ANTHROP\C

548 Market St., PMB 90375
San Francisco, CA 94104-5401

412-837-9797
anthropic.com

April 29, 2025

Transmitted via FederalRegister.gov

Re: Comment on the Framework for Artificial Intelligence Diffusion (RIN 0694-AJ90)

Anthropic strongly supports the Department of Commerce's "Framework for Artificial Intelligence Diffusion" interim final rule ("the Diffusion Framework"). The Diffusion Framework is essential to winning strategic competition with China and curbing the illicit smuggling of advanced semiconductors. We recommend that the Administration preserve the Diffusion Framework and take further action to strengthen it. Any changes to the Diffusion Framework should enhance its effectiveness at limiting China's access to advanced computing, and at ensuring that the coming AI infrastructure buildout takes place in America and by American firms.

We anticipate that extremely powerful AI systems will be built in the next few years, and that these systems will have transformative capabilities that can drive economic growth and strengthen national security. It's therefore crucial that the Administration ensure that America maintains control over the global distribution of AI infrastructure, including advanced semiconductor chips, to ensure these technologies cannot be weaponized by America's rivals.

The first Trump Administration recognized this dynamic and took aggressive action to restrict the export of advanced semiconductors and advanced semiconductor manufacturing equipment to China. These policies helped to establish America's current leadership at the AI frontier and to spur hundreds of billions of dollars of investment in domestic AI infrastructure.

We should continue to build upon the highly successful export control foundations laid by the first Trump Administration. The strategic window for export controls is now—not later. The first nation to develop powerful AI, which we estimate may emerge by 2027, will gain a decisive strategic advantage. Given the substantial performance gap between American-designed chips and inferior Chinese alternatives, export controls will yield their maximum strategic leverage only if implemented and maintained immediately. As of today, the latest offerings from Chinese semiconductor firms, while nominally competitive with American chips on a compute basis, face

substantial production and engineering challenges and require substantially more power to operate.[1]

While the Administration took an important step towards curbing China's ability to stockpile advanced semiconductors with its recent H20 ban, delaying implementation of the Diffusion Framework would provide China with a window to smuggle chips via third countries. Further, any changes to the Diffusion Framework should enhance its effectiveness at limiting China's access to advanced semiconductors. In this submission we suggest several actions to directly limit China's access to advanced semiconductors and to crack down on smuggling and circumvention of American export controls.

Maintaining America's AI dominance is contingent on America's continued compute advantage. The Diffusion Framework gives America the opportunity to solidify and grow its global lead in the development and deployment of powerful AI, by preventing China from accessing powerful AI infrastructure and enabling the United States to set the standards for the future of AI technology.

About Anthropic

Anthropic is a leading frontier AI model developer working to build reliable, interpretable, and steerable artificial intelligence systems. Our flagship AI assistant, Claude, represents the state of the art in Large Language Model (LLM) technology. As one of the ten most valuable private companies in the U.S., we conduct cutting-edge artificial intelligence research and deploy our research discoveries as products that serve millions of Americans. Our customers, ranging from Fortune 500 companies and U.S. government agencies to small businesses and consumers, use Claude as an AI co-pilot to enhance productivity on sophisticated tasks including software development, data analysis, and scientific research. In February 2025, we released Claude 3.7 Sonnet, which is by many performance benchmarks the most powerful and capable commercially-available AI system in the world.

Benefits and Risks of Powerful AI

Based on current research trajectories, we anticipate that powerful AI technology will be built during this administration, emerging as soon as late 2026 or 2027. In the five to ten years following these developments, American society will have the opportunity to harness major beneficial transformations resulting from the technology.[2] Powerful AI[3] systems will have the following properties:

---

[1] Patel, Dylan. "Huawei AI CloudMatrix 384 – China's Answer to Nvidia GB200 NVL72." SemiAnalysis, 16 Apr. 2025, available at: semianalysis.com/2025/04/16/huawei-ai-cloudmatrix-384-chinas-answer-to-nvidia-gb200-nvl72/.
[2] Amodei, Dario. "Machines of Loving Grace." Darioamodei.com, Oct. 2024, available at: darioamodei.com/machines-of-loving-grace.
[3] When we discuss "powerful AI," we are referring to systems that represent major advancement beyond today's AI model capabilities.

- Intellectual capabilities matching or exceeding that of Nobel Prize winners across most disciplines—including biology, computer science, mathematics, and engineering.
- The ability to navigate all interfaces available to a human doing digital work today, including the ability to process and generate text, audio, and video, the ability to autonomously control technology instruments like mice and keyboards, and the ability to access and browse the internet.
- The ability to autonomously reason through complex tasks over extended periods—hours, days, or even weeks—seeking clarification and feedback when needed,much like a highly capable employee would.
- The ability to interface with the physical world; controlling laboratory equipment, robotic systems, and manufacturing tools through digital connections.

