Articles
March 06, 2025
BIS Publishes Bold New Artificial Intelligence Diffusion Framework

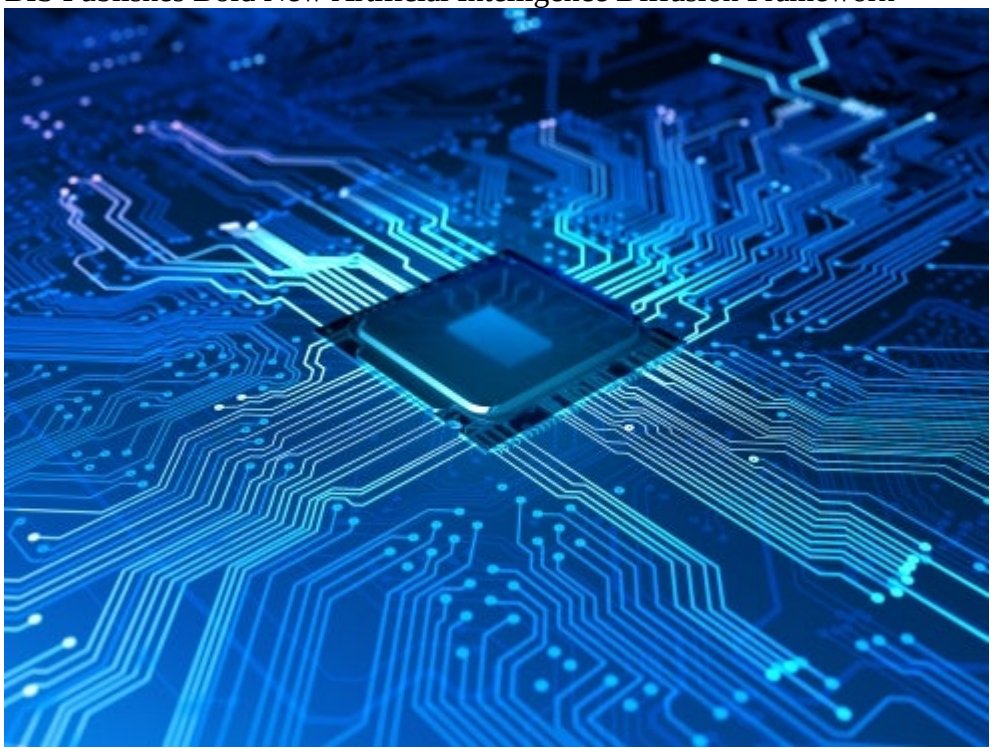

On January 13, 2025, the outgoing Biden administration announced a highly ambitious interim final rule (IFR) setting out a "Framework for Artificial Intelligence Diffusion" in the Export Administration Regulations (EAR).

Through the IFR, the U.S. Commerce Department's Bureau of Industry and Security (BIS) adds a new worldwide license requirement to the EAR for the export, reexport, or (in-country) transfer of advanced computing integrated circuits (ICs or chips), and model weights of closed-weight artificial intelligence (AI) models trained on more than 10^26 computation operations.

The IFR builds on prior controls on advanced computing ICs (aimed primarily at China, Hong Kong, and Macau) issued by BIS in October 2022 and updated in October 2023, April 2024, October 2024, and December 2024.[1] The new AI framework seeks to regulate the global diffusion of large clusters of advanced computing ICs and the most advanced AI models. BIS's stated objective in creating the AI framework is to balance continued U.S. technological leadership with the protection of the United States from malicious actors that may use such advances in ways that pose significant risks to U.S. national security and foreign policy interests (*e.g.*, to advance military end uses, develop weapons of mass destruction, conduct cyberattacks, or perpetrate human rights abuses).

Since this is an interim rule, BIS is soliciting public comments before issuing a final rule. **Comments on the IFR must be received by BIS no later than May 15, 2025.**

This Article provides a detailed analysis of the new provisions and technical requirements implemented by this new AI diffusion rule.

