SEMICONDUCTOR INDUSTRY ASSOCIATION

Comments of the
Semiconductor Industry Association

On

The Interim Final Rule Entitled
"Framework for Artificial Intelligence Diffusion"

90 Fed. Reg. 4544 (January 15, 2025)
RIN 0694-AJ90
Docket No. 250107-0007

Submitted April 15, 2025

The Semiconductor Industry Association (SIA) submits these comments in response to the request from the Bureau of Industry and Security (BIS) within the Department of Commerce (Commerce) in the interim final rule ("IFR" or "AI Diffusion Rule") entitled "Framework for Artificial Intelligence Diffusion," 90 Fed. Reg. 4544.

We understand export controls are an important tool for safeguarding national security, including by preventing adversaries and malign actors from leveraging American technologies to advance military capabilities and other interests counter to the United States. We also agree with President Trump that the United States must sustain and enhance America's global AI dominance, but "unnecessarily burdensome requirements…would stifle private sector innovation and threaten American technological leadership."[1]

We submit these comments 30 days ahead of the comment due date to ensure the Trump Administration has time to calibrate its approach informed by our inputs and perspectives. Time is of the essence. Absent any change to the IFR or its compliance date, exporters, re-exporters, and transferors must comply with the changes made in this IFR starting on May 15, 2025, the same day the comment period ends.

AI is pervasive in mainstream computing. The U.S. approach to regulating the diffusion of AI around the world will have profound impacts on SIA member companies and will influence whether America will maintain its global leadership on AI or cede global markets to alternative technologies and ecosystems. SIA strongly believes the IFR is flawed and warrants a significant rethink, as our forthcoming comments will elucidate.

Part I of these comments contains background comments about SIA and the role of semiconductors in artificial intelligence (AI). Part II outlines general comments on the IFR and articulates key principles around which regulatory action to control the export of

---

[1] The White House, "Fact Sheet: President Donald J. Trump Takes Action to Enhance America's AI Leadership," January 23, 2025. https://www.whitehouse.gov/fact-sheets/2025/01/fact-sheet-president-donald-j-trump-takes-action-to-enhance-americas-ai-leadership/

advanced AI chips should align. Part III of this submission articulates a set of non-exhaustive comments on specific provisions of the IFR to demonstrate its complexity and highlight some of its impracticality.

Part I – Background on SIA and Role of Semiconductors in AI

SIA has been the voice of the U.S. semiconductor industry for almost 50 years. Our member companies represent more than 99% of the U.S. semiconductor industry by revenue and are engaged in the full range of research, design, and manufacture of semiconductors – including both wafer fabrication and back-end assembly, test, and packaging of chips. Semiconductor technology was invented in America more than 65 years ago, and the U.S. remains the global leader in semiconductor technology and innovation, driving America's economic strength, national security, and global competitiveness in a range of downstream industries. More information about SIA and the semiconductor industry is available at https://www.semiconductors.org/.

Semiconductor technology comprises the computing, memory, and networking backbone that powers and enables AI systems. Advanced AI applications – from natural language processing to autonomous systems – benefit from high-performance chips called AI accelerators. AI accelerators include processing logic highly effective for AI training and inferencing workloads (e.g., GPUs, CPUs, ASICs) and often include high bandwidth memory (HBM) to accommodate the rapid movement of increasing data volumes, along with networking and optical connectors, among other critical components. AI servers are comprised of AI accelerators and a host of mature-node semiconductor components, including power chips, analog-to-digital converters (ADCs), digital-to-analog converters (DACs), and input-output controllers.

Together, the entire semiconductor supply chain, involving many U.S. semiconductor design, manufacturing, and manufacturing equipment companies, enables the production of AI systems. In short, without semiconductors there is no AI. U.S. economic security, national security, and sustained American leadership in AI leadership rest on the continued competitiveness and adoption of U.S. semiconductor technology globally.

