NOVEMBER 2025

The Architecture of AI Leadership

Enforcement, Innovation, and Global Trust

Charles Wessner and Shruti Sharma

As global competition in artificial intelligence (AI) intensifies, the United States faces a strategic dilemma: how to protect its most advanced technologies from misuse abroad without undermining the very ecosystem that sustains its leadership. In particular, the diversion of cutting-edge AI chips to China has exposed the limits of traditional export controls. Since the Biden administration placed export controls on advanced semiconductors and related equipment in October 2022, analysts **estimate** that over 1 million intentionally downgraded semiconductor chips by the U.S. firm Nvidia have reached China through gray markets. Meanwhile, the massive Chinese technology firm Huawei has allegedly **procured** more than 2 million chips from Taiwan Semiconductor Manufacturing Company (TSMC)—which makes over **90 percent** of the world's advanced semiconductors as of 2023—via front companies, highlighting the scale and sophistication of illicit procurement networks.

In response to this, in May, Senator Tom Cotton (R-AR) introduced the bipartisan Chip Security Act, **mandating** that high-performance AI chips, **classified** under Export Control Classification Numbers 3A090 and 4A090, must include "chip security mechanisms" before export. These may include export certificates, periodic location verification via server check-ins, and software-based authentication tools. Exporters would also be required to report any credible evidence of diversion, tampering, or unauthorized use.

Yet while the bill seeks to close post-shipment blind spots, it also raises deeper questions. Can embedded technical controls substitute for institutional enforcement? Will such measures preserve global trust in U.S. technology—or erode it? And most importantly, does this path strengthen or strain the very ecosystem that has kept the United States at the forefront of AI innovation?

CSIS | CENTER FOR STRATEGIC & INTERNATIONAL STUDIES

Understanding the Chip Security Act

The Chip Security Act establishes a two-phase compliance framework: an initial requirement mandating location verification within 180 days, followed by a feasibility assessment of additional security mechanisms—such as anti-tamper features or functionality throttling—led by the Departments of Commerce and Defense. If deemed appropriate, these secondary tools could be mandated within two years. The bill also grants the Secretary of Commerce authority to monitor post-export chip usage, collect data from exporters, and recommend adjustments to export rules based on observed compliance.

Amid mounting industry concerns, proponents of the Chip Security Act have emphasized that the bill's scope has always been narrowly focused on AI data center-grade graphics processing units (GPUs)—high-performance chips specifically designed for large-scale AI workloads in data centers. Unlike consumer or gaming GPUs, these chips are optimized for tasks such as training and inference of large language models, computer vision, and generative AI. These GPUs typically support massive parallel processing, high memory bandwidth, and specialized AI acceleration features. The draft bill explicitly excludes consumer hardware and edge devices and avoids language suggesting mandatory kill switches, instead prohibiting any feature that could degrade cybersecurity or disable functionality without due process. In theory, new cybersecurity safeguards require the secretary of commerce to ensure mandated mechanisms do not introduce exploitable vulnerabilities. Additionally, the bill defers more invasive enforcement mechanisms, such as performance throttling or telemetry, pending a feasibility study and post-implementation review.

> *Policies like the Chip Security Act aim to secure leadership by limiting adversary access, but if not carefully calibrated, they risk eroding the very trust that underpins U.S. dominance in the global semiconductor ecosystem.*

However, the draft leaves several critical issues unresolved. It does not specify how location will be verified if trusted execution environments are compromised, or what enforcement measures will apply if end users spoof telemetry signals. These gaps fuel technical skepticism and political resistance, particularly from stakeholders concerned about implementation feasibility, global trust, and institutional overreach.

Preserving U.S. Leadership Means Preserving Global Trust

More broadly, as the United States tightens controls on advanced AI chips, it must also contend with a shifting perception of U.S. technology and policy in global markets. Policies like the Chip Security Act aim to secure leadership by limiting adversary access, but if not carefully calibrated, they risk eroding the very trust that underpins U.S. dominance in the global semiconductor ecosystem. If implemented prematurely, chip-level controls could inadvertently undercut the administration's broader goal of promoting U.S. technology as trusted and interoperable abroad. Added power, thermal, or latency burdens may **affect** chip performance, while perceptions of intrusive or unreliable features could

weaken confidence in U.S. products. Rather than reinforcing trust, such measures risk creating openings for foreign suppliers to present themselves as more neutral alternatives.

