



NVIDIA GPU Export Restrictions: Current US Chip Controls 2026

Plan and build private AI compute with GPU Smith. We connect workload requirements with GPU hardware, networking, deployment and operating decisions so your infrastructure fits the work it must perform.

Published by GPU Smith · 2026-07-22 · [nvidia gpu export restrictions](#) · [us export controls on ai chips](#) · [nvidia h20 export ban](#) · [nvidia h200 china](#) · [commerce department entity list](#) · [nvidia blackwell china export](#) · [chip export policy 2026](#) · [ai chip license requirements](#)

Original article: <https://gpusmith.com/articles/en/nvidia-gpu-export-restrictions>



In this article

1. NVIDIA GPU Export Restrictions: Current US Chip Controls 2026
2. Introduction and Background
3. The Regulatory Architecture: How Export Controls on AI Chips Work
4. Which NVIDIA GPUs Are Currently Banned or Restricted in China
5. Fully Restricted Chips
6. China-Specific Compliant Variants and Their Fate
7. The H200: Conditionally Licensed, Barely Shipping
8. The Blackwell Loophole and Enforcement Gaps
9. How the Loophole Worked
10. Closing the Gap
11. The Entity List Pause
12. The China Side of the Equation: Beijing's Own Restrictions
13. The 2025-2026 Regulatory Timeline
14. From Total Ban to Conditional Access
15. Earlier Foundational Rules (2022-2024)
16. The Financial and Market Impact on NVIDIA
17. Direct Charges from the H20 License Requirement
18. Zero China Assumption Going Forward
19. Market Capitalization Context
20. Case Studies and Real-World Examples
21. The H20 Charge and Reversal (April-July 2025)
22. DeepSeek and the Question of Restricted Chip Access
23. The Super Micro Smuggling Prosecution
24. The Blackwell Subsidiary Loophole (2025-2026)
25. China's Market Share Reversal (2022-2026)
26. Data Analysis and Evidence
27. Implications and Future Directions
28. Frequently Asked Questions (FAQs)
29. Conclusion

NVIDIA GPU Export Restrictions: Current US Chip Controls 2026

Financial Times reporting cited by industry analysts projects Huawei's China AI chip revenue could reach \$12 billion in 2026 as its share climbs toward 50 to 60% ([wccftech.com](https://www.wccftech.com)). NVIDIA's own guidance now assumes zero China data center compute revenue going forward: its outlook for the second quarter of fiscal 2027 explicitly states the company is "not assuming any Data Center compute revenue from China" (nvidianews.nvidia.com), even as total company revenue reached a record \$81.6 billion in the first quarter of

fiscal 2027 (nvidianews.nvidia.com). Congress, meanwhile, is pushing the bipartisan Chip Security Act to mandate hardware-level location verification on exported chips and the GAIN AI Act to require export licensees to prioritize U.S. buyers, both signs that the legislative and executive branches remain at odds over how tightly to control the flow of the world's most valuable technology export (congress.gov).

Introduction and Background

NVIDIA GPU export restrictions are a set of U.S. federal controls that determine which of the company's AI accelerators can legally be sold to China and other restricted destinations, under what license conditions, and subject to what performance thresholds. The controls are administered primarily by BIS under the EAR, a body of regulation issued pursuant to the Export Control Reform Act of 2018 (ECRA) (bis.gov). Since October 2022, when BIS first imposed sweeping restrictions on advanced computing and semiconductor manufacturing equipment exports to China, the rules have been revised at least half a dozen times, oscillating between tightening and loosening as successive administrations balanced national security concerns against commercial interests and geopolitical leverage (mayerbrown.com).

At the center of these controls sits **NVIDIA**, the Santa Clara, California-based chipmaker that as of July 2026 carries a market capitalization of approximately \$5.02 trillion, making it the most valuable publicly traded company in the world (companiesmarketcap.com). NVIDIA's data center GPUs, most notably the Hopper generation (**H100**, **H200**) and the newer Blackwell generation (**B100**, **B200**, **GB200**), have become the dominant hardware substrate for training and running large language models and other frontier AI systems worldwide. Because these chips also carry potential military applications, including in weapons design, nuclear simulation, and command-and-control systems, the U.S. government treats their export the way it treats munitions: subject to a formal licensing regime that can range from automatic approval to outright prohibition depending on the destination country, the end user, and the chip's computational power.

The regulatory framework relies on technical thresholds rather than simple product names. BIS measures chip capability using metrics such as Total Processing Performance (TPP), a composite of computational throughput and interconnect bandwidth, and "total DRAM bandwidth," the speed at which a chip can move data to and from memory (federalregister.gov). A chip crossing a given TPP threshold requires a license for export to China regardless of its brand name; NVIDIA has repeatedly re-engineered chips, including the **A800**, **H800**, and **H20**, specifically to sit just below whatever threshold was in force at the time, only to see each variant eventually banned as thresholds tightened (cnbc.com).

The stakes extend well beyond one company's balance sheet. China represented \$17 billion in NVIDIA revenue in the fiscal year ending January 26, 2025, or roughly 13% of total sales at that time (reuters.com), a market too large for Washington to ignore but also one policymakers view as central to preventing China's People's Liberation Army (PLA) from closing the AI compute gap with the United States. This report explains the current state of NVIDIA GPU export restrictions as of July 2026: which chips are banned, which require licenses, how the rules evolved, what the data show about market share and enforcement, and where the policy is headed next.

The Regulatory Architecture: How Export Controls on AI Chips Work

Legal Authority and Core Mechanisms

BIS derives its authority to restrict semiconductor exports from ECRA, which allows the Commerce Department to control the export of dual-use items, technology, and software for reasons of national security and foreign policy ([bis.gov](https://www.bis.gov)). The regulations operate through several overlapping mechanisms. **Export Control Classification Numbers (ECCNs)** categorize products, and Category 3A090 covers the advanced integrated circuits at issue in the NVIDIA controls ([congress.gov](https://www.congress.gov)). The **Entity List**, formally Supplement No. 4 to Part 744 of the EAR, names specific organizations, universities, and companies for which "there is reasonable cause to believe" involvement in activities contrary to U.S. national security, triggering individualized licensing requirements ([bis.gov](https://www.bis.gov)). Decisions to add, remove, or modify Entity List designations flow through the interagency End-User Review Committee, composed of the Departments of Commerce, Defense, State, Energy, and, where appropriate, Treasury; additions require only a majority vote, while removals require unanimity ([bis.gov](https://www.bis.gov)).

The Foreign Direct Product Rule and Allied Coordination

U.S. export controls extend beyond goods physically shipped from American soil. The Foreign Direct Product Rule allows Washington to restrict foreign-made items if they were produced using U.S.-origin technology, equipment, or software, giving the EAR extraterritorial reach over chips manufactured overseas, including by Taiwan Semiconductor Manufacturing Company (TSMC), using U.S.-designed tools ([mayerbrown.com](https://www.mayerbrown.com)). Because chip fabrication is globally concentrated, the U.S. has also coordinated with allied governments controlling key steps in the supply chain. In January 2023, Japan and the Netherlands agreed to align their own export controls with Washington's, restricting companies such as **ASML Holding**, the Dutch maker of the extreme ultraviolet (EUV) lithography machines essential for cutting-edge chip fabrication, from selling advanced immersion lithography systems to China ([mayerbrown.com](https://www.mayerbrown.com)). ASML itself said at the time that "it is not possible to make any statement about the process going forward" regarding the long-term financial and industry-wide impact of the new controls ([mayerbrown.com](https://www.mayerbrown.com)). Because the European Union grants member states discretion to impose their own export authorizations on dual-use items for reasons of "public security," the Dutch government retained the legal latitude to restrict ASML unilaterally even outside a formal EU-wide measure ([mayerbrown.com](https://www.mayerbrown.com)). The Dutch government had already restricted ASML from shipping certain dual-use equipment to China since 2018 ([mayerbrown.com](https://www.mayerbrown.com)). Japan separately moved to amend a ministerial ordinance under its Foreign Exchange and Foreign Trade Act, setting equipment-export limits on Nikon Corp comparable to the Dutch restrictions on ASML ([mayerbrown.com](https://www.mayerbrown.com)).

