Center for AI Safety

ACTION FUND

The Center for AI Safety (CAIS) Action Fund welcomes this opportunity to comment on policy actions for inclusion in the new AI Action Plan ("Plan"). CAIS Action Fund is a nonpartisan advocacy organization dedicated to advancing policies that maintain U.S. leadership in AI and defend against AI-related threats to our national security.

Our response focuses on three interconnected pillars that should form the foundation of a robust AI Action Plan: implementing a comprehensive **nonproliferation** strategy to prevent AI chips from reaching rogue actors, **deterring** strategic competitors from pursuing destabilizing AI projects, and strengthening U.S. AI **competitiveness**.

Nonproliferation measures that prevent rogue actors from obtaining AI chips are critical for national security. The Administration should: 1) establish a comprehensive licensing and notification regime for export-controlled AI chips, 2) implement location verification technologies to combat chip smuggling, 3) invest in next-generation hardware security for AI chips, and 4) reposition the AI Safety Institute from NIST to BIS to focus on export control enforcement.

Deterrence is vital to prevent strategic rivals from pursuing destabilizing AI projects. The Administration should: 1) strengthen intelligence collection on foreign frontier AI capabilities, 2) develop offensive cyber capabilities to credibly deter rivals from pursuing high-risk AI development programs, and 3) implement stronger reporting requirements for developers of frontier AI models to enhance government visibility into emerging capabilities and potential risks.

Competitiveness in AI technologies will increasingly underpin America's economic and national security. The Administration should: 1) accelerate domestic AI chip production through a comprehensive strategy that could include targeted government subsidies, tax incentives, streamlined regulatory frameworks, and strategic tariffs; 2) streamline immigration pathways for top AI talent, and 3) secure guaranteed drone and robotics supply chains. By implementing these recommendations across all three pillars, the United States can protect America's economic and national security and enhance our global AI leadership.[1]

[1] This document is approved for public dissemination. The document contains no business-proprietary or confidential information. Document contents may be reused by the government in developing the AI Action Plan and associated documents without attribution.

TABLE OF CONTENTS

PROPOSALS

1. Implement a Nonproliferation Strategy to Secure the AI Chip Supply Chain

Nonproliferation measures are critical to prevent rogue actors from obtaining computational resources ("compute"). Advanced AI chips form the foundation of AI development and offer an ideal governance point because they are physical, trackable, and quantifiable.[2] Yet current control measures are failing to prevent widespread chip smuggling to strategic competitors and rogue actors, with sophisticated operations worth hundreds of millions of dollars evading controls.[3] These illicit networks operate with increasing sophistication, creating vulnerabilities that extend beyond great power competition to include potential acquisition by rogue states and non-state actors.[4]

The Plan should implement a robust nonproliferation strategy for AI chips spanning four key elements: 1) a comprehensive licensing and notification regime, 2) location verification technologies, 3) research for next-generation hardware security, and 4) strategic repositioning of the AI Safety Institute to strengthen enforcement capabilities. Together, these measures will help prevent critical computational components from falling into unauthorized hands while enabling legitimate AI development.

Recommendation 1.1: Implement a Comprehensive Licensing and Notification Regime for AI Chips

Chip smuggling has reached crisis levels, demanding a new tracking system to close intelligence gaps about the location and movement of advanced AI chips. The Plan should direct **BIS to establish a comprehensive licensing and notification regime for advanced AI chips**. BIS currently lacks reliable information about the locations of export-controlled AI chips.[5] Intelligence and market analysis reveal chip smuggling at an alarming scale, with estimates suggesting over 100,000 export-controlled GPUs—potentially as high as one

[2] Dan Hendrycks, *Compute Governance* in *Introduction to Compute Governance*, *AI Safety, Ethics, and Society* (2024); Girish Sastry et al., *Computing Power and the Governance of Artificial Intelligence* (Feb. 13, 2024).
[3] Erich Grunewald, *AI Chip Smuggling Is the Default, Not the Exception*, AI Policy Bulletin (Mar. 3, 2025).
[4] Dan Hendrycks, Eric Schmidt & Alexandr Wang, *Nonproliferation*, in *Superintelligence Strategy* (Mar. 2025).
[5] Tim Fist & Erich Grunewald, *Preventing AI Chip Smuggling to China*, Ctr. for a New Am. Sec. (Oct. 24, 2023).

million—smuggled into China last year alone.[6] Market data confirms this problem: Singapore accounted for 22% of Nvidia's revenue in a recent quarterly statement, despite Nvidia acknowledging most shipments ultimately went to users outside Singapore.[7] A robust tracking system would enable authorities to monitor these devices and quickly identify potential diversion attempts.

