

July 27, 2026

U.S. Trade Controls and National Security Considerations for Artificial Intelligence (AI) Investments

Navigating U.S. Export Control and National Security Risk When Investing in the AI Sector

By Billy Phalen, Shawn Cooley,
Antonia Tzinova, Sisi Liu, and Colin
Giuseppe Cox

At a Glance

- The U.S. Department of Commerce's Bureau of Industry and Security's (BIS) Foreign Direct Product Rule (FDPR) elevates risks for investments into non-U.S. AI-sector targets whose products, models, or infrastructure depend on U.S.-origin chips, chip design software, or manufacturing equipment.
- BIS is expected to expand export controls beyond just hardware to cover AI computing services and, potentially, model weights and capabilities, which could directly restrict U.S. parties from engaging with AI businesses with a significant nexus to China.
- Graphics Processing Units (GPU) and advanced-chip export licensing is currently governed by a patchwork of narrower rules following the rescission of the AI Diffusion Rule, with a reported draft framework from BIS that would tie licensing to compute cluster size.
- Any Chinese target carries elevated risk of being designated as a restricted party, which could sever access to U.S.-origin technology, U.S. customers, and U.S. investors if triggered.

BACKGROUND

Investor interest in the AI sector, spanning model developers, AI infrastructure and compute providers, and AI-enabled application companies, has grown alongside surging demand for AI products and services. Where a prospective investment target is based in or has significant ties to China, however, that opportunity is layered with a distinct set of U.S. export control and national security considerations that specifically target the AI sector.

Because many AI businesses depend, directly or indirectly, on U.S.-origin high performance chips, chip design software, cloud infrastructure, or customer relationships, they are more exposed to changes in U.S. export controls than a typical operating business. The degree of exposure varies by business model (e.g., a company that owns and operates GPU infrastructure faces different considerations than one that licenses a model or sells an AI-enabled application), but the considerations below are relevant, in varying degrees, to most AI-sector investments involving a Chinese target or user and each should be evaluated as part of an investor's pre-signing diligence process.

SUMMARY OF KEY RISK AREAS

Risk Area	Key Concern for AI-Sector Investors
Foreign Direct Product Rule (FDPR)	Chips, models, or infrastructure used by the target may fall within U.S. export control jurisdiction because they are the direct product of U.S.-origin chip design software and manufacturing equipment.
Export Controls on AI Computing Services & Model Controls	BIS rulemaking may extend export controls to AI computing services and, potentially, model weights and capabilities, directly affecting the target's core offering.
GPU & Advanced Chip Export Licensing Thresholds	A reported draft framework from BIS would tie chip export licensing to compute cluster size, potentially capping the scale at which a compute-intensive target can expand. The Remote Access Security Act currently pending in Congress is intended to close any perceived gap in regulation.
Restricted Party Designation Risk	Designation of the target, an affiliate, or a key supplier could cut off access to U.S.-origin technology, U.S. customers, and future U.S. investment.
Outbound Investment Screening	The investment itself may be a prohibited or notifiable transaction under the U.S. Department of the Treasury's Outbound Investment Security Program, intended to limit capital and managerial knowhow to Chinese companies, independent of export control exposure.

FOREIGN DIRECT PRODUCT RULE (FDPR)

The FDPR extends U.S. export control jurisdiction beyond U.S.-origin goods to cover foreign-made products that are the direct product of U.S.-origin technology, software, or production equipment (see 15 C.F.R. § 734.9).

- The FDPR is directly relevant to many AI-sector investments because a target's products, models, or infrastructure often depend on GPUs or other advanced chips that are themselves the product of U.S.-origin chip design software (Electronic Design Automation tools from Cadence or Synopsys) and U.S. semiconductor manufacturing equipment. This means the underlying compute a target relies on, whether owned, leased, or accessed through a cloud provider, could fall within U.S. export jurisdiction regardless of where the chips are fabricated or the servers are located.
- The degree of FDPR exposure will vary with a company's business model. Companies that own and operate significant compute infrastructure (e.g., data center operators, compute resellers) face the most direct exposure, since their ability to source, replace, and scale hardware is tied to chip-level export restrictions. Companies that primarily license software, sell AI-enabled applications, or access compute through a third-party cloud provider face more indirect exposure, but should still diligence whether their supply chain, cloud dependencies, or product architecture could implicate U.S.-origin technology.

