Arnold & Porter

January 27, 2025

Commerce Department Imposes Sweeping Global Restrictions on AI Technologies

Advisory

By

John P. Barker ,

Charles A. Blanchard ,

Ronald D. Lee ,

Soo-Mi Rhee ,

Andrew Varner ,

Nicholas L. Townsend ,

Junghyun Baek ,

Trevor G. Schmitt ,

Bell Johnson ,

Dustin Vesey

On January 15 and 16, 2025, the U.S. Department of Commerce's Bureau of Industry and Security (BIS) published two interim final rules (IFR) aimed at reducing the risk that countries of concern obtain the most advanced artificial intelligence (AI) models and advanced computing items. The first IFR is known as the Framework for AI Diffusion and may be referred to as the Diffusion rule. The second IFR may be referred to as the Presumption rule. Although the IFRs are already in effect, exporters, re-exporters, and transferors (in-country) are not required to comply with the rules' changes until various future dates. Most provisions under the Diffusion rule require compliance by May 15, 2025, although minor provisions, such as the one relating to security requirements to become certified under the new prongs of the validated end user program (described below), do not require compliance until January 15, 2026. The Presumption rule's changes, including the presumption itself, are set to require compliance on January 31, 2025.

The key elements of the IFRs are:

(1) Revisions to BIS' controls under the Export Administration Regulations (EAR) on advanced computing integrated circuits (ICs)

(2) The creation of a new rebuttable presumption that advanced ICs are covered by the revisions unless attested to by an "approved" or "authorized" company

(3) The addition of a new control on AI model weights for certain advanced closed-weight AI models

(4) The creation of three new license exceptions

(5) Changes to the data center validated end user program, including establishing two types of validated end user status

We discuss each of these elements in more detail below.

The IFRs aim to ensure that access to large clusters of advanced ICs is limited to destinations that pose comparatively low risks of diversion or misuse and model weights of advanced AI models are stored outside of the United States under strict security conditions. To accomplish these twin goals, the rules require a license to export, re-export, or transfer (in-country) advanced ICs and the model weights

of the most advanced AI models to any end user in any destination. License applications will be reviewed based on the sensitivity of the destination and end user, the quantity of compute power or performance of the AI model, and the security requirements agreed to by the recipient. The rules also include license exceptions for validated end users and particularly low-risk destinations.

New Controls on Advanced Computing Integrated Circuits

Training today's most advanced AI models requires large clusters of advanced computing ICs that are capable of processing large quantities of data. Prior BIS rules did not control exports to many destinations in the world and left open the possible diversion of controlled technologies to countries of concern. The Framework for AI Diffusion seeks to address these diversion risks by imposing a global license requirement for the exportation of advanced ICs while retaining license exceptions and validation opportunities to facilitate trusted transactions.

The Structure of the New Global Licensing Regime

The global licensing regime for advanced ICs created by the Diffusion rule divides the world into three tiers. The first tier (T1) comprises 19 destinations that maintain AI development ecosystems that are unlikely to threaten U.S. national security and have implemented measures to prevent diversion. Those destinations are Australia, Belgium, Canada, Denmark, Finland, France, Germany, Ireland, Italy, Japan, the Netherlands, New Zealand, Norway, South Korea, Spain, Sweden, Taiwan, the United Kingdom, and the United States. As discussed in more detail below, a number of new license exceptions are available for this tier, and a presumption of approval applies for applications to, or to an entity headquartered in, or whose ultimate parent company is headquartered in, these destinations, subject to remaining end use and end user restrictions.

The third tier (T3) comprises the group of destinations subject to U.S. arms embargoes (those classified as D:5 destinations in Supplement No. 1 to Part 740 of the EAR) and Macao, for whom there are already strict rules in place. The Diffusion rule does not change the rules governing those destinations; a presumption of denial applies for applications to, or to an entity headquartered in, or whose ultimate parent company is headquartered in, T3 destinations.