Part II – General Comments

SIA and our member companies stand ready to help inform the Trump Administration's approach to AI governance, particularly with respect to export controls. Unfortunately, the AI Diffusion Framework, embodied in the IFR, was developed and sent to the Office of Management and Budget for publication in the final days of the Biden Administration without formal industry consultation. It introduced a novel, complex, and economically significant export control framework for governing the global diffusion of U.S. semiconductor and AI technologies around the world. The approach promulgated through the IFR is poorly calibrated, uninformed by accurate data, market realities, and industry expertise, and risks undermining U.S. leadership.

The claim in the IFR's preamble that the rule will "address only the starkest risks identified at the frontier of AI development and [will] not affect the vast majority of AI technology" is simply incorrect.[2] Rather, the rule imposes sweeping U.S. government controls and a worldwide licensing regime with a country-by-country quota system on a vast array of mainstream computing products – regardless of use case or cluster size, and for both data center and consumer products. The controls also apply to more mature, less advanced chips, compute modules, and servers already deployed around the world and used by a diverse range of consumer industries, such as gaming, education, and agriculture, which have no nexus to national security.

Given its scope, the IFR threatens the long-term competitive position of U.S. industry. It risks ceding key markets and weakening U.S. semiconductor and AI leadership globally by incentivizing the adoption of alternate ecosystems for AI infrastructure and applications and positioning the U.S. as an unreliable supplier.

As of May 15, 2025, the heightened level of transaction monitoring, recordkeeping, and other compliance obligations imposed by this IFR will require degrees of oversight and bureaucracy for both industry and government that stand at odds with the policy outlined in President Trump's Executive Order 14179, "Removing Barriers to American Leadership in Artificial Intelligence."[3] It is also unclear whether BIS has the resources to implement and enforce the IFR, particularly at a time when the federal government is restructuring and prioritizing reductions in force.

Going forward, we ask that any approach to protecting U.S. national security and promoting U.S. leadership, particular with respect to U.S. semiconductor technologies that enable AI, should:

- Ensure the United States remains the global leader in the design and manufacture of the semiconductor technologies that enable AI;
- Avoid incentivizing the creation and promulgation of foreign alternatives to the U.S. AI and semiconductor technology stack, which drives global customers to design out American technologies;
- Reflect accurate data, economic and market realities, and a verifiable assessment of whether foreign items comparable in quality and in sufficient quantities are available to render U.S. export controls ineffective; and
- Be simple and straightforward for businesses to comply with and for BIS to enforce.

[2] Bureau of Industry and Security, "Export Control Framework for Artificial Intelligence Diffusion," 90 Fed. Reg. 4544 (Jan.15, 2025). https://www.federalregister.gov/documents/2025/01/15/2025-00636/framework-for-artificial-intelligence-diffusion

[3] The White House. "Removing Barriers to American Leadership In Artificial Intelligence," Executive Order 14179, January 23, 2025. https://www.whitehouse.gov/presidential-actions/2025/01/removing-barriers-to-american-leadership-in-artificial-intelligence/

SIA and our member companies hope to work in partnership with the U.S. government to achieve these goals with respect to AI diffusion, starting with a re-think of the IFR. We offer the following, non-exhaustive comments on the IFR.

<u>Part III – Specific Comments on Provisions related to Country Tiers and Related Country Caps in the IFR</u>

Comment III.A: The country tiers outlined in the IFR, including in paragraph (a) of Supplement No. 5 to Part 740, should align with the existing country groups in Supplement No. 1 to Part 740 of the Export Administration Regulations (EAR).

The IFR introduced new groupings of destination countries into three "tiers" that determine the application of export controls on AI-related technologies. Specifically, the rule creates License Exception Artificial Intelligence Authorization (AIA) in new § 740.27 for all transactions involving certain types of end users in 18 specific low-risk destinations (i.e., "Tier 1" countries) identified in Supplement No. 5 to Part 740. As noted in the IFR preamble, Tier 1 countries are destinations in which: "(1) the government has implemented measures to prevent diversion of advanced technologies, and (2) there is an ecosystem that will enable and encourage firms to use advanced AI models to advance the common national security and foreign policy interests of the United States and its allies and partners." While BIS makes no changes to export control restrictions on advanced semiconductors vis-a-vis Country Group D:5 destinations and Macau (i.e., "Tier 3" countries), the IFR creates a series of complex rules for all other country destinations not in Tier 1 (i.e., "Tier 2" countries) that condition access to items classified in 3A090.a, 4A090.a, and corresponding ".z" ECCNs based on certain factors (e.g., the volume of the transaction, the security measures agreed to by the recipient).