A useful conceptual framework is to envision powerful AI as equivalent to "a country of geniuses in a datacenter"—a concentration of intellectual capability that fundamentally transforms our understanding of what is possible. In his essay, *Machines of Loving Grace*, Anthropic CEO Dario Amodei noted the practical implications of these transformations.[4] Powerful AI is likely to yield major breakthroughs in healthcare, economic development, and government efficiency. We are already seeing aspects of these transformational benefits emerge with current capabilities. In just one commercial example, Anthropic's AI assistant Claude is being used in drug discovery to help get treatments to patients faster.

With these benefits also come risks. As dual-use technologies, frontier AI models will likely have significant national security implications. On one hand, this technology can provide the United States with strategic advantages when properly leveraged. On the other hand, the technology has the potential to enable adversaries to pursue harmful objectives. We are already seeing models approaching, and in some cases exceeding, undergraduate-level skills in cybersecurity and expert-level knowledge in some areas of biology.[5] These kinds of capabilities are on a developmental trajectory to change the face and nature of adversarial attacks. For example, in our evaluations of AI models' potential to make it easier for threat actors to acquire biological weapons, Anthropic's recent model, Claude 3.7 Sonnet, provided better advice in key steps of the weaponization pathway and made fewer mistakes in critical tasks compared to previous models.[6] As capabilities in these domains increase over the next 2-4 years,[7] the United States must take action to prevent this technology from ending up in the wrong hands.

---

[4] Amodei. "Machines of Loving Grace." (supra note 2)
[5] "Progress from Our Frontier Red Team." Anthropic, 19 Mar. 2025, available at: www.anthropic.com/news/strategic-warning-for-ai-risk-progress-and-insights-from-our-frontier-red-team.
[6] When observing end-to-end task success rates holistically, we find that Claude 3.7 Sonnet still makes several critical errors. As a result, the total amount of uplift Claude 3.7 Sonnet can provide in a given task is still limited. For more details see.Claude 3.7 Sonnet System Card (Feb. 24, 2025), available at: https://assets.anthropic.com/m/785e231869ea8b3b/original/claude-3-7-sonnet-system-card.pdf.
[7] Bengio, Y., et al. "International AI Safety Report." UK Government Publishing Service, Jan. 2025, available at: assets.publishing.service.gov.uk/media/679a0c48a77d250007d313ee/International_AI_Safety_Report_2025_accessible_f.pdf.

Denying China Access to High-End Chips is a Key Component to Securing America's AI Dominance

The first Trump Administration correctly diagnosed that AI will be central to strategic competition with China, and that the United States can and should use export controls to maintain and strengthen its AI leadership.[8] While the US still maintains a lead in AI development, recent advances of Chinese AI labs like DeepSeek have made inroads and underscore the importance of strong export controls on advanced chips.

As DeepSeek's founder, Liang Wenfeng, openly acknowledged: “the embargo on high-end chips" remains their primary constraint and that Chinese companies "have to consume twice the computing power to achieve the same results” as their American counterparts. Liang continued, “when combined with data efficiency gaps, this could mean Chinese labs need up to four times more computing power to achieve capabilities equivalent to American labs.”[9] As acknowledged by DeepSeek, compute is the major bottleneck for the substantial human, technological, and physical capital that China is bringing to bear on AI.

At this crucial moment in AI development, it would be a national security and economic mistake for the United States to weaken or eliminate the Diffusion Framework, as this could permanently cede AI leadership and innovation to China. Conversely, strengthening the Diffusion Framework and other export controls would extend and cement the United States’ AI advantage by preventing China access to the infrastructure it needs to catch-up to the US, while also enabling the US to build its existing lead. The Diffusion Framework is poised to further deny China access to a vital AI input while unleashing billions of dollars of investment in the United States’ domestic AI infrastructure.

Powerful chips are needed to create powerful AI

Our belief in AI scaling laws drives our perspective on the necessity of export controls to maintain the United States’ leadership on AI development. Artificial Intelligence system capabilities are the product of scaling laws: a given system’s performance is the function of its size, the amount of data it trains on, and the amount of compute it uses to train.[10] These scaling laws have characterized AI development for the past decade and continue to define it. The persistence of scaling laws has important policy implications for the underlying AI infrastructure

---

[8] Amodei, Dario. "On DeepSeek and Export Controls." Darioamodei.com, 2025, available at: www.darioamodei.com/post/on-deepseek-and-export-controls.