A relevant parallel can be seen in the defense sector, where allies such as Spain, Portugal, and Switzerland have recently walked back or reconsidered purchases of the F-35 fighter jet. Their reasons **include** concerns over U.S. control through embedded software, rising costs, and a desire for strategic autonomy, illustrating how overly restrictive terms can push even **close** partners to seek alternatives.

As efforts to tighten controls escalate, so do the incentives for global actors to diversify away from U.S.-linked supply chains. Therefore, the act may inadvertently create market openings for foreign chipmakers, particularly in Europe and Israel, who may position themselves as geopolitically neutral and free from U.S.-imposed compliance hooks. As concerns grow over embedded telemetry and U.S. oversight, governments and firms in Asia, the Middle East, and even parts of Europe may increasingly **seek** alternatives that offer greater supply chain sovereignty.

In this context, companies in countries like France, Germany, the Netherlands, and Israel stand to benefit, not necessarily because of technical superiority, but because their products may be viewed as less politically encumbered. The more U.S. chips are associated with embedded enforcement mechanisms, the more room there may be for international competitors to offer trusted alternatives, intentionally or not.

The controversy surrounding Nvidia's H20 chip in China **offers** an early example of how embedded security features can provoke diplomatic and commercial backlash. In July 2025, the Cyberspace Administration of China summoned Nvidia to address **alleged** "backdoor security risks" in its H20 chips, which were custom-designed for the Chinese market. Chinese regulators cited claims, originating from U.S. experts, that the chips could be used to track user locations or be remotely disabled.

> *The more U.S. chips are associated with embedded enforcement mechanisms, the more room there may be for international competitors to offer trusted alternatives, intentionally or not.*

The episode came just weeks after the Trump administration reversed an earlier restriction on H20 chip sales to China. Nvidia strongly denied the presence of any backdoors, but the controversy emphasized the diplomatic fragility that accompanies embedded security features, even when framed as compliance tools rather than control mechanisms.

This recent incident echoes a broader historical pattern, where security measures, however well-intentioned, risk triggering international backlash and long-term reputational damage. The Clipper chip episode of the 1990s serves as a historical warning about how security mandates can undermine trust in U.S. technology. The Clipper chip, a National Security Agency-backed initiative to **embed** encryption keys accessible to the government in consumer devices, triggered global backlash. European and Asian firms began **offering** alternatives to the Clipper chip marketed on privacy and sovereignty grounds, while U.S. tech companies saw reputational damage and commercial losses. The parallels

with the Chip Security Act are striking. A well-intentioned national security initiative can erode market confidence and provide competitors with a clear branding advantage: the ability to sell trust.

How Eroding Trust in the United States Helps China's Own Ambitions

In this environment, even narrowly targeted measures can generate outsized geopolitical consequences. China has **seized** on these dynamics to advance its domestic AI industry and portray U.S. controls as proof of the need for digital self-reliance. While the Chip Security Act targets only the most advanced data center–grade AI chips, the broader strategic and political effects are likely to extend beyond this scope. Beijing has already **used** earlier export controls to justify large-scale investment in domestic firms like Biren and Moore Threads, which now anchor its push for full-stack AI independence. Even limited tracking requirements can and will be **portrayed** by Beijing as evidence of U.S. overreach, reinforcing China's drive for indigenous alternatives and digital sovereignty and encouraging nonaligned countries to reconsider the neutrality of U.S. supply chains.

Within this shifting landscape, U.S. chipmakers such as Nvidia, Advanced Micro Devices, Inc. (AMD), and Intel have remained publicly cautious. While they have not openly opposed the Chip Security Act, they also have not endorsed its provisions. Their posture reflects broader industry sensitivity to proposals that embed compliance mechanisms directly into hardware. The primary concern is not the technical cost of implementation, but rather the strategic risks such features may **pose**, particularly in markets where sovereignty and supply chain neutrality are prioritized. For globally integrated firms, enforcement hooks that accompany the product risk recasting U.S. chips as tools of extraterritorial regulation rather than neutral infrastructure.