The second, middle, and biggest tier (T2), comprising all other destinations in the world, is the tier most impacted by the Diffusion rule. For these destinations, although transactions in limited quantities of controlled ICs will not require a license, as detailed further below, transactions involving larger quantities of controlled ICs will be subject to new specific destination allocations. However, in certain cases, T2 destinations may receive a higher allocation. Moreover, within those destinations, entities that have agreed to enact concrete, verifiable, robust security measures may apply to gain validated status to obtain larger quantities of ICs. A presumption of approval applies for applications to these destinations up to the enumerated per-destination allocation. A policy of denial will apply to applications after the allocation has been met.

For transactions that do not meet the requirements for use of a license exception (explained below), and for which the end user does not participate in one of the two new validated end user authorizations (also explained below), exporters and re-exporters must proceed with the traditional license application process. For those shipping to T2 destinations, those destinations will be subject to progressive quarterly maximum installed base allocations that total 790 million cumulative total processing performance (TPP) through calendar year 2027.[1] When the allocation has been met, applications will be reviewed under a policy of denial. To help the public stay abreast of destinations' progress toward reaching allocations, BIS will provide timely updates on its website on each destination's allocations as they are expended.

Items Impacted by the New Regime

The license requirement will apply to items listed under export control classification numbers (ECCN) 3A090.a, 4A090.a, and corresponding .z items that meet similar functional characteristics. These ECCNs include: advanced ICs with a certain level of TPP; computers, electronic assemblies, and components containing the same ICs; and associated software and technology.

In addition, the Diffusion rule expands the destination scope of the advanced computing foreign direct product rule (FDPR) that already exists. An FDPR provides that BIS can regulate the export, re-export, and transfer (in-country) of foreign-produced items if their production involves certain technology, software, or equipment subject to the EAR. Where the current FDPR for advanced ICs only covered certain destinations, it now covers any destination worldwide.

New Presumption for Exporters Relating to 3A090.a

To address potential evasion of the controls around 3A090.a, the Presumption rule creates a new presumption that any "applicable advanced logic IC" meets the parameters for the 3A090.a ECCN unless overcome in one of three ways.[2] If the presumption is not overcome, an "applicable advanced logic IC" is considered to be 3A090.a and designed or marketed for datacenters when a "front-end fabricator" (manufacturers of ICs designed by another party) or an outsourced semiconductor assembly and test (OSAT) company seeks

to export, re-export, or transfer (in-country) the relevant IC. This will in turn result in the application of the global licensing regime explained in the previous section and restrictions on license exception availabilities . Moreover, recently added Note 5 to the advanced computing FDPR reminds that a foreign-produced 3A090 item that is subject to the FDPR may, by way of the new presumption, also face additional regulatory requirements if the foreign-produced item is set for export, re-export, or transfer (in-country) by a front-end fabricator or OSAT company.

The first way to rebut the presumption is if the designer of the applicable IC is an "approved" or "authorized" IC designer (explained below). The second way is if the IC is packaged by a front-end fabricator outside of a T3 destination that makes certain attestations of the final packaged IC's performance capabilities. The third way is if the IC is packaged by an "approved" OSAT company that makes the same attestations. Attestations that must be made relate to the aggregated approximate transistor count of the final packaged IC and whether it contains high-bandwidth memory.

As part of instituting this new presumption, the Presumption rule creates a new framework for designating "approved"/"authorized" IC designers and "approved" OSAT companies. Approved IC designers and OSAT companies are ones BIS has determined, based on a holistic review, engage in trusted transactions. The rule provides an initial list of companies to be included in each list. Going forward, IC designers and OSAT companies that wish to be added to the lists of approved companies must submit a request in the form of an advisory opinion to BIS. "Authorized" IC designers, meanwhile, are determined differently. Prior to April 13, 2026, all IC designers that meet the following are deemed authorized: headquartered in Taiwan or an A:1 or A:5 destination in Supplement No. 1 to Part 740 of the EAR, neither located nor have an ultimate parent headquartered in a T3 destination, and have agreed to submit applicable information described in the rule to the front-end fabricator. After April 13, 2026, those IC designers that meet the previous requirements *and* have submitted an application to become an *approved* IC designer will be deemed authorized for not more than 180 days.

