

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

By Jacob Feldgoise, Kyle Miller, and Hanna Dohmen

Large-scale AI 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.

[Our analysis](#) suggests that location verification could be implemented for a feasible but non-trivial cost and would enhance U.S. export control enforcement efforts by generating leads and imposing costs on chip smugglers. These measures may also have a helpful deterrent effect. However, location verification has significant shortcomings that limit its effectiveness and ultimate net value; policymakers should understand these factors to make informed judgments about whether and how to pursue location verification.

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

Overall Policy Design

If policymakers pursue location verification, regulators should initially prescribe a flexible approach, where each responsible entity 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.

To mitigate the potential for noncompliance, regulators should centralize responsibility for reporting when a given chip is outside of its intended location.

- A single entity (i.e., BIS itself, the chip designers, or trusted third parties) should be responsible for knowing the intended location of each chip and verifying its location, even if the chip changes hands multiple times.

Best Option: Ping-Based Location Verification Systems Supplemented by Physical Inspections

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. Pairing the two provides the strongest detection and enforcement capability.
- 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.
- PLV generates leads by identifying chips that are not responding or are outside of their intended locations; it does not directly stop smuggling. PLV also cannot track chips that are in transit, storage, or otherwise not activated. Therefore, physical inspections should supplement PLV as a follow-on enforcement mechanism.

Table 1. Comparison of Location Verification Methods

	Ping-Based Geolocation	Physical Inspections	Topology- Based Geolocation	GPS-Based Geolocation	Inventory Management Tools	Mail-In Inspections	Video Inspections
Verifiable	✓ ^B	✓ ^B	✓ ^B	✓ ^B	✗	~	✗
Accurate	✓	✓	✓	✓	✗	✗	✗
Secure	✓ ^B	✓ ^B	✗	✗	✗	✗ ^E	✗
Repeatable	✓	~ ^D	✓	✓	✓	~ ^D	~ ^D
Low Cost	~ ^C	✗ ^C	✓	~ ^C	✓	~	~
Unintrusive ^A	✓	✗	✓	✓	✓	✗	✗

Notes:

(A) Intrusiveness is the degree of interference in the operation of data center servers.

(B) To be verifiable, in addition to estimating an AI chip’s physical location, these methods must also produce evidence—preferably involving cryptographic attestation—that establishes whether the chip being inspected is the one that investigators intend to inspect.

(C) For physical inspections, the cost primarily depends on the number of data centers and chips inspected; for ping-based geolocation, the cost primarily depends on the number of landmark servers; for GPS-based geolocation, the cost primarily depends on the type and amount of GPS hardware.

(D) While these partially repeatable location verification methods meet the “repeatable” criterion, it is infeasible for each method to physically inspect, mail in, or video inspect each chip many times per year.


(E) Mail-in inspections are not secure because the origin of the shipment can be obfuscated. For example, chips could be shipped from export-controlled countries within relatively short time periods using express shipping.

Source: CSET.

A PLV system measures the round-trip time (RTT) of a signal traveling from a chip to multiple trusted “landmark” servers and back, which is used to “trilaterate” (i.e., estimate) the chip’s location.

- The estimated accuracy of a PLV system depends on the exact PLV geolocation method, the number and location of landmark servers, and the amount of network processing latency when AI chips communicate with landmark servers.

Figure 1. Steps of a PLV System in Practice



Source: CSET illustration.

A global PLV system would likely need between 39 and 500 landmark servers. By positioning one landmark server in every major data center hub outside the United States, China, and Russia, PLV could likely achieve country-level accuracy for the vast majority of PLV-tracked chips.¹

- We find that 500 landmark servers, distributed globally, can very likely achieve accuracy of 1,000 miles or better (region-level accuracy). A low-effort PLV implementation would likely require at least 39 landmark servers; if distributed globally, these servers can very likely achieve accuracy of 2,120 miles or better (continent-level accuracy).
- However, our model does not calculate the overlap of multiple landmark servers, which could substantially improve PLV accuracy beyond the worst case.
- Best-case accuracy of a PLV system would very likely be no better than about 20 miles. Therefore, with 500 landmark servers, a PLV system will likely achieve region-level accuracy (worst case) to city-level accuracy (best case).

¹ A “major data center hub” is defined as an area containing a large concentration of data centers within a 19-mile radius. Country-level accuracy is defined as an error-circle radius of 116 miles or shorter.

Costs

A PLV system could be deployed for a feasible but non-trivial cost, 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 potential 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.
- Additional non-quantifiable, indirect costs, including follow-on enforcement, would require BIS to be well-resourced enough to act on the leads created by the PLV system.

Figure 2. Annual Cost Range Estimates by Landmark Server Procurement Model



Source: CSET Location Verification Cost Model (see “PLV Benefits, Costs, and Policy Design Choices” section in the full report).

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.

- BIS' 23% budget increase to \$235 million for FY2026 represents a positive development, but if BIS remains resource-constrained, the utility of a PLV system that produces additional leads (and potentially a high volume of false leads) will be partially undermined by the capacity challenges BIS already faces.

Limitations

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 from end users who would need to install a software or firmware update. Hardware-based implementations are not feasible on chips in circulation.

- Connectivity between chips and landmark servers: When an AI chip fails to respond to a PLV system, the system cannot inherently know whether it is due to diversion or legitimate reasons. Chips that remain offline—whether because they are in storage, in transit, or deliberately disconnected by bad actors—produce ambiguity in the system.
- Adversarial location obfuscation: Well-resourced adversaries could trick a PLV system into reporting that a smuggled chip is still in a permitted location. Even a PLV system designed with prudent security measures could be subverted, and as a result some chips that report being in their intended locations may actually have been diverted.
- 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.

Policy Design Choices

The true value of a PLV system depends on several policy design choices, including how the system is built, who operates it, and how broadly it applies. If a PLV system is pursued, we recommend the following:

- PLV should be implemented on AI chips with firmware-based signing, as it best balances robustness against adversarial attack, cost, and ease of implementation.
- Location verification should only be required for data center chips subject to the most stringent licensing requirements to align with current export control objectives.
- BIS-certified trusted third parties should operate the landmark and reporting servers in order to eliminate conflicts of interests and leverage private sector expertise.
- The program should aim to detect diversion globally (rather than only to countries of concern) as the additional costs are marginal.
- Policymakers also have to consider whether domestically deployed chips are tracked as well, which may require Congress to grant BIS additional authorities.

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.