Revised Controls on Advanced Computing Integrated Circuits (*i.e.*, ICs) and New Controls on AI Model Weights

Updated Advanced Computing Integrated Circuits (*i.e.*, Chips) Controls

The IFR expands existing controls on ECCNs 3A090.a and 4A090.a and the corresponding .z items.[2] It now requires a license from BIS to export, reexport, or in-country transfer advanced computing ICs under these export classifications *to any end user in any destination*.

New AI Model Weights Controls

The IFR adds a new ECCN 4E091, which specifies model weights of any closed weight model that has been trained on more than 10^26 computation operations. The ECCN is intended to cover the model weights for the most advanced AI models. BIS determined that a reasonable proxy for the performance of an AI model is the amount of compute (*i.e.*, the number of computational operations) used to train the model, which resulted in the baseline control of 10^26 computation operations. The new ECCN is subject to a worldwide license requirement for export, reexport, or in-country transfer.

ECCN 4E091 does not control open-weight AI models. However, a closed-weight model derived from an open-weight model is subject to control if the model has received additional training "operations" compared to the open-weight model that constitute more than 2.5 × 10^25 "operations" or more than 25% of the total training "operations," whichever is higher.

Additionally, ECCN 4E091 does not control closed AI models trained on fewer operations than the minimum operations needed to train an AI with performance equivalent to the most advanced open-weight AI model.

The AI model weight controls also do not apply to deemed exports and deemed reexports to "permanent regular employees" of entities that are headquartered in, or whose ultimate parents are headquartered in, Tier 1 countries (see below).

New Licensing Policy and License Exceptions

The licensing policy and license exception changes in the IFR rely on a three-tiered country approach as follows:

- **Tier 1.** Australia, Belgium, Canada, Denmark, Finland, France, Germany, Ireland, Italy, Japan, the Netherlands, New Zealand, Norway, Republic of Korea, Spain, Sweden, Taiwan, the United Kingdom, and the United States.
- **Tier 2.** All countries that are not in Tier 1 or Tier 3.
- **Tier 3.** Macau and "group D:5" countries (Afghanistan, Belarus, Cambodia, Central African Republic, China, Democratic Republic of Congo, Cuba, Eritrea, Haiti, Iran, Iraq, North Korea, Lebanon, Libya, Nicaragua, Russia, Somalia, Sudan, South Sudan, Syria, Venezuela, and Zimbabwe).

The licensing policy and license exception changes also rely on per-country total processing performance (TPP) quotas. These TPP quotas set forth the total processing power of all advanced computing ICs that may be exported to a given country during a given period. In the IFR, BIS has allocated a per-country TPP quota of 790,000,000 (equivalent to roughly 50,000 high-end AI GPUs). A country may increase its TPP quota by 100% (*i.e.*, to 1.58 billion TPP, or roughly 100,000 high-end GPUs) by providing certain government-to-government assurances with the United States.

Licensing Policy Changes

For exports that do not qualify for one of the license exceptions described below, the following licensing policies apply.

Advanced Computing ICs

License applications for exports to Tier 3 countries—or to entities headquartered, or whose ultimate parents are headquartered, in Tier 3 countries—will be reviewed by BIS under a presumption of denial.

License applications for exports to Tier 1 countries—or to entities headquartered, or whose ultimate parents are headquartered, in Tier 1 countries—will be reviewed by BIS under a presumption of approval.

For all other exports (*i.e.*, exports to Tier 2 countries or to entities that are headquartered, and whose ultimate parents are headquartered, in Tier 2 countries), the licensing policy will depend on whether the relevant Tier 2 country has exhausted its TPP quota.[3]

AI Model Weights

License applications for exports to entities that are not headquartered, or whose ultimate parents are not headquartered, in Tier 1 countries will be reviewed by BIS under a presumption of denial.

License Exception Artificial Intelligence Authorization (AIA)

The IFR includes a new license exception AIA that applies to both the advanced computing IC and AI model weights controls described above. License exception AIA is the *only* license exception that applies to the new AI model weights controls.

