The Cloud Loophole: Why Chip Sanctions Cannot Secure U.S. AI Leadership

The Economy Research Editorial[1,2]

[1]The Economy Research, 71 Lower Baggot Street, Dublin 2, Co. Dublin, D02 P593, Ireland

[2]Swiss Institute of Artificial Intelligence, Chaltenbodenstrasse 26, 8834 Schindellegi, Schwyz, Switzerland

Abstract

This article argues that cloud access now constitutes the core loophole in US sanctions on AI chips against China. While export controls may be effective at denying physical possession of sophisticated Nvidia chips, their ability to deny access to functionally similar computational power through rental of cloud infrastructure via foreign data centers, brokers, and increasingly globalized cloud operators is far weaker. The Remote Access Security Act correctly addresses a real logical loophole in the existing control regime by treating access to controlled chips via the cloud as an export. Yet, this article demonstrates that implementing these controls in practice is difficult. Compute does not behave like a tangible product in a controlled supply chain; it is networked, divisible, rentable, and routinely processed through third-country intermediaries, obscuring who is actually accessing it. True control over computing at the international level demands intrusive surveillance, alignment of legal regimes with allies, private sector compliance burdens, and extraterritorial enforcement that likely most US partners will reject. Comparison with rare earths highlights the inherent asymmetry. Raw materials generate leverage because there are physical choke-points in their extraction, processing, and production into magnets. Compute moves through a regime of contracts, accounts, servers, and data flows. Consequently, policy conclusions are not that controls on the cloud have no value but that they are not a substitute for maintaining superior technology, primarily through strengthening talent, research institutions, infrastructure, and their productive deployment.

1 Introduction - Cloud Access as the Missing Link in US Chip Sanctions

Since the outset of the semiconductor revolution, technological dominance in computing hardware has been viewed as a fundamental plank of national power.[1] Over the twentieth and early twenty-first centuries, the United States has always tried to sustain technological dominance in crucial fields, ranging from microprocessors to cutting-edge AI accelerators. This has been motivated by strategic and economic reasoning. Economically, this helps retain high-value industries and high-wage jobs. Strategically, technological dominance can offer leverage in peaceful competition and potential conflict.

The current competition with China over AI and advanced computing is only the latest page in a long history of technological rivalry among great powers. In the Cold War, the US and its allies in the Coordinating Committee for Multilateral Export Controls (CoCom) used restrictions to control not only weapons but also high-powered computers and semiconductor manufacturing equipment to set back the rival's technological development by denying them access to the most advanced Western science and engineering.[2] While some restrictions impeded Soviet technology, others were circumvented through third countries, smuggling networks, or domestic innovation.[3] These examples illustrated both the potential and the limitations of sanctions as tools of technological statecraft.

In the present moment, the logic of US Export controls on advanced AI chips is unambiguous. To American policymakers, artificial intelligence is a dual-use technology.[4] It can be harnessed to promote economic productivity, scientific discovery and consumer benefit; it also sustains military systems such as autonomous weapons, cyber weapons and intelligence collection. For this reason, it is argued, limiting China's access to the most advanced chips-one example being Nvidia's A100, H100 and Blackwell lines-can at least slow China's arrival at frontier AI techniques.[5] This, in turn, might buy the United States and its allies some additional time to hold onto their lead. The current sanctions regime, created in 2022 and since intensified, demonstrates the emerging consensus among policymakers in Washington that technological containment is a necessary, though not perfect, facet of strategic competition with China.[6]

But the reality of the global semiconductor and cloud computing ecosystem challenges that logic. The distinctions between hardware and services and between physical exports and virtual access to data, are getting fuzzy. US authorities can regulate the physical export of chips and the US has a hard time, if not an impossible time, trying to regulate access to cloud computing services across the borders of the Internet. There is now a new access point for the Chinese to access the world's computing power, which is largely beyond US control; we refer to that as the cloud loophole-the ability to rent out high-performance computing capacity from data centers outside of China.[7]

US export controls on AI chips are about blocking frontier AI development in China by prohibiting the sale of the most advanced chips to Chinese customers. This may sound reasonable; without the hardware, there is no way to use it. However, the reality is quite different and the world cloud computing market has irreversibly transformed the provisioning and consumption of access to advanced computing. Physical chips are no longer necessary; instead organizations buy time from vendors being able to rent GPU hours from a provider that is

either close or far away.[8] However, access to GPU-hours should not be mistaken for complete AI readiness: firm- and country-level adoption also relies on the quality of data, engineering talent, models-integration capacity, organizational restructuring, capital depth and conversion of compute into useful applications.

However, a number of authors identified an obvious inconsistency banning physical sales of equipment but not the ability to access the computing power by other means. A non-US Operator could buy Nvidia chips via legal means and provide computing services from a third-country data center to Chinese customers, resulting in a jurisdictional issue that US law could only address indirectly via licensing requirements, allied cooperation, and downstream compliance pressures. More recently, reports have emerged of one Chinese AI company renting space on 2,300 Blackwell chips, which are kept in Indonesia.[9] Others, such as Alibaba, ByteDance and Tencent, have signed large contracts with foreign suppliers to rent tens of thousands of Nvidia GPUs in external data centers.[10] Often, these arrangements involve a thick web of middlemen and shell corporations, making it hard to determine who the end user actually is.[11] The use of external facilities also gets around the export ban because the hardware is never physically inside China and, under US law, there has been no export.[12] Even regulators recognize this. According to the US government, 'cloud providers are not exporters when their hardware is used to provide cloud services to foreign customers.

Politicians have been aware of this issue with remote access. Members of the House Select Committee on the CCP have argued that the rental of cloud GPU computing is the Achilles' heel of the US export control system,[13] since offering computing capacity on an external network to China whether through chips or remote computing services, directly increases the Chinese Military and Intelligence Force. This was the reason for the House-passed 2026 Remote Access Security Act, or RASA. RASA would treat any access to computing power via cloud computing as an export, thus removing the loophole.[14] Congressmen have said that the measure is essential for national security, as the chips sourced from American providers are used in both military and civilian projects and that using cloud computing services, even when distant, provides the same benefit as having the chips in China, meaning the country could effectively sidestep the embargo on physical chips.

The ultimate enforcement challenge is also, arguably, the most intractable: the nature of cloud services. Cloud access is the logical missing link in US chip sanctions; however, turning that conclusion into an effective controls regime turns out to be more complex than the theoretical framework. The technical and political obstacles to policing cloud computing on a global level are so daunting that an embargo on cloud computing could be seen to be little more than an illusion. Unlike physical objects (such as rare earth metals) which have potential border-crossing points and can be examined and controlled at the border, computing power is untethered, borderless and integrally dynamic. Cloud computing infrastructure is highly distributed: not only do hyperscale providers such as Amazon, Google and Microsoft operate on a global scale, but a myriad of regional and niche providers ranging from Singapore to Frankfurt have already proliferated clusters of high-end GPUs.[15] Instead of using access control as a proxy for technological leadership, the US needs to focus its primary strategic effort in those areas where it still has lasting strength: talent, institutions, frontier research capacity, and the productive use of computing. This is the surest way to remain a leader.