The presumption's compliance date is set for January 31, 2025.

New Control on Model Weights

Model weights, like the keys to a car, are critical to the model's functionality and difficult to reproduce or reverse-engineer. The model weights of "closed-weight" models, or those models whose model weights are not publicly accessible (as compared to "open-weight models," whose model weights are publicly accessible), are one of the most well-guarded components of advanced AI models.

The Diffusion rule introduces a new ECCN, 4E091, to control certain advanced closed-weight model weights through a similar global licensing requirement. The licensing requirement does not apply to open-weight model weights. In addition, it only applies to those closed-weight models that are sufficiently advanced, defined in the IFR as those trained utilizing 10^26 or more "computational operations;"[3] less powerful closed-weight model weights will not require a license. Applications to export these most advanced closed-weight model weights will be subject to a presumption of denial to all destinations. However, applications to export to T1 destinations may benefit from the new license exception AIA (detailed below). The license requirement does not apply to "deemed exports or re-exports" for individuals employed by entities headquartered in or with an ultimate parent company headquartered in a T1 destination.

Alongside its creation of 4E091, the Diffusion rule introduces a new FDPR for the same, most advanced closed-weight model weights outlined above. As a result, closed-weight model weights that meet the parameter specifications outlined in 4E091 that are produced using ICs, servers, and other electronic equipment outlined in the rule that are themselves subject to the EAR are also subject to the licensing restriction when destined for any destination in the world.

New License Exceptions for Items Covered by the Rules

The IFRs create three new license exceptions for use when complying with the EAR: AIA (Artificial Intelligence Authorization), ACM (Advanced Compute Manufacturing), and LPP (Low Processing Performance).

The AIA exception may be used to authorize the export, re-export, or transfer (in-country) of eligible items (ECCNs 3A090.a, 4A090.a, corresponding .z items, and 4E091) to entities located in T1 destinations unless the entity is headquartered outside of or has a parent company headquartered outside of the T1 destination. To use the license exception, the exporter, re-exporter, or transferor must comply with certification and reporting requirements.

The ACM exception may be used to authorize the export, re-export, or transfer (in-country) of eligible items (ECCNs 3A090, 4A090, and corresponding .z items) to a "private sector end user" located in a T1 or T2 destination if the ultimate end use is the "development," "production," or storage of the items, provided the private sector end user is not headquartered in and does not have an ultimate parent company headquartered in a T3 destination. This exception may not be used if the end use is to train an AI model. Although exports under the ACM exception do not count against T2 destinations' allocations, exporters, re-exporters, and transferors must keep records for each

facility items are distributed to. The purpose of this exception is to minimize the impact of the rule on supply chains.

Pursuant to the Presumption rule, the AIA and ACM exceptions above may only be used for three specific commodities (ECCNs 3A090.a, 5A002.z.1.a, z.2.a, z.3.a, z.4.a, z.5.a, and 5A992.z.1) if the commodities are designed by an approved or authorized IC designer.

The LPP exception authorizes the export and re-export of up to 26,900,000 TPP of advanced computing ICs per calendar year to an individual ultimate consignee located in any T1 or T2 destination, provided its headquarters and the headquarters of its ultimate parent company are not located in a T3 destination. The TPP limit applies to aggregate shipments made to an ultimate consignee; before shipping, an exporter or re-exporter must obtain a certification from the ultimate consignee that it has not yet hit its annual TPP limit from all exporters and re-exporters. Use of the LPP license exception does not count toward a T2 destination's overall allocation under the Diffusion rule's new licensing regime. There are additional reporting requirements an ultimate consignee must abide under the LPP exception.

BIS also made changes to two other license exceptions — NAC (Notified Advanced Computing) and ACA (Advanced Computing Authorized) — to ensure they keep pace with the new regime and maintain robust reporting requirements.