The well-established country groups in Supplement No. 1 to Part 740 of the EAR, including Country Groups A, B, D, and E, are used across various export control programs, including in recent semiconductor technology-related controls.[4] These groupings reflect thoughtful assessments of national security risk, export control regime membership, foreign policy concerns, and economic partnerships. In particular, Country Groups A:5 and A:6 are comprised of "low-risk" countries broadly aligned with U.S. foreign policy interests. Adopting an entirely new set of destination countries introduces complex compliance challenges for industry, for countries, and for BIS administrators. Beyond this, more than twenty NATO allies (e.g. Poland and Czechia) are inexplicably included in Tier 2, rather than as trusted partners in Supplement No. 5 to Part 740.

Comment III.B: Country compute caps do not reflect market realities. They will artificially distort the global marketplace for advanced compute and incentivize end users in Tier 2 countries to seek alternatives to U.S. AI technology.

---

[4] For example, "Advanced Computing Rule," 87 Fed. Reg. 62186 (Oct. 13, 2022), "Semiconductor Manufacturing Equipment Controls," 88 Fed. Reg. 73458 (Oct. 25, 2023), "Validated End-User (VEU) Authorization," 89 FR 80080 (Oct. 2, 2024).

The IFR establishes pre-set country-specific quotas, or “caps,” of 790 million cumulative total processing performance (TPP) for export, reexports, and in-country transfers of Advanced Computing ICs to or within each Tier 2 country. BIS will track the fill rate of each country-specific TPP cumulative quota, and once a country-specific quota is reached, BIS will review license applications for that country with a presumption of denial. However, these caps do not account for the speed at which industry is moving in terms of consumption or use of compute. The country caps also do not account for items that are in a country temporarily, such as where a license is sought to permit a one-for-one replacement of failed items, such as in a situation where License Exception RPL for servicing and replacement of parts and equipment is not available.

The caps will also distort the marketplace involving otherwise acceptable transactions, potentially leading to anti-competitive behavior. The existence of a country cap will encourage applicants to file for licenses early and for larger quantities than they would under normal economic conditions to ensure their allotments are within the cap, as well as to block competition. For example, a single cloud hyperscaler can fulfill an entire country compute allotment, or “cap,” leaving no compute allotments for the rest of the provider market.

The Universal Validated End User (UVEU) framework outlined in the IFR also appears to be designed to create a global oligopoly – concentrating market power in the hands of only a few large cloud service providers, to the detriment of smaller U.S. cloud providers and technology companies. For UVEU holders that are headquartered in the United States, the rule also locks in place the amount of processing performance that can be added outside the United States and within any individual country without taking any steps to support the installation of additional computing capacity in the United States.

Further, the uniform nature of the country caps across all Tier 2 countries results in significant inequities for the allocation of compute power. Countries with substantially different populations and computing needs are allocated the same amount of computing power under the existing country cap (e.g., India vs. Luxembourg).

As a result of the above challenges, the rule will have the perverse effect of pushing foreign customers to opt for alternative compute ecosystems even if those alternatives are less advanced and secure than U.S. technology. Non-U.S. cloud providers are actively expanding in emerging markets, aiming to establish themselves as the reliable AI provider within regional marketplaces.[5] An overly complex and burdensome export

[5] Ann Cao, *South China Morning Post*, “Alibaba Cloud roles out expanded suite of AI models, development tools in overseas push,” January 21, 2025. https://www.scmp.com/tech/big-tech/article/3295689/alibaba-cloud-rolls-out-expanded-suite-ai-models-development-tools-overseas-push; Iris Deng, *South China Morning Post*, “Huawei unveils Arabic LLM, new data centre in Egypt as part of revenue diversification strategy,” May 21, 2024. https://www.scmp.com/tech/big-tech/article/3263533/huawei-unveils-arabic-llm-new-data-centre-egypt-part-revenue-diversification-strategy

framework, as embodied by the IFR, will only make these alternative AI solutions more attractive at the expense of U.S. technological leadership.