Country Tiers and Validated End User Status

The Biden administration published its January 2025 "Framework for Artificial Intelligence Diffusion," commonly called the AI Diffusion Rule, on January 15, 2025. Although the rule was effective on publication, its principal compliance requirements were scheduled for May 15, 2025; Commerce rescinded the rule on May 13, before that compliance date ([federalregister.gov](https://www.federalregister.gov)). The original rule text also imposed "Special Post Shipment Verification" reporting and recordkeeping requirements on computer exports to the most restricted country

group, a compliance mechanism distinct from the licensing thresholds themselves ([federalregister.gov](https://www.federalregister.gov)). Under that framework, Tier 1 countries (the United States and roughly 18 close allies including the United Kingdom, Canada, Germany, Japan, South Korea, Taiwan, Australia, and New Zealand) faced no chip procurement caps; Tier 2 countries (most of the rest of the world, including the Middle East, Latin America, and Southeast Asia) faced numeric caps of roughly 50,000 GPUs through 2027 absent special certification; and Tier 3 countries, roughly 20 states including China, Russia, Iran, North Korea, and Sudan, faced an outright prohibition on advanced GPU exports ([introl.com](https://www.introl.com)). AI safety company **Anthropic**, in a formal submission responding to the rule, described the same structure as "a three-tier system based on national security risk. Tier 1 includes close allies with few restrictions, Tier 2 includes most other countries with some limits, and Tier 3 includes adversarial nations with strict controls" ([anthropic.com](https://www.anthropic.com)), and separately noted that Tier 2 countries could purchase roughly \$40 million worth of chips, the equivalent of 1,700 NVIDIA H100-class accelerators, without a government license ([anthropic.com](https://www.anthropic.com)).

The Trump administration rescinded this framework on May 13, 2025, before its May 15 compliance deadline took effect, calling it a set of "burdensome new regulatory requirements" that "would have undermined U.S. diplomatic relations with dozens of countries by downgrading them to second-tier status" ([bis.gov](https://www.bis.gov)). In place of country tiers, organizations in Tier 2-type countries can now pursue **Validated End User (VEU)** status, which requires demonstrating regulatory compliance, accepting monitoring and audits, and committing to prevent diversion, in exchange for procurement allocations above baseline caps ([introl.com](https://www.introl.com)).

Transshipment Risk: Malaysia, Thailand, and Manufacturing Capacity Constraints

The rescission of the AI Diffusion Rule left a gap the Trump administration moved to address through country-specific measures instead of a global framework. Bloomberg reported in July 2025 that the administration was drafting new controls restricting AI chip shipments to **Malaysia** and **Thailand** specifically to curb suspected diversion to China ([datacenterdynamics.com](https://www.datacenterdynamics.com)), though that draft measure was described by people familiar with the matter as only "a first formal step" that remained "far from a comprehensive replacement" for the rescinded diffusion rule ([datacenterdynamics.com](https://www.datacenterdynamics.com)). Malaysia's Ministry of Investment, Trade and Industry (MITI) had already been investigating an unnamed Chinese company's potential use of NVIDIA chips barred from export to China even before it formalized its own transshipment rules ([datacenterdynamics.com](https://www.datacenterdynamics.com)), and formally announced its own rules on July 14, 2025, requiring companies to notify Malaysian authorities at least 30 days in advance of any export or transshipment of U.S.-origin AI chips, warning that violators "will face strict legal action" ([techcrunch.com](https://www.techcrunch.com)). The scale of the diversion problem this targeted was substantial: in an April 2025 policy submission written explicitly "in favor of the U.S. imposing more AI chip export rules" to prevent such smuggling ([techcrunch.com](https://www.techcrunch.com)), Anthropic wrote that "China has established sophisticated smuggling operations, with documented cases involving hundreds of millions of dollars worth of chips," and that smugglers had used methods including "hiding processors in prosthetic baby bumps and packing GPUs alongside live lobsters" to move hardware across borders ([anthropic.com](https://www.anthropic.com)). Anthropic separately projected that the resulting compute gap would compound over time, estimating that by 2027 countries relying on older, restricted chips could face AI training costs "ten times higher" than those with access to cutting-edge American technology, and noted that DeepSeek itself "likely used frontier chips for training their systems" despite the restrictions in place at the time ([anthropic.com](https://www.anthropic.com)). The Commerce Department had also drafted plans around the same time to revoke certain

operating authorizations for global chipmakers with China-based operations, a step that could complicate Samsung, SK Hynix, and TSMC's ability to use U.S. goods and technology at their Chinese factories (datacenterdynamics.com).

Export controls have also intersected with a genuine global capacity crunch. China's largest contract chipmaker, **Semiconductor Manufacturing International Corporation (SMIC)**, reported in May 2026 that overseas clients were increasingly shifting orders to Chinese foundries simply because capacity elsewhere was full: co-CEO Zhao Haijun said "these are among the few places with available production capacity, so we are seeing many overseas customers shift their orders to be manufactured in China," adding that SMIC's push into more advanced 7-nanometer manufacturing nonetheless continues to face direct constraints from U.S. export controls on semiconductor manufacturing equipment (reuters.com), underscoring that the allied equipment controls described above remain a binding constraint on China's most advanced domestic chip production even as its legacy-node capacity expands rapidly.

Which NVIDIA GPUs Are Currently Banned or Restricted in China

Fully Restricted Chips

As of July 2026, NVIDIA's most powerful chips remain formally barred from China under a presumption of denial. This category includes the **H100** and its variants along with the **A100** (both 40GB and 80GB configurations), both of which NVIDIA itself confirmed were subject to export restrictions in an October 2023 SEC filing following the second major round of BIS controls (cnbc.com). The restrictions also extend to NVIDIA's newest **Blackwell family**: **B100**, **B200**, and the **GB200** rack-scale system, all subject to the same presumption-of-denial posture for China and Macau (introl.com). The restrictions also reach performance-adjacent consumer hardware: the **RTX 4090** gaming GPU met the TPP threshold used in the October 2023 rules and was named explicitly in that same SEC disclosure (cnbc.com).

China-Specific Compliant Variants and Their Fate

NVIDIA has repeatedly engineered China-specific chip variants designed to sit just below whatever performance threshold was in force, a strategy that has succeeded only temporarily each time regulators moved the line:

- **A800**: A modified A100 with NVLink interconnect bandwidth reduced from 600 GB/s to 400 GB/s. Blocked by the October 2023 rules, which senior administration officials confirmed were designed to close loopholes that had emerged after the original 2022 restrictions (cnbc.com) (introl.com).
- **H800**: A modified H100 with interconnect bandwidth cut from 900 GB/s to roughly 300 GB/s. Also blocked in October 2023 (introl.com). Chinese AI startup **DeepSeek** later disclosed it used H800 chips to train its R1 and V3 models, though congressional investigators and U.S. officials have separately alleged the company also obtained restricted H100s through intermediaries (reuters.com).

- **H20**: A deeply stripped-down Hopper-generation chip. Initially permitted for export, it became subject to a U.S. license requirement for exports into the China market starting April 2025 ([sec.gov](#)).
- **L40 and L40S**: Data center GPUs also named in the October 2023 restrictions ([cnbc.com](#)). NVIDIA disclosed at the time that the restrictions also covered entire systems built around these chips, including its DGX and HGX server platforms ([cnbc.com](#)).