Advanced Computing ICs

License exception AIA is available for exports, reexports, and transfers of advanced computing ICs to entities located in Tier 1 countries so long as the entities (and their ultimate parent entities) are headquartered in a Tier 1 country. As discussed further below, to use the AIA license exception for advanced computing ICs, exporters, reexporters, and transferors must furnish the ECCN to the ultimate consignee and obtain a prior certification from the ultimate consignee that includes certain certification requirements[4] (Certification Requirements). Accordingly, Tier 1 entities will generally be able to obtain the most advanced ICs without a license as long as the "ultimate consignees" for the items certify compliance with the Certification Requirements, and exporters, reexporters, and transferors furnish the applicable ECCN(s) to the ultimate consignee.

Model Weights

License exception AIA is available for exports, reexports, and transfers of AI model weights controlled under new ECCN 4E091 to entities headquartered, or whose ultimate parent company is headquartered, in Tier 1 countries so long as these entities are not located within a Tier 3 country.

However, prior to use, the exporter, reexporter, or transferor must ensure that the controlled AI model weights will be stored in a facility that has instituted specific security measures set out in paragraphs 14, 15, and 18 of Supplement 10 of Part 748, which includes baseline security requirements for chips and data, AI-specific cybersecurity, and personnel security standards and practices (Security Measures).

Additional License Requirements for Foundries and OSAT Companies

In a supplemental interim final rule published on January 16, 2025, BIS modified license exception AIA so that certain exceptions are only available for items designed by "approved" IC designers who are likely to accurately report the ECCN of the items they request advanced foundries to fabricate.

Specifically, the January 16 IFR amends license exception AIA by adding a requirement for three of the eligible commodities for this exception: ECCNs 3A090.a; 5A002.z.1.a, z.2.a, z.3.a, z.4.a, z.5.a; and 5A992.z.1. These three ECCNs are eligible for license exception AIA only if they are designed by an "approved" or "authorized" IC designer, as described in newly issued supplement No. 6 to part 740 (list of "approved" IC designers) and new Note 1 to ECCN 3A090.a (criteria for "authorized" IC designers), respectively. BIS intends this new requirement to ensure that foundries and other entities seeking to use this exception can do so only if the items have been designed by an entity that presents a low risk of diversion.

As provided in new Note 1 to ECCN 3A090.a, "authorized" IC designers are those (1) that are in Taiwan or a location specified in Country Group A:1 or A:5, that are neither headquartered nor have an ultimate parent headquartered in Macau or a location specified in Country Group D:5 of supplement No. 1 to Part 740 and (2) for whom transactions are subject to the reporting requirements in § 743.9 of the EAR.

These additional requirements of license exception AIA are effective on January 16, 2025, but exporters, reexporters, and transferors (in-country) were not required to comply with the new requirements of this rule until January 31, 2025. "Authorized" IC designers must have submitted an application to become an "approved" IC designer by April 13, 2026, or risk losing their "authorized" status.

License Exception Advanced Compute Manufacturing (ACM)

License exception ACM is intended to limit the impact of the IFR on supply chains for otherwise-permitted advanced ICs. Accordingly, license exception ACM only applies when the reason for the export, reexport, or in-country transfer of controlled advanced ICs is the "development," "production," or storage (in a warehouse or other similar facility) of such advanced ICs.

This new license exception authorizes the export, reexport, or in-country transfer of advanced computing ICs (*i.e.*, ECCNs 3A090.a and 4A090.a and the corresponding .z items) to "private sector end-users" as long as they are not in Tier 3 countries and are not headquartered in, or have an ultimate parent headquartered in, a Tier 3 country. And the items ultimately developed, produced, or stored pursuant to the license exception must ultimately be destined for customers outside of Tier 3.[5]

As with license exception AIA, BIS modified license exception ACM in a supplemental interim final rule published on January 16, 2025, so that certain exceptions are only available for items designed by "approved" IC designers who are likely to accurately report the ECCN of the items they request advanced foundries to fabricate. Specifically, the January 15 IFR revised paragraph (b) of license exception ACM by adding a requirement for three of the eligible commodities for this exception: ECCNs 3A090.a; 5A002.z.1.a, z.2.a, z.3.a, z.4.a, z.5.a; and 5A992.z.1. These ECCNs are eligible for license exception ACM only if designed by an "approved" or "authorized" IC designer, as described in new supplement No. 6 to part 740 (list of "approved" IC designers) and new Note 1 to ECCN 3A090.a (criteria for "authorized" IC designers), respectively.