[9] Cao, H., and D. Papailiopoulos. "How Chinese Company DeepSeek Released a Top AI Reasoning Model despite US Sanctions." MIT Technology Review, 24 Jan. 2025, available at: www.technologyreview.com/2025/01/24/1110526/china-deepseek-top-ai-despite-sanctions/.

[10] Kaplan, J., et al. "Scaling Laws for Neural Language Models." arXiv, 23 Jan. 2020, available at: arxiv.org/pdf/2001.08361.

involved in model training and inference.[11][12] Compute, therefore, is an essential ingredient to leadership in AI.

The relationship between computational power and AI capabilities is firmly established. Since 2010, the amount of compute used for training notable AI systems has doubled roughly every six months, growing by a factor of 350 million over 13 years.[13] This massive growth has been made possible by steady advances in the chips used to power machine learning and AI training. Since 2006, the power of these chips has on average doubled every two years.[14]

The strategic benefits of the Diffusion Framework will compound with time. Under the Diffusion Framework, China and other countries subject to export controls will be frozen at the rule's current technology threshold. Concurrently, the United States and its partners will benefit from the continued improvements in chip capability and efficiency. As a result, by 2027, countries using older, export-compliant chips could face training costs that are ten times greater than those associated with current chips,[15] helping to ensure that the U.S. can extend its lead on AI development.

China is projected to invest $76 billion in its domestic semiconductor supply chain–more than any other nation–over the course of 2025 and 2026.[16] Yet, China has struggled to produce the types of advanced chips needed to train and run frontier AI models. In 2019, Huawei launched its Ascend series of chips as a direct competitor to NVIDIA's advanced semiconductor offerings. While Huawei has nominally been able to match–or even exceed–the capabilities of NVIDIA's integrated system offerings, it has been unable to do so on an individual chip basis. Huawei's newest Ascend series chip, the Ascend 910C, achieves only 50-percent of the per-watt compute performance of NVIDIA's GB200 released last year.[17] Further, the GB200[18] represents a 167% improvement over the per-watt compute performance of NVIDIA's circa-2020 A100 chip.[19] Comparatively, the Ascend 910C, slated to be released this coming May[20], represents an

---

[11] "Measuring AI Ability to Complete Long Tasks." METR, 19 Mar. 2025, available at: https://metr.org/blog/2025-03-19-measuring-ai-ability-to-complete-long-tasks/
[12] "Machine Learning Trends." EpochAI, 13 Jan. 2025, available at: epoch.ai/trends.
[13] Sastry, G., et al. "Computing Power and the Governance of Artificial Intelligence." arXiv, 2024, available at: arxiv.org/abs/2402.08797.
[14] Hobbhahn, M., and T. Besiroglu. "Trends in GPU Price-Performance." Epoch AI, June 2022, available at: epochai.org/blog/trends-in-gpu-price-performance.
[15] "AI Chips and Export Controls: Strategic Competition with China." Center for a New American Security, 2024, available at: s3.us-east-1.amazonaws.com/files.cnas.org/documents/CNAS-Report_AI-Trends_FinalC.pdf.
[16] "Global Fab Equipment Investment Expected to Reach $110 Billion in 2025." SEMI, 10 Dec. 2024, available at: www.semi.org/en/semi-press-release/global-fab-equipment-investment-expected-to-reach-110-billion-dollar-in-2025.
[17] Patel. "Huawei AI CloudMatrix 384 – China's Answer to Nvidia GB200 NVL72." (supra note 1).
[18] "NVIDIA Blackwell: The engine of the new industrial revolution." NVIDIA, 2025, available at: https://nvdam.widen.net/s/wwnsxrhm2w/blackwell-datasheet-3384703
[19] "A100 GPU's Offer Power, Performance, & Efficient Scalability." NVIDIA, 2022, available at: https://www.nvidia.com/content/dam/en-zz/Solutions/Data-Center/a100/pdf/nvidia-a100-datasheet-nvidia-us-2188504-web.pdf
[20] Lin, Liza and Huang, Raffael. "China's Huawei Develops New AI Chip, Seeking to Match Nvidia." The Wall Street Journal, 27 April 2025, available at: https://www.wsj.com/tech/chinas-huawei-develops-new-ai-chip-seeking-to-match-nvidia-8166f606

increase of 14%[21] in per-watt compute performance from the circa-2019 Ascend 910, Huawei's A100-equivalent[22].