The H200: Conditionally Licensed, Barely Shipping

The **H200**, NVIDIA's second most powerful shipping chip as of mid-2026, occupies a unique middle position: legally exportable under the January 2026 case-by-case rule, but subject to per-buyer caps, third-party testing, and volume restrictions. Roughly ten Chinese firms, including Alibaba, Tencent, and ByteDance, had been cleared to purchase the chip by mid-2026, each authorized for up to 75,000 units through NVIDIA directly or through authorized distributors Lenovo and Foxconn ([techtimes.com](#)) ([scmp.com](#)). On July 14, 2026, the approved buyer list expanded to include ZTE Kangxun Telecom, a unit of a company the U.S. government had itself designated a national security threat under the 2018 National Defense Authorization Act, a fact reported the same day Reuters revealed the firm's newly approved status ([kfgo.com](#)).

Both regulatory and commercial obstacles have slowed deliveries. Chinese customs authorities blocked H200 shipments from entering the country in January 2026, even after Washington approved them, with suppliers pausing production amid the confusion ([theguardian.com](#)). Beijing has separately restricted approved H200 use strictly to AI model training while directing companies toward domestic processors, chiefly Huawei's Ascend line, for routine inference workloads, and one proposal reportedly under consideration would require each H200 purchase to be bundled with a fixed ratio of Chinese-made chips ([techtimes.com](#)). SCMP reporting has separately noted that some industry experts had warned Chinese companies may not embrace the licensed H200 in large volumes regardless of approval, "citing security concerns and a push for domestic chip self-sufficiency" ([scmp.com](#)). A spokesperson for the Chinese embassy in Washington, Liu Chang, said China's position on U.S. chip export policy "has been consistent" in opposing what Beijing considers politically motivated trade restrictions ([scmp.com](#)), while Washington has separately clarified that Blackwell, NVIDIA's most advanced line, "remains strictly banned from direct export to China" even as H200 sales proceed ([scmp.com](#)).

The Blackwell Loophole and Enforcement Gaps

How the Loophole Worked

NVIDIA's Blackwell chips, along with AMD's MI350X-class parts, were subject to export license requirements for direct shipment into mainland China. However, the licensing trigger was based on shipping destination rather than the ultimate purchaser's corporate headquarters. Overseas subsidiaries of Chinese AI companies, incorporated in jurisdictions such as Malaysia, Singapore, or the United Arab Emirates, could purchase Blackwell-class chips without triggering a China-specific license requirement, because the goods were

technically destined outside mainland China ([techtimes.com](https://www.techtimes.com)). The Commerce Department created this opening in May 2025 when it announced it would not enforce the Biden-era AI Diffusion Rule's global licensing requirements ([reuters.com](https://www.reuters.com)).

Closing the Gap

BIS issued unexpected weekend guidance on May 31, 2026, to close the loophole, framing the move as a "clarification" of export license requirements that it said had technically been in place since 2023 ([reuters.com](https://www.reuters.com)). A paper circulating in Washington that prompted the guidance stated bluntly that "the floodgates have quietly opened" ([reuters.com](https://www.reuters.com)). One chip-industry source with deep supply chain knowledge estimated the volume affected was "in the hundreds of thousands" of units ([reuters.com](https://www.reuters.com)), and former State Department official Chris McGuire called the loophole "a HUGE problem," writing that "Chinese companies have been buying these chips, very likely at scale" ([reuters.com](https://www.reuters.com)).

Even the fix was incomplete. McGuire noted the new guidance left open a second gap: TSMC and other foundries were no longer explicitly required to perform extra due diligence to confirm that high-end AI chips they manufactured were not destined for Chinese front companies ([reuters.com](https://www.reuters.com)). The guidance also did not require data centers outside China that had already installed Blackwell-equipped servers to shut them down or stop offering computing services from them ([reuters.com](https://www.reuters.com)).

At the July 14, 2026 hearing, Representative Bill Huizenga pressed Kessler on why the May 31 guidance allowed Chinese companies to retain Blackwell chips already obtained through the loophole rather than requiring surrender or destruction, asking why the guidance said firms could keep chips "that they may have gotten through either smuggling or other loopholes" ([kfgo.com](https://www.kfgo.com)). Kessler responded that unlicensed holders should "voluntarily self-disclose" the violation, prompting Huizenga to reply, "That is just a frickin' circle that you talk in, and I'm sorry but this is unacceptable" ([kfgo.com](https://www.kfgo.com)).

The Entity List Pause

A separate enforcement gap concerns the Entity List itself. Representative Gregory Meeks, the senior Democrat on the House Foreign Affairs Committee, said the Commerce Department had not added any Chinese companies to the export control list since October, "which is the longest period in more than a decade" ([kfgo.com](https://www.kfgo.com)). Meeks charged that President Trump "has turned export controls into a bargaining chip in broader negotiations with China" and had "weakened existing safeguards, including by approving licenses for advanced AI chips destined for China" ([kfgo.com](https://www.kfgo.com)). More than 100 interagency-approved entities, including DeepSeek and memory chipmaker ChangXin Memory Technologies (CXMT), remain unpublished on the list despite committee approval, meaning U.S. suppliers face no formal legal trigger requiring a license before dealing with them ([kfgo.com](https://www.kfgo.com)). Kessler defended the department's posture at the hearing, saying it remained important to enforce the existing list of restricted Chinese companies even without new additions, and separately stated, "I don't want to replace the diffusion rule because I don't think the rule is worth replacing" while confirming "there will be future regulatory action in the area of chips and AI" ([kfgo.com](https://www.kfgo.com)). Philip Luck of the Center for Strategic and International Studies (CSIS) described the pause as an unprecedented enforcement gap, "longer than anything witnessed in the past decade," according to reporting on the matter ([techtimes.com](https://www.techtimes.com)).

The China Side of the Equation: Beijing's Own Restrictions

Export controls are not a one-way street. China has layered its own restrictions on top of U.S. rules, aimed at accelerating domestic chip self-sufficiency and reducing dependence on foreign hardware that could be cut off unpredictably.

In September 2025, China's internet regulator, the Cyberspace Administration of China, banned the country's largest technology companies from buying NVIDIA's AI chips ([reuters.com](https://www.reuters.com)). Two months later, in November 2025, Chinese regulators issued guidance requiring any new data center project that had received state funding to use only domestically made AI chips, ordering projects less than 30% complete to remove any installed foreign chips or cancel purchase plans altogether ([reuters.com](https://www.reuters.com)). AI data center projects in China have drawn more than \$100 billion in state funding since 2021, making the guidance's practical reach very large ([reuters.com](https://www.reuters.com)). Notably, the guidance covers not just banned chips like the B200 but also the licensed H20 and H200, meaning even chips the U.S. has approved for export may be excluded from major Chinese state-funded projects ([reuters.com](https://www.reuters.com)).

China's motives combine genuine security concerns with industrial policy. Beijing has separately summoned domestic tech firms to warn them against buying restricted chips even where legally permitted, while showcasing data centers built entirely on domestic silicon ([reuters.com](https://www.reuters.com)). This mirrors a 2023 precedent: China banned U.S. memory chipmaker Micron's products from critical infrastructure that year, a restriction that eventually led Micron to exit China's server chip market entirely in 2025 ([reuters.com](https://www.reuters.com)).