This system's foundation would be a new license requirement for all direct exporters of high-performance AI chips, regardless of destination. Under this framework, exporters would apply for licenses that identify the specific chips, their recipients, and intended end-use, with exports to trusted partners benefiting from a "presumption of approval" approach.[8] Entities with strong compliance records might qualify for streamlined processes but would still need to notify authorities of every resale or relocation. Reexporters and those conducting in-country transfers would follow a notification protocol to maintain the chain of custody information.[9] This approach creates an effective reporting mechanism by making data submission a condition of export authorization, similar to existing post-shipment verification reporting requirements for certain high-performance computers.[10] Because this system builds on familiar infrastructure, it can be implemented swiftly, enabling officials to track chips without stalling legitimate commerce.

With this licensing and notification regime in place, BIS could pilot a strategic chip inspection program using random sampling methods.[11] The program could operate through short-notice mail-in inspections to regional U.S. Commercial Service Offices, minimizing costs while leveraging existing maintenance processes in large data centers where most controlled chips are deployed. BIS personnel would inspect chips for ID matches and evidence of tampering, returning chips to owners within days. This system would make large-scale AI chip smuggling more difficult to sustain and provide early warning of smuggling activities.

---

[6] Grunewald, *AI Chip Smuggling is the Default*.
[7] Letter from John Moolenaar, Chairman, & Raja Krishnamoorthi, Ranking Member, H. Select Comm. on the Chinese Communist Party, to Michael Waltz, Nat'l Sec. Advisor (Jan. 29, 2025).
[8] Erich Grunewald & Michael Aird, *AI Chip Smuggling into China: Potential Paths, Quantities, and Countermeasures*, Inst. for AI Pol'y & Strategy (Oct. 4, 2023).
[9] Deric Cheng, *Evaluating an AI Chip Registration Policy*, Convergence Analysis (Apr. 8, 2024).
[10] Fist & Grunewald, *Preventing AI Chip Smuggling to China*.
[11] Fist & Grunewald, *Preventing AI Chip Smuggling to China*.

Recommendation 1.2: Require Location Verification for Advanced AI Chips

To maximize effectiveness, tracking capabilities should complement this licensing framework. The Plan should direct **BIS to require chipmakers to implement geolocation functionality** for advanced AI chips and require this functionality as an export license condition for advanced AI chips.[12] One promising approach to location verification measures round-trip communication time between chips and trusted landmark servers to determine a chip's location.[13] This method leverages existing cryptographic features in current-generation chips, such as Trusted Platform Modules, rendering it feasible to implement in the short term without requiring new chip designs.[14] Implementing location verification in this way would require establishing a network of landmark servers with verified locations strategically deployed across relevant countries, with increased density near restricted nations' borders to distinguish between permitted and restricted locations.[15] These landmarks measure delays in communication with AI chips, and convert these measurements into distance estimates using calibrated delay-to-distance mappings. By triangulating from multiple landmarks, this system establishes whether a chip remains within its authorized location. Each AI chip uses a secure digital identity with signed certificates issued during manufacturing for authentication with landmark servers, enabling encrypted communications and linking location data to specific chips. Research indicates that under typical network conditions, a delay-based geolocation system could verify location to within 100 kilometers—generally sufficient to determine whether a chip remains in its authorized country.[16]

The delay-based approach offers significant security advantages over alternatives. It inherently defends against common evasion tactics like false location reporting and VPN usage, as these typically increase network delays rather than help strategic competitors appear to be in permitted countries. Unlike GPS technology which can be spoofed for as little as $200, delay-based verification is significantly more resistant to manipulation.[17] Even sophisticated