EXPORT CONTROLS ON AI COMPUTING SERVICES AND MODEL ACCESS

BIS is reportedly considering expanding export controls related to AI, and most expect the final rulemaking to reach AI computing services and, potentially, model weights and capabilities — not just hardware.

- The current regulatory trajectory points toward controls on the provision of advanced AI computing services to or through Chinese entities, building on the AI diffusion framework developed since 2023. Companies that sell compute capacity, host or train models, or provide AI infrastructure as a service may sit squarely within the intended scope of these potential rules. If finalized in a more expansive form, such rules could prohibit U.S. persons from engaging with the target for covered AI activities without a BIS license, likely reviewed under a presumption of denial.
- Alternatively, BIS could focus controls more narrowly on model weights and AI capabilities themselves, requiring licenses to train, deploy, or transfer controlled models in or to Chinese facilities. The impact on a given target will depend heavily on its business model: a company whose offering does not involve transferring or hosting controlled models or capabilities may be less exposed than one whose core value proposition depends on doing so.

GPU AND ADVANCED CHIP EXPORT LICENSING THRESHOLDS

The regulatory framework governing quantity- and performance-based licensing thresholds for advanced GPUs and other AI-relevant chips is currently in flux, which creates both near-term uncertainty and the possibility of a more restrictive compute-based licensing regime in the future, including the potential expansion of restrictions to certain Southeast Asian countries (e.g., Singapore, Malaysia, and Thailand) that have faced explicit U.S. scrutiny over suspected chip diversion. This is most acute for targets that own or operate significant compute infrastructure but is also relevant to any AI company whose product or service depends on access to advanced chips through a cloud provider or hardware partner.

- The Biden-era “AI Diffusion Rule” (originally effective January 2025 but since rescinded) would have imposed a tiered, country-based licensing framework for advanced computing chips, but BIS rescinded the rule in May 2025 before its compliance date and did not announce any plans regarding what may replace it, while directing BIS enforcement staff not to enforce the rescinded rule in the interim. Since then, BIS has taken narrower, more targeted actions¹ rather than reissuing a comprehensive framework.
- More significantly for the growth trajectory of a compute-intensive AI target, BIS is reportedly drafting a new framework that would license GPU exports based on the compute capacity involved: shipments of up to 1,000 Nvidia GB300 GPUs or the equivalent would undergo “fairly simple review” with exemptions, while larger exports to companies building extensive computing clusters would require more pre-clearance and be subject to a stricter export licensing regime with conditions attached, potentially including disclosure of business models or allowing U.S. government access to facilities for site visits.²
- On May 31, 2026, [BIS issued guidance](#) clarifying that, under existing regulations, export licenses are required for advanced computing chips destined for any entity whose ultimate parent company is headquartered in China or Macau, regardless of where that entity is situated, and stated this standard has technically been in force since November 2023.
- Furthermore, in a [policy statement](#) of May 13, 2026, BIS advised that it had “determined that access to advanced computing integrated circuits (ICs) and commodities subject to the EAR for training AI models has the potential to enable military-intelligence and weapons of mass destruction (WMD) end uses in Country Group D:5 countries (including China) or Macau.” BIS provided examples of “activities [that] may trigger a license requirement under the catch-all controls of part 744 of the EAR when there is “knowledge” ... that the AI model will be used for a WMD or military-intelligence end use/user,” including:

¹ See the rule effective in January 2026 that shifted the license review standard for certain advanced computing chips exported to China and Macau from a presumption of denial to case-by-case review, but only for exporters meeting strict requirements, including various certifications as to sufficient supply in the United States, US third-party testing, and KYC/remote-access safeguards.

² See Bloomberg article [here](#); see Reuters article [here](#).

- Exports, reexports, or transfers (in-country) of advanced computing ICs and commodities subject to the EAR to any party, such as foreign Infrastructure as a Service (IaaS) providers (e.g., data center providers), when the exporter, reexporter, or transferor has “knowledge” that the IaaS provider will use these items to conduct training of AI models for or on behalf of parties headquartered in D:5 countries (including China) or Macau;
- Transfers (in-country), defined as a change in end use or end user, of advanced computing ICs and commodities subject to the EAR already in the possession of parties such as IaaS providers, if there is “knowledge” that the items will be used by the transferee to train AI models for or on behalf of parties headquartered in D:5 countries (including China) or Macau.
- A “U.S. person” provides any ‘support’ or performs any contract, service, or employment, when there is “knowledge” such activity will be used for or may assist the training of AI models for or on behalf of parties headquartered in D:5 countries (including China) or Macau.
- Separately, the Remote Access Security Act currently pending in Congress is intended to close any perceived gap in regulation.