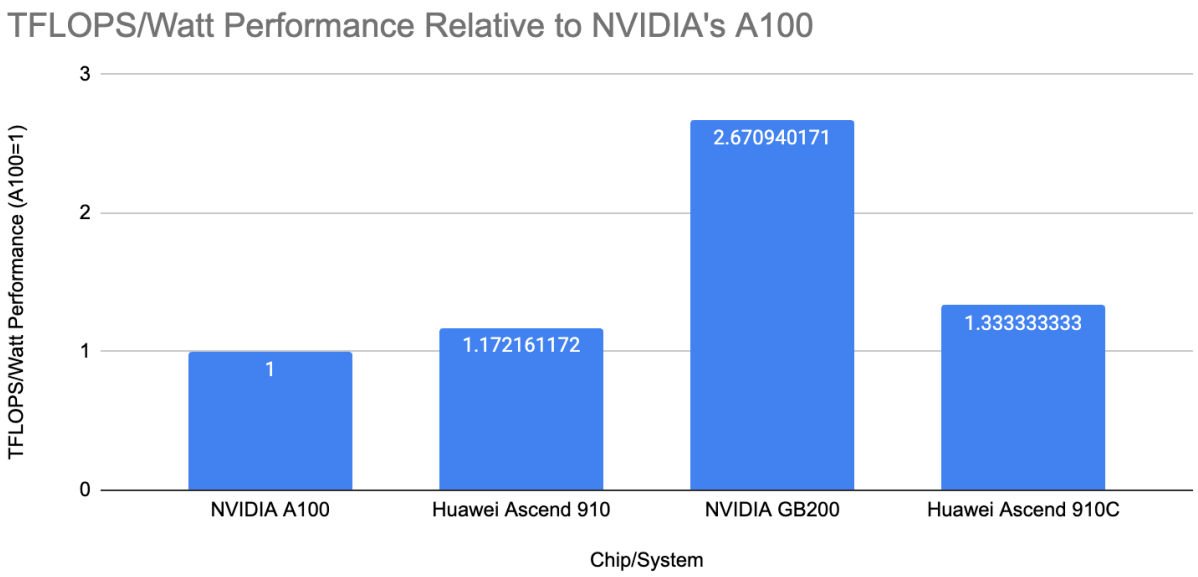

GPU Model	Year	FP16 TFLOPS	TDP (W)	TFLOP/Watt	Relative to A100
NVIDIA A100[23]	2020	312	400	0.78	1
Huawei Ascend 910[24]	2019	320	350	0.91	1.17
NVIDIA GB200[25]	2024	2500	1200	2.08	2.67
Huawei Ascend 910C[26]	2024	780	750	1.04	1.33

Subsequently, leading Chinese AI labs have been forced to stockpile and smuggle advanced American chips in order to remain close to the research and capability frontier.[27]

DeepSeek shows the value in denying China access to powerful chips

DeepSeek's R-1 model exemplifies China's reliance on high-end US chips. DeepSeek was among the earliest Chinese AI labs to recognize the importance of compute and acquired Asia's

[21] Patel. "Huawei AI CloudMatrix 384 – China's Answer to Nvidia GB200 NVL72." (supra note 1).
[22] Feldgoise, Jacob, and Hanna Dohmen. "Pushing the Limits: Huawei's AI Chip Tests U.S. Export Controls." Center for Security and Emerging Technology, 28 June 2024, available at: https://cset.georgetown.edu/publication/pushing-the-limits-huaweis-ai-chip-tests-u-s-export-controls/
[23] "A100." (supra note 19)
[24] Feldgroise and Dohmen. "Pushing the Limits." (supra note 22)
[25] "NVIDIA Blackwell." (supra note 18)
[26] Patel. "Huawei AI CloudMatrix 384 – China's Answer to Nvidia GB200 NVL72." (supra note 1).
[27] Patel, D. "DeepSeek Debates: Understanding China's AI Acceleration." SemiAnalysis, 31 Jan. 2025, available at: semianalysis.com/2025/01/31/deepseek-debates/.

first cluster of 10,000 NVIDIA A100 chips before they were subject to export controls.[28] DeepSeek has continued to focus on amassing and operating sizable fleets of cutting-edge compute.