---

[12] Asher Brass & Onnie Aarne, *Location Verification for AI Chips,* Institute for AI Policy and Strategy (May 6, 2024).
[13] Onni Aarne, Tim Fist & Caleb Withers, *Secure, Governable Chips: Using On-Chip Mechanisms to Manage National Security Risks from AI & Advanced Computing*, Ctr. for a New Am. Sec. (Jan. 8, 2024).
[14] Tim Fist, Tao Burga & Vivek Chilukuri, *Technology to Secure the AI Chip Supply Chain: A Primer*, Ctr. for a New Am. Sec. (Dec. 11, 2024).
[15] Brass & Aarne, *Location Verification for AI Chips*.
[16] *Id.*
[17] *Id.*

attacks using specialized infrastructure to manipulate delays face significant practical limitations and can be detected through secondary validation techniques such as network path analysis and ISP verification. In more advanced implementations, chips could be configured to deactivate if unable to verify they remain in authorized locations.

This location verification system would work hand-in-hand with the licensing and notification framework described earlier. Combined with a reporting requirement that tracks ownership and expected locations, suspicious movements become easier to detect and investigate. These integrated approaches would create a robust framework for securing the AI chip supply chain.

Recommendation 1.3: Develop Next-Generation Hardware-Enabled Mechanisms

While delay-based location verification represents a hardware-enabled mechanism (HEM) that can be implemented in the near term, the Plan should also direct **DARPA, IARPA, or another appropriate agency** to establish an R&D program to develop additional HEMs for more sophisticated governance of advanced AI chips in the future.[18] Modern AI chips such as the NVIDIA H100 already feature privacy-preserving corporate security measures like trusted execution environments, which can be adapted for national security purposes.

This R&D program should prioritize several promising HEMs, including AI chip interconnection controls to limit the use of chips in large-scale training clusters,[19] workload classification and measurement systems to enable verifiable measures of AI training runs,[20] offline licensing and remote attestation mechanisms that require regular validation from a remote trusted source for operational approval,[21] tamper-resistant enclosures with sensors to monitor electrical properties,[22] physical unclonable functions (PUFs) that leverage

[18] Fist, Burga & Chilukuri, *Technology to Secure the AI Chip Supply Chain*.
[19] Gabriel Kulp et al., *Hardware-Enabled Governance Mechanisms: Developing Technical Solutions to Exempt Items Otherwise Classified Under Export Control Classification Numbers 3A090 and 4A090*, RAND Corp. (Jan. 18, 2024).
[20] Yonadav Shavit, *What Does It Take to Catch a Chinchilla? Verifying Rules on Large-Scale Neural Network Training via Compute Monitoring* (Mar. 20, 2023); Lennart Heim et al., *Governing Through the Cloud: The Intermediary Role of Compute Providers in AI Regulation* (Mar. 13, 2024); Aarne, Fist & Withers, *Secure, Governable Chips*.
[21] James Petrie, *Near-Term Enforcement of AI Chip Export Controls Using A Firmware-Based Design for Offline Licensing* (May 28, 2024); Kulp et al., *Hardware-Enabled Governance Mechanisms*; Aarne, Fist & Withers, *Secure, Governable Chips*.
[22] Fist, Burga & Chilukuri, *Technology to Secure the AI Chip Supply Chain*.

manufacturing variations to create unique cryptographic identities,[23] hardware-level authentication systems that ensure chips only operate when properly authorized,[24] hardware-level IP protection and usage control systems that ensure model weights are only decryptable by authorized chips,[25] and secure hardware modules for workload verification that can track and report metrics relevant to safety regulations.[26] By investing in this research now, the United States can develop the technical foundations needed for next-generation export controls that address both immediate and long-term risks of AI chip proliferation.

Recommendation 1.4: Transform Export Control Enforcement by Repositioning the AI Safety Institute

Effective enforcement provides the critical foundation for any AI chip nonproliferation strategy. To transform our enforcement capabilities, the Plan should direct the **Department of Commerce to relocate the U.S. AI Safety Institute from NIST to BIS**, refocusing its mission on export control enforcement for AI chips and national security. This strategic reorganization would create a dedicated team of AI chip export control enforcement officers dedicated within BIS.