RESTRICTED PARTY DESIGNATION RISK

Any Chinese-affiliated company is at heightened risk of becoming a restricted party, and recent examples targeting major Chinese telecom OEMs demonstrate Commerce's willingness to sanction Chinese companies in the advanced-technology and AI space with limited warning.

- Outside counsel can assess in diligence whether a target or material customer/supplier of a target is more or less likely to be targeted based on factors such as its supply chain for chips or compute, its customer base, its role in the AI value chain, and any nexus to the Chinese government or military-linked entities.
- In a worst-case scenario, if the target, a key affiliate, or a critical supplier is added to the Entity List or designated a Chinese Military Company under the National Defense Authorization Act for Fiscal Year 2021 Section 1260H, U.S. persons would generally be barred from transacting with it without a license. For an AI company, this is a particularly acute risk: depending on the business model, it could cut off access to U.S.-origin chips or cloud infrastructure the business depends on to operate, sever ties to U.S. customers or partners, and potentially foreclose future U.S. investor participation, effectively stranding the investment.

OUTBOUND INVESTMENT SCREENING

Separate from U.S. export control considerations, the U.S. Outbound Investment Security Program (OISP) may independently prohibit or require notification of the investment itself. Effective since January 2025 under Executive Order 14105, the OISP restricts U.S. persons from making certain investments in Chinese entities engaged in specified activities in three sectors, including artificial intelligence, alongside semiconductors/microelectronics and quantum information technologies, with the intent to limit U.S. capital and managerial knowhow to Chinese companies operating in these sectors.

- Depending on the target's specific AI-related activities, the transaction may be a prohibited transaction (barred outright), a notifiable transaction (permitted, but requiring notice to Treasury), or excepted altogether — the analysis turns on defined terms (e.g., what constitutes “developing” a covered AI system) that require careful, fact-specific review.
- The regime is also in transition: the [Comprehensive Outbound Investment National Security Act of 2025](#), enacted in December 2025, codifies and expands the OISP, broadening covered countries and adding high-performance computing/supercomputing and hypersonic systems as covered sectors, with implementing regulations due no later than March 2027. Investors should structure diligence and deal timelines with this pending expansion in mind, in addition to complying with the current rules.

KEY TAKEAWAYS FOR PROSPECTIVE INVESTORS

What Investors Should Consider Now

- **Map hardware, software, and cloud provenance.** Diligence should trace the target's chip, model, and infrastructure dependencies, including any embedded U.S.-origin design software, manufacturing equipment, or cloud relationships, to assess FDPR exposure.
- **Stress-test the customer base.** Evaluate how much of the target's revenue and growth plan depends on U.S. hyperscaler, enterprise, or developer customers who may be restricted from, or reluctant to, transact with a Chinese AI company.
- **Build in regulatory flexibility.** Given the unsettled state of AI computing services controls, model weight restrictions, and chip export licensing thresholds, consider Material Adverse Change-type provisions or repricing mechanisms tied to adverse regulatory developments rather than underwriting to a single static regulatory assumption.
- **Assess Entity List and 1260H exposure.** Diligence the target's, its affiliates', and its key suppliers' nexus to the PRC government or military-linked entities, given the potential for swift, business-ending designations.

* * *

If you have questions concerning the contents of this alert, or would like more information, please speak to your regular contact at Weil or to the authors:

Authors

Billy Phalen (BO)	View Bio	william.phalen@weil.com	+1 617 772 8357
Shawn Cooley (D.C.)	View Bio	shawn.cooley@weil.com	+1 202 682 7103
Antonia Tzinova (D.C.)	View Bio	antonia.tzinova@weil.com	+1 202 682 7120
Sisi Liu (D.C.)	View Bio	sisi.liu@weil.com	+1 202 682 7176
Colin Giuseppe Cox (DA)	View Bio	colin.cox@weil.com	+1 214 746 7840

© 2026 Weil, Gotshal & Manges LLP. All rights reserved. Quotation with attribution is permitted. This publication provides general information and should not be used or taken as legal advice for specific situations that depend on the evaluation of precise factual circumstances. The views expressed in these articles reflect those of the authors and not necessarily the views of Weil, Gotshal & Manges LLP. If you would like to add a colleague to our mailing list, please [click here](#). If you need to change or remove your name from our mailing list, send an email to weil.alerts@weil.com.