The 2025-2026 Regulatory Timeline

From Total Ban to Conditional Access

The past 15 months have seen the fastest series of reversals in the history of U.S. AI chip export policy. **April 9, 2025:** NVIDIA disclosed via SEC filing that BIS informed the company a license would be required indefinitely to export H20 chips to China, citing risk the chip could be used in "a supercomputer in China" ([techcrunch.com](https://www.techcrunch.com)). **May 13, 2025:** Commerce rescinded the Biden-era AI Diffusion Rule ahead of its compliance deadline ([bis.gov](https://www.bis.gov)). **July 14-15, 2025:** Malaysia imposed its own 30-day advance notification rule on U.S. AI chip transshipments ([techcrunch.com](https://www.techcrunch.com)), and NVIDIA announced it was filing applications to resume H20 sales to China after receiving U.S. assurances the applications would be approved, a decision Commerce Secretary Howard Lutnick tied directly to a broader trade agreement covering rare earth minerals ("We put that in the trade deal with the magnets") ([reuters.com](https://www.reuters.com)).

August 2025: NVIDIA and AMD agreed to remit 15% of China chip sales revenue to the U.S. government in exchange for continued export licenses, an arrangement the Financial Times described as unprecedented, coming as Trump had separately threatened a 100% tariff on semiconductor imports unless companies built domestically ([cnbc.com](https://www.cnbc.com)). **December 8, 2025:** President Trump announced the administration would approve H200 sales to China. **January 13-14, 2026:** BIS published the final rule formalizing case-by-case review for


H200-class chips, and Trump signed a Section 232 proclamation imposing the 25% chip tariff following "the Secretary of Commerce's completion of a Section 232 investigation" that concluded existing semiconductor import levels "threaten to impair national security" ([whitehouse.gov](https://www.whitehouse.gov)). **January 17, 2026:** Chinese customs blocked incoming H200 shipments despite U.S. clearance, and NVIDIA's suppliers paused production (theguardian.com).

March 2026: Federal prosecutors unsealed charges against a Super Micro co-founder and two others for smuggling \$2.5 billion in NVIDIA-equipped servers to China ([cnn.com](https://www.cnn.com)). **May 31, 2026:** BIS issued emergency weekend guidance closing the Blackwell subsidiary loophole ([reuters.com](https://www.reuters.com)). **July 14, 2026:** Kessler's congressional testimony revealed the scale of the licensing-to-shipment gap and Commerce's Entity List pause ([kfgo.com](https://www.kfgo.com)).

Earlier Foundational Rules (2022-2024)

The current framework rests on a foundation laid across 2022 and 2023. In October 2022, BIS issued its first interim final rule imposing broad license requirements on advanced computing and semiconductor manufacturing equipment exports to China ([mayerbrown.com](https://www.mayerbrown.com)). On October 17, 2023, BIS closed loopholes exposed by the initial rule, explicitly banning the NVIDIA A800 and H800 chips that had been engineered to comply with the 2022 thresholds ([cnbc.com](https://www.cnbc.com)). Then-Commerce Secretary Gina Raimondo said the update was "specifically designed to control access to computing power, which will significantly slow the PRC's development of next-generation frontier model" ([cnbc.com](https://www.cnbc.com)), while officials cautioned the changes would only affect "a small fraction of chip exports to China" given that "China, even after the update of this rule, will import hundreds of billions of dollars of semiconductors from the United States" ([cnbc.com](https://www.cnbc.com)). In March 2025, the first Trump administration BIS action added 80 entities to the Entity List across China, the UAE, South Africa, and Iran, explicitly to "restrict the Chinese Communist Party's (CCP) ability to acquire and develop high-performance and exascale computing capabilities, as well as quantum technologies, for military applications," with 12 entities cited specifically for developing advanced AI, supercomputers, or high-performance AI chips for China-based end users tied to the military-industrial complex, including two entities added for selling to Huawei and its HiSilicon affiliate ([bis.gov](https://www.bis.gov)) ([bis.gov](https://www.bis.gov)). Then-Commerce Secretary Howard Lutnick framed the action as evidence the department was "taking decisive action to protect America," adding "we will not allow adversaries to exploit American technology to bolster their own militaries and threaten American lives" ([bis.gov](https://www.bis.gov)). Across 2025 as a whole, BIS added 65 Chinese entities to the list, 42 in March and 23 more in September, before the pause that began in October ([introl.com](https://www.introl.com)).

The Financial and Market Impact on NVIDIA



Direct Charges from the H20 License Requirement

The April 2025 H20 license requirement inflicted immediate, quantifiable financial damage. NVIDIA's own CFO commentary for the first quarter of fiscal 2026 disclosed a \$4.5 billion charge tied to H20 excess inventory and purchase obligations, plus an additional \$2.5 billion of H20 revenue the company was unable to ship in that quarter alone because of the new licensing requirement ([sec.gov](#)). Excluding that charge, non-GAAP gross margin for the quarter would have been 71.3% rather than the reported 61.0% ([sec.gov](#)). NVIDIA's outlook for the following quarter (Q2 FY2026) explicitly built in an anticipated \$8.0 billion loss in H20 revenue due to the export control limitations then in force ([nvidianews.nvidia.com](#)). When the H20 ban had first been proposed under looser terms in early 2025, NVIDIA had separately estimated the curbs could cut annual revenue by as much as \$15 billion ([reuters.com](#)). NVIDIA's stock fell approximately 6% in extended trading on the day the April 2025 restriction was disclosed ([techcrunch.com](#)).

Zero China Assumption Going Forward

By the first quarter of fiscal 2027 (reported May 2026), NVIDIA had stopped forecasting any China compute revenue altogether. The company reported record revenue of \$81.6 billion for the quarter, up 85% year over year, and record Data Center revenue of \$75.2 billion, up 92% year over year ([nvidianews.nvidia.com](#)). Its outlook for the subsequent quarter, targeting \$91.0 billion in revenue, states plainly that the company is "not assuming any Data Center compute revenue from China in its outlook" ([nvidianews.nvidia.com](#)). This posture matches disclosure NVIDIA made in its fiscal 2026 annual report (Form 10-K), stating the company was "effectively foreclosed from competing in China's data center computing/compute market" and that this foreclosure "helped our competitors build larger developer and customer ecosystems to challenge us worldwide" ([finance.yahoo.com](#)). NVIDIA's fiscal 2026 full-year China (including Hong Kong) revenue was \$19.67 billion, but the company said it did not anticipate meaningful further China revenue in the following quarter, adding that the government had granted it a license to ship H200 chips in February 2026 without yet knowing whether an import license would also be granted on the Chinese side ([finance.yahoo.com](#)).

Market Capitalization Context

Despite the China disruption, NVIDIA's overall business has continued to expand rapidly on the strength of global AI infrastructure demand. As of July 2026 the company's market capitalization stood at approximately \$5.02 trillion, the largest of any publicly traded company, with shares trading around \$207.29 and up nearly 21% over the preceding year (companiesmarketcap.com). NVIDIA's first-quarter fiscal 2027 non-GAAP gross margin reached 75.0%, roughly 14 percentage points above the H20-charge-depressed 61.0% margin reported a year earlier (nvidianews.nvidia.com). The scale of that broader growth illustrates why, even as China's contribution has shrunk toward zero, NVIDIA's aggregate financial position has remained historically strong.

Case Studies and Real-World Examples

The H20 Charge and Reversal (April-July 2025)

NVIDIA's H20 saga is the clearest illustration of how abruptly export policy can shift and how directly it flows through to a public company's financial statements. On April 9, 2025, BIS informed NVIDIA that the H20, a chip NVIDIA had specifically engineered to comply with earlier thresholds and legally sell into China, would now require an export license "indefinitely," citing the risk it could be used in a Chinese supercomputer (techcrunch.com). NPR had reported days earlier that CEO Jensen Huang may have secured a temporary reprieve during a dinner at President Trump's Mar-a-Lago resort by committing NVIDIA to invest in U.S. AI data center manufacturing, a commitment the company formalized the following day (techcrunch.com). Three months later, the reversal came just as suddenly: on July 15, 2025, NVIDIA announced it would resume filing license applications, explicitly tied by Commerce Secretary Lutnick to a rare earths trade agreement with China (reuters.com). The reversal drew immediate bipartisan pushback: Democratic Representative Raja Krishnamoorthi warned it would "hand our foreign adversaries our most advanced technologies," while Republican Committee Chair John Moolenaar noted the H20 had "played a significant role in the rise of PRC AI companies like DeepSeek" according to a bipartisan House investigation (reuters.com). Analyst Divyansh Kaushik of Beacon Global Strategies warned at the time that "if China is able to get a million H20 chips, it could significantly narrow, if not overtake, the U.S. lead in AI" (reuters.com).