As with the corresponding new requirement for license exception AIA, the additional requirement for license exception ACM is intended to ensure that foundries and other entities seeking to use license exception ACM can do so only if they are working with entities that pose a low risk of diversion. Exporters, reexporters, and transferors (in-country) are not required to comply with this additional requirement of license exception ACM beginning on January 31, 2025.

License Exception Low Processing Performance (LPP)

This new license exception authorizes the export or reexport of low amounts of compute (*i.e.*, up to 26,900,000 TPP, roughly equivalent to 1700 high-end AI GPUs) of advanced computing ICs per calendar year to any

individual ultimate consignee located outside of Tier 3, provided the ultimate consignee is not headquartered, and does not have an ultimate parent company headquartered, in a Tier 3 country. License exception LPP does not authorize in-country transfers or exports or reexports made through distributors (*i.e.*, the items must be sent directly to the ultimate consignee).[6]

Expansion of Existing License Exceptions Notified Advanced Computing (NAC) and Advanced Computing Authorized (ACA)

License exception NAC established an advanced *notification process* to BIS for authorizing certain exports/reexports to Tier 3 countries. The notification process involves providing detailed information to BIS and its interagency partners in order to evaluate the national security risk posed by such export of eligible ICs to these countries.

The IFR updated license exception NAC to include additional information that must be provided to BIS to assist in the interagency evaluation. The new information required includes: (1) all NAC and license approvals to the end-user in the past 12 months, (2) the memory bandwidth of the item(s) requested, and (3) whether the items are destined to be aggregated into a data center or computing cluster.

License exception ACA authorizes, without the BIS notification process of NAC:

- The export, reexport, and in-country transfer to any entity located outside of Tier 3, provided the entity is not headquartered, and does not have an ultimate parent company headquartered, in a Tier 3 country, **or**
- The in-country transfer to an entity located in a Tier 3 country of any item classified in ECCNs 3A090.a and 4A090.a and the corresponding .z items (except for items designed or marketed for use in a data center and meeting the parameters of 3A090.a).

The IFR expanded the scope of license exception ACA to include countries subject to the new "worldwide" advanced IC controls.

Updated Data Center Validated End-User (DC VEU) Authorization

Data center validated end-user (DC VEU) authorizations are for entities that have been qualified by BIS to receive eligible items (*i.e.*, items classified in ECCNs 3A090.a and 4A090.a and the corresponding .z items). Eligible items must be exported, reexported, or transferred (in-country) directly to the VEUs in eligible destinations (this authorization is not available for exports, reexports, or in-country transfers made through distributors).

For a data center to be eligible as a DC VEU, both its owner and operator must have VEU authorizations. To qualify as a VEU, the owner or operator of the data center must apply to BIS and go through an intensive application process that will be subject to interagency review. As part of the application, the owner or operator must attest to the Security Measures and must comply with all other guidelines in Supp. 10 to Part 748 (*e.g.*, vetting requirements, export restrictions, acceptable use policies, documentation, auditing, reporting, etc.). Information required includes, among other things, information about business activities and corporate relationships with government or military organizations in Tier 3 (*e.g.*, direct sales to or contracts with such entities).

The IFR split the DC VEU authorization into two types—universal validated end-users (UVEUs) and national validated end-users (NVEUs).

