Adam Boas
adamboas.com | anboas@gmail.com
February 17, 2026

Don’t Put Frontier AI Under ITAR: A Blunt Instrument for a Dual-Use Reality

There is a seductive simplicity to the idea: if frontier AI models and agentic runtimes can enable lethal autonomy, cyber operations, and strategic-scale influence, then treat them like a weapon. Pull major AI companies under the International Traffic in Arms Regulations (ITAR), lock it down, and stop pretending it is “just software.”

I am not convinced that works, and I am increasingly convinced it would backfire.

ITAR is built to control defense articles, defense services, and associated technical data on the U.S. Munitions List, administered by the State Department. It is intentionally broad and punitive, and “export” includes disclosing controlled technical data to foreign persons, including within the United States (a deemed export) [1]. That is a feature for missiles and fire-control systems. It is a bug for a technology ecosystem where capability is distributed across models, data, tooling, evaluation harnesses, chips, and the global talent base that builds and defends them.

The innovation cost would not be a side effect; it would be the main effect.

If you “ITARize” frontier AI in a way that captures general-purpose model development, you do not just throttle exports. You throttle basic operating mechanics:

- **Hiring and collaboration become export-control exercises.** In ITAR-land, foreign-person access is not a rounding error; it is a compliance event. Day-to-day AI work, including code review, model debugging, red-teaming, and vendor incident response, depends on cross-border and cross-national participation. Treating routine collaboration as a regulated export will push companies into U.S.-person enclaves and split teams along citizenship lines [2].
- **Research disclosure becomes legally radioactive.** AI is not a single item you can box up. Competitiveness depends on iterative publication, open benchmarks, shared safety practices, and rapid diffusion of defensive techniques. Over-controlling “technical data” risks converting a safety culture into a legal risk surface.
- **Compliance becomes a moat.** ITAR overhead does not fall evenly. It advantages incumbents with deep compliance departments and punishes startups and research labs where genuine innovation often occurs. You can accidentally build a policy that “secures” AI by concentrating it.

This is not hypothetical. Even under the Commerce Department’s Export Administration Regulations (EAR)-based approach, which is typically more flexible than ITAR, the U.S. has already seen backlash when proposed AI diffusion controls were viewed as innovation- and diplomacy-hostile. In May 2025, the Bureau of Industry and Security (BIS) described the Biden-era AI Diffusion Rule as innovation-stifling and diplomatically harmful, and rescinded it pending replacement [3].

If that friction triggered a course correction under EAR, it is hard to argue full ITAR treatment would be anything other than a strategic self-inflicted wound.

ITAR does not stop capability abroad; it mostly stops the U.S. from participating.

The uncomfortable truth is that frontier AI is already global. Export controls can slow diffusion of specific chokepoints, such as chips, select manufacturing tools, and some closed-weight artifacts, but cannot prevent capable states and well-resourced firms from developing comparable systems.

A maximalist ITAR approach would encourage:

- sovereign model programs outside U.S. jurisdiction,
- offshore safety and evaluation ecosystems that evolve without U.S. participation,
- and a bifurcated world where U.S. AI becomes compliance-heavy and slower-moving while others innovate around it.

That is not security. That is strategic isolation.

The U.S. government is already using a more targeted model. BIS has explored controlling model weights for the most advanced closed-weight AI models under EAR constructs. Congressional Research Service (CRS) coverage of the AI Diffusion Rule proposal discussed restrictions on model weights for advanced U.S. closed-weight dual-use AI models [4].

This is what scalable export control should look like: define a narrow set of controllable artifacts and chokepoints while avoiding a sweep of the domestic innovation base.

The real problem is not "AI exists"; it is "AI acts without governance".

The strongest ITARization argument is moral and operational: if AI can be used for lethal autonomy at scale, we have crossed a line, so treat enablers as munitions.

Even if you accept the premise, the conclusion does not follow.

The answer to "Pandora's box is open" is not "regulate the box." It is "regulate the interfaces where the box touches reality."

In practice, risk comes from systems that can:

- access sensitive data without enforced labeling and minimization,
- invoke tools with side effects without authorization gates,
- operate with unclear "on behalf of" semantics,
- and produce actions without provenance.

Those are governance failures. They are solvable with architecture, enforcement, and evidence. ITAR is an export regime, not an operational control plane.

That is where an agent control-plane approach matters: enforceable delegation, policy-gated action, and replayable operational evidence in high-consequence environments (the Agent Control Plane Reference Architecture, ACP-RA, is one proposed reference architecture for this).

A smarter alternative.

If the goal is reducing catastrophic misuse while preserving U.S. innovation, a more defensible approach is:

1. **Keep general-purpose frontier model development under EAR, not ITAR,** with narrowly defined controls on specific high-risk artifacts (for example: certain closed-weight model weights, specialized training/inference pipelines, and chip-enabled training services), updated as thresholds shift [4].
2. **Apply ITAR only to defense-unique implementations,** including fine-tunes, toolchains, integrations, and services specifically designed for military end use (targeting, weapons employment, mission systems integration). This aligns with ITAR's purpose: regulating defense articles/services and direct technical data [1].
3. **Require "no identity, no action" provenance for high-consequence agentic systems** in regulated environments (defense, critical infrastructure, classified or controlled domains). Make

provenance operationally mandatory: signed delegation chains, policy decisions at gateways, and replayable evidence.

This preserves what matters most: U.S. ability to innovate broadly, collaborate with allies, and lead on safety, while still drawing hard lines around defense-unique capabilities and sensitive artifacts. It also gives leaders a practical bridge between policy and execution: targeted export controls on one side, and operational governance on the other.

Conclusion.

Putting “AI companies” under ITAR feels strong because ITAR is strong. In this domain, strength is not measured by restrictiveness. It is measured by whether outcomes beat what adversaries can get for free.

A sweeping ITAR move would slow U.S. innovation, fragment talent, complicate allied collaboration, and push development offshore, while doing little to stop capable actors elsewhere.

If we want to keep the future from becoming science fiction, we should not start by regulating yesterday’s way. We should regulate the condition that actually creates harm: **uncontrolled autonomy without enforceable identity, authorization, and action provenance.**

Control the interfaces where autonomy touches reality - identity, authorization, and provenance - without freezing innovation.

References.

1. International Traffic in Arms Regulations (ITAR) | Wex | LII. https://www.law.cornell.edu/wex/international_traffic_in_arms_regulations_%28itar%29
2. International Traffic in Arms Regulations (ITAR) Compliance - Interim Policy | University of Arkansas Research Integrity and Compliance. https://rsic.uark.edu/exportcontrol/itar-compliance.php
3. Department of Commerce announces rescission of Biden-era AI Diffusion Rule | Bureau of Industry and Security. https://www.bis.gov/press-release/department-commerce-announces-rescission-biden-era-artificial-intelligence-diffusion-rule-strengthens
4. U.S. Export Controls and China: Advanced Semiconductors | CRS | Congress.gov. https://www.congress.gov/crs-product/R48642