Comment III.C: The requirements imposed by other EAR regulations, such as the "Foundry Due Diligence" rule, add further complexity and uncertainty to compliance with the IFR and lead to repeated counting against the country caps.

The U.S. government has promulgated multiple, consequential export control rules targeting semiconductor technologies in the past few years, creating a complex web of compliance burdens that are difficult for industry to navigate and understand. For example, days after publication of the IFR, BIS published a separate rule, "Implementation of Additional Due Diligence Measures for Advanced Computing Integrated Circuits",[6] or the "Foundry Dule Diligence Rule," which shifts export controls upstream by applying identical global license requirements to wafers and dies irrespective of their ultimate application or end use. The Foundry Due Diligence Rule establishes a presumption in Note 1 to 3A090.a. that all semiconductor wafers and dies dedicated to logic ICs produced using 16/14 nanometer or below technology or with non-planar transistor architecture are for data center uses and meet 3A090.a. control parameters, creating a cascading series of regulatory challenges.

For example, as a result of the Foundry Due Diligence Rule, the attribution of the TPP calculation applied to these wafers and dies as these items move through multi-stage, multi-country production processes results in their computing power being repeatedly counted against the country caps (See Annex I). This creates a compounding calculation such that the resulting cumulative TPP calculation for a country will have no association with the actual computational potential or national security concerns. Countries with foundries and Outsourced Semiconductor Assembly and Test (OSAT) operations would thus exceed TPP computational caps due to the massive volume of wafer and die movements.

Exporters inherently face additional uncertainty when attempting to comply with the regulatory requirements of both the IFR and Note 1 to ECCN 3A090.a. For example, for wafers and dies from designers not on BIS's "approved" list, Note 1 to ECCN 3A090.a stipulates that foundries and OSATs may not rely on the designer's specifications. As a result, foundries and OSATs will have no information to calculate the TPP – a calculation required both to ship the wafer (or provide TPP information BIS expects in the license application) and to calculate the applicable country compute cap under this IFR. Even if the foundry transfers the wafer to an OSAT for packaging, the OSAT cannot accurately determine the TPP to comply with the country compute cap in the IFR (See Annex II).

---

[6] "Implementation of Additional Due Diligence Measures for Advanced Computing Integrated Circuits," 90 Fed. Reg. 5289 (Jan. 16, 2025). https://www.federalregister.gov/documents/2025/01/16/2025-00711/implementation-of-additional-due-diligence-measures-for-advanced-computing-integrated-circuits.

Comment III.D: The scope of License Exception Advanced Compute Manufacturing (ACM) is too limited, increasing compliance burden for industry and license burden for the government.

The IFR creates new License Exception ACM under EAR § 740.28, which authorizes the export, reexport, and in-country transfer of Advanced Computing ICs and associated software and technology (ECCNs 3A090, 4A090, and corresponding .z items) to private sector end users located in a Tier 1 or Tier 2 destination, provided the private sector end user is not headquartered in, and does not have an ultimate parent company headquartered in, a Tier 3 jurisdiction. It also requires the transferred items' ultimate end use to be in the "development, production, or storage (in a warehouse or other similar facility)" of items covered by the license exception.

However, "maintenance (checking)," "repair," "overhaul," and "refurbish" are not elements of the EAR's definitions of "development" or "production" but are rather part of the EAR's definition of "use." As such, unless another license exception applies, companies will need to apply for licenses to conduct those activities, as they are not authorized under ACM – increasing the compliance burden for both industry and for the government to process unnecessary license applications.

Companies must also maintain unnecessary internal accounting records for exports, reexports, and transfers by and among corporate affiliates (e.g., parent, subsidiary, and sister companies) using License Exception ACM, unlike other license exceptions (e.g. License Exception ENC) which do not require such recordkeeping or reporting.

Comment III.E: The TPP Limit under License Exception Low Processing Performance (LPP) is already outdated due to the pace of technological advancements.