UVEUs

Entities headquartered, or whose ultimate parent is headquartered, in a Tier 1 country may apply for the UVEU authorization. UVEUs may receive advanced computing ICs without any TPP limit and may install such advanced computing ICs in any Tier 1 or Tier 2 country. However, UVEUs are subject to limitations on where they can geographically allocate their total AI computing power (*i.e.*, the AI computing power owned by the entity and all its subsidiary and parent entities), measured by the aggregate TPP of chips that meet or exceed the scope of ECCN 3A090.a. Specifically:

- A UVEU headquartered in a Tier 1 country may not transfer or install more than 25% of its total AI computing power to or in locations in Tier 2 countries and may not transfer or install more than 7% of its total AI computing power to or in any single country in Tier 2.
- A UVEU headquartered in the United States may not transfer or install more than 50% of its total AI computing power outside of the United States.

Approved UVEUs will be listed in the EAR, along with each address at which they are authorized to receive controlled advanced computing ICs. To add new data centers to its authorization, a UVEU must notify BIS 180 days prior to any exports, reexports, or in-country transfers to that location.

NVEUs

With certain limitations, entities that are not headquartered, and whose ultimate parent is not headquartered, in a Tier 3 country may apply for the NVEU authorization. The NVUE authorization is country specific and is subject to a per-company, per-country installed base allocation of TPP provided in a table in the IFR[7], measured by the aggregate TPP of chips that meet or exceed the scope of ECCN 3A090.a.

Approved NVEUs will be listed in the EAR.

New Foreign Direct Product (FDP) Rules

AI Model Weights FDP Rule

This new FDP rule applies to foreign-produced closed-weight AI models that meet or exceed the specifications of ECCN 4E091. Under the new FDP rule, such foreign-produced model weights remain subject to the EAR if they are produced by a complete plant or “major component” of a plant located outside the United States, when such plant or major component is subject to the EAR and specified in ECCNs 3A090.a and 4A090.a and the corresponding .z items. In other words, the new FDP rule controls foreign-produced AI models trained on EAR-controlled advanced computing ICs.

Expanded Advanced Computing FDP Rule

The existing advanced computing FDP rule required a license for certain foreign-produced advanced computing items that are the “direct product” of specified U.S. software or technology or a product of a complete plant or “major component” of a plant that is a “direct product” of specified U.S. software/technology, to or within certain countries of concern (China, Hong Kong, and Macau).

The IFR expands this FDP rule to apply whenever there is “knowledge” that the foreign-produced item is: (1) destined to any location worldwide or incorporated into any “part,” “component,” “computer,” or “equipment” not designated EAR99 destined to any location worldwide; or (2) “technology” “developed” by an entity headquartered, or whose ultimate parent company is headquartered, in a Tier 3 country for the “production” of a mask or an integrated circuit wafer or die.

Addition of AI "Red Flag"

BIS "red flags" are part of the agency's "Know Your Customer" guidance. The red flags are examples of situations, exporters, reexporters, or transferors one may encounter that may signal a potential violation of the EAR. BIS added a new AI "red flag 28," intended to help Infrastructure-as-a-Service (IaaS) cloud computing providers identify when training an advanced AI model for a customer that is a Tier 1 subsidiary of a non-Tier 1 parent company—and transferring the resulting model weights to that customer—creates a potential diversion concern.[8]

Important Dates and Deadlines Under the IFR

Effective date. The IFR took effect on January 13, 2025. However, compliance is not required until 120 days from the date of publication (*i.e.*, May 15, 2025). The Security Measures discussed above have a delayed compliance date of 365 days from publication (*i.e.*, January 15, 2026). This allows time for end users to ensure they have instituted the required security measures.

Comment due date. Comments on revisions and additions in this rule must be received by BIS no later than 120 days after the date of publication in the *Federal Register* (*i.e.*, May 15, 2025).

Takeaways and Recommendations

- Companies should carefully consider how the AI framework will affect their sales and/or sourcing of advanced computing equipment, data center implementation strategy, and development and use of AI models both now and into the future and consider submitting comments to BIS regarding such effects.
- The IFR is an interim rule issued by the prior administration. It remains to be seen how much the regulations contained in the IFR will change in the final rule issued by the Trump administration.
- Nevertheless, the IFR went into effect on January 13 and requires compliance by May 15, 2025. The IFR contains arguably the most complex export controls and license exceptions ever promulgated by BIS, so companies should start now to determine how to bring their export control compliance policies and procedures into compliance with the IFR.
- The 10^26 operations threshold for new ECCN 4E091 represents roughly double the training of the most advanced current AI models. Accordingly, existing AI models will not be affected. However, the advanced computing IC controls will likely have immediate practical effect.