Based on an analysis conducted by Epoch.AI ("Epoch"), existing export controls have seemingly slowed the progress of Chinese labs. Per Epoch's analysis, between 2019 and 2021, the amount of compute used to train Chinese models rapidly increased and was on trajectory to meet or even surpass the amount of compute used by American labs. Since late 2021, the rate of growth has dropped dramatically; American labs have continued to scale their training compute by an annual factor of five whereas Chinese labs have only been able to scale by an annual factor of three.[29][30]

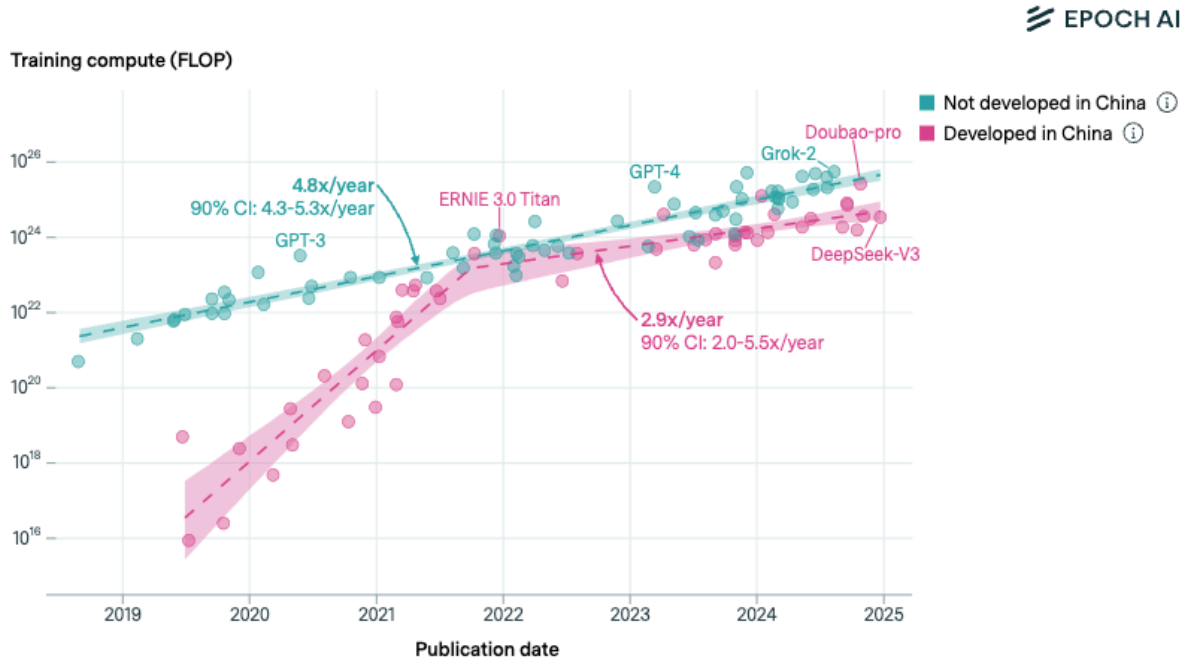

[31]

The compute gap is poised to grow despite China's billions of dollars of investments in its domestic semiconductor supply chain. Chinese semiconductor fabs continue to struggle to match the yields of TSMC; between 2020 and 2025, defect-free yields of Huawei's leading Ascend 910B chip remained flat at 20%.[32] While Chinese fabs will likely improve their yields, export restrictions on extreme ultraviolet lithography technology will significantly delay their

[28] Lin, A., and L. Heim. "DeepSeek's Lesson: America Needs Smarter Export Controls." RAND Corporation, 5 Feb. 2025, available at: www.rand.org/pubs/commentary/2025/02/deepseeks-lesson-america-needs-smarter-export-controls.html.
[29] "China Compute Trends." Epoch.ai, 2025, available at: epoch.ai/data-insights/china-compute-trends.
[30] Epoch expresses varying levels of confidence in its analysis of individual models, e.g., it assesses its analysis of the Doubao-Pro model as "speculative."
[31] "China Compute Trends." (supra note 29)
[32] Arcesati, Rebecca, and Gregory C. Allen. "DeepSeek, Huawei, Export Controls, and the Future of the US-China AI Race." Center for Strategic and International Studies, 2025, available at: www.csis.org/analysis/deepseek-huawei-export-controls-and-future-us-china-ai-race.

ability to move closer to the semiconductor frontier. As a result, even new offerings from Chinese semiconductor firms that are nominally competitive with American chips on a compute basis face substantial production and engineering challenges and require substantially more power to operate.[33]

Lastly, independent of these hardware challenges, the software needed to access and utilize Chinese AI chips lags far behind American equivalents. DeepSeek reportedly assessed that "it would be years" before Chinese hardware and software combinations would offer a viable alternative to their American counterparts.[34] While America's software advantage is less durable and harder to protect, it does serve to compound the effects of its compute advantage: matching the maturity and capability of American software could take years of dedicated effort and investment by Chinese firms, siphoning resources away from other software engineering challenges and opportunities.

The Danger of Offshoring AI Infrastructure

The US share of global semiconductor production has fallen from 40% in 1990 to just 12% today, with 90% of the world's leading-edge semiconductors now made outside the United States.[35] This decline represents a profound strategic vulnerability for a technology that underpins America's economy and critical national security systems.