Location Verification as a Compliance Mechanism in Strategic Export Control

As the geopolitical context of advanced semiconductors intensifies, governments seek new ways to enforce export controls beyond traditional licensing and inspection regimes. Location verification has emerged as a proposed compliance mechanism to bridge the gap between regulatory intent and real-world enforcement capacity. However, the debate over location-based compliance sits at the intersection of strategic technology control, cybersecurity, and supply chain governance.

A PRAGMATIC SOLUTION?

Proponents of the Chip Security Act **argue** that the bill offers a pragmatic solution to persistent enforcement blind spots without imposing intrusive surveillance or remote-control capabilities. They frame the proposed mechanisms, such as export certificates embedded in firmware and periodic Transport Layer Security (TLS)-based location checks, as minimally invasive tools to help the Bureau of Industry and Security (BIS) detect illicit diversion. TLS-based location verification means **verifying** that a chip or a system is connecting from an authorized geographic location by embedding geolocation metadata or cryptographic certificates within the TLS handshake process, without exposing sensitive user data. These mechanisms, supporters argue, do not **alter** chip performance or collect user data but instead generate investigative signals that can flag suspicious behavior. By requiring tamper-resistant verification, the act aims to raise the cost of black-market smuggling, particularly through shell companies and permissive jurisdictions.

Proponents also point to parallels in other sectors, such as enterprise networking and cloud access control, where similar compliance signals support visibility without undermining functionality. They

further emphasize that the bill allows for flexible implementation through driver-level authentication or software updates, avoiding the need for invasive firmware changes.

OR A FRAGILE WORK-AROUND?

Critics, however, see the act as a fragile technical work-around masquerading as a strategy. In their view, the bill underestimates adversaries' ability to spoof or circumvent location checks while overestimating the degree of control physical chip tracking provides. Once **deployed** in multi-tenant servers or integrated into broader systems, AI chips become nearly impossible to monitor meaningfully through geolocation alone; a chip in Singapore might serve a U.S. university or a Chinese military contractor, and location tells little about actual use.

Second, the **proposed** triangulation methods are highly susceptible to spoofing via VPNs, proxies, or artificial delays, particularly from state actors like China that control network infrastructure.

Third, **embedded** telemetry features also pose cybersecurity risks: "Calling home" opens attack surfaces that could be exploited for surveillance or sabotage. Further, requiring embedded geotracking in high-performance AI chips would **impose** significant nonrecurring engineering burdens.

Fourth, the act's six-month implementation deadline compounds these challenges. Major design and manufacturing changes typically take two to three years to move from concept through validation. Premature mandates could saddle U.S. firms with costly compliance burdens while delivering little real security.

While location verification is framed as a straightforward compliance tool, its practical **feasibility** is far less certain. Advanced chips themselves lack native connectivity, and many data centers operate behind firewalls, with intermittent or no external internet access, or are intentionally air-gapped for security reasons. Further, hyperscalers also rarely run all purchased chips **simultaneously**, meaning that the absence of a signal cannot be reliably equated with diversion or tampering. Hyperscalers are technology companies that operate vast, globally distributed data centers capable of rapidly expanding computing, storage, and networking capacity to meet demand. By designing their own hardware, software, and cloud infrastructure, they deliver large-scale digital services and power emerging technologies such as artificial intelligence. Leading examples include Amazon, Microsoft, and Google, whose hyperscale operations underpin much of today's cloud and AI ecosystem.

Moreover, the "trusted server network" provisioned by the Chip Security Act does not yet exist, and building such a global infrastructure would require significant international coordination and investment. Without these prerequisites, the proposal rests on assumptions of universal connectivity that do not reflect operational realities, making enforcement inconsistent at best and misleading at worst.

> *Hardware mandates cannot replace intelligence-driven enforcement. Strengthening U.S. capacity to track GPU end-use, expand counterproliferation tools, and share information with allies is essential to disrupt diversion networks.*

Designing a leading-edge AI GPU already costs **upwards** of $500 million at the 5-nanometer (nm) node, with projections exceeding $700 million for 3 nm and below. The 5 nm and 3 nm nodes **refer** to advanced semiconductor manufacturing generations that indicate how small and densely packed a chip's transistors are, with smaller nodes enabling higher performance and efficiency but at much greater design and production cost. To **implement** location verification at scale, companies like Nvidia would need two components: a firmware update enabling rapid location checks (estimated at under $1 million) and a global network of 100–500 trusted landmark servers, costing $2.5–12.5 million annually.