Endnotes

[1] October 2022 "Implementation of Additional Export Controls: Certain Advanced Computing and Semiconductor Manufacturing Items; Supercomputer and Semiconductor End Use; Entity List Modification," and expanded in October 2023 "Implementation of Additional Export Controls: Certain Advanced Computing Items; Supercomputer and Semiconductor End Use; Updates and Corrections," April 2024 "Implementation of Additional Export Controls: Certain Advanced Computing Items; Supercomputer and Semiconductor End Use; Updates and Corrections; and Export Controls on Semiconductor Manufacturing Items; Corrections and Clarifications," October 2024 "Expansion of Validated End User Authorization: Data Center Validated End User Authorization," and December 2024 "Foreign-Produced Direct Product Rule Additions, and Refinements to Controls for Advanced Computing and Semiconductor Manufacturing Items."

[2] ECCNs 3A001.z.1.a, z.2.a, z.3.a, z.4.a; 4A003.z.1.a, z.2.a; 4A004.z.1; 4A005.z.1; 5A004.z.1.a, z.2.a; and 5A992.z.1.

[3] If the Tier 2 country has not yet exhausted its TPP quota, the export license application will be reviewed by BIS under a presumption of approval. If the Tier 2 country has already exhausted its TPP quota, the export license application will be reviewed by BIS under a policy of denial (*i.e.*, will be denied).

[4] Such requirements include that the (1) items received will not be used to provide IaaS access sufficient to train an AI model classified under ECCN 4E091 to entities headquartered or located outside of, or whose ultimate parent company is head-quartered outside, the Tier 1 countries and (2) ultimate consignee will not export, reexport, or transfer the items to any end use or end user prohibited by EAR Part 744 or to an entity headquartered or located outside of, or whose ultimate parent company is headquartered outside of, the Tier 1 countries.

[5] Note that license exception ACM shipments do not count toward the per-country TPP quotas discussed above. Exporters, reexporters, and transferors using this license exception must maintain a system to account for the number of controlled items transferred into and out of each facility, and such accounting must be done at least every six months.

[6] Prior to using license exception LPP, the exporter or reexporter must obtain from the ultimate consignee a certification that it is not above the calendar-year limit, and then the exporter/reexporter must provide a copy of the certification to BIS within 30 days from the date on which the export or reexport of the items occurred. Any ultimate consignee that hits the TPP limit in a calendar year must notify BIS. In addition, exporters and reexporters must notify BIS of any export or reexport to a single ultimate consignee of TPP of more than 3,200,000.

[7] The following is the NVEU allocation table:

Quarter	Cumulative TPP per-company per-country allocation
2025 Q1	633,000,000
2025 Q2	949,500,000
2025 Q3	1,266,000,000
2025 Q4	1,582,500,000
2026 Q1	1,899,000,000
2026 Q2	2,690,250,000
2026 Q3	3,481,500,000
2026 Q4	4,272,750,000
2027 Q1-4	5,064,000,000

[8] Red flag 28 states, “You will be providing Infrastructure-as-a-Service (IaaS) products or services, or other computing products or services, to assist in training an AI model with model weights captured by ECCN 4E091 for an entity headquartered, or whose ultimate parent is headquartered, in any destination other than those listed in paragraph (a) of supplement no. 5 to part 740 of the EAR. Such assistance creates a substantial risk that such AI model weights, due to their digital nature, will be exported or reexported to a destination for which a license is required and, if a license is not obtained, that the IaaS provider will have aided and abetted in a violation of the EAR. In such cases, the IaaS provider should inquire if the customer intends to export the model and if so, apply for a license as required or inform the customer of their obligation to do so prior to export.”

Authors

Explore more in

Artificial Intelligence & Machine Learning