2 RASA and the Limits of Global Compute Enforcement

It appears fairly clear that Congress wants to achieve a consensus on US export Controls since they see current US export Controls as lacking. The RASA, or Remote Access Security Act, would allow US authorities to expand the scope of export law to cover cloud-based access to chips deemed a "controlled" item.[16] According to the RASA bill, they would be allowed to prompt cloud service providers to certify to the US authorities that any given foreign tenant of cloud service is not an entity restricted by the US as a foreign end-user, it would effectively treat computing access as export licensing. Although the RASA bill passed easily in the House in January 2026, making the concept popular, it would be a different matter to execute.

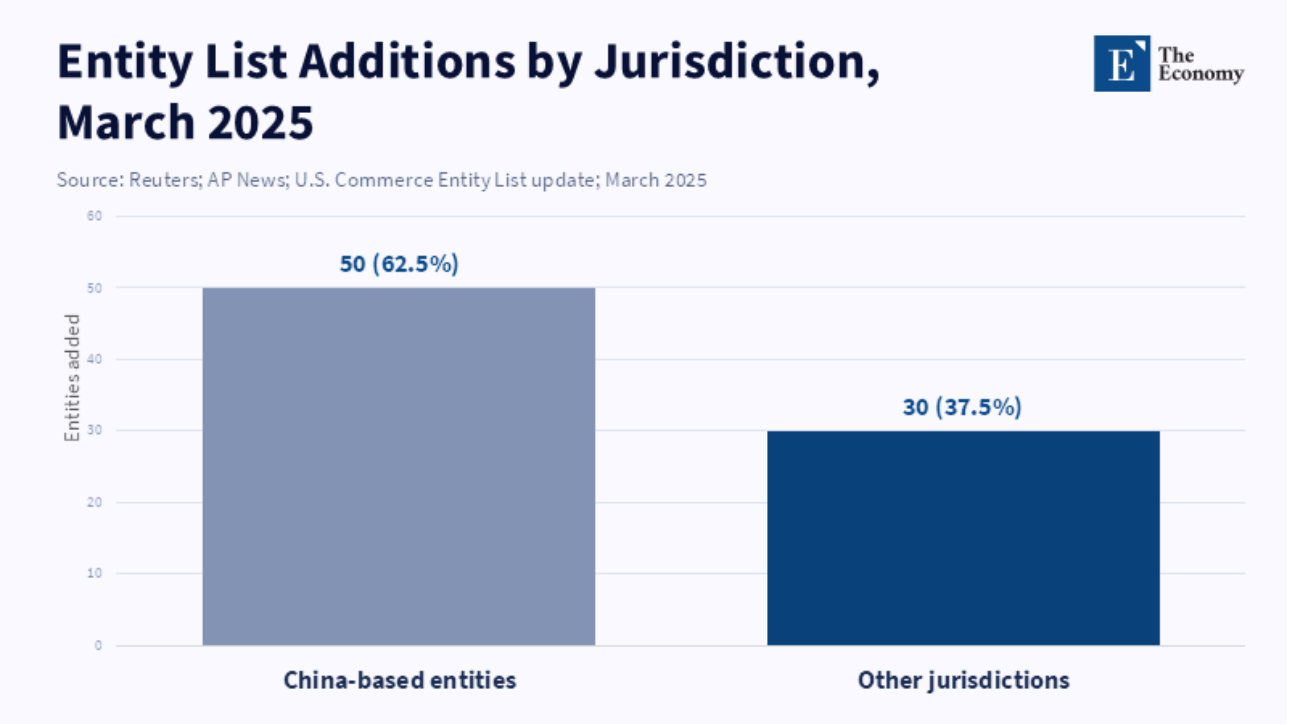

Figure 1: Entity-list expansion widens the enforcement perimeter, but it does not solve the deeper problem of identifying real end users.

Analysts are noting multiple issues: one issue involves determining who is actually using what remote GPU in what context. To do this, US authorities would essentially be asking US cloud service providers to vet each of their customers and the respective project to determine if it is an end-user who is a Chinese target. This would prove to be a near-impossible task for whoever attempts to circumvent the restriction, since the nature of digital service can often be masked or distorted at many levels. Chinese firms will be able to hide their identity by routing themselves through a third country, and through their shell corporations, splitting workloads amongst many different accounts, or disguising their IP addresses using VPNs;[17] no algorithm would ever accurately be able to discern how evasive their actions would be.

Highly advanced AI monitoring would struggle to distinguish valid traffic from disguised activity as users' actions would simply get more evasive. We already see this with the global financial industry which has tried for years with Anti Money Laundering and Know Your Customer policies to limit illegal access to the market but is still failing to meet that end [18] and, presumably, would also appear within the realm of the cloud compute market as well. The Commerce Department states that it can't keep out Chinese imports with current

inspectors and now it expects them to be able to discern between indirect foreign users accessing computing services.[19] Put differently, to effectively enforce a restriction on cloud-based computing, US cloud providers would essentially be performing a difficult detective job, bordering on the impossible. This would be no simple feat and without global, real-time, data sharing and identity verification, even the most resourceful agencies could not possibly police the entire cloud-based compute market.[20]