Beyond technical costs, firms would face ongoing management overhead, including export-specific compliance processes, audits, incident response playbooks, and coordination across legal, engineering, and sales teams. For globally integrated companies, these burdens compound operational complexity without offering clear commercial upside. More strategically, critics warn that embedding traceability features may **erode** trust in U.S. chips, prompting customers in key markets to shift toward non-U.S. suppliers viewed as more neutral or less vulnerable to embedded controls.

Enabling Effective Enforcement Through Institutional Strength

The Chip Security Act assumes that denying China access to advanced U.S. chips will significantly slow its frontier AI progress. Yet China's ecosystem—illustrated by the rise of labs like DeepSeek—shows that controls can delay but not decisively block development. DeepSeek has **built** competitive large language models using a mix of pre-ban chips, sanctioned imports, and smuggled hardware, highlighting persistent enforcement gaps. These flows, worth billions of dollars, have **blunted** the impact of 2022–2023 export controls. Addressing this challenge requires more than technical fixes; it demands **stronger** institutional capacity at BIS, which today lacks the scale, tools, and reach to monitor chips across complex global supply chains.

> *Sustaining U.S. advantage in a contested AI landscape will require not just smarter chips, but smarter institutions empowered to monitor, investigate, and disrupt illicit flows in real time.*

Hardware mandates cannot replace intelligence-driven **enforcement**. Strengthening U.S. capacity to track GPU end-use, expand counterproliferation tools, and share information with allies is essential to disrupt diversion networks. Smart monitoring of end-use patterns will deliver more durable security than embedding technical fixes into hardware.

While the Chip Security Act attempts to close this gap by embedding location verification features into hardware—an approach that shifts part of the enforcement burden to the chip itself—technical signals alone cannot substitute for strategic enforcement infrastructure. Sustaining U.S. advantage in a contested AI landscape will require not just smarter chips, but smarter institutions empowered to monitor, investigate, and disrupt illicit flows in real time.

As BIS leadership has **emphasized** in recent testimony, there is a need for enhancing legal authority, investigative capacity, and international coordination to bolster export control mechanisms. Strengthening these pillars would allow BIS to detect and disrupt illicit supply chains through targeted investigations, interagency intelligence sharing, and bilateral enforcement partnerships more effectively. Increased funding would support the expansion of BIS's Export Enforcement and Office of Technology Evaluation, investment in advanced supply chain analytics, and more rigorous end-use checks overseas.

> *Today, adversaries no longer need to import physical chips; they can simply rent high-performance compute from U.S. cloud providers to train frontier models.*

Additional investigative personnel, including attachés embedded with foreign governments, would improve monitoring of third-country transshipment routes and shell company networks. Stronger international coordination, through data-sharing agreements and joint **investigations** with customs and export control authorities in key diversion hubs such as the United Arab Emirates, Malaysia, and Hong Kong, would help close critical enforcement gaps that no embedded hardware feature can address.

Toward a More Effective Enforcement Strategy

To sustain leadership in AI and semiconductors, the United States must focus on a broader capability-driven strategy that **advances** the goals outlined in the administration's AI Action Plan: expanding global AI alliances, countering adversary influence, and securing the AI technology stack through coordinated governance and enforcement.

Strengthening BIS and Institutional Enforcement

To uphold the **effectiveness** of export controls in a rapidly evolving technological landscape, the United States must expand the institutional capacity of the BIS. BIS is tasked with enforcing export laws, yet it remains under-resourced relative to the scale of the challenge. Strengthening enforcement requires expanding staffing, deploying attachés to key transshipment hubs, and investing in advanced supply chain analytics to track chip diversion and end-use patterns in high-risk jurisdictions. These capabilities are **essential** for identifying illicit flows, closing enforcement gaps, and reinforcing the credibility of U.S. technology controls.