Other critical industries and supply chains have followed similar courses. Solar photovoltaic (PV) panels and lithium ion batteries were originally developed and commercialized in the United States. Yet, in 2025, Chinese firms control in excess of 90% of the global solar PV industry[36] and 80% of global lithium-ion battery production capacity.[37]

Without strong export controls, AI infrastructure–which includes the physical hardware and facilities needed to support AI development–will likely follow a similar trajectory as chips, PV panels, and batteries. There are stark limits to key inputs for the buildout of AI infrastructure, such as electric power grid components and AI chips such that it will not be possible to build both in the U.S. and overseas in the near term. The Diffusion Framework's domestic compute requirements and foreign compute caps are designed to ensure the AI infrastructure buildout takes place in America and by American firms.

---

[33] Patel, Dylan. "Huawei AI CloudMatrix 384" (supra note 1)
[34] Arcesati and Allen. "DeepSeek, Huawei" (supra note 32).
[35] "Semiconductors and National Defense: What Are the Stakes?" Center for Strategic and International Studies, 2024, available at: www.csis.org/analysis/semiconductors-and-national-defense-what-are-stakes.
[36] "Tsunami: Chinese Solar Company Insolvencies." PV Tech, 15 Feb. 2024, available at: www.pv-tech.org/tsunami-chinese-solar-company-insolvencies-pv-tech-bankability-report/.
[37] "Batteries and Secure Energy Transitions." International Energy Agency, 2024, available at: www.iea.org/reports/batteries-and-secure-energy-transitions.

AI infrastructure deployment is driven by economics

Domestically, AI infrastructure demands are projected to generate $500 billion in investment and the development of an additional 50 gigawatts of energy capacity by 2030.[38] Without appropriate controls, economic factors will drive this frontier AI infrastructure offshore. Some nations have AI industrial strategies that will offer hundreds of billions, if not more, of investment and energy resources to U.S. companies that build frontier AI infrastructure on their territories.[39]

The scaling laws and realities of the semiconductor industry suggest that within 3-5 years, America could find itself as a net consumer of compute—reliant on data centers in foreign countries to meet domestic inference and potentially even training needs—unless adequate policies are in place to ensure AI infrastructure is developed domestically. Currently, the compute available to American labs is constrained by supply.[40] Industry leaders have made public statements about "begging"[41] for chips or limiting the availability of certain AI tools due to chips "melting"[42] as a result of excess use. While additional semiconductor manufacturing capacity is coming online in the United States, demand for advanced chips will likely continue to outstrip supply in the coming years.[43]

These supply constraints and the framework's global export regime create clear incentives for other nations with AI ambitions: the only means of accessing large-scale cutting-edge compute is by investing in the United States' AI infrastructure and complying with its export controls.

The U.S. Must Take Decisive Action to Address Chip Smuggling

While existing country-specific export controls are foundational to maintaining America's compute advantage, they are increasingly coming under strain. Smugglers have employed creative methods to circumvent export controls, including hiding processors in prosthetic baby bumps[44] and packing GPUs alongside live lobsters.[45] Chinese firms continue to establish shell

---

[38] "How Data Centers and the Energy Sector Can Sate AI's Hunger for Power." McKinsey & Company, 17 Sept. 2024, available at: www.mckinsey.com/industries/private-capital/our-insights/how-data-centers-and-the-energy-sector-can-sate-ais-hunger-for-power.

[39] Pacheco, Filipe. "Race for AI Supremacy in Middle East Is Measured in Data Centers." Bloomberg, 11 Apr. 2024, available at: www.bloomberg.com/news/articles/2024-04-11/race-for-ai-supremacy-in-middle-east-is-measured-in-data-centers.

[40] Arcesati and Allen. "DeepSeek, Huawei" (supra note 32)

[41] Durkee, Alison. "Larry Ellison: Elon Musk Begged Nvidia's Jensen Huang for More GPUs at Fancy Sushi Dinner." Fortune, 16 Sept. 2024, available at: fortune.com/2024/09/16/larry-ellison-elon-musk-begged-nvidias-jensen-huang-more-gpus-fancy-sushi-dinner/.

[42] Altman, Sam [@sama]. "We need more GPUs. A lot more." Twitter, 27 Mar. 2025, available at: x.com/sama/status/1905296867145154688.

[43] "Prepare for the Coming AI Chip Shortage (Tech Report 2024)." Bain & Company, 2024, available at: www.bain.com/insights/prepare-for-the-coming-ai-chip-shortage-tech-report-2024/.

