Articles + Publications | September 27, 2024

Locke Lord QuickStudy: BIS Proposes New Reporting ?Requirements for the Development ?of Advanced AI Models & ?Computing Clusters

Locke Lord LLP

WRITTEN BY

Ryan Last

On September 9, 2024, U.S. Department of Commerce's Bureau of Industry and Security (BIS) issued a Notice of Proposed Rule Making ("NRPM"), which would mandate new reporting requirements for artificial intelligence ("AI") developers and cloud computing providers. The proposed rule would amend BIS's "Industrial Base Surveys—Data Collections" regulations by establishing reporting requirements for the development of advanced AI models and computing clusters under the Executive Order 14110 of October 30, 2023, "Safe, Secure, and Trustworthy Development and Use of Artificial Intelligence" ("EO 14110").

As background, the BIS is responsible for implementing and enforcing the U.S. Export Administration Regulations (the "EAR"), which regulates exports and deemed exports of "dual use goods," goods and technologies that have potential for both commercial and military application. Section 4.2(a)(i) of EO 14110, directs the Secretary of Commerce to require companies developing, or demonstrating an intent to develop, potential dual-use foundation AI models to provide certain information to the Federal Government on an ongoing basis. Section 4.2(a)(ii) of EO 14110 directs the Secretary of the U.S. Department of Commerce to require companies, individuals, or other organizations or entities that acquire, develop, or possess a potential large-scale computing cluster to report any such acquisition, development, or possession, including the existence and location of these clusters and the amount of total computing power available in each cluster.

As defined under EO 14110, a "dual-use foundation model" is "trained on broad data; generally uses self-supervision; contains at least tens of billions of parameters; is applicable across a wide range of contexts; and that exhibits, or could be easily modified to exhibit, high levels of performance at tasks that pose a serious risk to security, national economic security, national public health or safety, or any combination of those matters." The NRPM proposed reporting requirements would apply to dual-use foundation models that meet the technical standard issued by BIS, which BIS expects to update with industry technological advancements (pursuant to section 4.2(b) of EO 14110).

The NRPM aims to enhance government oversight of emerging and foundational models and large-scale computing clusters by collecting data and testing resiliency against misuse, with the intention to limit their export to mitigate national security risks.

Comments on this proposed rule should be submitted to the Federal rulemaking portal (www.regulations.gov) no later than October 11, 2024.

Key Requirements

- Covered Entities: U.S. persons and entities engaging in applicable activities, such as running AI models with more than 10^26 computational operations (e.g., integer or floating-point operations) or developing large-scale computing clusters, would be required to submit to BIS the information described below quarterly.
- Report Content: Entities must report on:
  - Developmental activities of dual-use AI models, including physical and cybersecurity measures taken to safeguard models.
  - Ownership and protection of AI model weights, crucial for ensuring control over the technology.
  - Results from "red-team" testing, aimed at evaluating the AI's ability to be misused, including potential use in cyberattacks or lowering barriers for dangerous weapons development (e.g., chemical, biological, radiological, or nuclear weapons). AI red-team testing 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.

Under 15 CFR 702.3 all information submitted to the BIS under this rule will be treated as confidential and afforded all the protections of section 705(d) of the Defense Production Act (50 U.S.C. 4501 *et seq*.).

- Compliance and Deadlines: If the NRPM passes in its current state, upon engaging in covered activities, U.S. persons will be required to provide detailed answers to BIS's questions within 30 calendar days of triggering reporting. Additionally, responses to follow-up questions would be required within seven days following request, in each case triggering potential civil and/or criminal penalties for non-compliance.
- Long-term Oversight: BIS is especially focused on gathering information to shape future regulatory actions, potentially tightening controls over AI models that could pose security risks to the U.S. or its defense industry.

Implications

For organizations involved in advanced AI and computing technologies, this NRPM signals a new era of government oversight. Compliance will require affected companies to review their AI safety measures, assess their model security, and ensure timely reporting of their AI activities. Given the focus on national security, these reports could prompt further regulatory action, so companies should be proactive in aligning their operations with the new requirements.

As of this date, BIS believes that there are no more than 15 U.S. persons that meet the reporting thresholds for models and computing clusters.

Conclusion

This paper is intended as a guide only and is not a substitute for specific legal or ?tax advice. ????Please ???reach out to the authors for any specific questions. We expect ?to continue to monitor the ????topics ???addressed in this paper and provide future ?client updates when useful.?

RELATED INDUSTRIES + PRACTICES

- Corporate
- International
- Sanctions + Trade Controls