- **Conduct an Independent Standards and Feasibility Review via NIST**

  Before mandating chip-level controls, the U.S. should ensure technical feasibility through an independent review led by the National Institute of Standards and Technology (NIST). As the government's standards body, NIST is best placed to validate solutions with industry and allies, reducing fragmentation and ensuring that any requirements are practical, secure, and widely trusted. A NIST-led process would also help separate credible, scalable technologies from unproven startup claims, grounding policy in independent science rather than vendor assurances. Anchoring technical standards in a transparent and collaborative framework is

essential for the United States to strengthen both the security and legitimacy of its export control system.

- **Promote Trusted AI Exports Through Voluntary Verification**

  Rather than relying solely on restrictive mandates, the United States should **encourage** voluntary end-user verification mechanisms that reward transparency and responsible behavior. This approach allows trusted firms to access U.S. technology under streamlined conditions, while reserving punitive measures for persistent noncompliance. Hence, by working with allies and partners to implement aligned standards, the U.S. can build a trusted AI export ecosystem that counters adversarial misuse without burdening legitimate innovation. This framework supports both commercial competitiveness and national security by fostering a network of reliable and accountable AI actors.

- **Shift Enforcement Focus to the Cloud**

  As AI development increasingly shifts from local infrastructure to scalable, on-demand cloud platforms, U.S. export controls must evolve accordingly. Today, adversaries no longer need to **import** physical chips; they can simply rent high-performance compute from U.S. cloud providers to train frontier models. This loophole undermines the intent of export restrictions and exposes U.S.-origin technology to strategic misuse. Without clear rules for cloud computing, efforts to constrain the diffusion of dual-use AI capabilities risk becoming obsolete.

  The goal of cloud restrictions is not to monitor, but to ensure accountability in how advanced compute is used. Just as banks are required to verify their customers to prevent illicit finance, cloud service providers should be required to verify and monitor high-risk AI workloads. Proposals such as the ENFORCE Act and H.R. 4683 **mark** important steps toward building an oversight architecture that matches the realities of modern compute delivery. These frameworks would enable regulators to detect and disrupt malicious actors without placing undue burden on trusted users.

  Done right, cloud restrictions can enhance, not inhibit, U.S. competitiveness. By establishing standards for responsible AI use in the cloud, the U.S. can build a coalition of trusted providers and democratic partners who agree on baseline safeguards.

- **Invest in Long-Term Advantage: Talent and Innovation**

  Finally, control alone is not sufficient to build lasting technological leadership; it must be matched by proactive investment in the capabilities that drive long-term advantage. Hence, to maintain leadership in AI, the United States must **invest** in the foundations of technological strength—world-class talent and universities—and fully support cutting-edge research. Export controls may impose short-term constraints on adversaries, but they cannot substitute sustained investments in innovation at competitive levels. This requires expanding funding for foundational and applied AI research, fostering public-private research and development collaboration, and building career pathways to develop and retain top technical talent in critical sectors. Just as vital is strengthening supply chain partnerships to ensure that U.S. technology remains a trusted, high-quality choice for allies and partners around the world.

> *Export controls may impose short-term constraints on adversaries, but they cannot substitute sustained investments in innovation at competitive levels.*

Conclusion

The Chip Security Act seeks to address real vulnerabilities in the current export control system, but embedding compliance into hardware is not a substitute for strategic enforcement. Many of the act's proposed technical mechanisms still require clarification—on feasibility, interoperability, and security—to ensure they do not introduce new risks or hinder competitiveness. As AI development moves to the cloud and global competition intensifies, U.S. leadership will hinge less on restrictive mechanisms and more on continued innovation backed by resilient institutions, credible governance, and trusted supply chains.

To stay ahead, the United States must modernize its enforcement architecture, bolster BIS capacity, extend oversight to virtual compute, and coordinate with allies to align standards. Export controls should enable innovation, not constrain it, and above all, not erode confidence and acceptance of U.S. high-tech manufacturing equipment and products. By investing in capability, not just control, the United States can preserve its edge in AI infrastructure while upholding the trust that sustains its global leadership. ■

Charles Wessner *is a senior adviser (non-resident) with Renewing American Innovation at the Center for Strategic and International Studies (CSIS) in Washington, D.C.* ***Shruti Sharma*** *is a research intern with Renewing American Innovation at CSIS.*

This work is made possible by general support to CSIS.