As the central hub for AI export control expertise, the repositioned Institute could lead a comprehensive enforcement strategy through several interconnected functions. The Institute should conduct more thorough end-use checks in high-risk regions where diversion frequently occurs and implement new cost-effective verification methods, including tamper-evident cameras in AI chip data centers to detect suspicious activity. Institute officers should also collaborate with BIS policy teams to identify necessary export control expansions, such as controls on advanced AI chips with integrated HBM like the NVIDIA H20, which enable deployment of state-of-the-art reasoning models that could pose risks if misused.

The Institute should enforce verified decommissioning protocols for non-functional or outdated AI chips. Similar to disposal requirements for nuclear or chemical materials, these protocols would ensure controlled hardware is properly destroyed or disabled, preventing supposedly "retired" chips from being repurposed for unauthorized use. Enforcement officers could verify proper decommissioning through on-site inspections. The Institute should also

[23] Aarne, Fist & Withers, *Secure, Governable Chips*.
[24] *Id.*
[25] Fist, Burga & Chilukuri, *Technology to Secure the AI Chip Supply Chain*.
[26] *Id.*

enforce harsher penalties for export violations. The largest administrative penalty in BIS history—$300 million—pales in comparison to the global annual AI chip market worth tens of billions of dollars with substantial profit margins.[27] Given these economics, companies may treat violations as acceptable business risks rather than serious legal infractions. The Institute should enforce meaningful penalties for violations to levels that genuinely threaten company profits and hold companies responsible for "knowing" violations when they fail to investigate clear signs of diversion.[28]

Beyond its direct enforcement responsibilities, the Institute should coordinate a whole-of-government approach to enforcing export controls. This should include strengthening intelligence community support for AI hardware tracking, revitalizing Cold War-era capabilities that have atrophied but remain essential for tracking sophisticated procurement efforts and measuring the effectiveness of our control regimes.[29] By centralizing enforcement expertise in the repositioned Institute, the Plan would create a more robust nonproliferation framework for advanced AI technologies.

2. Deter Strategic Competitors From Pursuing Destabilizing AI Projects

The United States must **deter** rival states from pursuing an AI-enabled strategic monopoly prevent potentially catastrophic outcomes.[30] Frontier AI systems could enable strategic competitors to develop weapons of mass destruction, launch sophisticated cyberattacks against critical infrastructure, or gain decisive military advantages. No nation will passively accept a rival's pursuit of strategic AI monopoly that could threaten national survival. Without early detection and effective countermeasures, a rival's race toward decisive AI advantage could force impossible choices between acceptance of strategic defeat or dangerous escalation.

To operationalize a credible **deterrence** framework, the Plan should adopt a three-pronged approach: 1) strengthening intelligence on strategic competitors' AI capabilities,

---

[27] Gregory C. Allen, *Understanding the Biden Administration's Updated Export Controls*, Ctr. for Strategic & Int'l Stud. (Dec. 11, 2024).

[28] Majority Staff of S. Permanent Subcomm. on Investigations, 118th Cong., *The U.S. Technology Fueling Russia's War in Ukraine: Examining the Bureau of Industry and Security's Enforcement of Semiconductor Export Controls* (Dec. 18, 2024).

[29] Gregory C. Allen, *DeepSeek, Huawei, Export Controls, and the Future of the U.S.-China AI Race*, Ctr. for Strategic & Int'l Stud. (Mar. 7, 2025).

[30] Hendrycks, Schmidt & Wang, *Deterrence with Mutual Assured AI Malfunction (MAIM)* in *Superintelligence Strategy*.

2) developing counter-AI operations to deter high-risk foreign projects, and 3) expanding reporting requirements for U.S. frontier AI developers to enhance government visibility into emerging capabilities. These measures would help the United States credibly deter our adversaries while minimizing the risk of unintended escalation.

Recommendation 2.1: Strengthen Intelligence Collection on Foreign Frontier AI Development

Rapidly evolving AI technologies make robust intelligence on strategic competitors' capabilities essential for national security. Without this insight, the United States faces risks of technological surprise that could undermine our strategic position. Our goal should be to **increase collection and analysis on foreign AI capabilities**.[31] To achieve that goal, the Plan should direct intelligence agencies to double personnel devoted to studying foreign AI capabilities.[32] This investment in human capital should be accompanied by structural changes in prioritization. The DNI to elevate AI intelligence in the National Intelligence Priorities Framework and establish dedicated AI analyst teams to assess how strategic competitors integrate AI into military and intelligence operations.[33] Enhanced intelligence on foreign AI development will strengthen decision-making across government. With deeper understanding of competitors' AI progress, U.S. leaders can make more informed decisions about investments and coordinate effectively with allies to maintain democratic leadership in these critical technologies.