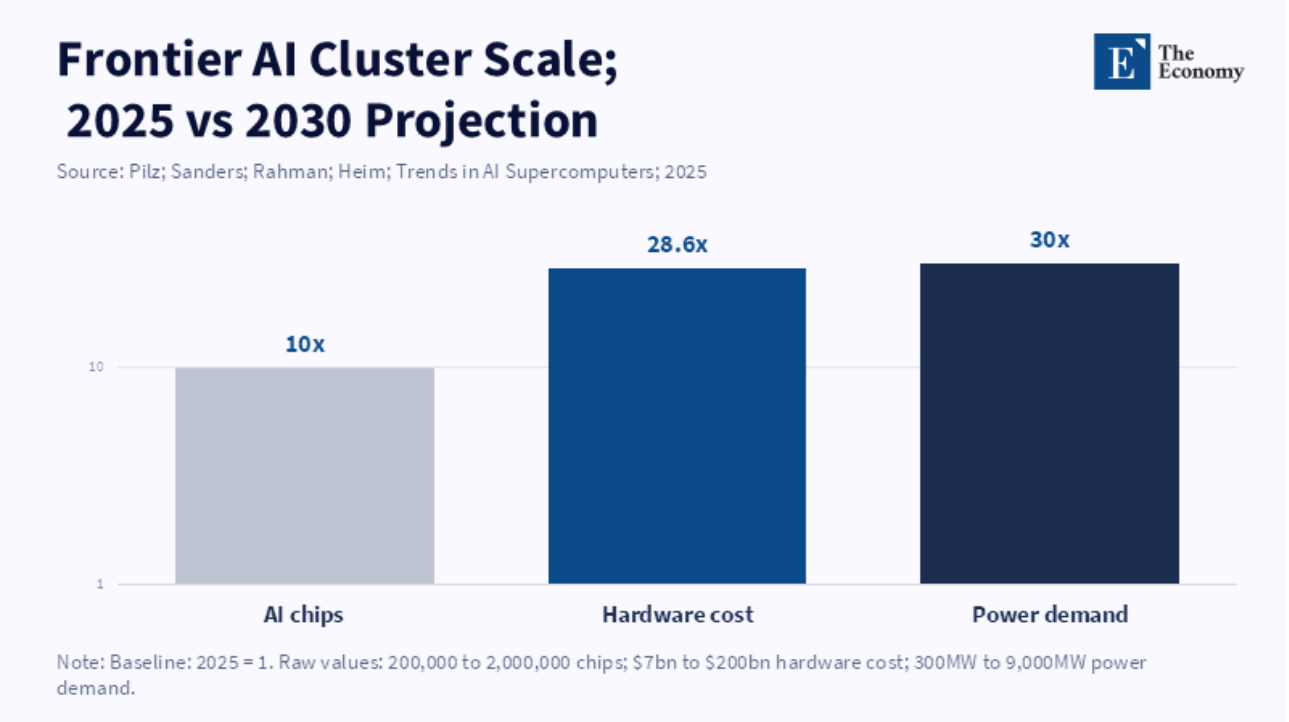

Figure 2: The scale of frontier AI infrastructure is rising faster than conventional compliance systems can realistically monitor.

Another issue with forcing cloud service providers to identify the user who would access what would involve cost. Critics are concerned that forcing stringent know-your-customer (KYC) requirements on US cloud providers would cause significant increases in operational costs and decrease profits; this could then possibly push US cloud providers to limit sales to listed Chinese entities due to business and political pressure. However, restricting US cloud providers to only entity-listed firms will not have the effect they desired. Instead, restricted parties can always find ways to get around it through technicality, such as having the Chinese AI firms employ Chinese citizens outside of the United States and pay them with non-US-sourced cash since the US cloud provider would have no reason to believe they are indeed a foreign end-user unless specifically restricted. In the absence of a worldwide, real-time registry of who can use which computer, the government would force private US companies to effectively self-regulate, which would not work.

A third issue revolves around the availability of controlled compute outside of the jurisdiction of the United States.[21] One could procure Nvidia chips in South Korea and sell computing services to Chinese clients using South Korean-based servers since US law does not apply. A Chinese national could also potentially obtain access to a data center in a country outside of US jurisdiction and provide computing services to China from that foreign location. However, with the RASA, enforcing the restriction abroad would require all other nations to enact their own laws governing US restrictions,[22] and many close allies within the EU, South Korea, and

Japan have never enacted such strict export control policies and refuse to implement them within their respective countries. If South Korean data centers are selling GPUs to Chinese end-users, US officials have no control over such practice without reintroducing them to the US Commerce Department. Again, thereby invading their national sovereignty.

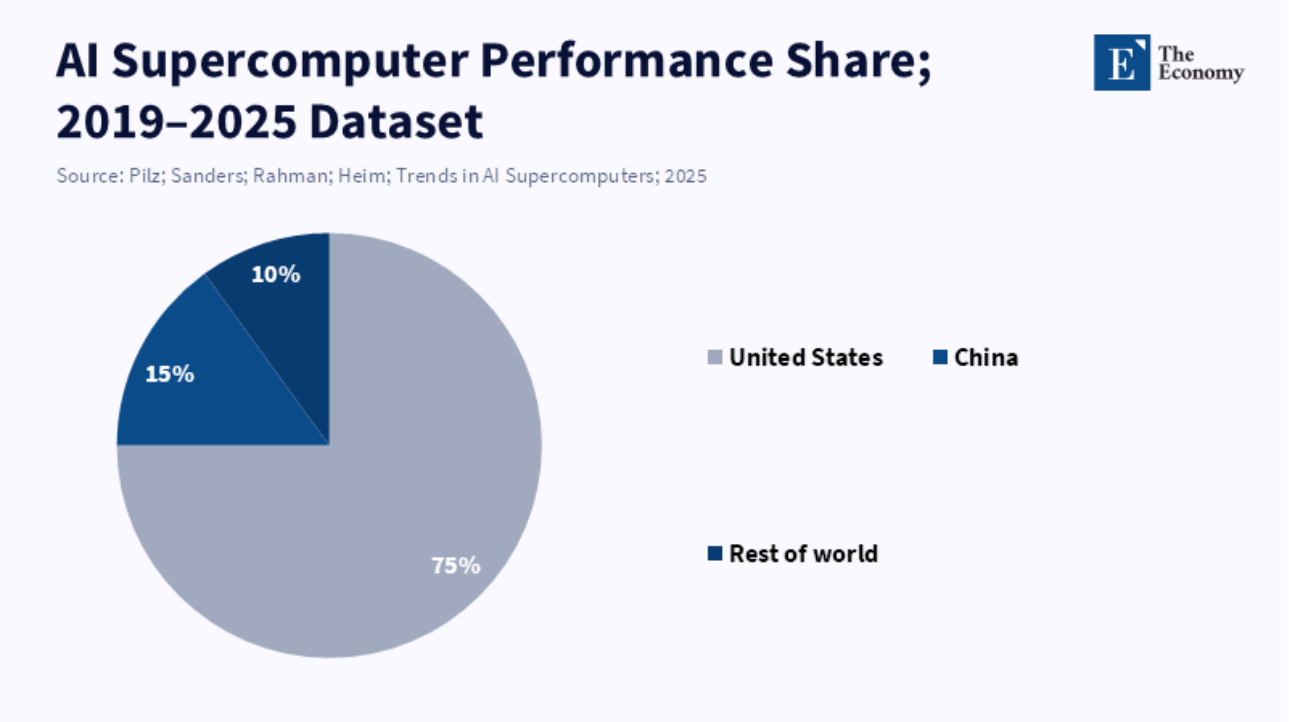

Figure 3: Compute power remains heavily U.S.-concentrated, but concentration is not the same as full jurisdictional control.