[44] Zhong, R. "Chinese Woman Fakes Pregnancy, Tries to Evade Customs with over 200 Intel Alder Lake CPUs." WCCF Tech, 12 Dec. 2022, available at: wccftech.com/chinese-woman-fakes-pregnancy-tries-to-evade-customs-over-200-intel-alder-lake-cpus/.

[45] Humphries, M. "Crustacean Cargo: Hong Kong Drivers Smuggle Nvidia GPUs with Live Lobsters." PC Magazine, 8 Nov. 2023, available at: www.pcmag.com/news/crustacean-cargo-hong-kong-drivers-smuggle-nvidia-gpus-with-live-lobsters.

companies in third countries at a rapid pace to evade export controls, often exploiting the capacity constraints of and lack of harmonization between American regulators and law enforcement agencies.[46]

One major smuggling operation involved sending 53,000 banned American chips to China worth $12 million through a South Korean company between August 2020 and August 2023.[47] In another recent case, Singaporean prosecutors alleged that firms based in the country orchestrated a $390 million scheme to transfer servers with NVIDIA chips to DeepSeek.[48] These are all pattern and practice of a larger smuggling ecosystem in which multiple illicit operations regularly engage in transactions valued at more than $100 million,[49] and which demonstrates the critical importance of U.S. computing power to Chinese AI infrastructure and capabilities, despite claims to the contrary.

A global licensing requirement is crucial to ensuring that controlled items are not diverted to destinations or end users of concern, as evidenced by China's frequent use of shell companies in third countries to acquire and subsequently smuggle controlled semiconductors. The Diffusion Framework curbs the ability of these shell companies to exploit current gaps in the United States' export controls.

The Diffusion Rule Should Not Be Paused for Changes

Should the Administration decide to make changes to the Diffusion Framework, it is critical that the rule remain in place during this process and compliance deadlines not be altered. This is especially critical as Chinese firms have engaged in aggressive stockpiling ahead of the Diffusion Framework's May 15, 2025 implementation date, and further delay would only invite more stockpiling and ultimately weaken the effectiveness of the Diffusion Framework at a critical moment – when powerful AI capabilities could be only 18-36 months away.

China has an established history of evading export controls

Even with the Administration's recent ban of the H20 chip, China is likely to continue attempting to circumvent export controls via third countries.

---

[46] "The Urge to Merge: Streamlining US Sanction Lists Targeting China." Rhodium Group, Oct. 2024, available at: rhg.com/research/the-urge-to-merge-streamlining-us-sanction-lists-targeting-china/.
[47] Crider, M. "Chip Smuggling Operation That Sent 53,000 Banned American Chips to China Gets Busted." Tom's Hardware, 18 Feb. 2024, available at: www.tomshardware.com/tech-industry/semiconductors/chip-smuggling-operation-that-sent-53000-banned-american-chips-to-china-gets-busted-dollar12-million-worth-of-chips-funneled-through-south-korean-company.
[48] "Singapore prosecutors says US servers fraud case involves $390 million." Reuters, 13 Mar. 2025, available at https://www.reuters.com/world/asia-pacific/singapore-prosecutors-says-us-servers-fraud-case-involves-390-million-2025-03-13/.
[49] Yang, W. "Nvidia AI Chip Smuggling to China Becomes an Industry." The Information, 11 Jan. 2025, available at: www.theinformation.com/articles/nvidia-ai-chip-smuggling-to-china-becomes-an-industry.

China imported semiconductor manufacturing equipment valued at $27.4 billion in 2023, a 46.48% increase compared to 2022.[50] Between January and August 2023, China imported $3.2 billion worth of semiconductor manufacturing machines from the Netherlands, a 96.1% increase over the $1.7 billion recorded over the same period in 2022.[51] Chinese firms engaged in similar stockpiling with the NVIDIA H20 chip prior to the imposition of licensing requirements by the Trump Administration. Prior to the chip's ban, Chinese firms had placed $16 billion in orders for H20 chips (1.3 million units)[52] on top of the existing $12 billion (1 million units) shipped to China in 2024.[53]

America's continued compute advantage depends on strong, adaptive export controls like the Diffusion Framework and their vigorous enforcement. China has invested billions of dollars towards developing a domestic semiconductor industry and fostering the world-class technical talent needed to compete with leading American labs. That America's compute advantage has endured in the face of such a massive, multi-faceted effort is testament to the need for robust, adaptive export controls like the Diffusion Framework and their vigorous enforcement. Any pause or extension of the Diffusion Framework would almost certainly result in further stockpiling and undermine any successor rule or modification prior to its taking effect.