Recommendation 2.2: Expand Offensive Cyber Capabilities to Credibly Deter Strategic Competitors' Destabilizing AI Projects

Intelligence collection provides a necessary foundation for national security, but it must be paired with an ability to act. For deterrence to function effectively, the United States must develop and signal its capability to target destabilizing AI projects. While espionage identifies threats, operational capabilities make deterrence more credible. AI development centers present ideal targets for cyber operations due to their inherent vulnerabilities because intense

[31] Special Competitive Studies Project, *Intelligence Innovation: Repositioning for Future Technology Competition* (Apr. 2024).
[32] CSIS Technology and Intelligence Task Force, *Maintaining the Intelligence Edge: Reimagining and Reinventing Intelligence through Innovation* (Jan. 13, 2021).
[33] Special Competitive Studies Project, *Intelligence Innovation*.

computational requirements create dependencies on vulnerable cooling and power systems. By developing targeted cyber capabilities to exploit these vulnerabilities, the United States can ensure strategic competitors understand that destabilizing AI monopoly attempts will face effective countermeasures with minimal collateral damage or escalation risks.

Operationalizing this approach requires institutional commitment and technical expertise. The Plan should direct U.S. Cyber Command to develop a comprehensive suite of capabilities for selectively impeding strategic competitors' high-risk AI projects. The Plan should direct **U.S. Cyber Command to develop a comprehensive suite of capabilities for selectively impeding strategic competitors' high-risk AI projects**.[34] Cyber operators might target the computational infrastructure supporting AI development by exploiting vulnerabilities in cooling systems, power supplies, and software that manages GPU failures during critical training runs. A clear governance framework should complement these technical capabilities. The Administration should establish and publicize standards for determining when an adversary's AI project represents a sufficiently serious threat to merit consideration of counter-AI operations.[35] This framework should distinguish between destabilizing AI projects and acceptable use cases, creating a basis for proportionate responses while reducing risks of unintended escalation.

Rather than simply threatening cyber intervention, the United States should leverage this capability to demand transparency from strategic competitors. This could include inspection protocols similar in spirit to the Open Skies Treaty, which employed unarmed overflights to demonstrate that neither side was hiding missile deployments. This approach balances security with stability in global AI development by preventing unnecessary disruption to everyday AI services and reducing the risk of indiscriminate sabotage.

Recommendation 2.3: Strengthen Reporting Requirements for Frontier AI Developers

Effective deterrence requires symmetric application of transparency principles. While we develop capabilities to monitor competitors' AI projects, we must ensure our own developers maintain appropriate transparency with the government. The federal government

[34] *Cf.* Special Competitive Studies Project, *Artificial General Intelligence* (2025) (calling for the NSC to "establish an AGI attack framework").

[35] Gary Corn & Eric Talbot Jensen, *Attacking Big Data: Strategic Competition in the Race for AI through Cyber Sabotage*, Lieber Inst. West Point (Feb. 8, 2024).

currently has limited insight into frontier AI systems being developed by U.S. companies, creating dangerous national security blind spots. China likely has better visibility into AI capabilities through its comprehensive monitoring of domestic companies and intensive intelligence collection targeting U.S. firms. This information deficit undermines our deterrence posture by preventing our government from accurately calibrating our responses to foreign threats.

Strong reporting requirements can close this critical information gap. The Plan should direct **BIS to strengthen its proposed rule establishing reporting requirements for developers of frontier AI models.**[36] The new rule should establish a comprehensive reporting framework requiring companies to document ongoing and planned development activities, cybersecurity measures, model ownership structures, results of safety evaluations, and mitigation strategies for identified risks. BIS should extend reporting requirements beyond quarterly submissions to include prompt disclosure of unforeseen system behaviors that may pose security risks, particularly those that could lower barriers to developing chemical, biological, radiological, or nuclear weapons or facilitate sophisticated cyberattacks against critical infrastructure.[37] At minimum, this reporting should cover: 1) training compute used (measured in floating-point operations), 2) key benchmark performance metrics for capability and national security evaluations, and 3) implementation of risk management frameworks and security protocols.[38]

Beyond formal corporate reporting, individual researchers need secure channels to report concerns. The Plan should direct **Commerce to consider creating an AI whistleblower hotline** for AI company employees to securely report dangerous capabilities or misrepresented safety information. This multi-layered reporting system would provide government with crucial early warnings about emerging AI risks, enabling more targeted regulatory interventions and fostering improved industry-government coordination on safety measures.