Finally, extraterritorial enforcement will not be palatable in foreign countries. One report stated the Carnegie Institute expressed that nations would object to US demands to enact KYC protocols on their servers and foreign governments, particularly in countries with massive investments in their own data center infrastructure such as those in Southeast Asia, would argue this oversteps US authority. Furthermore, nations such as Indonesia, Malaysia, and Thailand have had no problem with their respective data centers using the latest available technology and implementing their own policies, so if they are asked to go above and beyond what they normally would, by enacting US political policies on their own soil they are almost certainly to refuse to cooperate. In the last few years, there has been a global push towards modernizing existing infrastructure with the new developments in technology, particularly focusing on expanding into developing countries. They see this as crucial to economic growth and development; the US ordering these countries to enact its specific policies is, in their view, simply overstepping and attempting to dictate economic policy to other countries.

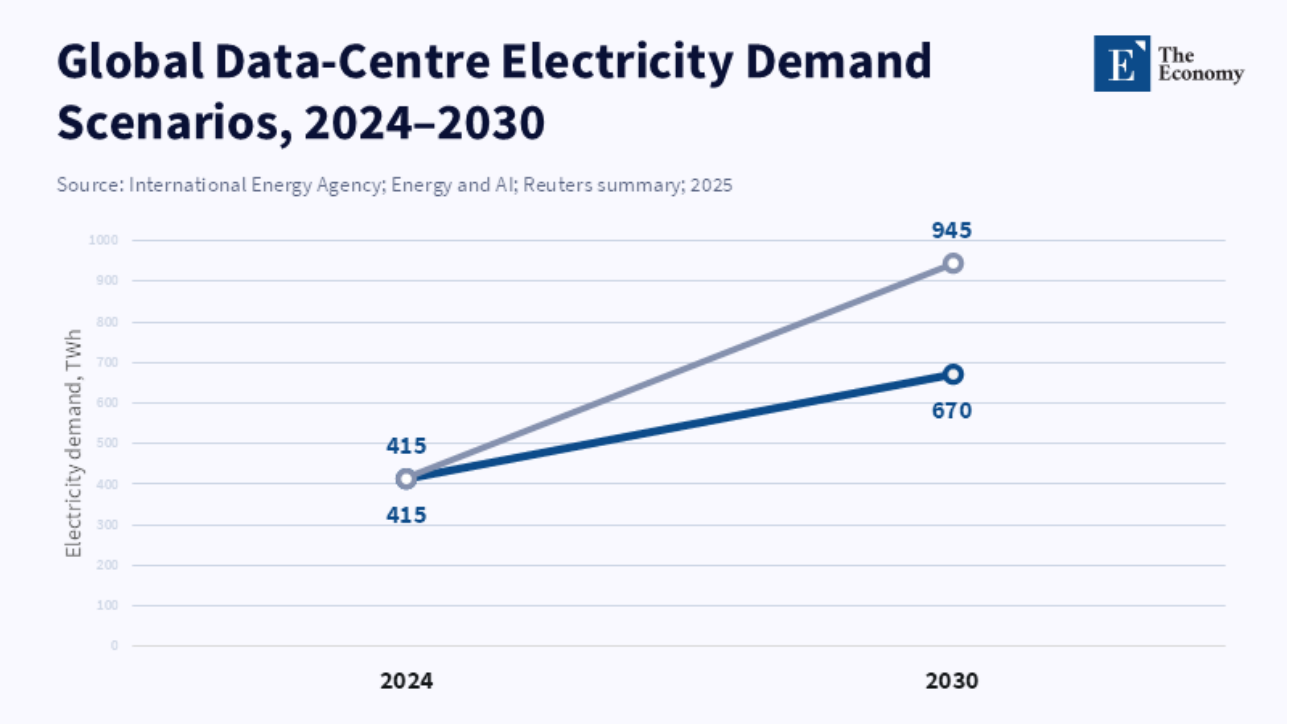

Figure 4: Cloud access is remote for users, but data-centre growth is physical, energy-intensive, and politically located.

If a US corporation were to ask for servers in Indonesia, Thailand, or Malaysia to go above and beyond their normal protocol and attempt to implement political restrictions, this would create enormous diplomatic strains between the nations and lead to either losses in revenue for the data center providers, or cause the nation to simply decide it would not cooperate with a restriction that is simply not economically viable for their market. The US has already heard these arguments from other nations in response to chip export controls, and such complaints would only be amplified by controls on cloud-based computing. Nations in Southeast Asia which have become a center for global AI development may simply choose a different route to achieve compute power by trading with a country that is not subject to the US restrictions. In the end, this would only cost US companies revenue and clout and isolate the Chinese economy. The GAO also commented that the Commerce Department is currently too understaffed to implement hardware controls and applying cloud controls would spread their limited resources too thin and further compromise existing export controls.[23]

The European Union has a different story as while they sometimes cooperate with US policy, they have consistently rejected US extraterritorial sanctions that affect their own economies.[24] This has created a complex and fragmented enforcement mechanism that will be unlikely to allow for a single unified standard to apply across countries.

In essence, while an approach to control access to cloud computing like in the RASA seems politically popular among lawmakers in Congress and conceptually simple in principle, a practical execution would prove nightmarish. US Law simply could not contend with the real-world circumstances of the global marketplace: shell companies, encrypted communications, and an enthusiastic alternative market. Ultimately, restricting access to cloud-based computing would simply prove ineffective and as one commentator succinctly stated, Chinese companies would continue to find ways around the regulations while the US continues to attempt to

enforce them at an ever-increasing radius until the efforts are clearly in vain and the sanctions become mere theater.

3 Rare Earths and Cloud Compute: Physical Chokepoints, Networked Loopholes

The fundamental differences between rare-earth and AI computing power explain precisely why a compute ban can only be a pale shadow of the effects of rare-earth controls.[25] The market factors giving China so much leverage with rare earths virtually complete, easily observable domination of the physical supply chain of a vital commodity simply do not exist for computing power. Most of all, China's power lies not just in its mines, but in the processing of rare-earth and its manufacture of permanent magnets, in which its position in the world is many times more dominant.[26]