Hardening Export Controls to Widen the U.S.'s lead

We applaud the Administration's recent efforts to control the H20 chips as an important measure to slow China's AI progress. In addition, to prevent advanced AI models and AI infrastructure from being acquired by adversaries, we strongly recommend the administration strengthen export controls on computational resources and implement appropriate export restrictions on certain model weights. Amongst other things, we recommend the Administration:

- Consider adjustments to the tiering system to allow countries that have robust and established data center security and technology control regimes the ability to move tiers and obtain a greater number of chips. One avenue the Administration may consider is requiring countries in Tier 2 to sign government-to-government agreements outlining measures to prevent smuggling as a mechanism for obtaining more chips. As a prerequisite for hosting data centers with more than 50,000 chips from U.S. companies, the U.S. should mandate that countries at high-risk for chip smuggling comply with a

---

[50] "China Imports Massive Lithography Equipment Ahead of Restrictions." Digitimes, 24 Jan. 2024, available at: www.digitimes.com/news/a20240124VL203/china-equipment-lithography-netherlands.html.
[51] "The US and Allies Tightened the Embargo: China Imported Massive Amount of Netherlands Chip Manufacturing Equipment." FPT Semiconductor, 2023, available at: fpt-semiconductor.com/blogs/the-us-and-allies-tightened-the-embargo-china-imported-massive-amount-of-netherlands-chip-manufacturing-equipment/.
[52] Liu, Qianer. "Nvidia Faces Dilemma After Chinese Firms Rush to Order $16 Billion in New AI Chips." The Information, 2 Apr. 2025, available at: www.theinformation.com/articles/nvidia-faces-dilemma-chinese-firms-rush-order-16-billion-new-ai-chips.
[53] Potkin, Fanny, and Che Pan. "Nvidia's H20 Chip Orders Jump as Chinese Firms Adopt DeepSeek's AI Models." Reuters, 25 Feb. 2025, available at: www.reuters.com/technology/artificial-intelligence/nvidias-h20-chip-orders-jump-chinese-firms-adopt-deepseeks-ai-models-sources-say-2025-02-25/.

government-to-government agreement that 1) requires them to align their export control systems with the U.S., 2) takes security measures to address chip smuggling or remote access compute to China, and 3) stops their companies from working with the Chinese military. These requirements also could be aligned with the Trump administration's America First Investment Policy. The Department of Commerce's January 2025 Interim Final Rule on the Framework for Artificial Intelligence Diffusion (the "Diffusion Rule") already contains the possibility for such agreements, laying a foundation for further policy development.

- Closely examine and reduce the 1,700 H100 no-license required threshold for orders to Tier 2 countries in the Diffusion Rule. Currently, the Diffusion Rule allows advanced chip orders from Tier 2 countries for less than 1,700 H100s —an approximately $40 million order—to proceed without review. These orders do not count against the Rule's caps, regardless of the purchaser. While these thresholds address legitimate commercial purposes, we believe that they also pose smuggling risks. We recommend that the Administration consider reducing the number of H100s that Tier 2 countries can purchase without review to further mitigate smuggling risks. Determining the optimal lower threshold would require comprehensive analysis balancing smuggling prevention against commercial facilitation. To determine a revised figure, we recommend the determination be made by the four members of the End-User Review Committee.

- Increase funding as well as analytic and monitoring resources for the Bureau of Industry and Security (BIS) for export enforcement. Export controls are only effective with proper enforcement. A thorough assessment of BIS's current enforcement capabilities and the potential benefits of additional resources would significantly enhance the overall effectiveness of these controls.

Conclusion

The Diffusion Framework represents a critical policy tool for maintaining America's strategic advantage in artificial intelligence. As powerful AI systems approach breakthrough capabilities by 2027, the strategic imperative to preserve America's compute advantage over China has never been more urgent. Export controls on advanced semiconductors are not merely technical regulations—they are foundational to America's national security and economic prosperity in an era defined by AI innovation.

The first Trump Administration correctly diagnosed the centrality of AI to strategic competition with China and the Diffusion Framework builds upon this foundation by preventing AI infrastructure offshoring, addressing smuggling vulnerabilities, and creating incentives for domestic investment. Any pause or weakening of these measures would provide China with an opportunity to stockpile advanced semiconductors and accelerate its efforts to close the compute gap just as transformative AI capabilities are emerging.

America stands at a pivotal moment in technological history. The development of AI systems with capabilities rivaling Nobel Prize winners across multiple disciplines will transform our economy, national security, and society. By strengthening—not weakening—the Diffusion Framework and vigorously enforcing export controls, America can ensure that these transformative technologies are developed domestically, by American firms, and in alignment with American values and interests. Our continued leadership in AI depends on maintaining and expanding our compute advantage through decisive policy action today.