Changes to the Data Center Validated End User Program

Data centers are specialized facilities designed to train AI models. In 2024, BIS amended the EAR to allow companies that operate data centers abroad to apply for and receive validated end user (VEU) status. The Diffusion rule bifurcates that approval process into a "universal VEU" (UVEU) and a "national VEU" (NVEU). Each type of VEU authorization would allow the company operating the data center to acquire controlled technology more easily.

UVEU is only available to companies headquartered in, or whose ultimate parent company is headquartered in, a T1 destination. UVEUs, once authorized, will be empowered to use the status to build data centers comprised of advanced ICs subject to U.S. export controls around the world, except in T3 destinations. UVEUs will be subject to limitations on where they can geographically allocate their computing power, with no more than 25% (of all computing power owned by the entity and all its subsidiary and parent entities) located outside of T1 destinations and no more than 7% located in any single non-T1 destination. U.S.-headquartered UVEUs will be prohibited from allocating more than 50% of total AI computing power outside of the United States.

With certain limitations, companies headquartered in, or whose ultimate parent company is headquartered in, both T1 and T2 destinations may apply for the NVEU authorization. NVEUs will benefit from a per-company, per-destination installed base allocation of TPP that is calculated separately from the destination's overall allocation under the regime. The TPP allocations, which are currently set through 2027, are designed to allow powerful AI models to be trained, while restricting NVEUs' computing access in such a way that they stay one generation of development behind the most advanced AI models.

Applicants applying for VEU status must go through an intensive application process and certify they will follow certain guidelines that impose potentially significant and onerous restrictions. Some examples of such restrictions include the following:

- The VEU and all subsidiary and parent entities must be free of ties to military end users or military-intelligence end users (including research and development agreements and joint activities).
- The VEU and all subsidiary and parent entities must adhere to outbound investment rules set forth in 31 C.F.R. Part 850, even if they are a non-U.S. person.
- The VEU must demonstrate that it has eliminated supply chain dependencies on certain advanced semiconductors and networking equipment produced by any entities headquartered in T3 destinations as well as equipment and services listed by the U.S. Federal Communications Commission.
- The VEU must notify the government of all cooperative activities with any entities headquartered in a T3 destination or any individuals on the U.S. Department of Commerce's Entity List or the U.S. Department of the Treasury's Specially Designated Nationals List.
- The VEU must perform ongoing monitoring, evaluation, and end user due diligence and maintain related records.

Conclusion

Since 2022, the Biden administration pursued a dual-track strategy to promote U.S. AI superiority: hinder development of advanced technologies by countries of concern through the use of export controls while promoting U.S. and ally development through industrial policy. These new IFRs cap off the prior administration's dual-track approach. On the first front, the rules consolidate previous piecemeal

restrictions into one global licensing regime that covers both the inputs and outputs of the most advanced AI models: ICs and model weights. On the second front, the Diffusion rule's bifurcated VEU program advances the availability of AI technologies in T1 and T2 destinations as a way to combat the market base and attractiveness of T3 capabilities in those destinations.

Going forward, it will be prudent for individuals and businesses seeking to export, re-export, or transfer (in-country) advanced ICs and other AI-related technologies to keep abreast of their licensing, disclosure, and reporting obligations. This is especially important to keep in mind because, in addition to the recent change of administrations (which may portend changes to AI policies and regulations), the comment periods for these rules, as well as the December 2, 2024 foreign direct product IFR, are open for approximately two more months. Comments received by BIS may prompt new rules or changes to prior rules at any time. Please contact any author of this Advisory, our interdisciplinary AI team, or your Arnold & Porter relationship attorney if you have any questions or to seek further guidance or advice.

[1] Total processing performance measures the number of operations a chip can perform per second.

[2] "Applicable advanced logic IC" is defined as one using the 16/14 nanometer node or below or using a non-planar transistor architecture.

[3] The threshold of 10^26 computational operations is recognized as the presumptive threshold for advanced systems under President Biden's October 2023 Executive Order on the Safe, Secure, and Trustworthy Development and Use of Artificial Intelligence. It is a higher threshold of computational operations than the one included in the U.S. Department of the Treasury's recently published rule covering outbound investments in China. See our Advisory on the outbound rule for more information.