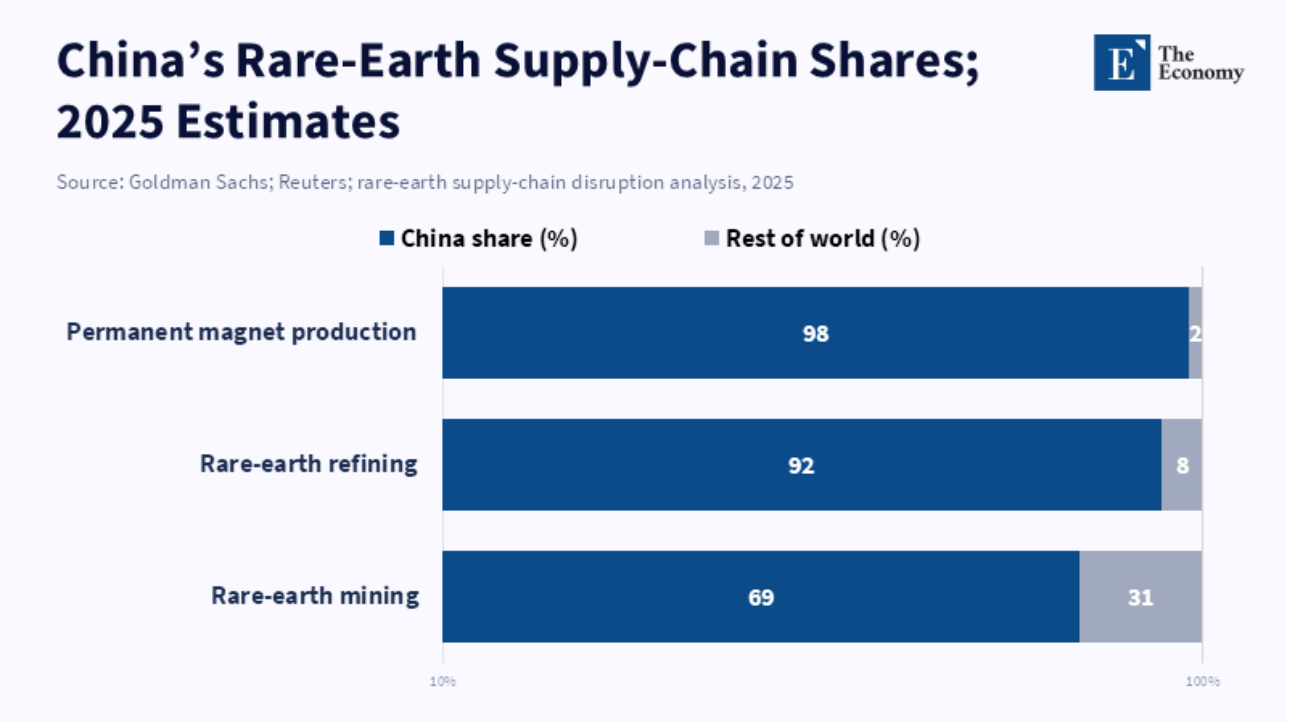

Figure 5: Rare-earth leverage is strongest where supply is physical, concentrated, and observable.

In 2010, when Sino-Japanese relations turned sour, China cut off all rare-earth exports, forcing its high-tech firms to grind to a halt.[27] Again in 2025, when the US tightened export controls to cover Chinese affiliates, Beijing retaliated with rare-earth export restrictions.[28] China controls some 93 percent of global production of magnets, a key component for US defense electronics.[29] Physical commodities with few sources and fixed capacities are immediately responsive to embargoes. A ban on rare earths can cause immediate, crippling scarcity for a rival.

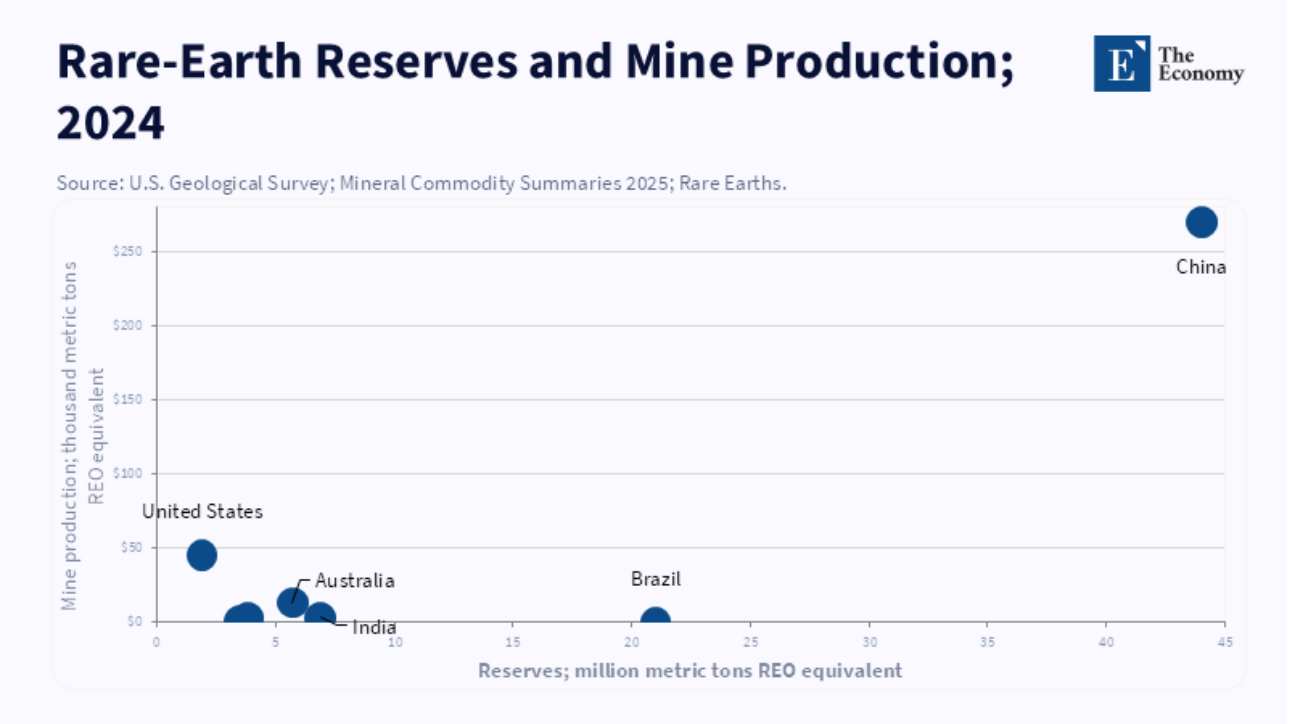

Figure 6: Reserves are geographically broader than production, which is why physical supply-chain control can still create immediate scarcity.

In contrast, no such bottleneck exists with compute: the vast bulk of it is effectively intangible services that flow over the internet and no more physical than email. A data center can spin up and scale any amount of compute capacity it desires. If China cannot buy the silicon it needs, it can rent it from a cloud service; after all, a GPU in Singapore, Tokyo, or the US performs a Chinese algorithm just as well as it would perform one that it directly owns. Compute is networked, interchangeable, and users can redirect their purchasing to meet their capacity needs and budget: no choke point can be targeted with a shutoff switch, because such capacity exists as an interconnected web of global centers where computation can be performed for anyone.

The scale factors already make market distortions more likely. A global cloud GPU market is a natural competitor; specialized providers and the largest of the hyperscale clouds across the globe all compete fiercely for the computing power business by constructing their own AI hardware clusters. Massive investments in data centers worldwide continue.[30] As an example, in 2025 alone, global data center construction spending was over $40 billion, and much of this investment is centered in newly expanding hubs like India, Southeast Asia and the Middle East. In Asia, Singapore now stands among the leaders for AI data center investment, while Vietnam and Indonesia are rapidly following, driven by a mix of governmental inducements, affordable power, and client base competition both locally and internationally.