The IFR under § 740.29 created a new License Exception LPP to authorize the export and reexport (but not in-country transfer) of Advanced Computing ICs up to 26.9 million TPP per calendar year to an individual consignee located in a Tier 1 or Tier 2 destination that is not headquartered in, and does not have an ultimate parent headquartered in, a Tier 3 destination.

Both the number of systems in a data center cluster for foundational models and the TPP of individual components of those systems and clusters are rapidly increasing. For example, some SIA member companies have seen a nearly fourfold increase in their products' TPP over a two-year period and they expect a doubling or tripling of TPP to continue into future products. Capping the TPP limit under License Exception LPP fails to reflect the spirit and purpose of the exception, i.e., to avoid impacting ordinary sales of 3A090.a, 4A090.a, and related .z items for data centers that are not creating "frontier models." The threshold contemplated in the IFR is likely to be quickly outdated and fails to acknowledge such a reality.

Comment III.F: The "ultimate consignee" definition introduced in LPP conflicts with existing definitions of the term in the EAR, complicating compliance.

The IFR introduces a new definition for "ultimate consignee" in §740.29 for license exception LPP, namely as "the ultimate parent entity that has ultimate ownership" of the eligible commodities. The EAR separately defines "ultimate consignee" under §748.5 and §772.1 of the EAR as "the principal party in interest located abroad who receives the exported or reexported items. The ultimate consignee is not a forwarding agent or other intermediary but may be the end-user."

The term "ultimate consignee" is a common term throughout the EAR and in export activities. The introduction of a conflicting definition for the same term for purposes of a single license exception adds unnecessary complexity to the regulations and creates confusion for export control compliance personnel.

Comment III.G: The exclusion of sales through distributors or for in-country transfers from the LPP License Exception, even when the OEM knows the end-user, unnecessarily increases compliance burden.

New License Exception LPP "authorizes the export and reexport of low amounts of compute that do not present significant national security risks" but is "not available for exports or reexports made through distributors or for in-country transfers." This limitation in LPP is intended to ensure authorized and trusted parties shipping under the license exception do not lose visibility of the identity of the ultimate end-user.

But because the products controlled by this IFR are complex, the OEM supplying the items for a data center must understand the intended end use, the networking requirements, and the energy consumption to ensure that the correct products are provided. This typically means that the OEM knows the end user before supplying an item, even if the items will first be sold through one or more distributors. Such a scenario illustrates that it is likely not BIS's intent to limit distribution models in this way when the ultimate consignee (i.e., end user) is known to the exporter. However, this limitation as drafted in the IFR creates needless burden on industry as well as BIS in the form of avoidable resulting licenses. This is yet another example of how the provisions in the IFR are overcomplicated and unnecessarily increase the licensing burden on companies.

* * *

The complexity, breadth, and challenges of the IFR, as highlighted above, place an artificial cap on the growth of U.S. exports, incentivize the development and diffusion of non-U.S. alternative AI compute ecosystems, and will lead to the design-out of U.S. semiconductor technologies around the world. This outcome undermines the policy of the United States "to sustain and enhance America's global AI dominance in order to promote human flourishing, economic competitiveness, and national security."[7]

The specific comments in this submission are meant to highlight particularly complex provisions in the IFR and the undue compliance burden it places on both industry and the government. They are, however, non-exhaustive. Other aspects of the rule that warrant closer review and modification include simplifying License Exceptions Advanced Computing Authorized (ACA) and Notified Advanced Computing (NAC), defining the term "headquartered," and establishing a clear process for companies to confirm their controlling HQ location for export control compliance purposes.

Finally, we note that while semiconductor hardware has become the primary – if not singular – target of recent export control efforts aimed at restricting AI, there are other means through which entities can leverage AI compute power without having physical access to controlled semiconductor hardware, and which fall outside semiconductor companies' ability to control.

SIA and its member companies stand ready to partner with BIS and other agencies in providing support and feedback on this IFR and other export control policies with respect to semiconductors.

Uploaded to www.regulations.gov.   ID – BIS-2025-0001-0001

[7] "Removing Barriers to American Leadership In Artificial Intelligence," Executive Order 14179.