DeepSeek and the Question of Restricted Chip Access

DeepSeek's January 2025 release of its R1 reasoning model, which the Hangzhou-based startup claimed matched frontier U.S. models at a fraction of the training cost, became a pivotal case study shaping subsequent export policy debates. A senior State Department official told Reuters in June 2025 that DeepSeek was "willingly" providing support to China's military and intelligence operations and had used Southeast Asian shell companies to attempt to access high-end U.S. semiconductors barred from export to China, alleging DeepSeek had access to "large volumes" of NVIDIA's restricted H100 chips (reuters.com). NVIDIA disputed the specific chip claim, stating its own review indicated DeepSeek used "lawfully acquired H800 products, not H100" (reuters.com), and separately stated it does "not support parties that have violated U.S. export controls or are on the U.S. entity lists," while acknowledging that "with the current export controls, we are effectively out

of the China data center market, which is now served only by competitors such as Huawei" ([reuters.com](https://www.reuters.com)). The episode directly motivated the Chip Security Act now moving through Congress, which its sponsors describe as a legislative response to the House Select Committee on China's finding that DeepSeek trained its model on restricted chips that should never have reached it ([thecyberexpress.com](https://www.thecyberexpress.com)).

The Super Micro Smuggling Prosecution

In March 2026, the U.S. Attorney's Office for the Southern District of New York unsealed charges against Super Micro Computer co-founder Yih-Shyan "Wally" Liaw and two others, Ruei-Tsang "Steven" Chang and Ting-Wei "Willy" Sun, for conspiring to divert approximately \$2.5 billion worth of NVIDIA AI chip-equipped servers to China ([cnn.com](https://www.cnn.com)). Prosecutors alleged the men routed orders through a pass-through company in Southeast Asia to obscure China as the final destination, used a logistics company to repackage servers into unmarked boxes, and, to deceive compliance auditors, staged "dummy" non-working replica servers at a rented warehouse while the real hardware shipped separately, going so far as to use hair dryers to remove and replace shipping labels ([cnn.com](https://www.cnn.com)). U.S. Attorney Jay Clayton said "diversion schemes like those disrupted today generate billions of dollars in ill-gotten gains and pose a direct threat to U.S. national security" ([cnn.com](https://www.cnn.com)). The company itself, in a statement, said the conduct alleged "is a contravention of the Company's policies and compliance controls, including efforts to circumvent applicable export control laws and regulations," adding that "Supermicro maintains a robust compliance program and is committed to full adherence to all applicable U.S. export and re-export control laws and regulations" ([cnn.com](https://www.cnn.com)). The case rippled to Taiwan, where the Keelung District Prosecutors' Office opened the island's first formal criminal probe into AI chip smuggling in May 2026, alleging fraudulent shipping declarations were used to route Super Micro servers containing NVIDIA chips to China, Hong Kong, and Macau (finance.yahoo.com). Jensen Huang, arriving in Taipei days after the Taiwan probe was announced, told reporters, "We insist our partners are compliant. We hope that they will enhance and improve their regulation compliance and prevent that from happening in the future" (finance.yahoo.com). In the same conversation, Huang confirmed the H200 had "been licensed to ship to China," calling it "terrific to be able to serve that market" while noting that, months after approval, "not a single H200 has been delivered" to a Chinese customer (finance.yahoo.com).

The Blackwell Subsidiary Loophole (2025-2026)

This case, detailed above, stands as perhaps the most consequential enforcement failure of the current export control regime. Over roughly a full year, between the May 2025 non-enforcement decision and the May 31, 2026 closure, Chinese company subsidiaries incorporated outside mainland China were able to purchase NVIDIA's most advanced Blackwell chips without triggering license requirements. The magnitude, estimated in the hundreds of thousands of units by an industry source with supply-chain visibility ([reuters.com](https://www.reuters.com)), dwarfs the volumes involved in either the DeepSeek or Super Micro cases and demonstrates how a single definitional gap, destination-based rather than ownership-based licensing triggers, can undermine an otherwise stringent regulatory framework.

China's Market Share Reversal (2022-2026)

The cumulative effect of four years of export restrictions and China's own countermeasures is visible in market share data. According to IDC data reviewed by Reuters, Chinese GPU and AI chip makers captured nearly 41% of China's AI accelerator server market in 2025, up from a market that NVIDIA dominated with roughly 95% share as recently as 2022 ([reuters.com](#)). Total shipments of AI accelerator cards by NVIDIA, AMD, and Chinese chipmakers reached approximately four million units in China in 2025 ([reuters.com](#)). Huawei Technologies alone shipped approximately 812,000 AI chips in 2025, roughly half of all domestically branded shipments, ahead of Alibaba's T-Head unit (approximately 265,000 cards), Baidu's Kunlunxin, and Cambricon (each around 116,000 cards) ([reuters.com](#)).

Data Analysis and Evidence

Table 1 below summarizes the current export status of NVIDIA's principal AI accelerator product lines as of July 2026, drawing on BIS's own regulatory thresholds and multiple corroborating trade press reports.

NVIDIA GPU / Product	Generation	China Export Status (July 2026)	Key Regulatory Milestone
A100 (40GB/80GB)	Ampere	Fully restricted (presumption of denial)	Named in NVIDIA's own October 2023 SEC restriction filing (cnbc.com)
A800	Ampere (China-compliant variant)	Fully restricted	Blocked by October 2023 rules (cnbc.com)
H100	Hopper	Fully restricted	Named in NVIDIA's own October 2023 SEC restriction filing (cnbc.com)
H800	Hopper (China-compliant variant)	Fully restricted	Blocked by October 2023 rules; linked to DeepSeek training (cnbc.com)
H20	Hopper (China-compliant variant)	Licensed, case-by-case, 15% revenue share to U.S. government	Banned April 2025, relicensed July 2025 (sec.gov)
H200	Hopper (top-tier)	Conditionally licensable, case-by-case, when the rule's specified thresholds and other conditions are met	Case-by-case rule effective January 15, 2026; separately, a 25% Section 232 duty applies to certain covered chip imports into the United States, not exports to China (federalregister.gov) (whitehouse.gov)
B100 / B200 / GB200 (Blackwell)	Blackwell	Fully restricted (presumption of denial); subsidiary loophole closed May 2026	Loophole closure guidance issued May 31, 2026 (reuters.com)
L40 / L40S	Ada Lovelace	Fully restricted	Named in October 2023 SEC filing (cnbc.com)

NVIDIA GPU / Product	Generation	China Export Status (July 2026)	Key Regulatory Milestone
RTX 4090	Ada Lovelace (consumer)	Fully restricted	Named in October 2023 SEC filing on restrictions (cnbc.com)

Table 1 shows that regulatory treatment is not reliably determined by a “China-specific” product label. The January 2026 pathway is limited to qualifying transactions, while applications outside it remain subject to a presumption of denial; prior changes to controls also show that product and licensing treatment can change quickly.

Table 2 below tracks the financial and market share metrics that quantify the scale of disruption from these restrictions.