An additional billion-plus in data center investment in just Southeast Asia is predicted, and the overall market for cloud computing will pass $30 billion by 2030. Naturally, an overwhelming glut of computing capacity has caused price convergence. Experts predict that the soaring costs for Nvidia GPUs due to scarcity will only persist briefly; soon prices will fall due to supply gluts. One report suggested some Nvidia H100 cloud instances could be found in early 2026 at $2.50-$3.50 per hour, a fraction of the 7−12 typically seen at more famous hyperscale clouds. These lower costs make AI hardware accessible to smaller firms, but will also allow

both major Chinese tech players and smaller startups to secure capacity via multi-year contracts with non-US cloud service providers, bypassing the export restrictions.

Ultimately, supply and demand will set a reasonable market price for computing hours for anyone, even if no one customer is able to guarantee artificially inflated costs. Chinese companies can, and do, purchase GPU hours at far lower prices through non-US cloud providers and, if necessary, through domestic or regional cloud services at costs lower than US hyperscale providers. The growing market for alternative cloud providers and specialized brokers means that Chinese companies can and will tap capacity virtually anywhere in the world, undermining price-based sanctions and Artificial Scarcity.

Chinese AI developers can generally obtain computation at the market rate. Unlike rare earth exports where China can punish foreign clients by increasing costs, cloud providers typically do not distinguish between customers on the basis of nationality alone. There is no evidence that Chinese users are being charged a premium on foreign GPU instance rents; users abroad pay a price for their compute time that, while dependent on provider scale and server availability, is largely independent of the server location and in some cases lower than US hyperscale cloud providers. Nevertheless, cloud GPU rentals naturally impose a market rate upon usage and nothing short of US complete ownership of the entire cloud computing sector could change that.

Most significant, perhaps, is the observability question. Governments can see ore moving through customs checkpoints and track shipments through GPS. But the bits flying across cloud networks are effectively invisible to governments and have no passport. A petabyte of training data is not a physical package stamped at customs and numerous AI inference requests carry no visible markers. Under RASA, all leverage points are at the sale of the physical hardware, not the utilization of compute power beyond that sale. Once the hardware is elsewhere, or even once it has simply been activated within a server abroad and used, it is no longer covered by export logs and export controls are effectively useless. A boycott on a port of physical goods like rare earths or semiconductors simply doesn't exist for bits in the global internet cloud.

So, in sum, the controllability of rare earth and advanced computer power is fundamentally different. For rare earths, China possesses a virtually complete, single leverage point of market control; for compute there exists not a single controller, but rather a network of hundreds of service centers globally ready to offer compute services to anyone. An export ban on hardware makes some sense. US policymakers would be misguided to use control of computer access as leverage against the Chinese mineral advantage. While an attack on rare earths has direct impact on China, controls on the bits traveling globally have their primary impact on US companies.

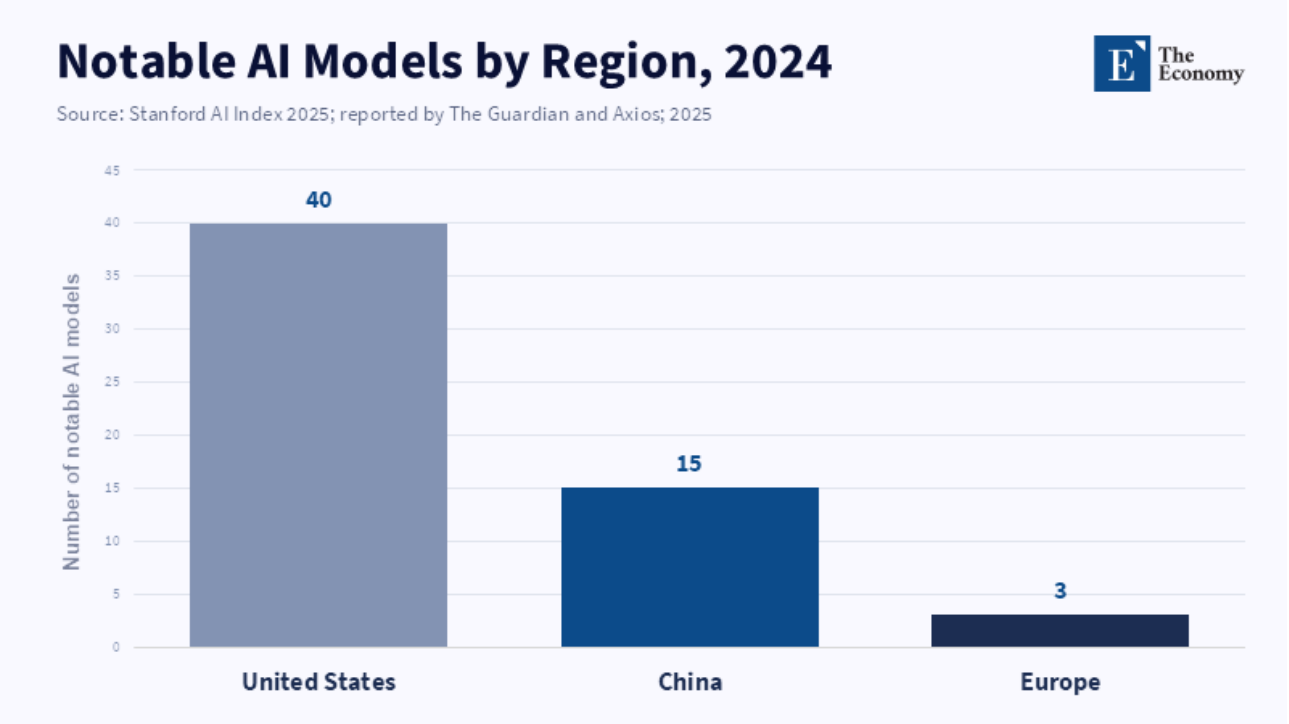

Figure 7: Compute access matters, but frontier model production still depends on deeper institutional and research capacity.

There are distinct strategic calculations for each form of control. While previously the US had exported raw mineral resources in exchange for permissions from China to produce its raw materials, the current relationship means China has the ability to punish the US through a lack of rare-earths minerals in return. Not so with compute: while the US may attempt to create obstacles for China accessing compute, it would have nothing analogous to the mineral market it faces from China. A ban on computing would, at best, result in the US harming its own businesses by denying them a lucrative market, whereas market dynamics of GPU pricing make China much less likely to accede to US demands in the first place when it can buy from the cheapest supplier regardless. Carnegie shows this only incentivizes China to find other means to reduce reliance on US chips, such as domestic chip fabrication.

