CHAMBER OF PROGRESS

October 9, 2024

Department of Commerce
Under Secretary of Commerce for Industry and Security

Bureau of Industry and Security
14th St NW & Constitution Ave. NW
Washington, DC 20230

Re: RIN 0694-AJ55, Establishment of Reporting Requirements for the Development of Advanced Artificial Intelligence Models and Computing Clusters

On behalf of Chamber of Progress—a tech industry association supporting public policies to build a more inclusive country in which all people benefit from technological advances—we appreciate the opportunity to share this response to the U.S. Bureau of Industry and Security's (BIS) Request for Comment on the Establishment of Reporting Requirements for the Development of Advanced Artificial Intelligence Models and Computing Clusters.

The United States is the global leader in artificial intelligence (AI). As BIS considers its approach to AI, particularly in dual-use models, it should prioritize preserving America as the epicenter of AI innovation. While we understand the need for oversight in guarding U.S. national security, overly prescriptive frameworks stifle innovation by creating high barriers for startups and smaller firms, limiting their ability to compete and innovate. We encourage BIS to pursue a balanced regulatory approach that fosters a competitive and dynamic AI ecosystem while at the same time addressing specific national security risks without imposing undue burdens on the AI development community.

Ensuring a Balanced Approach: "Let a Thousand Flowers Bloom":

Ensuring a diverse array of foundation models, both open-source and closed-source, is essential for fostering competition in the technology sector. We therefore caution against policies that favor one type over the other, and instead urge the adoption of a balanced framework that allows both to thrive. As

currently drafted, **Section 702.7's emphasis on the need for "physical and cybersecurity protections"[1] to safeguard the integrity of AI model training places open-source developers at a disadvantage**. "Ownership and possession" of open-source models are not always clear due to their collaborative and decentralized nature thus it is not obvious how major open source models would or could comply. . Similarly, the **model weights reporting requirement[2]**, while manageable for closed-source models with centralized ownership, could impose unique and unclear burdens on open-source developers, potentially stifling innovation and chilling critical investment across AI development generally.

Broad Rules Hinder Innovation:

Concerningly, the proposed rule incorrectly assumes all dual-use models carry similar risks. For instance, the definition of **"dual-use foundation models" includes models with "at least tens of billions of parameters" and broad applicability across multiple contexts.[3]** In practice, this would impose uniform reporting requirements without distinguishing between relatively higher- and lower-risk models. Research shows that such stringent regulations like this disproportionately harm smaller players, locking them out of the market and stifling innovation.[4] As noted by the *Center for a New American Security*, growing secrecy and costs in AI development could further delay the release of competitive models by smaller labs, ultimately consolidating market power among a few large corporations.[5] This risks creating a regulatory environment that discourages the development of next-generation AI technologies, thereby diminishing the U.S.'s competitive edge in AI innovation. Moreover, overly broad regulations, as highlighted by research from the Harvard Kennedy School, could hinder technical progress by imposing undue burdens across the board instead of focusing on specific threats.[6]

---

[1] *See* § 702.7 (Special requirements for on-going reporting regarding the development of advanced artificial intelligence models and computing clusters.)
[2] *Id.*
[3] *See* § 702.7 (Definition of "Dual-use foundation model")
[4] *See* https://www.nber.org/papers/w28381
[5] *See* https://www.cnas.org/publications/commentary/response-to-ntia-request-for-comment-dual-use-foundation-artificial-intelligence-models-with-widely-available-model-weights
[6] *See* https://www.hks.harvard.edu/centers/mrcbg/publications/dual-imperative-innovation-and-regulation-ai-era

Relatedly, given AI technology's rapid progress, the **collection thresholds[7] may quickly become outdated**, further stifling innovation. Specifically, the requirement around reporting of computing clusters exceeding specific computational thresholds, along with detailed technical information about those clusters, may bring many more companies, particularly small under-resourced startups, under the reporting requirements and impose high compliance costs that smaller entities are unable to pay. Training compute for AI models has been doubling roughly every 6 months, further underscoring the fact that these thresholds may no longer be relevant as model developments technology advances.[8] While ensuring AI development is aligned with US national security concerns is paramount, this approach could discourage mission-critical investment in next-generation AI technologies by increasing operational costs and exposing proprietary information to undue scrutiny.

Focusing on Specific Risks, Not Blanket Reporting

The government's stated aim is to ensure that "dual-use foundation models operate in a safe and reliable manner." However, in doing so, BIS should consider any specific risks posed by AI models, rather than adopting broad mandates. Section 702.7(b)(2) requires companies to report on physical and cybersecurity protections, ownership, and performance testing of their models, including red-team[9] testing results. While these are valuable tools, the broad application of these requirements could impose unnecessary compliance burdens on companies whose foundation models or computer clusters do not pose significant threats to US national security interests. Furthermore, the red-teaming requirements may lead to the disclosure of proprietary information, which could be exploited by competitors or hostile actors, undermining America's leadership. Instead, BIS should shift its focus on clear, identifiable risks, rather than applying broad mandates across all dual-use foundation models. Specific areas of interest

---

[7] *See* § 702.7 "Special requirements for on-going reporting regarding the development of advanced artificial intelligence models and computing clusters."

[8] *See* https://arxiv.org/html/2405.10799v2 ("Since the emergence of the Deep Learning Era around 2010, training compute has been increasing at a much faster rate, doubling roughly every 6 months (increasing by about 4× per year).")

[9] *See* § 702.7 "AI red-teaming means a structured testing effort to find flaws and vulnerabilities in an AI system, often in a controlled environment and in collaboration with developers of AI. In the context of AI, red-teaming is most often performed by dedicated "red teams" that adopt adversarial methods to identify flaws and vulnerabilities, such as harmful or discriminatory outputs from an AI system, unforeseen or undesirable system behaviors, limitations, or potential risks associated with the misuse of the system."

might be where dual-use models pose specific cybersecurity vulnerabilities or misuse by malicious actors.

We commend the BIS for taking the initiative to establish rules around dual-use foundation models and computing clusters that align with the aim of keeping Americans safe. This reflects a commitment to responsible AI development and ensuring the U.S. remains a leader in technological innovation. However, we urge BIS to reconsider aspects of the proposed rule that could stifle competition and innovation and thus fail any cost-benefit analysis. A balanced approach that addresses specific risks without imposing undue burdens on companies developing foundation models is essential to maintaining the U.S.'s competitive edge in AI and drive overall economic growth for all Americans.

We appreciate the opportunity to provide feedback and look forward to continued engagement on this issue.

Sincerely,

Andres Calzada
Policy Fellow
Chamber of Progress