Metric	Value	Period / As of Date	Source
H20 inventory and purchase obligation charge	\$4.5 billion	Q1 FY2026 (quarter ended April 27, 2025)	NVIDIA SEC filing (sec.gov)
H20 revenue unable to be shipped	\$2.5 billion	Q1 FY2026	NVIDIA SEC filing (sec.gov)
Projected H20 revenue loss (subsequent quarter)	\$8.0 billion	Q2 FY2026 outlook	NVIDIA earnings release (nvidianews.nvidia.com)
NVIDIA China (incl. Hong Kong) full-year revenue	\$19.67 billion	Fiscal year 2026	NVIDIA 10-K, cited via Yahoo Finance (finance.yahoo.com)
Revenue share paid to U.S. government on China chip sales	15%	Since August 2025	Financial Times via CNBC (cnbc.com)
Tariff on H200-class chip imports	25%	Since January 14, 2026	White House fact sheet (whitehouse.gov)
Approved H200 export license value	Approximately \$10 billion	As of July 2026	Congressional testimony via Tech Times (techtimes.com)
Chinese firm H200 orders vs. NVIDIA inventory	2 million+ ordered vs. ~700,000 available	2026	Tech Times, citing NBC News (techtimes.com)
Estimated H200/H100 volume permitted under the 50% cap	~1 million H200s, potentially ~1 million H100s	2026 estimate	CFR analysis (cfr.org)
Potential increase in China's 2026 installed AI compute if cap volume is reached	250%	2026 estimate	CFR analysis (cfr.org)
NVIDIA China AI accelerator market share	95% (2022) to 0% (claimed, 2026)	2022 vs. 2026	Jensen Huang interview via Yahoo Finance (finance.yahoo.com)
NVIDIA China AI accelerator server share (independent data)	55%	2025	IDC data via Reuters (reuters.com)

Metric	Value	Period / As of Date	Source
Chinese domestic vendor combined share	41%	2025	IDC data via Reuters (reuters.com)
Huawei projected China AI chip revenue	\$12 billion	2026 (forecast)	Financial Times via Wccftch (wccftch.com)
China AI chip market projected size (domestic suppliers, 86% share)	\$67 billion (total), \$21 billion (2026 domestic)	2030 forecast	Morgan Stanley via Financial Times/Wccftch (wccftch.com)
NVIDIA market capitalization	\$5.02 trillion	July 2026	CompaniesMarketCap (companiesmarketcap.com)
NVIDIA Q1 FY2027 total revenue	\$81.6 billion	Quarter ended April 26, 2026	NVIDIA earnings release (nvidianews.nvidia.com)
Chinese state funding drawn by AI data centers since 2021	Over \$100 billion	Since 2021	Reuters review of government tenders (reuters.com)
Tier 2 country no-license compute threshold under the rescinded AI Diffusion Rule	~1,700 H100-class chips (~\$40 million)	Effective January-May 2025	Anthropic policy submission (anthropic.com)
U.S. share of global semiconductor production	Fallen from 40% (1990) to 12% (2025)	1990 vs. 2025	Anthropic policy submission (anthropic.com)

Table 2 illustrates a striking divergence between NVIDIA's company-wide financial trajectory and its China-specific fortunes. While total revenue and market capitalization have both grown substantially through 2026, the China-specific figures show a business unit that has effectively been reduced to negligible planned contribution, with the company explicitly building its forward guidance around zero China data center compute revenue. The disparity between the roughly \$10 billion in approved H200 export licenses and the "trivial" actual shipments Kessler described to Congress captures the central tension of the current policy: paper permissiveness has not translated into commercial reality. The table's CFR-derived estimate that a fully utilized 50% volume cap could still authorize roughly 1 million H200s to China, enough to increase the country's 2026 installed AI compute by 250%, illustrates that even the "restrictive" case-by-case regime leaves substantial headroom for China's compute buildout if licenses are approved at scale.

Implications and Future Directions

Several structural forces will shape how NVIDIA GPU export restrictions evolve over the remainder of 2026 and beyond. First, the legislative branch and the executive branch remain in open tension. The **GAIN AI Act**, introduced by Senator Jim Banks with bipartisan cosponsors including Senators Warren, Cotton, Schumer, McCormick, and Coons, would require exporters seeking a license to sell advanced chips to countries of concern to certify that U.S. customers had a documented "right of first refusal," including a public notice period of at least 15 business days before any China-bound shipment could proceed (congress.gov). The bill separately defines a "backlog of requests" as any documented but unfulfilled U.S. purchase order that a chipmaker cannot meet within commercially standard delivery timeframes, and defines a "trusted United States person" as one specifically designated eligible for streamlined treatment under the Act, thresholds designed to

prevent exporters from claiming sufficient domestic supply while U.S. orders go unfilled ([congress.gov](https://www.congress.gov)). Critics, including CSIS and Brookings analysts, argue the bill's public notice requirement would depress the global value of U.S. AI chips and undermine competitiveness ([csis.org](https://www.csis.org)) ([brookings.edu](https://www.brookings.edu)). CFR's own analysis of the January 2026 rule notes that its 50 percent volume-cap formula "borrows heavily from concepts" already present in the GAIN AI Act, suggesting the executive and legislative approaches are converging on similar mechanics even as they clash over their scope ([cfr.org](https://www.cfr.org)). CFR further cautions that extending the same 50 percent formula to next-generation Blackwell-class GB300 chips could eventually authorize the sale of up to 2.5 million units to China, and separately flags that the rule's Know Your Customer verification requirements are difficult to apply credibly to the most likely large-volume buyers, since Tencent, one of China's largest cloud providers, has itself been designated a "Chinese Military Company" by the U.S. Department of Defense ([cfr.org](https://www.cfr.org)) ([cfr.org](https://www.cfr.org)). CFR's own conclusion is unambiguous: "the most effective policy remains the simplest one: denying the export of all AI chips to China" ([cfr.org](https://www.cfr.org)), a position that puts think-tank analysis sharply at odds with the administration's current case-by-case approach.

Separately, the **Chip Security Act** (H.R. 3447) was introduced in the House and referred to the House Foreign Affairs Committee on May 15, 2025; it remains at introduced status. If enacted, it would require the Commerce Secretary to mandate location-verification security mechanisms embedded in covered chips within 180 days of enactment, alongside mandatory reporting when license holders learn of diversion to unauthorized users ([thecyberexpress.com](https://www.thecyberexpress.com)). The bill's covered products include chips classified under ECCNs 3A090, 3A001.z, 4A090, and 4A003.z, the same classifications that cover NVIDIA's H100-class accelerators ([thecyberexpress.com](https://www.thecyberexpress.com)). White House AI czar David Sacks has publicly criticized the bill, prompting House Foreign Affairs Committee Chairman Brian Mast to note that the criticism echoed talking points he had heard directly from NVIDIA ([thecyberexpress.com](https://www.thecyberexpress.com)). Industry groups including the Information Technology and Innovation Council have separately warned that a government chip-tracking mandate risks pushing non-U.S. customers toward alternative suppliers ([thecyberexpress.com](https://www.thecyberexpress.com)).

Second, NVIDIA's own public position has grown more openly critical of the policy it must operate under. Jensen Huang has repeatedly called U.S. export controls a "failure," arguing at Taiwan's Computex trade show in May 2025 that the restrictions cut NVIDIA's China market share from 95% to 50% and motivated Beijing "to make its own chips faster" ([cnbc.com](https://www.cnbc.com)). Huang also told lawmakers in Washington that same spring, "China is right behind us. We are very close. Remember this is a long-term, infinite race," while separately praising Huawei's capabilities: "They're incredible in computing and network technology, all these essential capabilities to advance AI. They have made enormous progress in the last several years" ([cnbc.com](https://www.cnbc.com)). By April 2026, Huang told the Special Competitive Studies Project's "Memos to the President" series that share had fallen further, to zero, and that "conceding an entire market the size of China probably doesn't make a lot of strategic sense," adding that "I think that has already largely backfired" (finance.yahoo.com). This dynamic places Huang in the position of publicly lobbying against a policy set by the same administration he has also cultivated close ties with, including accompanying Trump on a Middle East trip and receiving public praise from the president ([cnbc.com](https://www.cnbc.com)).