---

[36] *Establishment of Reporting Requirements for the Development of Advanced Artificial Intelligence Models and Computing Clusters*, 89 Fed. Reg. 73,612 (proposed Sept. 11, 2024) (to be codified at 15 C.F.R. pt. 702).
[37] Noam Kolt et al., *Responsible Reporting for Frontier AI Development* (Apr. 3, 2024).
[38] Helen Toner & Timothy Fist, *Regulating the AI Frontier: Design Choices and Constraints*, Ctr. for Sec. & Emerging Tech. (Oct. 26, 2023).

3. Strengthen U.S. AI Competitiveness

Competitiveness in AI technologies will be increasingly important for America's economic security as AI capabilities rapidly advance. The United States must secure critical supply chains while accelerating AI innovation to maintain our edge over strategic competitors. Ensuring our competitiveness requires addressing three priorities that form the foundation of AI leadership: 1) accelerating domestic AI chip production to reduce vulnerabilities from Taiwan dependence, 2) implementing streamlined immigration pathways for top AI talent who drive innovation, and 3) securing guaranteed drone and robotics supply chains to support military applications of AI technologies on future battlefields. These interconnected measures will strengthen America's economic position moving forward.

Recommendation 3.1: Accelerate Domestic AI Chip Production

Securing a domestic supply of advanced AI chips has become critical for national security and technological leadership. The strategic importance of domestic production will intensify as AI systems deliver measurable economic value and national power becomes directly linked to AI chip access. In times of crisis, only domestic manufacturing can guarantee this critical supply. The United States currently leads in chip design but relies heavily on overseas manufacturing, particularly in Taiwan—which China has threatened to annex by force.[39] This dependence creates serious vulnerabilities in our AI supply chain that demand urgent action. Analysts estimate a double-digit probability of Chinese invasion of Taiwan within the next decade, which would severely disrupt global AI chip supplies and grant China a decisive advantage in AI capabilities.

The Administration should **adopt a comprehensive strategy to secure domestic AI chip production** that could include multiple complementary approaches. Domestic production entails higher costs, but targeted government subsidies could bridge this gap through government spending or tax incentives. The strategy could also include streamlining regulatory frameworks for energy-intensive manufacturing facilities. Strategic tariffs on Taiwanese semiconductors could further incentivize reshoring of critical manufacturing capabilities. By strengthening domestic capabilities in AI chip production through these

[39] Hendrycks, Schmidt & Wang, *Economic Security* in *Superintelligence Strategy*.

coordinated measures, the United States can enhance its competitive position and build resilience against foreseeable supply chain disruptions.

Recommendation 3.2: Implement Streamlined Immigration Pathways for Top AI Talent

Human capital forms the foundation of America's technological edge in AI. Yet this foundation is eroding as immigration barriers impede our ability to attract and retain top talent. The United States' AI leadership depends significantly on exceptional scientists from abroad, similar to how immigrant scientists contributed to the Manhattan Project. Recent data shows this talent pipeline is at risk: 60% of non-citizen AI PhDs working in the United States report significant immigration difficulties, with many indicating they are more likely to leave. [40] This talent leakage threatens to undermine our AI leadership.

While the U.S. maintains outdated policies, other countries have implemented aggressive talent attraction programs. Canada now processes skilled worker permits in as little as two weeks, and the UK has exempted PhD-level occupations from visa caps and introduced a Global Talent visa to attract STEM talent.[41] In contrast, America's high-skilled immigration policies remain largely unchanged for decades, creating a clear disadvantage in the global competition for AI expertise. The Administration should prioritize specialized immigration pathways for AI scientists—distinct from broader immigration reforms or southern border policy—to focus on an area increasingly vital to our national security and economic strength.