Price and scale again: a shortage of rare earth minerals could push their costs to unheard-of levels for every consumer of them; for GPUs, however, a similar lack of supply would likely not push up prices for so much as several dollars per hour, with the cost of a high-end Nvidia AI chip expected to fall by 2026 to below that amount if supplied by third-party cloud providers. If the US blockades its markets, the rest of the world can simply buy GPUs at any price available to it and pass them on to Chinese customers. Supply will not be scarce; it will just be cheaper for them, while US firms lose out on a major revenue stream. This fact, coupled with already low hardware costs across the board, makes the cost penalty too small for China to consider in a broader sense of paying premiums.

In short, the nature of advanced computing means that export controls along the lines of those used for rare earths are largely impracticable. The cloud is a global, decentralized network with no singular geographic bottleneck or border crossing through which data must physically travel. Any efforts to stop the transmission of computing capacity would require stopping the internet itself.

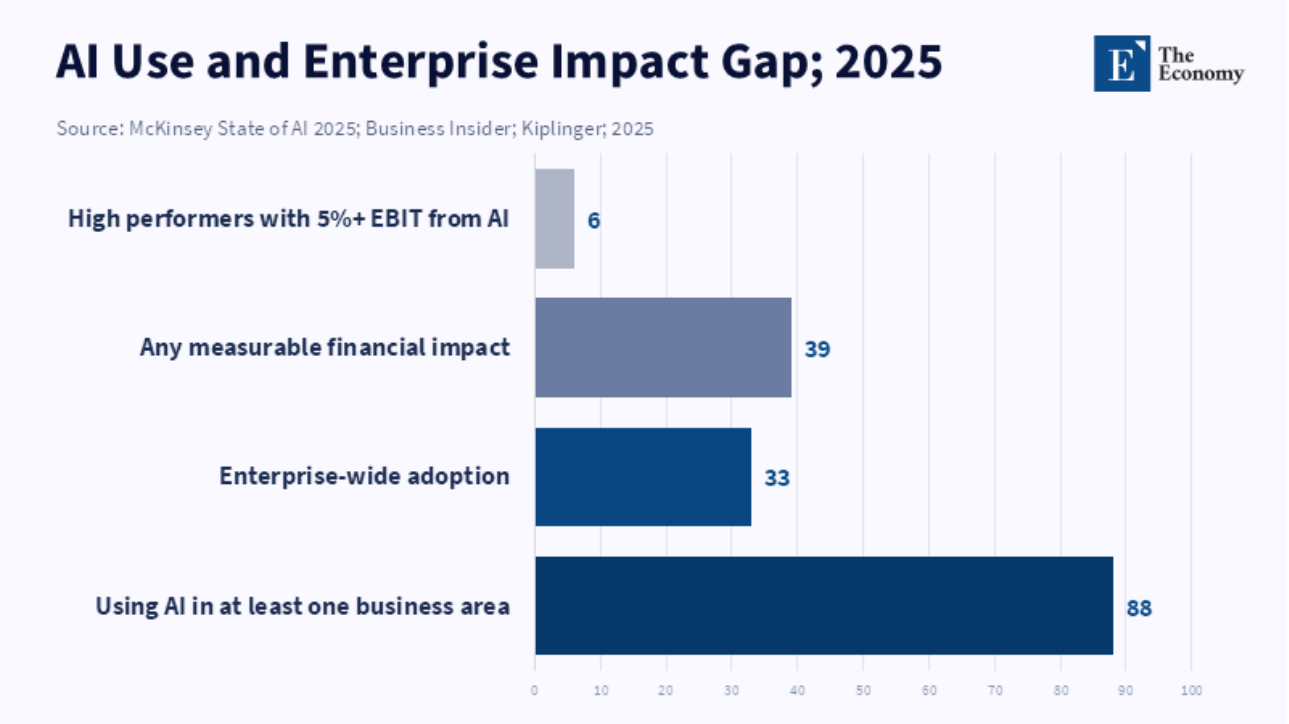

Figure 8: The strategic question is not who can access AI, but who can convert access into measurable organizational capability

4 Conclusion -Talent, Institutions, and Strategic AI Use Beyond Access Control

The US correctly identified cloud computing as a loophole within the constraints that it has placed on its AI chips, although evidence points to it being a theoretical observation rather than a practical maneuver. Blocking China from access to cloud data centers worldwide might supplement the chip embargo, but would only make an infinitesimal difference. Chinese users reataine multiple oathways- legal, informal and illicit- to obtain GPU power, whether via third-country clouds, shell corporations or domestic capacity buildup. While policy has driven up the costs of cutting-edge chips, China continues to build domestic champions and invest in building domestic substitutes, as the explosive growth of such companies as Huawei, Baidu and Sensetime attests. For policymakers seeking to control compute resources, embracing a model of global regulatory policing would be necessary, a feat that no single nation has ever succeeded in doing at any substantial scale.

The empirical data support the conclusion as well. GPU-based AI clusters are deployed worldwide with substantial investments already occurring in Southeast Asia, Europe and the US in particular. The US and China are no longer the sole investors in large-scale AI infrastructure. Southeast Asia is emerging as a key region for data centers, and Europe is advancing public-private AI infrastructure initiatives, such as a 20 billion euro EU drive for AI gigafactories. Though US-based giants are now rapidly globalizing, local and regional providers with a detailed understanding of their clients are as much of a challenge as they are an opportunity, so the ability to restrict access to computing power is naturally limited by the sheer breadth of suppliers and platforms. And global cloud competition means that compute costs are falling everywhere-no ”artificial scarcity premium” is available to Chinese developers and any lingering hardware supply, given chip smuggling and accelerated

domestic production, has probably already moved beyond the reach of US supply. Ultimately, the sanctions' nullification is not due to leaks but the fundamental nature of computing. The US should not mistake the absence of leaks for success in technology competition.

There is still further reason for skepticism. According to experts, hardware flow restrictions often divert attention from the most crucial element: what is being built with the computing power? The US's largest advancements in AI are the result of massive investments in data, algorithms and research personnel-factors which export controls cannot constrain as readily. Although US policy does indeed make top-tier chips more expensive, China continues to cultivate domestic champions and invest in developing alternatives, as seen in the meteoric rise of companies like Huawei, Baidu and Sensetime.

The US's genuine strength lies in its innovation ecosystem: a recent AI Index report showed that the US still far outspends China on AI research and development and produces more major breakthroughs. In 2024 alone, the US released roughly 40 noteworthy large AI models, compared to about 15 from China, with private investment in AI research in the US dwarfing that in China by an order of magnitude.[31] It also holds the world's largest cohort of leading AI companies and researchers.[32] Such a position is not guaranteed to be immutable or easily replicable, but relies on long-term policy support: "America leads the world in advanced semiconductor design, period. We're a couple of years ahead of China," Commerce Secretary Raimondo observed, emphasizing that it's imperative to accelerate, not decelerate, American capacity.