Third, China's own domestic chip industry is on a trajectory that could make the entire U.S. licensing debate increasingly moot for large segments of the Chinese market regardless of what Washington decides. Morgan Stanley forecasts that China's AI chip market will reach \$67 billion by 2030, with 86% of that supplied by Chinese vendors rather than foreign chipmakers ([wccftech.com](https://www.wccftech.com)). Huawei's roadmap already extends through

its Ascend 950, 960, and 970 chip generations planned through 2028, with the newest Ascend 950PR chip supporting the same MXFP4 instruction set used by DeepSeek's latest models (wccftech.com), and both NVIDIA and Huawei provided simultaneous "Day-0" support for DeepSeek's V4 model release, a milestone the trade press marked as a significant validation of China's parallel AI hardware ecosystem (wccftech.com). Compounding this trend, SMIC's own capacity data show Chinese foundries' share of global legacy-node semiconductor capacity, in the 22 to 40 nanometer range, is expected to reach 37% in 2026 and 41% in 2027, up from 32% in 2025 (reuters.com), meaning China's chip self-sufficiency push extends well beyond AI accelerators into the broader semiconductor supply chain.

Fourth, enforcement capacity itself remains a bottleneck regardless of how the underlying rules are written. The Entity List pause, the Blackwell subsidiary loophole, the Super Micro and Taiwan smuggling cases, the Malaysia and Thailand transshipment restrictions, and the persistent gap between billions of dollars in approved licenses and "trivial" actual shipments collectively suggest that BIS's administrative and investigative capacity has not kept pace with the volume and sophistication of transactions it is tasked with policing. If enacted, the Chip Security Act's location-verification mandate would shift enforcement from paperwork-based licensing toward physical, continuously auditable hardware controls; whether it could be implemented within its proposed 180-day window while withstanding industry pushback remains uncertain.

Frequently Asked Questions (FAQs)

What are NVIDIA GPU export restrictions? They are U.S. federal controls, administered by BIS under the EAR, that require licenses for exporting NVIDIA's advanced AI accelerators to China and certain other restricted destinations, based on technical performance thresholds such as Total Processing Performance and DRAM bandwidth rather than simple product names (federalregister.gov).

Which NVIDIA GPUs are banned in China as of July 2026? The H100, A100, A800, H800, L40, L40S, RTX 4090, and the entire Blackwell family (B100, B200, GB200) remain subject to a presumption of denial for direct export to China (cnbc.com) (introl.com). The H20 and H200 are conditionally licensable on a case-by-case basis, subject to volume caps, testing requirements, and a 25% tariff for the H200 (federalregister.gov) (whitehouse.gov).

What is the NVIDIA H20 export ban and is it still in effect? The H20 export ban was imposed on April 9, 2025, requiring a license "indefinitely" and costing NVIDIA a \$4.5 billion charge in a single quarter (sec.gov). It was reversed in July 2025, and as of July 2026 the H20 is legally exportable under case-by-case license review, though subject to the 15% revenue-sharing arrangement with the U.S. government (reuters.com) (cnbc.com).

What does the Commerce Department's Entity List have to do with AI chips? Entity List designation triggers individualized, generally denied, licensing requirements for named organizations. As of July 2026, Commerce has not added new Chinese entities since October 2025 despite more than 100 additional entities, including DeepSeek and CXMT, awaiting formal listing, an enforcement gap lawmakers on both sides of the aisle have criticized (kfgo.com).

Does NVIDIA support the export restrictions? No. CEO Jensen Huang has publicly called the restrictions a "failure" that has ceded market share to Chinese competitors and accelerated China's domestic chip development rather than slowing it ([cnbc.com](https://www.cnbc.com)).

Do the restrictions require an export license for every sale to China? Not under the January 2026 rule alone. That rule changes review from a presumption of denial to case-by-case only for specified advanced-computing commodities exported from the United States to end users in China or Macau that are commercially available in the United States, have Total Processing Performance below 21,000 and total DRAM bandwidth below 6,500 GB/s, and meet the rule's other conditions. Other licensing requirements and end-user or end-use controls may still apply ([federalregister.gov](https://www.federalregister.gov)).

Is there a 25% tariff on NVIDIA chips sold to China? Yes. President Trump imposed a 25% Section 232 tariff on certain advanced computing chips, including the NVIDIA H200 and AMD MI325X, effective January 14, 2026, though the tariff does not apply to chips imported to support U.S. domestic manufacturing buildout ([whitehouse.gov](https://www.whitehouse.gov)).

Can China simply buy banned NVIDIA chips through other countries? U.S. regulators have identified this "transshipment" or diversion risk as a central enforcement challenge, evidenced by the Blackwell subsidiary loophole, the Super Micro and Taiwan smuggling cases, Malaysia's and Thailand's new transshipment rules, and DeepSeek's alleged use of Southeast Asian shell companies; Reuters reporting indicates banned chips like the B200 and H200 remain available in China through grey-market channels despite the controls ([reuters.com](https://www.reuters.com)).

Has China banned its own companies from buying NVIDIA chips? Yes, in part. China's internet regulator barred major tech firms from purchasing NVIDIA chips in September 2025, and in November 2025 Beijing mandated that any state-funded data center project use only domestic AI chips, a directive covering even the licensed H20 and H200 ([reuters.com](https://www.reuters.com)).

How does the US government's AI chip export licensing process actually work? Exporters must file a license application with BIS specifying the item, destination, and end user; for chips below the case-by-case thresholds, applicants must additionally certify sufficient domestic U.S. supply, absence of diversion to prohibited end uses, adequate Know Your Customer screening, and pre-export third-party lab testing to confirm the chip's technical specifications match the application ([federalregister.gov](https://www.federalregister.gov)).

Conclusion

NVIDIA GPU export restrictions as of July 2026 represent a regulatory framework caught between competing imperatives that have not been reconciled: the national security case for keeping advanced AI compute out of Chinese military hands, the commercial case for capturing revenue from the world's second-largest technology market before Chinese domestic suppliers close the gap entirely, and the enforcement reality that licensing rules on paper have proven far easier to write than to police. Current treatment is transaction-specific rather than reliably determined by a product name: commodities below both the 21,000 TPP and 6,500 GB/s total-DRAM-bandwidth thresholds may receive case-by-case review for qualifying U.S. exports to China or Macau, while other applications remain subject to a presumption of denial. The qualifying pathway includes volume-cap, certification, and third-party-testing requirements; the separate Section 232 duty applies to covered chip

imports into the United States, not exports to China ([federalregister.gov](https://www.federalregister.gov)) ([whitehouse.gov](https://www.whitehouse.gov)). China's own countermeasures, from customs blockages to outright procurement bans on state-funded projects, have compounded the effect, pushing NVIDIA's China market share from a reported 95% in 2022 toward figures its own CEO has characterized as effectively zero, even as independent IDC data suggests the erosion, while severe, has not yet been total.