Although comprehensive immigration reform requires congressional action, the Administration can take immediate steps to address talent challenges. **Modernizing the O-1 "extraordinary ability" visa criteria** would better accommodate AI researchers and entrepreneurs without requiring new legislation. [42] Current criteria for demonstrating "extraordinary ability" fail to recognize achievements specific to emerging fields like AI, disadvantaging those with unconventional backgrounds but exceptional capabilities. By updating these criteria, more top AI talent could qualify for this uncapped visa category

[40] Catherine Aiken, James Dunham & Remco Zwetsloot, *Immigration Pathways and Plans of AI Talent*, Ctr. for Sec. & Emerging Tech. (Sept. 2020).
[41] Tina Huang & Zachary Arnold, *Immigration Policy and the Global Competition for AI Talent*, Ctr. for Sec. & Emerging Tech. (June 2020).
[42] *Id.*

without requiring legislative change. Through these targeted reforms, the United States can improve its ability to attract and retain the world's best AI talent.

Recommendation 3.3: Secure Guaranteed Drone and Robotics Supply Chains

America's battlefield capabilities face significant risk due to dependency on foreign components for drones and advanced robotics. Chinese manufacturers dominate the global market for key drone and robotic components, creating vulnerabilities that could impair military readiness during conflict. This overreliance on foreign supply lines puts our strategic position at risk, particularly as autonomous systems become central to battlefield operations.

The Plan should prioritize **securing reliable drone and robotics supply chains** to ensure American forces have uninterrupted access to these critical technologies. History demonstrates the dangers of technological dependence—even pioneering states can lose their advantage when unable to manufacture and deploy new capabilities at scale. Just as Britain's early lead in tank warfare was ultimately overcome by Germany's superior industrial integration, America must ensure its production capacity matches its technological capabilities. To address these challenges, the Administration should **implement incentives for production** of essential robotic and drone components and **establish robust supply chain security partnerships with trusted allies**. These actions would reduce vulnerabilities while maintaining access to global innovation. Even cutting-edge AI advancements provide limited military value without the physical platforms needed to deploy them effectively.

The rapid proliferation of autonomous systems creates additional security concerns of unintended escalation. The sheer volume and autonomy of robotics and drones can drive conflicts into unintended terrain if they approach disputed lines or misread ambiguous signals. To mitigate these risks, the Plan should encourage the development of **confidence-building measures between major powers**. These could include crisis communication channels, information exchange protocols, and operational guidelines for autonomous system activities in contested areas. Human oversight remains essential as AI integration into military command and control systems expands. AI can enhance battlefield decision-making by processing vast quantities of data, but requires meaningful human oversight of key military decisions. The Administration should develop **frameworks that ensure human approval of escalatory actions** while allowing AI to support lower-level tactical operations.

Strategic advantage in modern warfare depends on both technological innovation and industrial capacity. By securing robotics and drone supply chains and establishing mechanisms to prevent unintended escalation, the United States can strengthen its position in this critical domain. Large-scale production and responsible deployment of unmanned systems have become strategic imperatives for maintaining military advantage in the 21st century.

CONCLUSION

The stakes could not be higher as AI reshapes the global security landscape. Our response has outlined a coordinated strategy built on three pillars that together create a robust framework for American leadership in AI.

Our **nonproliferation** recommendations establish a sophisticated control framework for AI chips, from licensing and geolocation verification to next-generation hardware security mechanisms and enhanced enforcement capabilities. This approach prevents critical computing resources from falling into unauthorized hands while maintaining legitimate access for responsible users.

For **deterrence** to be effective, the Plan must pair enhanced intelligence capabilities with operational tools that can counter destabilizing AI projects. The proposed approach combines strengthened intelligence collection with targeted cyber capabilities and improved domestic reporting requirements, creating credible consequences for dangerous AI development while illuminating potential risks.

Securing America's **competitiveness** requires investments in both hardware and human capital. By developing domestic AI chip manufacturing capabilities, modernizing immigration pathways for AI talent, and securing supply chains for AI-enabled systems, the United States can reduce dangerous dependencies while maintaining our technological edge. By taking decisive action across all three pillars, the Administration can protect America's economic and national security and enhance our global AI leadership.