

Quick Highlights from “Tracking AI Chips: Assessing Location Verification as an Export Control Enforcement Tool”

By Jacob Feldgoise, Kyle Miller, and Hanna Dohmen

Chip smuggling undermines U.S. export controls that are designed to prevent adversaries from developing and deploying advanced AI capabilities. Location verification—which would help regulators determine where AI chips are located—has emerged as a key proposal to strengthen export control enforcement efforts.

If policymakers pursue location verification, regulators should initially prescribe a flexible location verification approach, where each company can implement their preferred location verification method to allow industry to develop, test, and evaluate solutions. After a testing period, the Bureau of Industry and Security (BIS) can mandate one or more approaches that it deems most suitable, or it can maintain a flexible approach.

A ping-based location verification (PLV) system—supplemented by physical inspections to investigate leads—would offer the most robust and feasible approach to location verification.

- We evaluated seven location verification methods based on four criteria: verifiability, accuracy, security, and repeatability. Of these methods, PLV and physical inspections were the only two that satisfied all four criteria.
- PLV—which estimates chip locations by measuring the time it takes for a signal to travel from a chip to multiple trusted “landmark” servers and back—is more robust than other geolocation mechanisms (i.e., topology-based and GPS-based), inventory management tools, mail-in inspections, and video inspections.
- Relative to physical inspections, PLV is more cost-effective and scales better in the long term as more chips are deployed. Across more than ten million simulated scenarios, physical inspections never achieved higher detections-per-cost than a PLV system, whether the landmark servers are owned or rented.
- A global PLV system would likely require between 39 and 500 landmark servers. With 500 landmark servers, a PLV system would likely achieve region-level accuracy (worst case) to city-level accuracy (best case).
- By positioning one landmark server in every major data center hub—defined as an area containing a large concentration of data centers within a 19-mile radius—outside the United States, China, and Russia, PLV would likely achieve country-level accuracy (a radius of 116 miles) for the vast majority of PLV-tracked chips.

A PLV system could be deployed for a feasible but non-trivial cost, but costs vary widely depending on implementation.

- We estimate that a PLV system would cost roughly \$3.1 million to \$28.8 million annually if the landmark servers are owned and \$2.6 million to \$72.4 million annually if they are rented.
- The potentially government-borne costs would represent 0.5% to 7.2% of BIS's FY2026 budget if landmark servers are owned or 0.3% to 25.7% if they are rented.
- Before implementing a PLV system, policymakers should ensure that BIS has sufficient resources and staffing to be able to effectively use the information that such a system would provide.

Any PLV system will have significant limitations. Most importantly, a PLV system can only detect diversion after it has already occurred. We identify five additional limitations:

- Enabling a PLV system on AI chips in circulation: While PLV capability could be added to AI chips that have already been manufactured and are in circulation, doing so would require compliance by end users who would need to install a software or firmware update.
- Connectivity between chips and landmark servers: When a controlled AI chip fails to respond to a PLV system, the system cannot know whether it is due to diversion or legitimate reasons, introducing ambiguity into the system.
- Adversarial location obfuscation: Adversaries could trick a PLV system into reporting that a smuggled chip is still in a permitted location.
- Approximate geolocation of AI chips: Since the best-case accuracy of a PLV system would very likely be no better than about 20 miles, the system cannot reliably distinguish locations between nearby data centers operated by different entities. This may be particularly problematic in areas with a high density of data centers, as well as border regions near restricted end locations such as China.
- End-user obfuscation: Because PLV systems verify locations rather than end-user identities, the system cannot detect diversion to an unauthorized user where the chip's physical location remains unchanged.

Download the report: <https://cset.georgetown.edu/publication/tracking-ai-chips/>

For questions and feedback, or to contact the authors, please write to us at cset@georgetown.edu.