The financial consequences have been real but contained within a company whose overall scale has grown enormously regardless: a \$4.5 billion inventory charge and billions more in unshippable revenue in 2025 sit alongside record company-wide quarterly revenue exceeding \$81 billion in 2026 and a market capitalization above \$5 trillion. What has changed most durably may not be any single rule but the strategic posture on both sides: Washington increasingly treats chip exports as a bargaining chip in broader trade negotiations, tying H2O relicensing explicitly to rare earth supply agreements, while Beijing accelerates a self-sufficiency drive that Huawei's rising market share and SMIC's expanding legacy-node capacity both suggest is beginning to bear fruit independent of whatever Washington ultimately decides to permit. For companies, investors, and policymakers tracking this space, the operative lesson of the past 15 months is that the rules governing NVIDIA GPU exports to China are unlikely to remain stable for long, and that enforcement gaps, whether in the Entity List, in destination-based licensing loopholes, or in third-country transshipment through Malaysia, Thailand, and Singapore, will likely matter as much as the formal thresholds themselves in determining how much U.S. AI compute actually reaches Chinese end users going forward.

Source links

Links cited in this article, in order of first appearance. Full addresses are provided for readers using extracted text.

1. <https://wccftech.com/huawei-biggest-winner-in-china-ai-market-after-nvidia-pullout-60-percent-ai-share-2026/amp/>
2. <https://nvidianews.nvidia.com/news/nvidia-announces-financial-results-for-first-quarter-fiscal-2027>
3. <https://www.congress.gov/bill/119th-congress/senate-bill/3150/text>
4. <https://www.bis.gov/press-release/commerce-further-restricts-chinas-artificial-intelligence-advanced-computing-capabilities>
5. <https://www.mayerbrown.com/en/insights/publications/2023/02/japan-and-the-netherlands-agree-to-new-restrictions-on-exports-of-chipmaking-equipment-to-china>
6. <https://companiesmarketcap.com/nvidia/marketcap/>
7. <https://www.federalregister.gov/documents/2026/01/15/2026-00789/revision-to-license-review-policy-for-advanced-computing-commodities>
8. <https://www.cnbc.com/2023/10/17/us-bans-export-of-more-ai-chips-including-nvidia-h800-to-china.html>
9. <https://www.reuters.com/technology/nvidia-resume-h20-gpu-sales-china-2025-07-15/>
10. <https://www.federalregister.gov/documents/2025/01/15/2025-00636/framework-for-artificial-intelligence-diffusion>
11. <https://introl.com/blog/ai-export-controls-navigating-chip-restrictions-globally-2025>
12. <https://www.anthropic.com/news/securing-america-s-compute-advantage-anthropic-s-position-on-the-diffusion-rule>
13. <https://www.bis.gov/press-release/departement-commerce-announces-rescission-biden-era-artificial-intelligence-diffusion-rule-strengthens>
14. <https://www.datacenterdynamics.com/en/news/trump-administration-planning-to-restrict-ai-chip-exports-to-malaysia-and-thailand-report/>

15. <https://techcrunch.com/2025/07/14/malaysia-will-require-trade-permits-for-u-s-ai-chips/>
16. <https://techcrunch.com/2025/04/30/anthropic-suggests-tweaks-to-proposed-u-s-ai-chip-export-controls/>
17. <https://www.reuters.com/world/china/chinas-smic-says-foreign-clients-shifting-orders-back-china-2026-05-15/>
18. <https://www.reuters.com/world/china/deepseek-aids-chinas-military-evaded-export-controls-us-official-says-2025-06-23/>
19. <https://www.sec.gov/Archives/edgar/data/1045810/000104581025000115/q1fy26cfocommentary.htm>
20. <https://www.techtimes.com/articles/320544/20260715/nvidia-h200-shipments-china-called-trivial-blackwell-loophole-draws-fire.htm>
21. <https://www.scmp.com/news/china/diplomacy/article/3360582/us-says-nvidias-h200-exports-china-remain-trivial-despite-approvals>
22. <https://kfgo.com/2026/07/14/us-official-says-shipments-of-h200-chips-to-china-have-begun/>
23. <https://www.theguardian.com/technology/2026/jan/17/china-blocks-nvidia-h200-ai-chips-that-us-government-cleared-for-export-report>
24. <https://www.reuters.com/world/china/us-takes-step-halt-nvidia-ai-chip-shipments-chinese-firms-outside-china-2026-05-31/>
25. <https://www.reuters.com/world/china/china-bans-foreign-ai-chips-state-funded-data-centres-sources-say-2025-11-05/>
26. <https://techcrunch.com/2025/04/15/nvidia-h20-chip-exports-hit-with-license-requirement-by-us-government/>
27. <https://www.cnn.com/2025/08/10/nvidia-amd-15percent-of-china-chip-sales-revenues-to-us-ft-reports.html>
28. <https://www.whitehouse.gov/fact-sheets/2026/01/fact-sheet-president-donald-j-trump-takes-action-on-certain-advanced-computing-chips-to-protect-americas-economic-and-national-security/>
29. <https://www.cnn.com/2026/03/19/politics/super-micro-computer-founder-charged-ai-chips-china>
30. <https://nvidianews.nvidia.com/news/nvidia-announces-financial-results-for-first-quarter-fiscal-2026>
31. <https://finance.yahoo.com/sectors/technology/article/nvidia-ceo-jensen-huang-says-company-now-has-zero-market-share-in-china-150805330.html>
32. <https://thecyberexpress.com/advanced-ai-chip-chip-security-act/>
33. <https://finance.yahoo.com/sectors/technology/articles/2-5-billion-supermicro-smuggling-142029534.html>
34. <https://www.reuters.com/world/china/chinese-chipmakers-claim-nearly-half-of-local-market-nvidias-lead-shrinks-idc-2026-04-01/>
35. <https://www.whitehouse.gov/presidential-actions/2026/01/adjusting-imports-of-semiconductors-semiconductor-manufacturing-equipment-and-their-derivative-products-into-the-united-states/>
36. <https://www.cfr.org/articles/new-ai-chip-export-policy-china-strategically-incoherent-and-unenforceable>
37. <https://www.csis.org/blogs/perspectives-innovation/gain-ai-act-will-undermine-global-competitiveness-us-ai-chipmakers>
38. <https://www.brookings.edu/articles/why-the-gain-ai-act-would-undermine-us-ai-preeminence/>
39. <https://www.cnn.com/amp/2025/05/21/nvidia-ceo-jensen-huang-slams-us-chip-restrictions-as-a-failure.html>

THE PUBLISHER

About GPU Smith

GPU Smith provides independent engineering and research for private AI infrastructure. We help technical buyers and operators reason about GPU workload sizing, hardware selection, procurement, deployment and inference operations, from the compute system to the networking, power and cooling requirements around it.

Start with the workload

Useful infrastructure decisions begin with the models, data, throughput, latency and operating constraints a team actually has. GPU Smith helps connect those requirements with system design and deployment choices. The goal is a privately controlled AI environment whose capacity and operational demands are understood before a purchasing decision.

Hardware and supplier research

Our public reference library covers GPUs, complete systems and networking components. The server-vendor directory and manufacturing research help teams investigate suppliers, compare options and follow sources behind technical claims. We distinguish manufacturer specifications from measured benchmarks and advertised capabilities from tested configurations. Reference pages are research resources, not a live inventory listing or a binding equipment quote.

Deployment and operations

GPU Smith's engineering scope includes infrastructure integration, networking, power, cooling and the practical operation of inference workloads. Our published articles explain the tradeoffs behind deployment, procurement and ongoing operations so teams can ask better questions and document decisions.

Work with GPU Smith

Explore the [hardware library](#), [GPU references](#), [networking references](#) and [vendor research](#). [Contact GPU Smith](#) with your workloads and deployment constraints to discuss a project and confirm current scope, pricing and availability.

Inclusion of a manufacturer or product in our research does not imply a partnership, certification or endorsement.

Website: <https://gpusmith.com>

This article is educational. Verify consequential medical, legal, financial, regulatory and technical decisions with appropriate professionals and the cited primary sources. Company services are described separately from the article's research findings.