Annex I: Hypothetical manufacturing flow illustrating compounding calculation of TPP

As wafers and dies move through multi-stage, multi-country production processes, their TPP is repeatedly counted, resulting in an inflated TPP of 2.7M across all countries (sum of each country's TPP) for an actual computational capacity of 0.9M TPP.

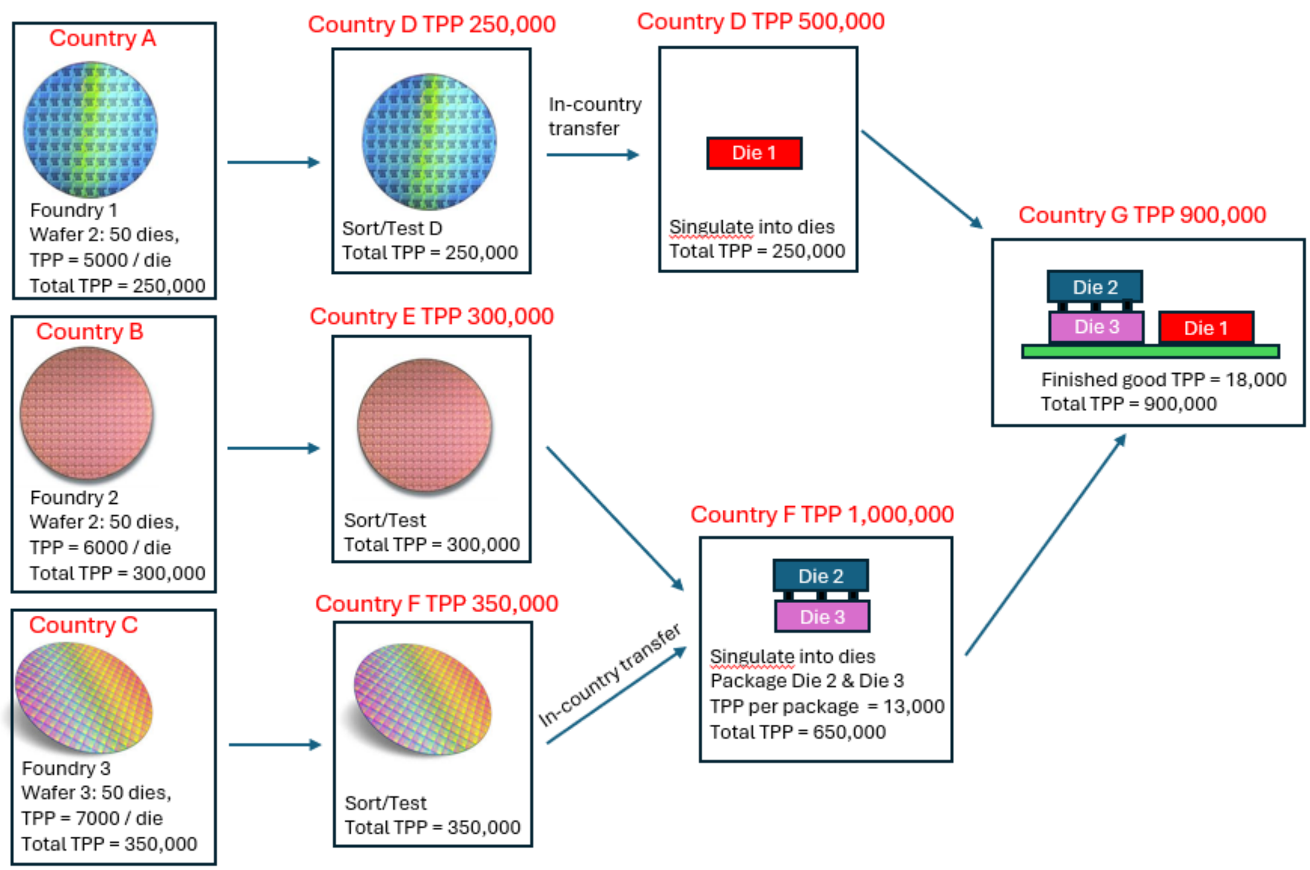

<u>Annex II: Illustration showing unknown TPP and ECCN in manufacturing flow</u>