelerator designed for the PRC market.
- Compute balance impact is severe.** With a continued ban on advanced AI chip exports to China in 2026 (and no smuggling), the US would hold roughly a 31× compute advantage. Approving B30A exports shrinks that lead to <4×, and broad US vendor sales could flip the advantage to China.
- Domestic Chinese supply cannot substitute near term.** Huawei and others cannot yet produce Blackwell-class chips or computing clusters due to U.S. export controls on semiconductor manufacturing equipment and high bandwidth memory (HBM) chokepoints. B30A sales thus unlock capacity Beijing otherwise lacks, both for domestic AI data centers, and for its ability to compete with the US in global markets.

FIGURE 1: B30A Export Impacts



THE B30A IS NOT “HALF AS GOOD”

While industry sources have described the B30A as half as powerful as Blackwell chips designed for U.S. companies,⁴ the risks of selling China the B30A are not halved. That frames the issue at the chip level, when the decisive variable is cluster-level capability per dollar. If the B30A delivers roughly half B300 performance at half price, the effective price-performance ratio is similar, meaning Chinese labs can scale out to parity.

In practice, PRC state backing removes the modest 10 to 20 percent cost penalty that recent research from IFP estimates for matching US clusters, nullifying the “half as good” premise.

Selling a less powerful AI chip capable of stacking with more AI chips isn’t like selling degraded fighter jets to a foreign nation. Whereas downgraded versions of aircraft, such as F-16s, are strategically outmatched on the battlefield by cutting-edge fighter jets, a chip like the B30A has all of the same capabilities as cutting-edge chips when purchased in quantity. A more accurate metaphor might be the export of liquid natural gas (LNG) if the U.S. government were to require shipments in tanks half as large. The end result is a small tax on foreign purchases of the U.S. product with no downgraded impact on the product’s utility.

STRATEGIC RISKS OF PERMITTING B30A EXPORTS

Right now, the U.S. maintains a global edge in the development of advanced artificial intelligence, but China is closing the performance gap.⁵ Both nations have different strategic advantages, including in energy, data, access to talent, and compute. To date, the United States has maintained a decisive advantage in access to compute, supporting the U.S. lead in developing AI. Sale of the B30A chip to China threatens to erase that critical advantage.

- **Erosion of US aggregate compute lead.** The net effect is faster PRC access to frontier-scale training and deployment, reducing the US margin for error in next-gen model races and enabling catch-up on agentic, multimodal, and autonomy-relevant systems.
- **Global market displacement.** B30A-equipped Chinese cloud providers can undercut or match US cloud offerings in third markets, siphoning workloads, data processing, and the downstream AI services stack.
- **Defense and state diffusion.** Once in-country, chips flow—formally or informally—across PRC state-owned enterprises and the defense ecosystem, compounding military-civil fusion.

CHECKING FALSE CLAIMS REGARDING B30A EXPORTS

- ➡ **“We’ll crowd out Huawei.”** Huawei cannot meet near-term domestic/global demand at scale due to semiconductor manufacturing equipment (SME) and HBM chokepoints. PRC policy can also create artificial demand for domestic chips regardless of US

⁴Iyer, Ryan. [Nvidia said to be developing new, more powerful AI chip for sale in China](#). TechCrunch. August 19, 2025.

⁵ Yolanda Gil and Raymond Perrault. [The 2025 AI Index Report](#). Stanford University. April 2025.

exports. Allowing B30A exports thus fails to durably suppress indigenization while immediately boosting China's capabilities. Exporting large numbers of B30As could even enable the PRC to export the Huawei Ascend chips they can produce, allowing the PRC to compete with the U.S. on global hardware diffusion.

- **“Limited quantities are safe.”** Even a conservative forecast of exports measurably shrinks the US compute edge. If ceilings loosen—or if chips leak to larger buyers—the advantage compresses below 4× and could invert. Compute is additive; leaks compound.
- **“It's compliant with thresholds.”** Reported B30A specs exceed current thresholds by over 18×. The B30A exceeds the performance of the current top AI chip available to China by 12×. While industry sources have described the chip as being designed to fall within government thresholds, independent research shows that the B30A would leapfrog current chips on the market and exceed acceptable performance under current export control standards.

POLICY OPTIONS AND RECOMMENDED COURSE

The most prudent course of action is to deny export licenses for the B30A entirely. Doing so would preserve the United States' widening lead in advanced AI compute capacity through 2026 and 2027, ensuring that American labs and allied partners maintain the decisive infrastructure needed for frontier model development. Approving these exports, by contrast, would grant Chinese firms near-parity in compute performance and undermine the strategic advantage that current export controls have created.

If policymakers determine that some form of limited export must be considered, any such relaxation should be strictly confined to non-frontier accelerators that fall below current export thresholds not only in technical specifications, but also in overall price-performance capability. This distinction is critical: a chip that meets compliance on paper can still, when deployed at scale, erode U.S. aggregate compute advantages. Any narrow authorization should include explicit annual unit caps, rigorous end-use monitoring, and automatic revocation clauses in the event of diversion or leakage to high-risk entities.

Finally, to lock in America's lead regardless of future licensing debates, the Administration should complement chip-level restrictions with upstream measures. Strengthening enforcement on SME exports and HBM components would limit China's ability to domestically replicate Blackwell-class performance. Together, these steps would maintain a consistent and enforceable framework that protects U.S. national security, sustains the domestic innovation ecosystem, and preserves the credibility of U.S. commitments as President Trump seeks a trade deal with the PRC.