In practice, what this implies is that US Efforts should be refocused to foster human and institutional capacity. It is countries which are fertile ground for deep expertise and robust ecosystems that have historically been at the forefront of technology. Educators and universities should strengthen the advanced human-capital pipeline by investing in STEM and AI education programs that train the next generation of researchers and engineers; recent initiatives by the NSF and university/industry partnerships are positive signs that should be continued and expanded to reflect the scale of the challenge. Industry and academia must cooperate on the creation of open standards and continue to push the envelope on AI capabilities; broad-based alliances and support for open-source projects will sustain US leadership, as will continued investment in domestic AI infrastructure and workforce development programs, grants and the recruitment of international talent. It's not for nothing that the US has been a magnet for international researchers and why immigration policy should be shaped with AI talent in mind. The new national AI strategy aptly calls for prioritizing innovation, infrastructure and talent and for an industry-driven focus on reskilling.[33] These are the dimensions of the competition in which the US still possesses true power.

By contrast, focusing on a cloud computing ban is likely an illusion. Blocking Chinese access to remote GPUs will only result in diplomatic annoyance and diplomatic pronouncements, potentially pushing partners away from the US toward either Chinese tech or their own native alternatives. In a connected globalized economy, "winning" the AI race has never come from more embargos but from advancing the cutting edge via smart education, the nurturing of research ecosystems and the forging of partnerships around common standards – these are ultimately what matters in terms of human capital and technological advancement. In summary, compute access controls are an ineffective strategy to maintain US dominance in AI and attempts to seal off GPU hours over the global cloud will result in costly inefficiencies. The US must play to its strengths and

focus on talent, innovation and building institutions to sustain its edge while China catches up. The empirical evidence is clear that China's progress thus far has resulted from the development of domestic industrial bases and talent pools rather than clandestine chip transfers. Policymakers must heed this lesson: invest in people and ideas and maximize computational utility rather than using a policy of embargoes as a means to limit the US itself.

References

[1] Miller, C. (2022) *Chip War: The Fight for the World's Most Critical Technology*. Scribner.

[2] Yasuhara, Y. (1991) 'The Myth of Free Trade: The Origins of COCOM 1945–1950', *The Japanese Journal of American Studies*.

[3] Mastanduno, M. (1992) *Economic Containment: CoCom and the Politics of East-West Trade*. Cornell University Press.

[4] National Security Commission on Artificial Intelligence (2021) *Final Report*. NSCAI.

[5] U.S. Bureau of Industry and Security (2023) *Commerce Strengthens Restrictions on Advanced Computing Semiconductors, Semiconductor Manufacturing Equipment, and Supercomputing Items to Countries of Concern*. U.S. Department of Commerce.

[6] U.S. Bureau of Industry and Security (2022) *Commerce Implements New Export Controls on Advanced Computing and Semiconductor Manufacturing Items to the People's Republic of China*. U.S. Department of Commerce.

[7] Carnegie Endowment for International Peace / Noah Tan (2026) *The Geopolitical Debates Over Controlling Cloud Compute*. Carnegie Endowment for International Peace.

[8] Heim, L., Fist, T., Egan, J., Huang, S., Zekany, S., Trager, R., Osborne, M.A. and Zilberman, N. (2024) *Governing Through the Cloud: The Intermediary Role of Compute Providers in AI Regulation*. arXiv.

[9] Carnegie Endowment for International Peace / Noah Tan (2026) *The Geopolitical Debates Over Controlling Cloud Compute*. Carnegie Endowment for International Peace.

[10] Carnegie Endowment for International Peace / Noah Tan (2026) *The Geopolitical Debates Over Controlling Cloud Compute*. Carnegie Endowment for International Peace.

[11] Gupta, R., Walker, L. and Reddie, A.W. (2024) *Whack-a-Chip: The Futility of Hardware-Centric Export Controls*. arXiv.

[12] U.S. Bureau of Industry and Security (2022) *Export Administration Regulations and Advanced Computing Controls*. U.S. Department of Commerce.

[13] U.S. House Select Committee on the Chinese Communist Party (2026) *Remote Access Security Act State-*

ment and Committee Materials. U.S. House of Representatives.

[14] U.S. Congress (2026) *Remote Access Security Act*. U.S. House of Representatives.

[15] Pilz, K.F., Sanders, J., Rahman, R. and Heim, L. (2025) *Trends in AI Supercomputers*. arXiv.

[16] U.S. Congress (2026) *Remote Access Security Act*. U.S. House of Representatives.

[17] Gupta, R., Walker, L. and Reddie, A.W. (2024) *Whack-a-Chip: The Futility of Hardware-Centric Export Controls*. arXiv.

[18] Financial Action Task Force (2023) *International Standards on Combating Money Laundering and the Financing of Terrorism and Proliferation*. FATF.

[19] U.S. Senate Permanent Subcommittee on Investigations (2024) *The Failure of Export Controls to Prevent Advanced U.S. Technology from Reaching Adversaries*. U.S. Senate.

[20] Egan, J. and Heim, L. (2023) *Oversight for Frontier AI through a Know-Your-Customer Scheme for Compute Providers*. arXiv.

[21] Center for Strategic and International Studies (2024) *Understanding U.S. Allies' Current Legal Authority to Implement AI and Semiconductor Export Controls*. CSIS.

[22] Center for Strategic and International Studies (2024) *Understanding U.S. Allies' Current Legal Authority to Implement AI and Semiconductor Export Controls*. CSIS.

[23] U.S. Senate Permanent Subcommittee on Investigations (2024) *The Failure of Export Controls to Prevent Advanced U.S. Technology from Reaching Adversaries*. U.S. Senate.

[24] European Commission (2021) *Communication on the European Economic and Financial System: Fostering Openness, Strength and Resilience*. European Commission.

[25] International Energy Agency (2025) *Global Critical Minerals Outlook 2025*. IEA.

[26] U.S. Geological Survey (2025) *Mineral Commodity Summaries 2025: Rare Earths*. U.S. Department of the Interior.

[27] Bradsher, K. (2010) 'Amid Tension, China Blocks Vital Exports to Japan', *The New York Times*.

[28] Reuters (2025) 'China Tightens Rare Earth Export Controls, Targets Defence and Semiconductor Users', *Reuters*.

[29] International Energy Agency (2025) *Global Critical Minerals Outlook 2025*. IEA.

[30] International Energy Agency (2025) *Energy and AI*. IEA.

[31] Stanford Institute for Human-Centered Artificial Intelligence (2025) *Artificial Intelligence Index Report 2025*. Stanford University.

[32] MacroPolo (2024) *The Global AI Talent Tracker 2.0*. Paulson Institute / MacroPolo.

[33] White House (2025) *Winning the Race: America's AI Action Plan*. The White House.